

JOSEPH SAVERI LAW FIRM, LLP

Joseph R. Saveri (SBN 130064)
601 California Street, Suite 1505
San Francisco, CA 94108
Telephone: (415) 500-6800
Facsimile: (415) 395-9940
Email: jsaveri@saverilawfirm.com

Matthew Butterick (SBN 250953)
1920 Hillhurst Avenue, 406
Los Angeles, CA 90027
Telephone: (323) 968-2632
Facsimile: (415) 395-9940
Email: mb@buttericklaw.com

*Counsel for Individual and Representative
Plaintiffs and the Proposed Class*

[Additional Counsel Listed on Signature Page]

**UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN FRANCISCO DIVISION**

Sarah Andersen, an individual; Kelly McKernan, an individual; Karla Ortiz, an individual; H. Southworth pka Hawke Southworth, an individual; Grzegorz Rutkowski, an individual; Gregory Manchess, an individual; Gerald Brom, an individual; Jingna Zhang, an individual; Julia Kaye, an individual; Adam Ellis, an individual;

Individual and Representative Plaintiffs,

v.

Stability AI Ltd., a UK corporation;
Stability AI, Inc., a Delaware corporation;
DeviantArt, Inc., a Delaware corporation;
Midjourney, Inc., a Delaware corporation;
Runway AI, Inc., a Delaware corporation;

Defendants.

Case No. 3:23-cv-00201-WHO

THIRD AMENDED COMPLAINT

CLASS ACTION

DEMAND FOR JURY TRIAL

REDACTED

TABLE OF CONTENTS

Page(s)

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

- I. AI IMAGE PRODUCTS ARE TRAINED ON VAST NUMBERS OF COPYRIGHTED IMAGES WITHOUT CONSENT, CREDIT, OR COMPENSATION AND VIOLATE THE RIGHTS OF MILLIONS OF ARTISTS 1
- II. JURISDICTION AND VENUE.....4
- III. PLAINTIFFS5
- IV. DEFENDANTS 7
- V. AGENTS AND CO-CONSPIRATORS8
- VI. CLASS ALLEGATIONS9
 - A. CLASS DEFINITIONS 9
 - B. NUMEROSITY 11
 - C. TYPICALITY 11
 - D. COMMONALITY & PREDOMINANCE 11
 - E. ADEQUACY 12
 - F. OTHER CLASS CONSIDERATIONS 12
- VII. ARTISTS AND THEIR WORKS 13
- VIII. THE SOURCE OF THE TRAINING DATASETS: LAION..... 15
- IX. HOW AI IMAGE PRODUCTS WORK: CLIP-GUIDED DIFFUSION.....22
- X. PROTECTED EXPRESSION FROM TRAINING IMAGES IS COPIED, COMPRESSED, STORED, AND INTERPOLATED BY DIFFUSION MODELS29
- XI. EXAMPLES OF TEXT PROMPTS USING PLAINTIFF NAMES IN AI IMAGE PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY.....37
- XII. EXAMPLES OF IMAGE PROMPTS USING PLAINTIFF IMAGES IN AI IMAGE PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY.....43
- XIII. USER AND LICENSEE ACTIVITY51

1 XIV. DEFINITIONS FOR THE CAUSES OF ACTION.....53
2 XV. CAUSES OF ACTION AGAINST STABILITY.....54
3 XVI. CAUSES OF ACTION AGAINST MIDJOURNEY.....62
4 XVII. CAUSES OF ACTION AGAINST RUNWAY82
5 XVIII. CAUSES OF ACTION AGAINST DEVIANTART.....84
6 XIX. JURY TRIAL DEMANDED.....93
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

1 Artists and Plaintiffs Sarah Andersen, Kelly McKernan, Karla Ortiz, Hawke Southworth,
2 Grzegorz Rutkowski, Gregory Manchess, Gerald Brom, Jingna Zhang, Julia Kaye, and Adam Ellis
3 (“Plaintiffs”), on behalf of themselves and all others similarly situated, bring this class action
4 complaint against defendants Runway AI, Inc. (“Runway”); Stability AI Ltd. and Stability AI, Inc.
5 (collectively “Stability”); Midjourney, Inc. (“Midjourney”); and DeviantArt, Inc. (“DeviantArt”)
6 (all collectively “Defendants”). Plaintiffs allege various violations of the Copyright Act (17 U.S.C.
7 § 501); and violations of the Lanham Act (15 U.S.C. § 1125(a)(1)) as to Defendant Midjourney.

8 **I. AI IMAGE PRODUCTS ARE TRAINED ON VAST NUMBERS OF COPYRIGHTED**
9 **IMAGES WITHOUT CONSENT, CREDIT, OR COMPENSATION AND VIOLATE THE**
10 **RIGHTS OF MILLIONS OF ARTISTS**

11 1. An *AI image product* is a software product designed to output images through so-
12 called artificial intelligence techniques. But “artificial intelligence” is a misnomer. The AI image
13 products at issue in this complaint are all built around the same asset: human intelligence and
14 creative expression, in the form of billions of artworks copied from the internet. An AI image
15 product simply divorces these artworks from the artists and attaches a new price tag. The profits
16 from the misappropriation of these works can then flow directly into Defendants’ pockets. But the
17 artists who provided the intelligence and creativity—including Plaintiffs—were not asked for their
18 consent. They were not given any credit. And they have not received one cent in compensation.

19 2. Under the hood, AI image products are powered by one or more *machine-learning*
20 *models* (or simply *models*). Models are not programmed directly in the manner of traditional
21 software, but rather *trained*. Training a model first requires amassing a huge corpus of data, called a
22 *dataset*. The models at issue in this complaint were trained on datasets containing billions of images
23 paired with descriptive captions. In this complaint, each image–caption pair is called a *training*
24 *image*. During training of the model, the training images in the dataset are directly copied in full
25 and then completely ingested by the model, meaning that protected expression from every training
26 image enters the model. As it copies and ingests billions of training images, the model progressively
27 develops the ability to generate outputs that mimic the protected expression copied from the dataset.
28 The outputs of a model are derived entirely and exclusively from what it has extracted from the

1 dataset.

2 3. Users elicit images from AI image products through *prompting*. Early versions of AI
3 image products only accepted text prompts—that is, short textual descriptions of an image. But
4 Defendants have progressively evolved their AI image products to also accept images as prompts to
5 allow users to describe their desired result more easily and precisely. Whether based on text, image,
6 or a combination, a prompt is converted into a numerical descriptor called a *CLIP embedding*. This
7 CLIP embedding is then used to guide the AI image product to generate an image resembling
8 whatever is described by the prompt.

9 4. Defendants Runway and Stability have collaborated on the training and deployment
10 of a series of models called Stable Diffusion. In September 2022, Stability CEO Emad Mostaque
11 described it thus: “Stable Diffusion is the model itself. It’s a collaboration that we did with a whole
12 bunch of people ... We took 100,000 gigabytes of images and compressed it to a two-gigabyte file
13 that can recreate any of those [images] and iterations of those.”¹ To train multiple versions of Stable
14 Diffusion, Runway and Stability have each downloaded copies of billions of copyrighted images
15 without permission—including those belonging to Plaintiffs. Runway and Stability induce others to
16 download Stable Diffusion by distributing it for free through public websites. Runway and Stability
17 also make Stable Diffusion available to end users by building it into AI image products that they
18 market and sell. In August 2022, Mostaque wrote “Ironically [the] main funding of stability except
19 me is ... artists” (ellipsis in original) followed by “Lol” (internet slang for “laughing out loud”).²

20 5. Defendant Midjourney has trained multiple models of its own and deployed them
21 commercially as part of its AI image product. Midjourney has downloaded copies of billions of
22 copyrighted images without permission to train its models—including those belonging to Plaintiffs.
23 It has also copied the Stable Diffusion model and deployed it commercially as part of its AI image
24 product. It sells subscription access to its AI image product.

25 _____
26 ¹ Narratives Podcast, Sept 2022 ([https://narrativespodcast.com/2022/09/19/112-emad-mostaque-ai-
27 alignment-and-stable-diffusion/](https://narrativespodcast.com/2022/09/19/112-emad-mostaque-ai-alignment-and-stable-diffusion/))

28 ² <https://discord.com/channels/729741769192767510/730095596861521970/1008530914525061190>

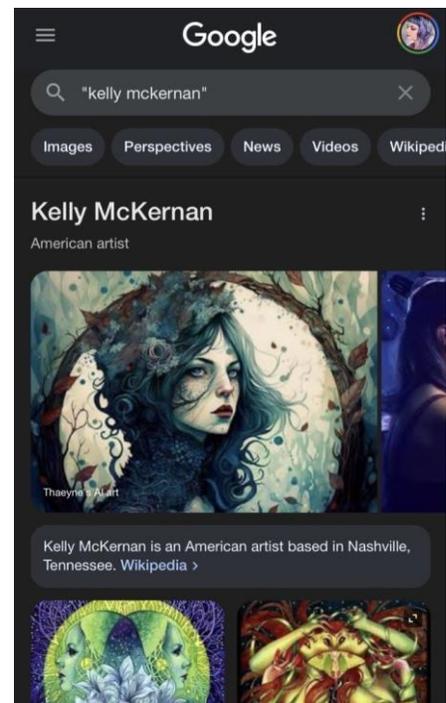
1 6. Defendant DeviantArt has copied the Stable Diffusion 1.4 model and deployed it
2 commercially as part of a subscription-based AI image product. This model was also trained on
3 billions of training images, many of which were harvested from the artist portfolios hosted on
4 DeviantArt’s own website—including certain Plaintiffs.

5 7. In developing its AI image product, DeviantArt also created various infringing
6 copies of Plaintiff Jingna Zhang’s registered copyrighted works.

7 8. Though the Defendants claim to be selling access to AI image products, what they’re
8 really selling is copyright infringement as a service. The scale of this misappropriation is staggering
9 and unprecedented, with violations of law happening at every phase: the gathering and copying of
10 the dataset, the training and deployment of the model, and the output images.

11 9. Worst of all, the Defendants hold out their AI image products as being able to create
12 substantially similar substitutes for the very works they were
13 trained on—either specific training images, or images that
14 imitate the trade dress of particular artists—including
15 Plaintiffs. This is already damaging the market for Plaintiffs’
16 artwork and labor, and the art market more broadly.

17 Midjourney, for instance, has repeatedly promoted the use of
18 artist names—including Plaintiffs’ names—within text
19 prompts as a means of getting better results. Runway,
20 Stability, and Midjourney also encourage the use of images—
21 including images made by Plaintiffs—as a means of
22 prompting their AI image products. Recently, Plaintiff Kelly
23 McKernan was astonished to find that the top internet search
24 result for their name is now an AI-generated image made with Midjourney, prompted with Mx.
25 McKernan’s name.³ Without intervention, this is the grim future that awaits many other artists.



26
27
28 ³ <https://thaeyne.com/2022/12/11/image-in-the-style-of-kelly-mckernan-1/>

1 substantial part of the events giving rise to Plaintiffs’ claims and the interstate trade and commerce
2 involved and affected by Defendants’ illegal conduct occurred in this Division.

3 **III. PLAINTIFFS**

4 15. Sarah Andersen lives in Oregon and owns the copyrights in training images shown
5 under her name in **Exhibit A** (showing Plaintiff images in LAION-5B) and **Exhibit B** (showing
6 Plaintiff images in LAION-400M). Ms. Andersen’s copyright registrations for these works are
7 included in **Exhibit C**, which contains copies of Plaintiffs’ copyright registrations.

8 16. Kelly McKernan lives in Tennessee and owns the copyrights in training images
9 shown under their name in **Exhibit A**. Mx. McKernan is a member of DeviantArt under the name
10 ‘kellymckernan’, which is found in the text caption of some of their works.

11 17. Karla Ortiz lives in California and owns the copyrights in training images shown
12 under her name in **Exhibit A**. Ms. Ortiz’s copyright registrations for certain works are included in
13 **Exhibit C**.⁵

14 18. H. Southworth PKA Hawke Southworth lives in Oregon and owns the copyrights in
15 training images shown under his name in **Exhibit A**. Mr. Southworth is a member of DeviantArt
16 under the name ‘Hauket’, which is found in the text captions of some of his works.

17 19. Grzegorz Rutkowski lives in Poland and owns the copyrights in training images
18 shown under his name in **Exhibit A**. Mr. Rutkowski is a member of DeviantArt under the name
19 ‘88grzes’, which is found in the text captions of some of his works.

20 20. Gregory Manchess lives in Kentucky and owns the copyrights in training images
21 shown under his name in **Exhibit A**. Mr. Manchess’s copyright registrations for these works are
22 included in **Exhibit C**.

23 21. Gerald Brom lives in Georgia and owns the copyrights in training images shown
24 under his name in **Exhibit A**. Mr. Brom’s copyright registrations for these works are included in
25 **Exhibit C**.

26 _____
27 ⁵ Ms. Ortiz registered these copyrights after the initial complaint in this action was filed. She does
28 not seek to assert copyright infringement claims against Stability, Midjourney, or DeviantArt.

1 22. Jingna Zhang lives in the State of Washington and owns the copyrights in training
2 images shown under her name in **Exhibit A** and **Exhibit B**. Ms. Zhang’s copyright registrations for
3 these works are included in **Exhibit C**. Ms. Zhang is a member of DeviantArt under the name
4 ‘Zemotion’, which is found in the text captions of some of her works.

5 23. Julia Kaye lives in California and owns the copyrights in training images shown
6 under her name in **Exhibit A**. Ms. Kaye’s copyright registrations for these works are included in
7 **Exhibit C**.

8 24. Adam Ellis lives in the state of New York and owns the copyrights in training images
9 shown under his name in **Exhibit A**. Mr. Ellis’s copyright registrations for certain works are
10 included in **Exhibit C**.

11 25. The images shown in **Exhibit A** and **Exhibit B** are offered as a representative sample
12 of works by Plaintiffs that appear in the LAION datasets, not an exhaustive or complete list.
13 Plaintiffs confirmed that these particular images were in the LAION-5B and LAION-400M datasets
14 respectively by searching for their own names on two websites that allow searching of the LAION
15 datasets: <https://haveibeen trained.com> and <https://rom1504.github.io/clip-retrieval/>.

16 26. [REDACTED]

17 [REDACTED]

18 [REDACTED]

19 [REDACTED]

20 [REDACTED]

21 [REDACTED]

22 [REDACTED]

23 [REDACTED]

24 [REDACTED]

25 27. The images shown in Exhibit E are offered as a representative sample of works by
26 Plaintiffs that appear in the DataComp xLarge dataset, not an exhaustive or complete list. The
27 exhibits shown in Exhibit E are offered as a representative sample of works by Plaintiffs that appear
28

1 in the DataComp xLarge dataset, not an exhaustive or complete list. Plaintiffs confirmed the
2 presence of these particular images in the DataComp xLarge dataset by searching a publicly
3 available version of the dataset hosted on Hugging Face.
4 https://huggingface.co/datasets/mlfoundations/datacomp_xlarge. On information and belief, the
5 version of DataComp xLarge used by Midjourney to train its models contains these same images, as
6 the dataset is distributed as a standardized collection and Midjourney has not produced its actual
7 training data copies for this data.

8 28. The images shown in Exhibit F are offered as a representative sample of works by
9 Plaintiff Jingna Zhang that DeviantArt copied in developing DreamUp. Plaintiffs confirmed that
10 specific images belonging to Plaintiff Zhang are contained in a dataset DeviantArt produced to
11 Plaintiffs in discovery in this action.

12 29. On information and belief, all of Plaintiffs' works that were registered as part of the
13 collections in **Exhibit C** and were online and were scraped into one of these datasets.

14 30. The registrations shown in **Exhibit C** are only a partial list of registered copyrights
15 owned by Plaintiffs.

16 31. Given the size of the LAION datasets and the search methodology of
17 <https://haveibeenentrained.com> and <https://rom1504.github.io/clip-retrieval/>, it is possible that there
18 are more examples of Plaintiffs' works that have yet to be identified. It is also possible that a
19 particular Plaintiff's work may be included in both datasets even if their works have so far only
20 been identified in one.

21 **IV. DEFENDANTS**

22 32. Defendant Stability AI Ltd. is a UK corporation with its principal place of business
23 at 88 Notting Hill Gate, London, England, W11 3HP. Stability was founded in 2020 by Mohammad
24 Emad Mostaque, a former hedge-fund manager. Mostaque is currently the Chief Executive Officer
25 of Stability AI. Stability AI also employs Robin Rombach, formerly a member of the CompVis
26 research group at Ludwig Maximilian University in Munich, where he was a principal developer of
27 the technology underlying Stable Diffusion.

1 with respect to the acts, violations, and common course of conduct alleged herein.

2 **VI. CLASS ALLEGATIONS**

3 **A. Class Definitions**

4 34. The “Class Period” as defined in this Complaint begins on at least January 13, 2020
5 and runs through the present. Because Plaintiffs do not yet know when the unlawful conduct alleged
6 herein began, but believe, on information and belief, that the conduct likely began earlier than the
7 date listed above, Plaintiffs reserve the right to amend the Class Period to comport with the facts
8 and evidence uncovered during further investigation or through discovery.

9 35. Plaintiffs bring this action for damages and injunctive relief on behalf of themselves
10 and all others similarly situated as a class action pursuant to Rules 23(a), 23(b)(2), and 23(b)(3) of
11 the Federal Rules of Civil Procedure, on behalf of the following Classes:

12 **“Injunctive Relief Class” under Rule 23(b)(2):**

13 All persons or entities nationalized or domiciled in the United States
14 that are legal or beneficial owners of a copyright interest in any work
15 that was used to train any version or component of an AI image product
that was offered directly or incorporated into another product by one or
more Defendants during the Class Period.

16 **“Damages Class” under Rule 23(b)(3):**

17 All persons or entities nationalized or domiciled in the United States
18 that are legal or beneficial owners of a copyright interest in any work
19 that was used to train any version or component of an AI image product
that was offered directly or incorporated into another product by one or
more Defendants during the Class Period.

20 **“LAION-5B Damages Subclass” Under Rule 23(b)(3)**

21 All persons or entities nationalized or domiciled in the United States
22 that are legal or beneficial owners of a registered copyright in any work
23 in the LAION-5B dataset that was used to train any version or
24 component of an AI image product that was offered directly or
25 incorporated into another product by one or more Defendants during
the Class Period.

26 **“LAION-400M Damages Subclass” Under Rule 23(b)(3)**

1 All persons or entities nationalized or domiciled in the United States that are legal or beneficial
2 owners of a registered copyright in any work in the LAION-400M dataset that was used to train any
3 version or component of an AI image product that was offered directly or incorporated into another
4 product by one or more Defendants during the Class Period.

5 **“Midjourney Named Artist Class” under Rule 23(b)(3):**

6 All persons or entities who appear on the Midjourney Names List and
7 whose names were invoked within prompts of the Midjourney Image
8 Product during the Class Period.

8 [REDACTED]
9 [REDACTED]
10 [REDACTED]
11 [REDACTED]
12 [REDACTED]

13 **“DataComp xLarge Damages Subclass” under Rule 23(b)(3):**

14 All persons or entities nationalized or domiciled in the United States
15 that are legal or beneficial owners of a registered copyright in any
16 work in the DataComp xLarge dataset that was used to train any version
or component of an AI image product that was offered directly or
incorporated into another product by one or more Defendants during
the Class Period.

17 **“DeviantArt DreamUp Damages Subclass” under Rule 23(b)(3):**

18 All persons or entities nationalized or domiciled in the United States
19 that are legal or beneficial owners of a registered copyright in any
20 work DeviantArt copied without authorization in developing
21 DreamUp during the relevant time period, which was registered with
22 the United States Copyright Office within five years of the work’s
publication and which was registered with the United States
Copyright Office before being reproduced by DeviantArt as part of
that development, or within three months of publication.

23 These “Class Definitions” specifically exclude the following person or entities:

- 24 a. Any of the Defendants named herein;
- 25 b. Any of the Defendants’ co-conspirators;
- 26 c. Any of Defendants’ parent companies, subsidiaries, and affiliates;
- 27 d. Any of Defendants’ officers, directors, management, employees, subsidiaries,
28 affiliates, or agents;

1 e. All governmental entities; and

2 f. The judges and chambers staff in this case, as well as any members of their
3 immediate families.

4 **B. Numerosity**

5 36. Plaintiffs do not know the exact number of Class members, because such information
6 is in the exclusive control of Defendants. Plaintiffs are informed and believe that there are at least
7 thousands of Class members geographically dispersed throughout the United States such that
8 joinder of all Class members in the prosecution of this action is impracticable.

9 **C. Typicality**

10 37. Plaintiffs' claims are typical of the claims of their fellow Class members because
11 Plaintiffs' claims arise out of the same course of conduct from which their injuries result. Plaintiffs
12 and all Class members own copyrights in the Works. Plaintiffs and the Class created or owned
13 Works that were published on the internet by themselves or others. The Works were used to train
14 various AI Image Products without permission. Plaintiffs and absent Class members were damaged
15 by this and other wrongful conduct of Defendants as alleged herein. Damages and the other relief
16 sought herein are common to all members of the Class.

17 **D. Commonality & Predominance**

18 38. Numerous questions of law or fact common to the entire Class arise from
19 Defendants' conduct—including, but not limited to those identified below.

20 39. **Direct Copyright Infringement:** Whether Defendants violated the copyrights of
21 Plaintiffs and the Class when they downloaded and stored copies of the Works; Whether Defendants
22 violated the copyrights of Plaintiffs and the Class when they used copies of the Works to develop AI
23 Image Products.

24 40. **Vicarious Copyright Infringement:** Whether Defendants violated the copyrights of
25 Plaintiffs and the Class when they distributed their AI Image Products in order to induce, materially
26 contribute, or otherwise encourage users and licensees of their AI Image Products to directly
27 infringe Plaintiffs and Class members' works.
28

1 41. **DMCA Violations:** Whether Defendants violated the rights of Plaintiffs and the
2 Class by falsely attributing CMI to the models, and also making copies of Plaintiffs and Class
3 members' Works with CMI removed or altered.

4 42. **Lanham Act Violations:** Whether Defendants misappropriated Plaintiffs and Class
5 members' trade dress and distinctive look and feel in violation of the Lanham Act.

6 43. **Anticipated Defenses:** Whether any affirmative defense excuses Defendants'
7 conduct, including whether some or all of Defendants' conduct is allowed under fair use.

8 44. These and other questions of law and fact are common to the Class and predominate
9 over any questions affecting the Class members individually.

10 **E. Adequacy**

11 45. Plaintiffs will fairly and adequately represent the interests of the Class because they
12 have experienced the same harms as the Class and have no conflicts with any other members of the
13 Class. Furthermore, Plaintiffs have retained sophisticated and competent counsel ("Class Counsel")
14 who are experienced in prosecuting federal and state class actions throughout the United States and
15 other complex litigation and have extensive experience advising clients and litigating intellectual
16 property, competition, contract, and privacy matters.

17 **F. Other Class Considerations**

18 46. Defendants have acted on grounds generally applicable to the Class, thereby making
19 final injunctive relief appropriate with respect to the Class as a whole.

20 47. This class action is superior to alternatives, if any, for the fair and efficient
21 adjudication of this controversy. Prosecuting the claims pleaded herein as a class action will
22 eliminate the possibility of repetitive litigation. There will be no material difficulty in the
23 management of this action as a class action.

24 48. The prosecution of separate actions by individual Class members would create the
25 risk of inconsistent or varying adjudications, establishing incompatible standards of conduct for
26 Defendants.

VII. ARTISTS AND THEIR WORKS

49. Plaintiffs are artists who have created recognized and influential contemporary artwork.

50. Creating successful artwork that is recognized and appreciated, let alone for an artist to become financially successful, requires immense dedication, energy, and creativity. An artist may become well-known for a variety of reasons. But at the core, each artist is known for the value of their particular expression. As such, it is important for artists to protect their works from being copied or used without their permission.

51. Copyright law protects artists’ works from infringement by creating *exclusive* rights of artists to make copies of their works, to make derivative works of their copyrighted works, and to distribute such copies and derivative works, which protects the interests of artists and preserves the incentives for humans to produce art. As recognized by no higher source than the United States Constitution itself, copyright law is intended to “promote the Progress of Science and useful Arts.” By offering artists protection, they can be rewarded for their efforts. In accordance with copyright law, many artists, including certain Plaintiffs, register copyrights in their works.

52. Artists also protect their works in other ways. One of the most common ways is to affix a distinctive mark, watermark, signature, website URL or other identifying mark to their work. These marks ensure that artists receive credit and recognition for the artwork they have created and serves as a means to manage the use of their work.

53. A few examples of Plaintiffs’ distinctive marks are below:

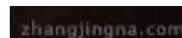
Ex. A at p. 5 (Plaintiff Brom’s signature—lower right corner):



Ex. A at p. 13 (Plaintiff Brom’s personal website URL—lower left corner):



Ex. A at p. 62 (Plaintiff Zhang’s personal website URL—lower-right corner):



Ex. A at p. 79 (Plaintiff Ortiz’s signature—lower-right corner):

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28



Ex. A at p. 86 (Plaintiff McKernan’s mark—lower right corner):



Ex. A at p. 103 (Plaintiff Andersen’s signature—lower right corner):



54. Artists and their livelihoods are facing a new threat, however. AI image products have begun to proliferate. As described more fully herein, these machine-learning models are trained on billions of artworks, many of which are protected either by being registered under the copyright laws, or otherwise designated as protected by, for example, bearing an artist’s distinctive mark.

55. Machine-learning models depend on massive quantities of high-quality data that is digitally copied without authorization to train the model. The quality of the dataset a model is trained on determines the quality of the model itself. The models within AI image products are no different. As a result, the images these models are trained on have an enormous impact on the quality of the models and the outputs they can produce. There are many images that are in the public domain, *i.e.*, images that are not subject to copyright. The developers of many AI image products, however, made the decision to greatly expand the datasets their models are trained on by including billions of protected works. They made this choice because many desirable works are not in the public domain, *i.e.*, are not subject to unfettered use by anyone without consent, credit, or compensation.

56. These AI image products compete with the artists whose very works comprise the raw material for the models within the AI image products. Rather than pay or commission an artist to create a particular work, pay an artist for a print or copy of a particular work, or pay for the artist’s permission to use the work, now users and licensees of these AI image products merely have to prompt the AI image product using terms—such as an artist’s name, or titles of their works—to generate an image indistinguishable from one the artist might’ve created themselves. Worse, certain

1 users of AI image products have at times abused this “feature” to harass and annoy the artists
2 themselves by making knockoff versions of their artwork and publicizing it.

3 **VIII. THE SOURCE OF THE TRAINING DATASETS: LAION**

4 57. LAION (acronym for “Large-Scale Artificial Intelligence Open Network”) is an
5 organization based in Hamburg, Germany. According to its website, LAION is led by Christoph
6 Schuhmann. LAION’s stated goal is “to make large-scale machine learning models, datasets and
7 related code available to the general public.” All of LAION’s projects are made available for free.
8 Other members of LAION’s current team include Stability engineers Robin Rombach and
9 Katherine Crowson, and Google engineer Romain Beaumont.

10 58. LAION’s most well-known projects are the datasets of training images it has
11 released for training machine-learning models, which are now widely used in the AI industry.

12 59. In August 2021, LAION released LAION-400M, a dataset of 400 million training
13 images assembled from images accessible on the public internet. At the time, LAION-400M was
14 the largest freely available dataset of its kind. LAION distributes the LAION-400M dataset to the
15 public through its own website and elsewhere. Information about LAION-400M is available in an
16 accompanying paper by Schuhmann, Beaumont, and others titled “LAION-400M: Open Dataset of
17 CLIP-Filtered 400 Million Image-Text Pairs,” released in November 2021 (hereafter, the “LAION-
18 400M Paper”).⁶

19 60. When one downloads the LAION-400M dataset, one gets a list of metadata records,
20 one for each training image. Each record includes the URL of the image, the image caption, the
21 similarity of the caption and image (as measured by the proximity of their respective CLIP
22 embeddings), a NSFW flag (indicating whether the CLIP embedding of the image suggests that it
23 contains so-called “not safe for work” content), and the width and height of the image.

24 61. The actual images referenced in the LAION-400M dataset records are not included
25 with the dataset. Anyone who wishes to use LAION-400M for training their own machine-learning
26

27 _____
28 ⁶ <https://arxiv.org/abs/2111.02114>

1 model must first acquire copies of the actual images from their URLs. To facilitate the copying of
2 these images, LAION provides a software tool called ‘img2dataset’⁷ that takes the metadata records
3 as input and makes copies of the referenced images from the URLs in each metadata record, thereby
4 creating local copies.

5 62. Training a model with the LAION-400M dataset cannot begin without first using
6 ‘img2dataset’ or another similar tool to download the images in the dataset. Thus, every person or
7 entity that has trained a model on LAION-400M has necessarily made one or more copies of
8 images belonging to Plaintiffs as shown in **Exhibit B** (Plaintiff images in LAION-400M), either by
9 using ‘img2dataset’ or another tool. These Plaintiffs never authorized any of these LAION dataset
10 users to copy their images or use them for training any models.

11 63. One of the entities that has made unauthorized copies of the LAION-400M training
12 images is LAION itself. According to the LAION-400M Paper, LAION made the dataset by starting
13 with Common Crawl metadata records. Common Crawl is a corpus of 250 billion web pages copied
14 from the public web, including assets like Plaintiffs’ images (<https://commoncrawl.org/>). The
15 metadata records contain web URLs. According to the LAION-400M Paper, LAION created
16 training images by first “pars[ing] through [the metadata records] from Common Crawl and
17 pars[ing] out all HTML IMG tags containing an alt-text attribute [that is, a text caption].” Then,
18 LAION “download[ed] the raw images from the parsed URLs”. *See* LAION-400M Paper at 3. To
19 ensure that the training images in the dataset had reasonably accurate captions, LAION used a CLIP
20 model to calculate the CLIP embeddings for the image and text of each image–text pair. These two
21 CLIP embeddings were compared to measure how well the text described the image. Image–text
22 pairs with low CLIP-similarity scores were omitted from the dataset.

23 64. Sometime after the release of LAION-400M in August 2021, Stability funded
24 LAION’s creation of a similar dataset, but much larger. In March 2022, Stability CEO Mostaque
25 called himself “the biggest backer of LAION.”⁸ In August 2022, Stability CEO Mostaque said, “I
26

27 ⁷ <https://github.com/rom1504/img2dataset>

28 ⁸ <https://discord.com/channels/662267976984297473/938713143759216720/954674533942591510>

1 funded LAION, underlying dataset for ... stable diffusion.”⁹ (After the initial complaint in this
2 action was filed, Mostaque changed his story, saying “We actually have/had no influence on
3 CompVis or LAION, did not funded [sic] either.”¹⁰)

4 65. In October 2022, LAION released LAION-5B, a dataset of 5.85 billion training
5 images—more than 14 times bigger than LAION-400M. Information about LAION-5B is available
6 in an accompanying paper called “LAION-5B: An open large-scale dataset for training next
7 generation image-text models,” by Schuhmann, Beaumont, Crowson, and others (hereafter, the
8 “LAION-5B Paper”).¹¹ According to the LAION-5B Paper, LAION-400M is a subset of LAION-
9 5B, meaning every image in LAION-400M is also in LAION-5B.

10 66. Much like the LAION-400M dataset, when one downloads the LAION-5B dataset,
11 one gets a list of metadata records, one for each training image. Each record includes the URL of
12 the image, the image caption, the similarity of the caption and image (as measured by their
13 respective CLIP embeddings), the width and height of the image, and other metadata fields. Each
14 record also includes a watermark detection score, which indicates the likelihood an image possesses
15 a distinctive mark of a copyright holder, *e.g.*, the artist.

16 67. Just like the LAION-400M dataset, the actual images referenced in the LAION-5B
17 dataset records are not included with the dataset. Anyone who wishes to use LAION-5B for training
18 their own machine-learning model must first acquire copies of the actual images from their URLs.
19 To facilitate the copying of these images, LAION provides a software tool called ‘img2dataset’¹²
20 that takes the metadata records as input and makes copies of the referenced images from the URLs
21 in each record, thereby creating local copies.

22 68. Training a model with the LAION-5B dataset cannot begin without first using
23 ‘img2dataset’ or another similar tool to download the images in the dataset. Thus, every person or
24

25 _____
26 ⁹ <https://twitter.com/EMostaque/status/1559332564787240962>

27 ¹⁰ <https://twitter.com/EMostaque/status/1629516125150011394>

28 ¹¹ <https://arxiv.org/abs/2210.08402>

¹² <https://github.com/rom1504/img2dataset>

1 entity that has trained a model on LAION-5B has necessarily made one or more copies of images
2 belonging to Plaintiffs as shown in the **Exhibit A** (Plaintiff images in LAION-5B), either by using
3 ‘img2dataset’ or another tool. These Plaintiffs never authorized any of these LAION dataset users to
4 copy their images or use them for training any models.

5 69. One of the entities that has made unauthorized copies of the LAION-5B training
6 images is LAION itself. According to the LAION-5B Paper, the LAION-5B dataset was built in
7 much the same way as the LAION-400M dataset. First, image–text pairs were assembled by
8 starting with “Common Crawl’s ... metadata files,” extracting URLs of images with captions, and
9 “download[ing] the raw images from the parsed URLs.” *See* LAION-5B Paper at 5. LAION then
10 used a CLIP model to test the CLIP similarity of the image and text, omitting pairs with low
11 similarity.

12 70. The authors of the LAION-5B Paper also included a boldface warning:

13 [W]e strongly recommend that LAION-5B should only be used for
14 academic research purposes in its current form. We advise against any
15 applications in deployed systems without carefully investigating
behavior and possible biases of models trained on LAION-5B.

16 71. Despite this warning that LAION-5B “should only be used for academic research
17 purposes,” all the Defendants—Runway, Stability, Midjourney, and DeviantArt—have trained,
18 distributed, promoted, or deployed commercial systems that rely on LAION-5B, directly in
19 contravention of this warning.

20 72. Though LAION-400M and LAION-5B are often used to train diffusion models, they
21 are used to train other models as well. For instance, Stability sponsored LAION’s creation of
22 OpenCLIP, a CLIP model trained on a subset of images from LAION-5B called LAION-2B
23 consisting of training images with English-language captions.¹³ On information and belief, because
24 all the Plaintiff images in **Exhibit A** (Plaintiff images in LAION-5B) have English-language
25 captions, they are also part of LAION-2B and were thus used to train OpenCLIP. To create
26
27

28 ¹³ <https://huggingface.co/laion/CLIP-ViT-H-14-laion2B-s32B-b79K>

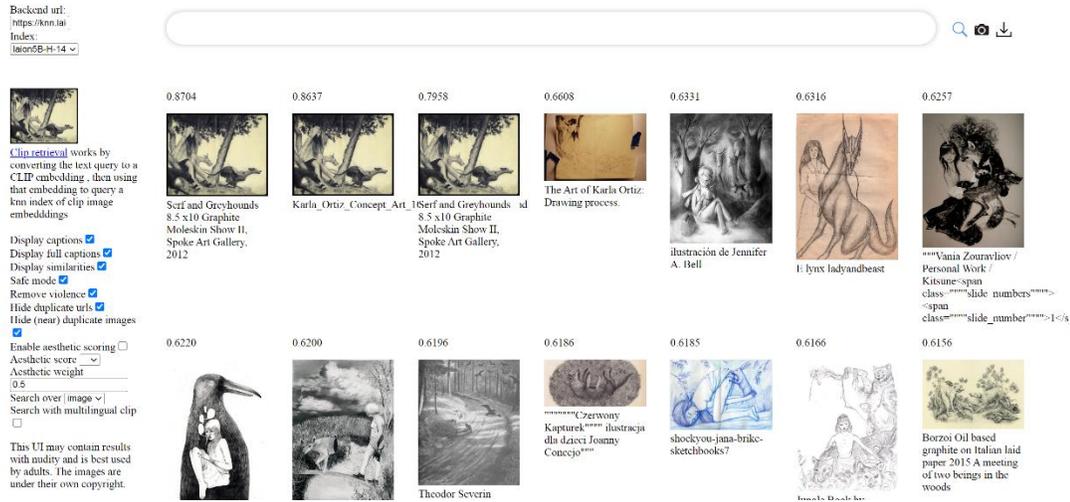
1 OpenCLIP, LAION necessarily had to create one or more copies of these images. Plaintiffs never
2 gave their permission to LAION to copy their images or use them to train OpenCLIP.

3 73. It is possible to search whether a specific image is included in the LAION dataset
4 through the use of the websites <https://haveibeen trained.com> and [https://rom1504.github.io/clip-](https://rom1504.github.io/clip-retrieval/)
5 [retrieval/](https://rom1504.github.io/clip-retrieval/). These websites use CLIP embeddings to search the LAION datasets to discover whether
6 particular images are included. Based on the size of the datasets, however, and the search
7 methodologies, exact or exhaustive results are not guaranteed for every example of a particular
8 artist's work.

9 74. Below is an image hosted on plaintiff Karla Ortiz's website:



19 75. A search of the LAION dataset for this image by Ms. Ortiz on
20 <https://rom1504.github.io/clip-retrieval/> may generate the below results:

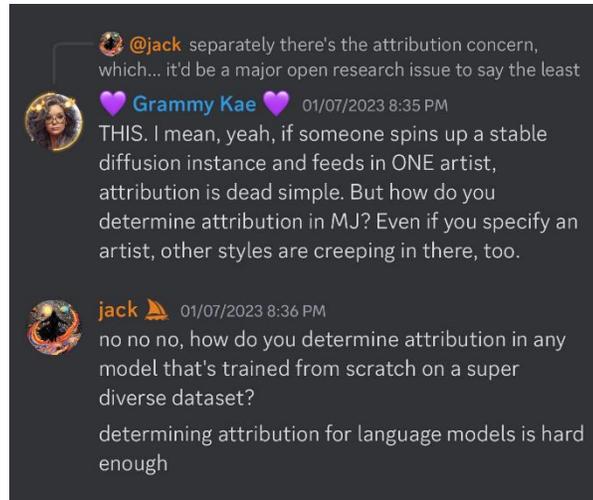


76. The caption shown for the first image in the search results of the LAION dataset is exactly the same as the caption shown on Plaintiff Ortiz’s website: “Serf and Greyhounds / 8.5 x 10 Graphite / Moleskin Show II, Spoke Art Gallery, 2012”.

77. Further, inspection of the metadata given by <https://rom1504.github.io/clip-retrieval/> indicates that the image was scraped from the following URL: https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKft7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uy2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg, which indicates the source of the image was Ms. Ortiz’s own website. In other words, the metadata indicates that this image was scraped from Ms. Ortiz’s personal website for inclusion into the LAION datasets.

78. In this way, the captions included in the training images also function as copyright-management information. Much as a music publisher may search on a streaming platform for unlawful uses of their work in order to conduct a DMCA strike, artists utilize keywords (for example, their name) as search terms on <https://haveibeentrained.com> to identify whether their works have been scraped and used as training material for AI image products.

1 79. The developers of AI image products know the datasets upon which their models are
2 trained contain copyrighted material. As admitted by Midjourney engineer Jack Gallagher on
3 Midjourney’s Discord server, Midjourney knew that attribution was a difficult issue:



16 80. Stability CEO Emad Mostaque has publicly acknowledged the importance of using
17 licensed training images, saying that future versions of Stable Diffusion would be based on “fully
18 licensed” training images.¹⁴ But so far, Stability has taken no steps to negotiate suitable licenses.
19 Neither has Runway. Neither has Midjourney. They all just use LAION datasets—with no consent,
20 no credit, and no compensation to the artists.

21 81. In July 2023, the topic of AI training reached the U.S. Senate.¹⁵ During a hearing
22 convened by a subcommittee of the Senate Judiciary Committee, Sen. Mazie Hirono quizzed Ben
23 Brooks, a representative from Stability, about Stability’s position on licensing training data. Sen.
24 Hirono asked directly, “So basically you don’t pay for the data that you put into your—to train your
25 models?” Mr. Brooks replied, “There is no arrangement in place.” Sen. Hirono then turned to
26 Plaintiff Karla Ortiz, who was testifying on the same panel. “So you have Ms. Ortiz, who says that
27 that is wrong. Is that correct, Ms. Ortiz?” Ms. Ortiz replied, “A hundred percent, Senator.”

28

¹⁴ @EMostaque, Twitter (Dec. 15, 2022, 8:03 AM),
<https://twitter.com/EMostaque/status/1603390169192833027>.

¹⁵ <https://youtu.be/uoCJun7gkBA?t=3578>

IX. HOW AI IMAGE PRODUCTS WORK: CLIP-GUIDED DIFFUSION

82. *CLIP-guided diffusion* is a technique that uses two machine-learning models in cooperation. The first is called a *diffusion model*, which generates the image over a sequence of steps. The second is called a *CLIP model*, which converts the user prompt into a form that can be used to nudge the diffusion model closer to a suitable result at each step in the generation process.

83. Stable Diffusion is a model that generates images using CLIP-guided diffusion. Therefore, AI image products based on Stable Diffusion—including those offered by Runway, Stability, and DeviantArt—rely on CLIP-guided diffusion as well. In June 2022, a Midjourney developer confirmed on their public message board that “we use clip guided diffusion” too.¹⁶ Thus, the description below applies to all the AI image products at issue in this complaint.

84. In 2015, the diffusion technique for training a machine-learning model was proposed by a team of researchers led by Jascha Sohl-Dickstein at Stanford University and introduced in their paper “Deep Unsupervised Learning Using Nonequilibrium Thermodynamics.”¹⁷ Though the technique can be applied to any kind of data, the description below uses digital images as an example.

85. Diffusion proceeds in two phases. To begin the training phase, initial copies are made of many training images. For each training image, progressively more noise is added over a series of steps. At each step, the model records how the addition of noise changes the image. By the last step, the image has been “diffused” into essentially random noise. A simplified version of this process is shown in the diagram below:¹⁸

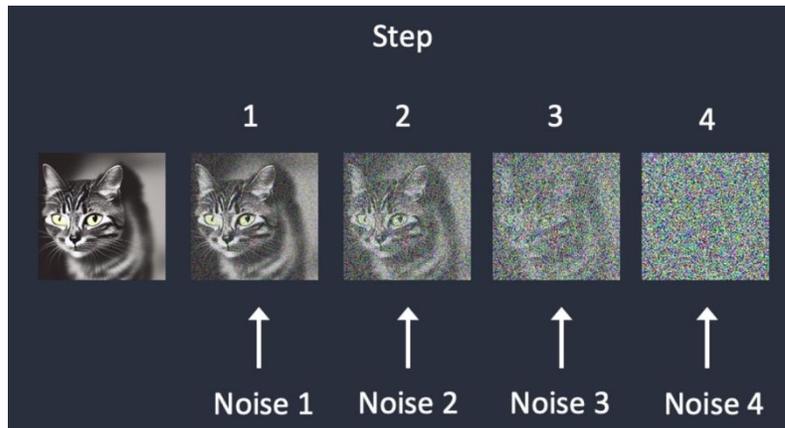
¹⁶

<https://discord.com/channels/662267976984297473/938713143759216720/982136076888068156>

¹⁷ <https://arxiv.org/abs/1503.03585>

¹⁸ Noising & denoising illustrations from https://stable-diffusion-art.com/how-stable-diffusion-work/#How_training_is_done

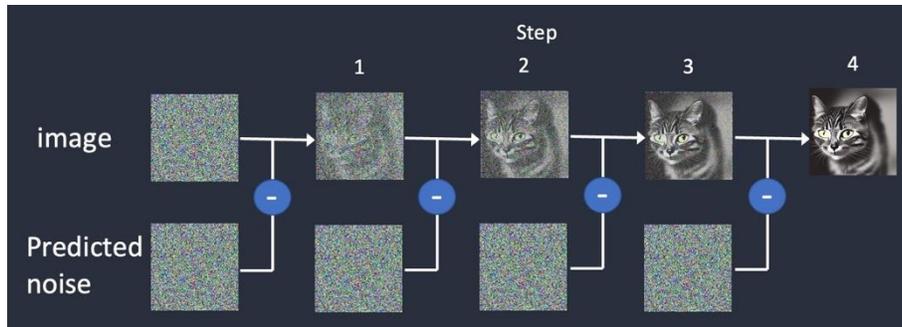
86.



87. The diagram also illustrates that many intermediate copies of a training image are necessarily made during the training process, with increasing amounts of noise added.

88. After the diffusion model is trained, it can perform the second phase of diffusion, which is like the first but reversed. Having recorded the process of turning a certain image into noise over many steps, the model can then run the sequence backwards. Starting with a patch of random noise, the model applies the steps in reverse order. As it progressively removes noise (or “denoises”) the data, the model is eventually able to reveal that image, as illustrated below:

89.



90. In sum, diffusion is a way for a machine-learning model to calculate how to reconstruct a copy of its training images. For each training image, a diffusion model finds the sequence of denoising steps to reconstruct that specific image. Then it stores this sequence of steps. In practice, this training would be repeated for many images—likely millions or billions. A diffusion

1 model is then able to reconstruct copies of each training image based on this denoising process.
2 Furthermore, being able to reconstruct copies of the training images is not an incidental side effect.
3 The primary objective of a diffusion model is to reconstruct copies of its training images with
4 maximum accuracy and fidelity. Or, in the words of prominent machine-learning researcher
5 Nicholas Carlini, who has studied the behavior of diffusion models: “diffusion models are explicitly
6 trained to reconstruct the training set.”¹⁹

7 91. In December 2020, the diffusion technique was improved by a team of researchers at
8 UC Berkeley led by Jonathan Ho. These ideas were introduced in their paper “Denoising Diffusion
9 Probabilistic Models.”²⁰

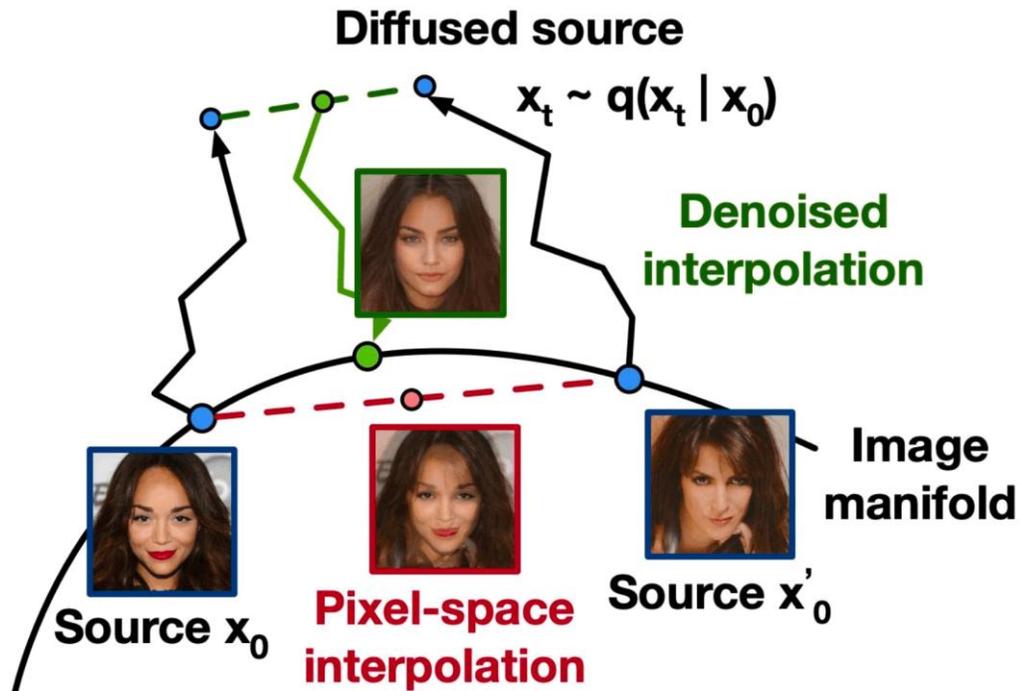
10 92. Ho showed how diffused images could be interpolated—meaning, blended
11 mathematically—to produce new derivative images. Rather than combine two images pixel by
12 pixel—which gives unappealing results—Ho showed how protected expression from training
13 images can be stored in the diffusion model and then interpolated to generate another image.

14 93. The diagram below, taken from Ho’s paper, shows how this process works, and
15 demonstrates the difference in results between interpolating via pixels and interpolating via
16 diffusion and denoising.

17
18
19
20
21
22
23
24
25
26
27 ¹⁹ “Extracting Training Data from Diffusion Models,” p. 12, available at
<https://arxiv.org/abs/2301.13188>

28 ²⁰ Available from <https://arxiv.org/abs/2006.11239>

94.



95. In the diagram, two photos are being blended: the photo on the left labeled “Source x_0 ,” and the photo on the right labeled “Source x'_0 .”

96. The image in the red frame has been interpolated pixel by pixel and is thus labeled “pixel-space interpolation.” This pixel-space interpolation simply looks like two translucent face images stacked on top of each other, not a single convincing face.

97. The image in the green frame, labeled “denoised interpolation,” has been generated differently. In that case, the two source images have been converted into diffused images (illustrated by the crooked black arrows pointing upward toward the label “Diffused source”). Once these diffused images have been interpolated (represented by the green dotted line), the newly interpolated diffused image (represented by the smaller green dot) has been denoised into pixels (a process represented by the crooked green arrow pointing downward to a larger green dot). This process yields the image in the green frame.

98. Compared to the pixel-space interpolation, the difference is apparent: the denoised interpolation looks like a single convincing human face, not an overlay of two faces. An enlarged

1 detail of the two interpolated images is shown below:

2 99.



10 100. Despite the difference in results, these two modes of interpolation are similar in that
11 they both blend protected expression from the source images but using different techniques.

12 101. In April 2022, the diffusion technique was further improved by a team of researchers
13 led by Robin Rombach at Ludwig Maximilian University of Munich. These ideas were introduced
14 in his paper “High-Resolution Image Synthesis with Latent Diffusion Models.”²¹ (Rombach is
15 currently employed by Stability as a developer of Stable Diffusion.)

16 102. Rombach’s diffusion technique offered one key improvement over previous efforts.
17 Rombach devised a way to supplement the denoising process with extra information, a technique
18 called *conditioning*. One common tool for conditioning is a *prompt*, which is text or image provided
19 by a user that might describe elements of the image, *e.g.*— “a dog wearing a baseball cap while
20 eating ice cream.” As the denoising process proceeds, the conditioning data is used to nudge the
21 denoising process closer to the desired result.

22 103. Though in principle, conditioning data can come from a variety of sources, in the AI
23 image products at issue in this complaint, the conditioning data is provided by a *CLIP model*.

24 104. In 2021, researchers from OpenAI introduced the idea of a CLIP model in a paper
25 called “Learning Transferable Visual Models From Natural Language Supervision.”²² A CLIP
26

27 ²¹ <https://arxiv.org/abs/2112.10752>

28 ²² <https://arxiv.org/abs/2103.00020>

1 model quantifies the semantic correlation between images and captions.

2 105. “CLIP” stands for “contrastive language–image pretraining.” This connotes the idea
3 that during training, a CLIP model learns to correlate images and captions by ingesting protected
4 expression from training images along with their text captions. Whereas a diffusion model learns to
5 generate actual images, the CLIP model learns to correlate images and captions. An image is
6 meaningless to a CLIP model without its accompanying text caption. These images and their text
7 captions are colloquially known as an “image-text pair” or “text-image pair” (or in this complaint, a
8 training image).

9 106. These captions are often generated by the artists themselves. For example, when an
10 artist uploads an image to their personal website, they may include a caption that describes the
11 image and also identifies themselves as a way of managing the image’s use. Examples of training
12 images showing both image and caption are included in **Exhibit A** (Plaintiff images in LAION-5B)
13 and **Exhibit B** (Plaintiff images in LAION-400M) attached hereto.

14 107. Like a diffusion model, a CLIP model is trained by copying and ingesting a huge
15 number of training images—on the scale of hundreds of millions or billions. Though a diffusion
16 model cooperates with a CLIP model in CLIP-guided diffusion, the two models are trained
17 separately. They may be trained on the same training dataset. But this is not required.

18 108. For instance, by copying and ingesting a diverse set of images of dogs that have
19 “dog” in the caption, the CLIP model will learn to correlate the word “dog” more strongly with
20 images containing what humans perceive as dogs, and less with images of other things. The success
21 of CLIP training depends on the training images having accurate captions. If all images of dogs are
22 labeled “cat,” then the CLIP model will make incorrect correlations. The CLIP model has no visual
23 or other knowledge of the world that allows it to make these correlations. It is entirely dependent on
24 the captions.

25 109. Consistent with this behavior, a CLIP model that is exposed to training images with a
26 certain artist’s name in the caption—say, Plaintiff Grzegorz Rutkowski—will learn to associate the
27 work of Mr. Rutkowski with the caption “Grzegorz Rutkowski.”
28

1 110. The CLIP model is able to do this by converting both images and text captions into a
2 common intermediate format called a *CLIP embedding*. The embedding is a list of numbers
3 representing a point in a geometric space. To use an analogy, a CLIP embedding is like an x–y
4 coordinate in the two-dimensional plane, but with many more dimensions. To find out how well a
5 particular image matches a particular caption, one converts both into their respective CLIP
6 embeddings and measures the proximity of the CLIP embeddings within this geometric space.
7 When the CLIP embeddings are closer together, it means there is a stronger semantic correlation
8 between the image and the caption. In general, there is no human-intelligible meaning of the
9 numbers in a CLIP embedding. They are only intelligible to the CLIP model.

10 111. In an AI image product, the role of a CLIP model is to guide the diffusion model
11 toward the user’s intended result. For example, if a user includes the word “dog” in the text prompt,
12 first the CLIP model converts the prompt into a CLIP embedding. Then, as the diffusion model
13 iteratively denoises the image, the CLIP embedding is used as conditioning data, nudging the
14 image-denoising process toward a more dog-like result. Similarly, if a user includes the name of
15 Plaintiff “Grzegorz Rutkowski” in the text prompt, then as the diffusion model iteratively denoises
16 the image, the CLIP embedding for the phrase “Grzegorz Rutkowski” nudges the image-denoising
17 process toward a more Rutkowski-like result.

18 112. Certain words and phrases have stronger correlations within CLIP models. For
19 example, artist names are particularly influential when included in a prompt. Indeed, users of AI
20 image products quite often use an artist’s name to get a particular result. Defendants actively
21 promote such use.

22 113. As mentioned at the beginning of this section, this combination of a CLIP model and
23 diffusion model is called *CLIP-guided diffusion* and is used by all the AI image products at issue in
24 this complaint.

25 114. Because a CLIP embedding can be generated from either text or an image, an AI
26 image product that relies on CLIP-guided diffusion can be prompted with either text or image, since
27 either can be converted into a CLIP embedding. Image prompts, however, tend to produce more
28

1 precise and descriptive CLIP embeddings. Thus, image prompting has become an increasingly
2 prominent feature in AI image products, because it allows finer control of the prompting process.

3 115. To recap, an AI image product that works based on CLIP-guided diffusion contains
4 two models that cooperate: a CLIP model and a diffusion model. Initially, the CLIP model is trained
5 on a dataset of training images and learns to relate the semantic meaning of images and associated
6 text through an intermediate format called a CLIP embedding. The diffusion model is also trained on
7 a dataset of training images and learns how to take a patch of noise and “denoise” it to reveal an
8 image. These trained models are deployed as part of the AI image product. When a user submits a
9 prompt to the AI image product—either text, image, or a combination—the CLIP model converts
10 this prompt to an embedding. The embedding is then used as conditioning data as the diffusion
11 model progressively generates the image through denoising. The image that emerges at the end of
12 the denoising process is presented to the user as the output.

13 **X. PROTECTED EXPRESSION FROM TRAINING IMAGES IS COPIED, COMPRESSED, 14 STORED, AND INTERPOLATED BY DIFFUSION MODELS**

15 116. As mentioned above, training a diffusion model requires vast numbers of training
16 images—often billions. When the training process is complete, a diffusion model is then able to
17 reconstruct copies of each training image. Furthermore, being able to reconstruct copies of the
18 training images is not an incidental side effect. The primary objective of a diffusion model is to
19 reconstruct copies of its training images with maximum accuracy and fidelity.

20 117. Consistent with this, a machine-learning model—including a diffusion model—can
21 be conceptualized as an evolution of the database. As described by François Chollet, Google
22 machine-learning researcher and author of the book *Deep Learning with Python*, “Deep learning
23 takes data points and turns them into a query-able structure that **enables retrieval and**
24 **interpolation between the points**. You could think of it as a continuous generalization of database
25 technology. ... Because it is analogous to a database, the usefulness of a deep learning system
26 depends entirely on the data points it was constructed with. **You get back what you put in (or**
27
28

1 **interpolations of the same).**²³ (Emphases added.)

2 118. With the phrase “continuous generalization of database technology,” Chollet is
3 contrasting a traditional database, which stores its data in discrete records, with a machine-learning
4 model, which treats its data as sitting on a continuous geometric surface, called a *manifold*. The
5 manifold is a mathematical construct discovered by the model during training and represents the
6 “information space” of the training data. By representing training data in a continuous rather than
7 discrete manner, a machine-learning model permits flexible operations of data, such as measuring
8 the proximity of data points, and as Chollet notes, “retrieval and interpolation” of data points.
9 Furthermore, because the representations of the training data on the manifold are simplified
10 compared to their original form, the model essentially uses the manifold to accomplish compression
11 of the training dataset.

12 119. Though the framing of machine learning as a form of data compression has been
13 resisted by some, research shows an ever-stronger connection between the two, and between
14 diffusion models and compression in particular. In November 2023, a team of machine-learning
15 researchers led by Yaodong Yu at UC Berkeley published a paper called “White-Box Transformers
16 via Sparse Rate Reduction: Compression Is All There Is?”²⁴ (Below, the “Yu Paper”.) In their paper,
17 the authors describe in detail a strong mathematical and experimental correspondence between
18 diffusion models and data compression, and conclude by saying (italics in original, bold emphasis
19 added):

20 [W]e hope that this work ... help[s] clarify the ultimate capabilities
21 of modern artificial intelligence (AI) systems ... Just as with all
22 other natural phenomena or technical innovations that were once
23 “black boxes” to people, significant confusion and anxiety is arising
24 in society about the potential or implications of emerging new AI
25 systems, including ... large image generation models such as
26 Midjourney ... From the perspective of this work ... **these large
models are unlikely to do anything beyond purely mechanical
data compression (encoding) and interpolation (decoding).** That
is, this work suggests that for these existing large AI models,
however magical and mysterious they might appear to be:
Compression is all there is.

27 ²³ <https://twitter.com/fchollet/status/1563153087514419206>

28 ²⁴ <https://arxiv.org/abs/2311.13110>

1 See Yu Paper at 53.

2 120. In public statements, Stability CEO Emad Mostaque and Stability itself have
3 repeatedly and consistently characterized Stable Diffusion as a compressed copy of its training
4 images. Some examples are listed below (emphasis added):

5 121. In August 2022, Mostaque described Stable Diffusion in a recorded interview: “What
6 happens is you take 250 thousand gigabytes of images and you **compress it down** to X gigabytes.
7 We’ll share the details soon. But it’s surprisingly small.”

8 122. In August 2022, Mostaque described Stable Diffusion in another recorded interview:
9 “It’s worth taking a step back and thinking about how crazy insane this is: we took a hundred
10 terabytes of data—a hundred thousand thousand megabytes of images—2 billion of them—and we
11 **squished it down** to a 2–4 gigabyte file. And that file can create everything that you’ve seen. That’s
12 insane, right? That’s about **as compressed as you can get.**”

13 123. In August 2022, Stability said in its launch announcement for Stable Diffusion that it
14 “is the culmination of many hours of collective effort to create a single file that **compresses the**
15 **visual information** of humanity into a few gigabytes.”²⁵

16 124. In September 2022, Mostaque said in a podcast interview: “Stable Diffusion is the
17 model itself. ... We took 100,000 gigabytes of images and **compressed it** to a two-gigabyte file that
18 can recreate any of those and iterations of those.”²⁶

19 125. In January 2023, Mostaque said in a podcast interview: “We took 100,000 gigabytes
20 of image-label pairs—2 billion images—and created a 1.6 gigabyte file ... that **basically**
21 **compresses the visual information** of a snapshot of the internet.”²⁷

22 126. In February 2023, Mostaque said in a podcast interview: “We’ve created the **most**
23 **efficient compression** in the world.”²⁸

24
25
26 ²⁵ <https://stability.ai/blog/stable-diffusion-public-release>

27 ²⁶ <https://narrativespodcast.com/2022/09/19/112-emad-mostaque-ai-alignment-and-stable-diffusion/>

28 ²⁷ <https://www.youtube.com/watch?v=jgTv2W0mUP0>

²⁸ <https://sarahguo.com/blog/emadmostaque>

1 127. In May 2023, Mostaque said to a tech journalist that Stable Diffusion is “a hundred
2 thousand gigabytes of **images compressed** to a two-gigabyte file.”²⁹

3 128. Though the estimated quantity of training images and size of the model has varied in
4 these statements by Mostaque and Stability, the core message has been consistent: Stable Diffusion
5 is a “compressed” version of its training images that can be used to “recreate any of those [images]
6 and iterations of those.”

7 129. The subject of whether diffusion models store copies of protected expression from
8 their training images is an active area of research in the AI field. So far, the answer is: yes, they do,
9 and as the models get larger, so does their capacity to store such copies (that is, “memorize”).

10 130. This topic was explored in a January 2023 paper called “Extracting Training Data
11 from Diffusion Models” by Nicholas Carlini of Google and others.³⁰ (Below, the “Carlini Paper”.)
12 Carlini is one of the world’s leading AI researchers. He often studies the security of machine-
13 learning models, in particular issues related to the privacy and security of training data after it has
14 been ingested into the model.

15 131. According to Carlini, “[t]he appeal of generative diffusion models is rooted in their
16 ability to synthesize novel images that are ostensibly unlike anything in the training set.” But
17 Carlini notes that “diffusion models are explicitly trained to reconstruct the training set.” *See* Carlini
18 Paper at 12. Based on further experiments, Carlini concludes “that state-of-the-art diffusion models
19 *do* memorize and regenerate individual training examples.” *See* Carlini Paper at 1.

20 132. Carlini’s experiment involved supplying text prompts to Stable Diffusion 1.4 to see if
21 the prompts could elicit images essentially identical to those found in the training dataset. In many
22 instances Carlini was able to coax Stable Diffusion 1.4 to output copies of its training images. An
23 example from the paper is shown below, comparing certain training images with images output by
24 Stable Diffusion. In each case, the image in the “Original” line is a training image; the image
25

26
27 ²⁹ <https://www.zdnet.com/article/why-open-source-is-essential-to-allaying-ai-fears-according-to-stability-ai-founder/>

28 ³⁰ <https://arxiv.org/abs/2301.13188>

1 directly below in the “Generated” line is an image output from Stable Diffusion. As Carlini notes,
2 the generated images are nearly identical to the training images:

3 133.



8 134. Based on these tests with Stable Diffusion and another diffusion model, Carlini
9 concludes that storage of copies of training images “is pervasive in large diffusion models—and that
10 ... extraction [of these stored copies] is feasible.” *See* Carlini Paper at 7. Carlini concludes that
11 “diffusion models memorize and regenerate individual training images ... and more useful diffusion
12 models memorize more than weaker diffusion models. This suggests that the vulnerability of
13 generative image models may grow over time.” *See* Carlini Paper at 15.

14 135. Carlini also poses a question: “[d]o large-scale models work by generating novel
15 output, or do they just copy and interpolate between individual training examples?” *Id.* He
16 concludes that “because our attacks [*i.e.*, attempts to elicit stored copies of training images]
17 succeed, this question remains open.” *Id.* François Chollet has taken an even stronger position,
18 saying that “It’s accurate that generative art models create new content by recombining images from
19 their training data.”³¹ Yaodong Yu concurs, stating “large image generation models ... are unlikely
20 to do anything beyond purely mechanical data compression (encoding) and interpolation
21 (decoding).” *See* Yu Paper at 53.

22 136. Carlini notes that a limitation of his experiment is that it relied on a very strict
23 “definition of ‘memorization’: whether diffusion models can be induced to generate” essentially
24 identical copies of certain training images “when prompted with appropriate instructions.” *See*
25 Carlini Paper at 4. Carlini says later: “[o]ur work highlights the difficulty in defining *memorization*
26

27 _____
28 ³¹ <https://twitter.com/fchollet/status/1600230516934209536>

1 ... a more comprehensive analysis will be necessary to accurately capture more nuanced definitions
2 of memorization that allow for more human-aligned notions of data copying.” See Carlini Paper at
3 15.

4 137. On information and belief, had Carlini adopted a more “human-aligned” standard of
5 visual correspondence—say, output images that were merely substantially similar to training
6 images—his experiment would’ve yielded many more successful results.

7 138. In May 2023, researcher Ryan Webster extended Carlini’s research in this direction
8 in a paper called “A Reproducible Extraction of Training Images from Diffusion Models.”³²
9 Webster found that by using a less strict technique for detecting stored copies of training images,
10 more instances of stored copies could be discovered. Webster tested several diffusion models,
11 including Stable Diffusion 2.0 and version 4 of the Midjourney Model, and found stored copies of
12 training images within all of them.

13 139. Carlini’s paper tested Stable Diffusion version 1.4, which had fewer than a billion
14 parameters. (A *parameter* is a single numerical value that a model learns during training, and
15 models with more parameters are considered “larger” than those with fewer.) But in July 2023,
16 Stability released Stable Diffusion XL 1.0, which has 3.5 billion parameters. On information and
17 belief, based on Carlini’s theory that larger models are more likely to “memorize and regenerate
18 individual training images,” a model like Stable Diffusion XL 1.0 is even more likely to exhibit this
19 behavior than the version 1.4 that Carlini tested.

20 140. More broadly, over time, AI image products are tending to adopt models with more
21 parameters, and thus, according to Carlini’s theory, these models are only getting better at storing
22 copies of training images and regenerating them in whole or in part.

23 141. A related topic was explored in a July 2023 paper called “Measuring the Success of
24 Diffusion Models at Imitating Human Artists”³³ by Stephen Casper of MIT and others. (Below, the
25 “Casper Paper”.)

26
27 ³² <https://arxiv.org/abs/2305.08694>

28 ³³ <https://arxiv.org/pdf/2307.04028.pdf>.

1 142. Starting with a list of 70 artist names, Casper supplied prompts to Stable Diffusion
2 version 1.5 in the form of “artwork from [name of artist]” to produce output images. Casper then
3 passed these images into a CLIP model to see whether it could correctly predict the artist being
4 imitated.

5 143. Casper found that the CLIP model “classified 81.0% of the generated images as
6 works made by artists whose names were used to generated them ... Overall, these results suggest
7 that Stable Diffusion has a broad-ranging ability to imitate the style of individual artists.” *See*
8 Casper Paper at 3.

9 144. In sum—based on work by leading AI researchers, AI image products are only
10 getting better and better at storing copies of training images and can even produce images
11 indistinguishable from those created by a specific artist in the training dataset.

12 145. Carlini’s paper shows that diffusion models—and Stable Diffusion in particular—
13 have the ability to store copies of protected expression from training images and later regenerate it.
14 Moreover, diffusion models have an increasing propensity to do so as they get larger, leading to a
15 stronger inference that these models generate output merely by “copy[ing] and interpolat[ing]
16 between individual training” images. *See* Carlini Paper at 15.

17 146. Casper’s paper shows another effect of this propensity to store copies of protected
18 expression from training images: diffusion models—and Stable Diffusion in particular—are
19 exceptionally good at creating convincing images resembling the work of specific artists if the
20 artist’s name is provided in the prompt. *See* Casper Paper at 3.

21 147. Stable Diffusion is held out as an “open-source” program. But in the hands of
22 Runway and Stability, the term “open source” is more of a marketing and competitive gimmick than
23 a substantive virtue, intended mostly to ensure the widest distribution of Stable Diffusion, and the
24 economic opportunities that result. David Widder and others strongly criticized this corruption of
25 the traditional meaning of “open source” in an August 2023 paper called “Open (for Business): Big
26
27
28

1 Tech, Concentrated Power, and the Political Economy of Open AI”³⁴ (emphasis added below):

2 As a rule, ‘open’ refers to systems that offer transparency,
 3 reusability, and extensibility—they can be scrutinized, reused, and
 4 built on. ... we also find that marketing around openness and
 5 investment in (somewhat) open AI systems is being leveraged by
 6 powerful companies to bolster their positions in the face of growing
 7 interest in AI regulation. And that some companies have moved to
 8 embrace ‘open’ AI as a mechanism to entrench dominance, using the
 9 rhetoric of ‘open’ AI to expand market power, and investing in
 10 ‘open’ AI efforts in ways that allow them to set standards of
 development while benefiting from the free labor of open source
 contributors ... **Companies like ... Stability AI offer open source
 AI models to their customers and the public. Their business
 models rely not on licensing proprietary models themselves, but
 on charging for extra features and services on top of open
 models,** features such as API access, model training on custom data,
 and security and technical support as a paid service to clients ...

11 See Widder at 11.

12 148. If one downloads the Stable Diffusion 2.0 model from Stability via GitHub,³⁵ one
 13 does not get everything one needs to operate Stable Diffusion 2.0, let alone train a comparable
 14 model from scratch. Rather, one gets a set of scripts (mostly written in Python) and configuration
 15 files for generating images using a weights file (not included, some assembly required). A *weights*
 16 file is a binary file that encodes all the information that the model captured during training by
 17 copying protected expression from training images. An example of a Stable Diffusion weights file is
 18 available at <https://huggingface.co/stabilityai/stable-diffusion-2/blob/main/768-v-ema.ckpt>. This is a
 19 mass of binary data that is meaningful when accessed via the Stable Diffusion scripts, but otherwise
 20 not intelligible to humans. In that way, the weights file has a status similar to that of a videogame
 21 cartridge or DVD—it contains protected expression that can be perceived only with the aid of a
 22 device (in this case, the accompanying software scripts), which can only be seen when interacted
 23 with using the appropriate device, for example, a videogame console or a DVD player. In other
 24 words, simply because certain code for a particular model is labeled “open source” does not mean
 25 one can meaningfully interrogate the model or experiment with it. All the value of the model is

26 _____
 27 ³⁴ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543807

28 ³⁵ <https://github.com/Stability-AI/stablediffusion>

1 encapsulated in the weights file, and a weights file is unintelligible to human beings.

2 149. For this reason, the most direct method of interrogating a model is to generate output
3 images with an AI image product containing that model.

4 150. The behaviors of diffusion models described in the Carlini Paper and the Casper
5 Paper can be observed in the output of the AI image products offered by Stability, Runway, and
6 Midjourney when prompted with Plaintiffs' names and images, as shown in the next sections.

7 **XI. EXAMPLES OF TEXT PROMPTS USING PLAINTIFF NAMES IN AI IMAGE**
8 **PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY**

9 151. What follows are examples from the current versions of image products offered by
10 Stability, Runway, and Midjourney (as of November 2023) showing the results of text prompts
11 invoking the names of certain Plaintiffs.

12 152. As mentioned below, Plaintiffs have found at least one instance where a Defendant
13 (Stability) has apparently adjusted the behavior of its AI image product to make prompting with
14 Plaintiffs' names more difficult, possibly in response to filing the initial complaint in this action. At
15 this juncture, it is impossible for Plaintiffs to know the full scope of measures that Defendants may
16 have adopted in their AI image products to frustrate Plaintiffs' investigation of the claims in this
17 complaint.

18 153. On information and belief, each Defendant is able to control the output of their
19 specific AI image products, including prohibiting the use of certain keywords in prompts. In this
20 way, Defendants can exercise control over the behavior of their AI image products.

21 154. **Stability:** Stability makes the current version of Stable Diffusion available through
22 an online AI image product called DreamStudio (<https://beta.dreamstudio.ai>). DreamStudio was
23 used to make the text-prompt examples below. The version of the Stable Diffusion model made
24 available in DreamStudio is called Stable Diffusion XL 1.0. This version of Stable Diffusion was
25 trained on the Plaintiff works in **Exhibit A**.

26 155. Plaintiffs Sarah Andersen, Kelly McKernan, and Karla Ortiz cannot show samples of
27 DreamStudio images with their names in the text prompts because their names have been blocked
28

1 since the initial complaint in this action was filed. Using one of these three names in DreamStudio
2 produces the error “Something isn’t quite right with your prompts.” Before Plaintiffs Andersen,
3 McKernan and Ortiz filed their initial complaint, however, their names could be used as prompts to
4 generate images.

5 156. On information and belief, Stability has blocked these names deliberately within its
6 DreamStudio app as a response to Ms. Andersen, Mx. McKernan, and Ms. Ortiz, that is intended to
7 thwart investigation of their claims against Stability.

8 157. Plaintiff Grzegorz Rutkowski cannot show samples of Stable Diffusion images with
9 his name in prompts. Due to the massive popularity of his name in Stable Diffusion prompts—one
10 report estimated Mr. Rutkowski’s name had been invoked “over 400,000 times”³⁶—Mr. Rutkowski
11 was also removed as a possible prompt.

12 158. Still, despite Stability’s attempt to inhibit the use of artist names in prompts,
13 demonstrative output can still be elicited. The text-prompt examples for Stable Diffusion that appear
14 below rely on the names of three Plaintiffs who have not been blocked: Gregory Manchess, Gerald
15 Brom, and Jingna Zhang.

16 159. Each of these artists has a distinctive artistic style that can be seen in the examples of
17 their work included in **Exhibit A**:

- 18 a. Gregory Manchess is known for his classic oil paintings distinguished by
19 their handcrafted brushwork, calligraphic style, and realistic themes.
- 20 b. Gerald Brom is known for his gritty, dark, fantasy images that combine
21 classical realism, gothic, and countercultural aesthetics.
- 22 c. Jingna Zhang is known for her painting and romantic photography, with
23 special attention to color, movement, and props. (Many of the subjects of Ms.
24 Zhang’s photographs happen to be Asian, a fact that is incidental to her work,
25

26
27 ³⁶ See “Greg Rutkowski Was Removed From Stable Diffusion, But AI Artists Brought Him Back,”
28 <https://decrypt.co/150575/greg-rutkowski-removed-from-stable-diffusion-but-brought-back-by-ai-artists>

1 but affects diffusion models in a peculiar way.)

2 160. The examples below appear in **Exhibit D: Stability text prompts**. To reveal the
3 effect of an artist's name on a text prompt, each artist's name is combined with a single word
4 representing a generic subject. For these examples, the subjects "chef" and "teacher" have been
5 used. Any difference between the output for each text prompt must therefore be attributable to the
6 influence of the artist's name on the prompt.

7 161. The first set of Stability text prompts consists of "*chef*," "*gregory manchess chef*,"
8 "*gerald brom chef*," and "*jingna zhang chef*." **Exhibit D**, pp. 1–4. The initial "*chef*" prompt shows
9 what Stable Diffusion produces by default, without an artist name in the prompt. The Manchess-
10 inspired chefs are rendered in the calligraphic brushwork characteristic of Mr. Manchess's oil
11 paintings. The Brom-inspired chefs have a gothic and countercultural air, including one serving a
12 shrunken skull, like a typical artwork by Mr. Brom. The Zhang-inspired chefs are Asian and
13 rendered photographically, like many of Ms. Zhang's images. In each case, the addition of the
14 artist's name causes the generic term "chef" to be rendered in a manner characteristic of the artist.

15 162. The next set of Stability text prompts consists of "*teacher*," "*gregory manchess*
16 *teacher*," "*gerald brom teacher*," and "*jingna zhang teacher*." **Exhibit D**, pp. 5–8. The initial
17 "*teacher*" prompt shows what Stable Diffusion produces by default, without an artist name in the
18 prompt. The Manchess-inspired teachers are rendered in calligraphic brushwork with realistic
19 settings, as frequently found in the artwork of Mr. Manchess. The Brom-inspired teachers are
20 demonic, and feature images of weapons and skulls, like a characteristic artwork by Mr. Brom. The
21 Zhang-inspired teachers are Asian and rendered photographically, like many of Ms. Zhang's
22 images. As with "*chef*," the addition of the artist's name causes the generic term "teacher" to be
23 rendered in a manner characteristic of the artist. Furthermore, the changes provoked by the insertion
24 of the artist's name are comparable for both "*chef*" and "*teacher*."

25 163. **Runway:** Runway makes a text-to-image generator available via its online AI image
26 product called AI Magic Tools (<https://app.runwayml.com/>). This online app was used to make the
27 text-prompt examples below. On information and belief, Runway's AI Magic Tools app uses Stable
28

1 Diffusion 1.5, because Runway trained that version of Stable Diffusion, and trained it on the
2 Plaintiff works in **Exhibit A**.

3 164. The examples below appear in **Exhibit E: Runway text prompts**.

4 165. The text-prompt examples for Runway use the same three Plaintiffs as before—
5 Gregory Manchess, Gerald Brom, and Jingna Zhang—and also Kelly McKernan and Sarah
6 Andersen, who have not been blocked in Runway prompts.

7 166. Like the other three artists, Mx. McKernan and Ms. Andersen also have distinctive
8 styles:

- 9 a. Kelly McKernan is known for their colorful, flowing, Art Nouveau-inspired
10 images that frequently feature female-presenting subjects and intricate
11 backgrounds.
- 12 b. Sarah Andersen is known for her black & white comic, “Sarah’s Scribbles,”
13 whose main character is a young woman with dark hair, large eyes, and a
14 striped shirt.

15 167. The first set of Runway text prompts consists of “*chef*,” “*gregory manchess chef*,”
16 “*gerald brom chef*,” “*jingna zhang chef*,” “*kelly mckernan chef*,” and “*sarah andersen chef*.”
17 **Exhibit E**, pp. 1–4. The initial “*chef*” prompt shows what the Runway image product produces by
18 default, without an artist name in the prompt. Much like the Stability results, the Manchess-inspired
19 chefs are rendered in the calligraphic brushwork characteristic of Mr. Manchess’s oil paintings. The
20 Brom-inspired chefs have gothic and countercultural air, with skulls hanging in the background of
21 one image, typical of artworks by Mr. Brom. The Zhang-inspired chefs are Asian and rendered
22 photographically, like many of Ms. Zhang’s subjects. The McKernan-inspired chefs feature colorful
23 female-presenting faces with elaborate hair and decorations, commonly found in Mx. McKernan’s
24 work. The Andersen-inspired chefs are all obvious variations on the main character of Ms.
25 Andersen’s celebrated comic “Sarah’s Scribbles,” and two of the images even include the panels
26 common in Ms. Andersen’s work.

27 168. The next set of Runway text prompts consists of “*teacher*,” “*gregory manchess*
28

1 *teacher*,” “*gerald brom teacher*,” “*jingna zhang teacher*,” “*kelly mckernan teacher*,” and “*sarah*
2 *andersen teacher*.” **Exhibit E**, pp. 5–8. The initial “*teacher*” prompt shows what the Runway image
3 product produces by default, without an artist name in the prompt. Much like the Stability results,
4 the Manchess-inspired teachers are rendered in calligraphic brushwork with realistic settings,
5 common in the work of Mr. Manchess. The Brom-inspired teachers are fantastic, otherworldly
6 creatures, resembling those often found in Mr. Brom’s work. The Zhang-inspired teachers are Asian
7 and rendered photographically, like many of Ms. Zhang’s subjects. The McKernan-inspired teachers
8 feature colorful female-presenting faces with elaborate hair and decorations, as often seen in Mx.
9 McKernan’s work. The Andersen-inspired teachers are all obvious variations on the main character
10 of Ms. Andersen’s comic “Sarah’s Scribbles,” and two of the images even include the panels
11 common in Ms. Andersen’s work.

12 169. **Midjourney**: Midjourney makes the current version of its AI image product
13 available through an online discussion service called Discord (<https://discord.com>). Midjourney’s
14 AI image product is also called “Midjourney.” Thus, for clarity below, the Midjourney AI image
15 product as presented via Discord will be called the **Midjourney Image Product**. The text-prompt
16 examples below were made using the Midjourney Image Product. Because the Midjourney Image
17 Product incorporates both Stable Diffusion as an underlying model as well as a custom model
18 trained by Midjourney, this latter model will be called the **Midjourney Model**.

19 170. The version of the Midjourney Model made available in the current Midjourney
20 Image Product is version 5.2. On information and belief, this version of the Midjourney Model was
21 trained on the Plaintiff works in **Exhibit A**. This is the default model used by the Midjourney Image
22 Product. Midjourney Model 5.2 was used to make the examples below.

23 171. The examples below appear in **Exhibit F: Midjourney text prompts**.

24 172. The text-prompt examples for Midjourney use the same five Plaintiffs as before—
25 Gregory Manchess, Gerald Brom, Jingna Zhang, Kelly McKernan, and Sarah Andersen, who have
26 not been blocked in Midjourney prompts.

27 173. By default, the Midjourney Model layers onto every output image what it calls the
28

1 “Midjourney default aesthetic.” Because this “default aesthetic” is an overbaked visual style
2 supplied by designers at Midjourney, not the training images, it has been turned off in the examples
3 below by appending “--style raw” to each prompt listed, thereby more accurately revealing the style
4 changes that arise from changing the artist names.

5 174. The first set of Midjourney text prompts consists of “*chef*,” “*gregory manchess chef*,”
6 “*gerald brom chef*,” “*jingna zhang chef*,” “*kelly mckernan chef*,” and “*sarah andersen chef*.”

7 **Exhibit F**, pp. 1–6. The initial “*chef*” prompt shows what the Midjourney Model produces by
8 default, without an artist name in the prompt. Much like the Stability results, the Manchess-inspired
9 chefs are rendered in the calligraphic brushwork characteristic of Mr. Manchess’s oil paintings. The
10 Brom-inspired chefs have gothic and countercultural air, including several fantastic creatures,
11 typical of artwork by Mr. Brom. The Zhang-inspired chefs are Asian and rendered photographically,
12 like many of Ms. Zhang’s works. The McKernan-inspired chefs feature colorful female-presenting
13 faces with elaborate hair and decorations, commonly found in Mx. McKernan’s work. The
14 Andersen-inspired chefs are all obvious variations on the main character of Ms. Andersen’s
15 celebrated comic “Sarah’s Scribbles.”

16 175. The next set of Midjourney text prompts consists of “*teacher*,” “*gregory manchess*
17 *teacher*,” “*gerald brom teacher*,” “*jingna zhang teacher*,” “*kelly mckernan teacher*,” and “*sarah*
18 *andersen teacher*.” **Exhibit F**, pp. 7–12. The initial “*teacher*” prompt shows what the Midjourney
19 Model produces by default, without an artist name in the prompt. Much like the Stability results, the
20 Manchess-inspired teachers are rendered in calligraphic brushwork with realistic settings, common
21 in the work of Mr. Manchess. The Brom-inspired teachers are demonic, and feature images of
22 skulls, common motifs in Mr. Brom’s work. The Zhang-inspired teachers are Asian and rendered
23 photographically, like many of Ms. Zhang’s images. The McKernan-inspired teachers feature
24 colorful female-presenting faces with elaborate hair and decorations, as often seen in Mx.
25 McKernan’s work. The Andersen-inspired teachers are all obvious variations on the main character
26 of Ms. Andersen’s comic “Sarah’s Scribbles.”

27 176. In sum, Stability’s diffusion model (Stable Diffusion XL 1.0), Runway’s diffusion
28

1 model (inferred to be Stable Diffusion 1.5), and the Midjourney Model version 5.2 demonstrate
2 behavior similar to that described in the Casper research paper: by adding a certain artist name to a
3 prompt, one can consistently elicit characteristic elements of that artist's body of work in the output
4 images, allowing the creation of unlimited output images that are substantially similar to, and could
5 be mistaken for those of the original artist. These results are consistent between prompts and
6 between models. This strongly suggests that the Stable Diffusion XL, the Runway Model, and the
7 Midjourney Model store copies of protected expression after copying and ingesting training images.

8 **XII. EXAMPLES OF IMAGE PROMPTS USING PLAINTIFF IMAGES**
9 **IN AI IMAGE PRODUCTS OFFERED BY STABILITY, RUNWAY, AND**
10 **MIDJOURNEY**

11 177. What follows are examples from the current versions of AI image products offered
12 by Stability, Runway, and Midjourney (as of November 2023) showing the results of prompting
13 these AI image products with the works of certain Plaintiffs that appear in the LAION-5B dataset.

14 178. These examples demonstrate an even more precise way of prompting an AI image
15 product: through image prompts. As explained above, systems based on CLIP-guided diffusion use
16 a CLIP model to convert each text prompt to a numerical descriptor called a CLIP embedding,
17 which in turn guides the diffusion process as the image emerges. When AI image products were first
18 released, users would provide a text prompt as input, which was converted into a CLIP embedding.

19 179. But a CLIP model can also produce a CLIP embedding from an image. Thus, an
20 image prompt for an AI image product works the same way as a text prompt, but with an image
21 rather than text as the initial user input that produces the CLIP embedding. The CLIP embedding
22 does not directly represent text or image data.

23 180. **Stability:** The first set of image-prompt examples were made with a Stability AI
24 image product called Reimagine XL, released in May 2023. Reimagine XL is built atop the Stable
25 Diffusion XL model that was also used for the text-prompt examples in the previous section. The
26 difference is that Reimagine XL accepts image prompts rather than text prompts. As Stability
27
28

1 explains³⁷ (emphasis added below)—

2 The classical text-to-image Stable Diffusion XL model is trained to
3 be conditioned on text inputs. [Reimagine XL] replaces the original
4 text encoder with an image encoder. So instead of generating images
5 based on text input, images are generated from an image. ... This
6 approach produces similar-looking images with different details and
7 compositions. Unlike the image-to-image algorithm, **the source
8 image is first fully encoded, so the generator does not use a single
9 pixel from the original one!**

10 181. The emphasized text is key: “not ... a single pixel” from the input image is being
11 passed into the model, just a higher-level numerical description of the image in the form of a CLIP
12 embedding. Stability emphasizes that every image output by Reimagine XL is freshly generated
13 with its own “details and composition” and promises output images that are merely “similar.”

14 182. Carlini’s research indicated that large diffusion models like Stable Diffusion XL have
15 a greater propensity for storing copies of protected expression from training images.

16 183. **Exhibit G: Stability image prompts** contains examples of prompting Reimagine
17 XL with training images from **Exhibit A** (Plaintiff images in LAION-5B). These training images
18 were made by plaintiffs Gerald Brom, Gregory Manchess, Grzegorz Rutkowski, Hawke
19 Southworth, Jingna Zhang, Karla Ortiz, Kelly McKernan, and Sarah Andersen.

20 184. On each page of this exhibit, the original training image is positioned in the upper
21 left; the other three images are output images. In every case, the output images are not merely
22 similar to the training image, but substantially similar—in some cases startlingly so. On information
23 and belief, because Stability says “not ... a single pixel” from the input image is being passed into
24 Stable Diffusion XL (via the Reimagine XL image product), it would not be possible for Stable
25 Diffusion XL to produce output images substantially similar to the training images unless it had
26 stored copies of protected expression from those training images, and the CLIP embedding
27 generated from the image prompt was eliciting the output of this copied expression.

28 185. **Runway:** The next set of image-prompt examples were made with Runway’s AI

³⁷ <https://clipdrop.co/stable-diffusion-reimagine>

1 Magic Tools using its “Image Variation” feature. On information and belief, based on the output of
2 this tool, it works in a manner similar to Stability’s Reimagine XL: none of the pixels of the input
3 image are retained, but rather a CLIP embedding is generated from the input image, which guides
4 the subsequent diffusion process.

5 186. **Exhibit H: Runway image prompts** contains examples of prompting Runway’s
6 Image Variation tool with training images from Exhibit A (Plaintiff images in LAION-5B). These
7 training images were made by plaintiffs Gerald Brom, Gregory Manchess, Grzegorz Rutkowski,
8 Hawke Southworth, Jingna Zhang, Karla Ortiz, and Kelly McKernan.

9 187. On each page of this Exhibit, the original training image is positioned in the upper
10 left; the other three images are output images. In every case, the output images are not merely
11 similar to the training image, but substantially similar. On information and belief, it would not be
12 possible for the Runway Image Variation tool to produce output images substantially similar to the
13 training images unless it had stored copies of protected expression from those training images, and
14 the CLIP embedding generated from the image prompt was eliciting the output of this copied
15 expression.

16 188. **Midjourney:** The next set of image-prompt examples were made with Midjourney
17 using its image-prompting feature. This feature was released by Midjourney one day after the initial
18 complaint in this action was filed. According to Midjourney CEO David Holz, this feature does not
19 copy pixels from the input, but rather “looks at the ‘concepts’ and ‘vibes’ of your images and
20 merges them together into novel interpretations.”³⁸ According to Midjourney’s documentation,
21 using an image as a prompt merely “influence[s] a Job’s composition, style, and colors.”³⁹

22 189. **Exhibit I: Midjourney image prompts** contains examples of prompting
23 Midjourney with training images from **Exhibit A** (Plaintiff images in LAION-5B). These training
24 images were made by Plaintiffs Gerald Brom, Gregory Manchess, Grzegorz Rutkowski, Hawke
25

26 38

27 <https://discord.com/channels/662267976984297473/952771221915840552/1064031587735445546>

28 ³⁹ <https://docs.midjourney.com/docs/image-prompts>

1 Southworth, Jingna Zhang, Karla Ortiz, and Kelly McKernan. Midjourney requires that an image
2 prompt be accompanied by a text prompt, so in these cases, each image prompt was supplemented
3 by the artist’s name. Each prompt was also supplemented with the command ‘--iw 2’ to ensure that
4 the image portion of the prompt was treated as the primary part (where “iw” means “image
5 weight”), thereby maximizing “the ‘concepts’ and ‘vibes’” derived from the image.

6 190. On each page of this Exhibit, the original training image is positioned in the upper
7 left; the other three images are output images. In every case, the output images are not merely
8 similar to the training image, but substantially similar. On information and belief, since Midjourney
9 says only “‘concepts’ and ‘vibes’” are being taken from the input image, it would not be possible
10 for the Midjourney Model to produce output images substantially similar to the training images
11 unless it had stored copies of protected expression from those training images, and the CLIP
12 embedding generated from the image prompt was eliciting the output of this copied expression.

13 191. In sum, the models offered by Stability, Runway, and Midjourney demonstrate
14 behavior apparently similar to that described in the Carlini paper: by using a sufficiently precise
15 CLIP embedding as conditioning, one can consistently elicit protected expression from a training
16 image, allowing the creation of unlimited output images that could be mistaken for copies of the
17 training images. These results are consistent between prompts and between models.

18 192. Taken together, these examples of text prompting and image prompting strongly
19 imply that diffusion models like the ones shown above store copies of protected expression from
20 copying and ingesting training images.

21 193. Further, because the makers of these AI image products allow users and licensees to
22 generate copies based on uploaded images and promote their models’ ability to do so, the
23 proliferation of CLIP models invite further infringement.

24 194. The models also create visually similar copies based on original work with
25 copyright-management information removed or altered.

26 195. In this example (**Ex. G** at p. 2), the original image is in the top-left quadrant. Plaintiff
27 Brom’s CMI in the form of the URL for his personal website is clearly visible. None of the visually
28

1 similar copies of the original image generated by the Stable Diffusion XL model contained the
2 original CMI.



1 196. In this example (**Ex. G** at p. 13), the original image is in the top left. Plaintiff
2 Zhang’s CMI in the form of the URL for her personal website is clearly visible in the bottom right
3 corner. The Stable Diffusion model again generated visually similar copies of the work without its
4 CMI.



1 197. In this example (**Ex. G** at p. 15), the original image is on the left. Plaintiff Ortiz's
2 CMI in the form of her signature is visible in the bottom right. The Stable Diffusion model
3 generated visually similar copies, including the one depicted which plainly showed Plaintiff Ortiz's
4 CMI altered on the bottom-left corner of the image.



1 198. In this example (**Ex. H** at p. 2), the original image is on the left. Plaintiff Brom's
2 CMI in the form of the URL for his personal website is clearly visible. The Runway model
3 generated visually similar copies, including the one depicted with Plaintiff Brom's CMI removed.
4



17 199. In this example (**Ex. H** at p. 12), the original image is on the left. Plaintiff Zhang's
18 CMI in the form of the URL for her personal website is clearly visible in the bottom-left of the
19 image. The Runway model generated visually similar copies, including the one depicted to the right,
20 with Plaintiff Zhang's CMI removed.
21
22
23
24
25
26
27
28



200. In each of these examples, the copies generated by the AI image product could not have been generated but for copying the original image which included CMI by operation of the processes described herein.

XIII. USER AND LICENSEE ACTIVITY

201. What is described herein is not hypothetical. Individuals have and are using the AI image products to create images that mimic and imitate Plaintiffs and Class members' work. Further, users and licensees, with assistance from Defendants, track and update the specific artists (including Plaintiffs) which the AI image products are able to mimic or imitate.

202. Midjourney Image Product users and licensees maintain a spreadsheet that features community-created images and tracks the specific artists that the Midjourney Image Product can successfully mimic or imitate. The list includes thousands of artists and tracks user- and licensee-generated works using artist names as prompts for various versions of the Midjourney Model. Plaintiffs listed in the Midjourney community-generated spreadsheet include Sarah Andersen, Gerald Brom, Gregory Manchess, Kelly McKernan, and Jingna Zhang.

203. An AI image product user posted onto the internet a purported study representing "[a] collection of images from Midjourney that demonstrate the recognized artist styles and their

1 results on Midjourney (v4 model).”⁴⁰ The purpose of the post was, in the words of the user, to “help
2 decide the style to invoke when prompting the AI to create your desired images. The user used the
3 following prompt into Midjourney “Art by <artist name>.” Included in the study were several
4 Plaintiffs including Sarah Andersen, Kelly McKernan and Jingna Zhang.

5 204. That same individual later posted an “artist study” wherein the individual used the
6 prompt “art by<artist>” with a “negative prompt” of “blurry, soft, low quality”; the artist’s name as
7 the only change in the prompt into Stable Diffusion XL.⁴¹ As demonstrated by the user, “[t]he
8 prompt was straight forward ‘art by <artist>’ which would get the SDXL mode [sic] to emulate the
9 style and creations of that artists [sic].” Included in this artist study were several of the Plaintiffs
10 including Sarah Andersen, Gerald Brom, Kelly McKernan, Karla Ortiz and Jingna Zhang.

11 205. Another example involves users conducting “Artist Style Studies” using “Stable
12 Diffusion V1”.⁴² The user input a series of six prompts: prompt 1-3 included “a portrait of a
13 character in a scenic environment by [artist]” and prompts 4-6 included “a building in a stunning
14 landscape by [artist].” This particular study includes over 1,781 artists to date and includes a form to
15 recommend other artist names to input. Plaintiffs who have been included in this study include
16 Plaintiffs Andersen, Brom, McKernan, Ortiz, and Zhang.

17 206. Similar artist studies exist for the Runway Models, including compilations of artists
18 whose names were confirmed to be recognized by Stable Diffusion 1.5 and other Runway Models,
19 *i.e.*, artists who the Runway Models were capable of mimicking or imitating.

20 207. As these example images show, users of AI image products are employing machine-
21 learning models to create output that is indistinguishable from works created by Plaintiffs and Class
22 Members. Users are doing so with the intent of emulating the artist’s work without any of the
23 compensation or credit that would typically be required if an individual wants to commission an
24 artist to create artwork.

25
26 ⁴⁰ <https://weirdwonderfulai.art/resources/artist-styles-on-midjourney-v4/>

27 ⁴¹ <https://weirdwonderfulai.art/resources/stable-diffusion-xl-sd-xl-artist-study/>

28 ⁴² <https://proximacentaurib.notion.site/e28a4f8d97724f14a784a538b8589e7d?v=42948fd8f45c4d47a0edfc4b78937474>

1 208. This is not done without assistance by Defendants. Defendants each materially assist
2 by distributing the models themselves. As described herein, Defendants also encourage the use of
3 specific artist names—including Plaintiffs—as text prompts in order to adduce artwork
4 indistinguishable from Plaintiffs from the AI Image Products.

5 **XIV. DEFINITIONS FOR THE CAUSES OF ACTION**

6 209. The term **Statutory Copy** denotes the definition of *copies* in 17 U.S.C. § 101 of the
7 U.S. Copyright Act: “material objects ... in which a work is fixed by any method ... and from
8 which the work can be perceived, reproduced, or otherwise communicated, either directly or with
9 the aid of a machine or device.”

10 210. The term **Statutory Derivative Work** denotes the definition of *derivative work* in 17
11 U.S.C. § 101 of the U.S. Copyright Act: “a work based upon one or more preexisting works, such as
12 ... [an] abridgment, condensation, or any other form in which a work may be recast, transformed, or
13 adapted.”

14 211. The term **LAION-5B Works** denotes the works in **Exhibit A**, all of which are part of
15 the LAION-5B dataset. Each Plaintiff is the author of one or more of the LAION-5B Works. The
16 Plaintiffs hold the exclusive rights to their respective LAION-5B Works under 17 U.S.C. § 106,
17 including the rights to make Statutory Copies, prepare Statutory Derivative Works, and distribute
18 both Statutory Copies and Statutory Derivative Works.

19 212. The term **LAION-5B Registered Works** denotes the subset of works in the LAION-
20 5B Works that are covered by registered copyrights.

21 213. The term **LAION-5B Registered Plaintiffs** denotes the subset of Plaintiffs who hold
22 copyrights in these LAION-5B Registered Works that were registered before the filing of the initial
23 complaint in this action, namely Sarah Andersen, Jingna Zhang, Gerald Brom, Gregory Manchess,
24 Julia Kaye, and Adam Ellis.

25 214. [REDACTED]
26 [REDACTED]
27 [REDACTED]
28 [REDACTED]

1 215. [REDACTED]
2 [REDACTED]
3 [REDACTED]
4 [REDACTED]

5 216. The term **DataComp xLarge Registered Works** denotes the subset of works
6 Midjourney downloaded that are covered by registered copyrights.

7 217. The term **DataComp xLarge Registered Plaintiffs** denotes the subset of Plaintiffs
8 who hold copyrights in these **DataComp xLarge Registered Works** that were registered before the
9 filing of the initial complaint in this action, including Adam Ellis, Gerald Brom, Gregory Manchess,
10 Grzegorz Rutkowski, Jingna Zhang, Julia Kaye, Karla Ortiz, and Sarah Andersen.

11 218. The term **DeviantArt DreamUp Registered Works** denotes the subset of works
12 DeviantArt used as part of its development of DreamUp that are covered by registered copyrights.

13 219. The term **DeviantArt DreamUp Registered Plaintiffs** denotes the subset of
14 Plaintiffs who hold copyrights in these **DeviantArt DreamUp Registered Works** that were
15 registered before the filing of the initial complaint in this action, including Jingna Zhang.

16
17 **XV. CAUSES OF ACTION AGAINST STABILITY**

18 220. Between April and November 2022, Stability trained an image model called Stable
19 Diffusion 2.0. According to Stability, “The model developers used the following dataset for training
20 the model: LAION-5B and subsets 5B” (*See* <https://huggingface.co/stabilityai/stable-diffusion-2#training>).
21

22 221. Between November 2022 and July 2023, Stability trained an image model called
23 Stable Diffusion XL 1.0. On information and belief, Stable Diffusion XL is also trained on LAION-
24 5B, because Stability has funded LAION and used LAION datasets for all its previous models. AI
25
26
27
28

1 chipmaker NVIDIA features Stable Diffusion XL as one of its “AI Foundation Models.” On its
2 information page for Stable Diffusion XL, the training dataset is listed as LAION-5B.⁴³

3 222. Because LAION-5B is an openly accessible dataset, Stability knew the LAION-5B
4 dataset contained copyrighted works, including those of the LAION-5B Registered Plaintiffs.
5 Additionally, because the LAION-5B dataset contains detection scores for watermarks and because
6 CMI is ubiquitous in art, Stability also knew that the LAION-5B dataset contained copyrighted
7 works with CMI affixed to them.

8 223. The LAION-5B Registered Works are included in the LAION-5B dataset. Therefore,
9 Stability used the LAION-5B Registered Works for training. Below, the term **Stability Models**
10 refers to all models trained by Stability on the LAION-5B Registered Works, including Stable
11 Diffusion 2.0 and Stable Diffusion XL 1.0.

12 224. [REDACTED]

13 [REDACTED]

14 [REDACTED]

15 225. Since the filing of Plaintiffs Andersen, McKernan, and Ortiz’s initial complaint,
16 Stability has changed the behavior of the Stability models. Plaintiffs Andersen, McKernan and
17 Ortiz’s names can no longer be used as prompts. Plaintiff Rutkowski’s name has similarly been
18 blocked as a prompt. As demonstrated herein, however, each of their names have been used in the
19 past with Stability Models to generate work that mimicked their works.

20 **COUNT ONE**

21 **Direct copyright infringement of the LAION-5B Registered Works**
22 **by training the Stability Models, including Stable Diffusion 2.0 and Stable Diffusion XL 1.0**
23 **against Stability on behalf of the LAION-5B Registered Plaintiffs and Damages Subclass**

24 226. The preceding factual allegations are incorporated by reference.

25 227. The LAION-5B Registered Plaintiffs never authorized Stability to use their
26 respective LAION-5B Registered Works in any way. Nevertheless, Stability repeatedly violated the

27 _____
28 ⁴³ See <https://catalog.ngc.nvidia.com/orgs/nvidia/teams/ai-foundation/models/sdxl/overview>

1 exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs and continues to do
2 so today.

3 228. The LAION-5B dataset contains only URLs of training images, not the actual
4 training images. Therefore, anyone who wishes to use LAION-5B for training their own machine-
5 learning model must first acquire copies of the actual training images from their URLs using the
6 img2dataset or other similar tool. Consistent with this, in preparation for training the Stability
7 Models, Stability made one or more Statutory Copies of the LAION-5B Registered Works so they
8 could be fed to each Stability Model as training data. The Statutory Copies made of each registered
9 work were substantially similar to that registered work.

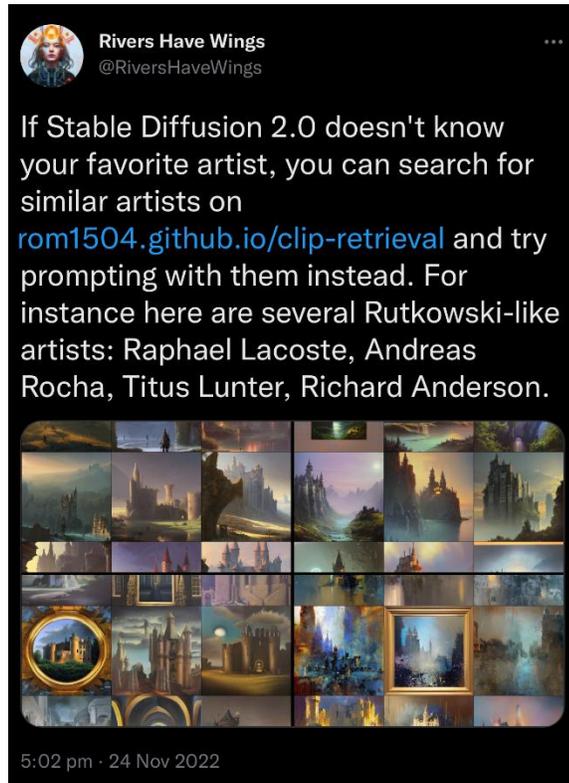
10 229. During the training of each Stability Model, Stability made a series of intermediate
11 Statutory Copies of the LAION-5B Registered Works. For instance, diffusion models are trained by
12 creating “noised” copies of training images, as described herein, all of which qualify as Statutory
13 Copies. The intermediate Statutory Copies of each registered work that Stability made during
14 training of the Stability Models were substantially similar to that registered work.

15 230. By the end of training, Stable Diffusion XL 1.0 was capable of reproducing protected
16 expression from each of the LAION-5B Registered Works that was in each case substantially
17 similar to that registered work, as shown in **Exhibit D: Stability text prompts** and **Exhibit G:**
18 **Stability image prompts**. Therefore, Stable Diffusion XL 1.0 qualifies as an infringing Statutory
19 Copy of the LAION-5B Registered Works. Because Stable Diffusion XL 1.0 represented a
20 transformation of the LAION-5B Registered Works into an alternative form, Stable Diffusion XL
21 1.0 also qualifies as an infringing Statutory Derivative Work.

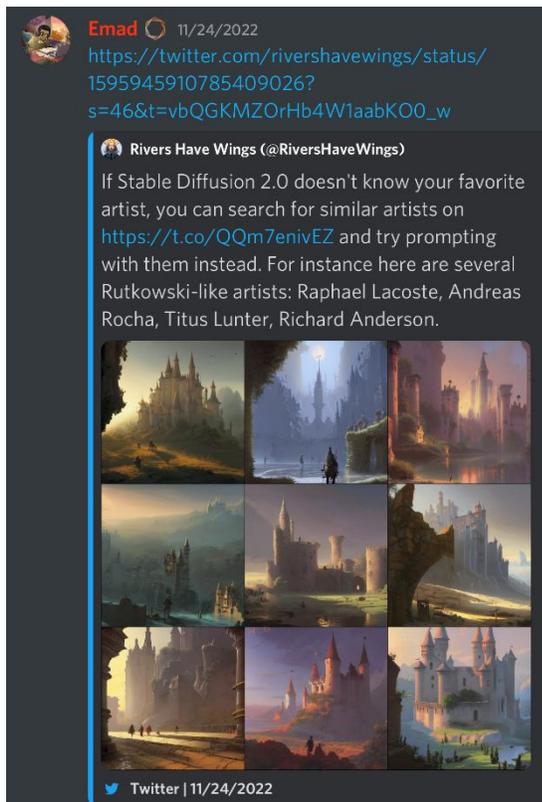
22 231. Executives and high-level employees of Stability know that one of the most
23 attractive features of the Stability models is its ability to mimic and copy artists’ works, including
24 Plaintiffs. As such, they routinely advertise the Stability Models’ ability to mimic artwork.

25 232. For example, once Stability prohibited the use of Plaintiff Rutkowski’s name as a
26 prompt, Stability’s employees and executives encouraged the use of similar artist names in lieu of
27 Plaintiff Rutkowski’s in order to achieve similar results.

1 233. For example, Katherine Crowson, a principal researcher at Stability AI tweeted the
2 following on November 24, 2022:



1 234. Emad Mostaque, Stability’s CEO retweeted Crowson’s advice:



15

16 235. Stability also maintains a Discord channel where executives routinely offered

17 resources to users including encouragement to use artist names as prompts:

18

19

20

21

22

23

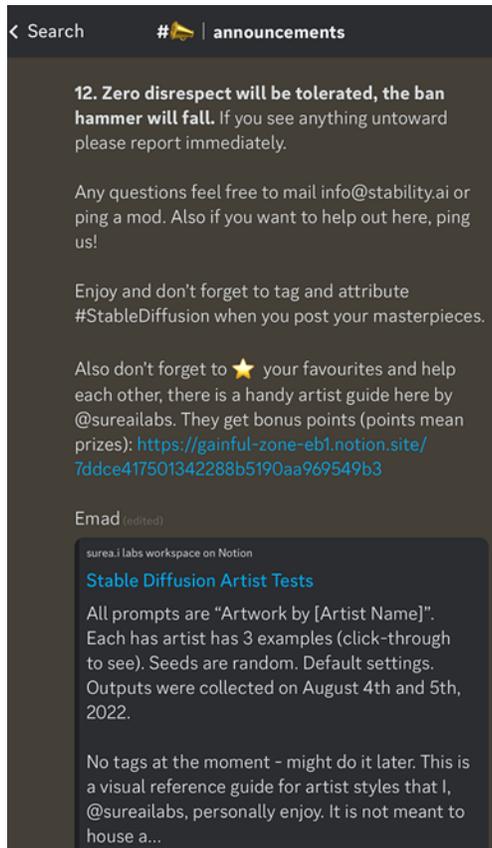
24

25

26

27

28



236. On information and belief, the other Stability Models exhibit the same properties, because they were trained on the same LAION-5B dataset.

237. Since November 2022, Stability has distributed Stable Diffusion 2.0 to the public, for instance via websites like GitHub (*See, e.g.*, <https://github.com/Stability-AI/stablediffusion>) and Hugging Face (*See, e.g.*, <https://huggingface.co/stabilityai/stable-diffusion-2>). Since July 2023, Stability has distributed Stable Diffusion XL 1.0 to the public, for instance via websites like GitHub (*See, e.g.*, <https://github.com/Stability-AI/generative-models>) and Hugging Face (*See, e.g.*, <https://huggingface.co/stabilityai/stable-diffusion-xl-base-1.0>). In so doing, Stability infringed the exclusive distribution rights of the LAION-5B Registered Plaintiffs.

238. The LAION-5B Registered Plaintiffs have been and continue to be injured by Stability's multiple acts of direct copyright infringement. These Plaintiffs are entitled to statutory damages, actual damages, restitution of profits, and other remedies provided by law.

COUNT TWO

Inducement of copyright infringement by distributing Stable Diffusion 2.0 and Stable Diffusion XL 1.0 for free against Stability on behalf of the LAION-5B Registered Plaintiffs and Damages Subclass

239. The preceding factual allegations are incorporated by reference.

240. Stability distributes Stable Diffusion 2.0 and Stable Diffusion XL 1.0 under the MIT License, which allows anyone to download, use, and deploy the Stability Models for free, for instance, via websites like GitHub (*See* <https://github.com/Stability-AI/stablediffusion>) and Hugging Face (*See* <https://huggingface.co/stabilityai/stable-diffusion-2>).

241. Stable Diffusion 2.0 and Stable Diffusion XL 1.0 violate the exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs. Therefore, anyone who in fact downloads, uses, or deploys Stable Diffusion 2.0 or Stable Diffusion XL 1.0 is engaged in infringing activity.

242. Stability has made a material contribution to this infringing activity by creating Stable Diffusion 2.0 and Stable Diffusion XL 1.0 and then distributing them for free.

243. Stability intends to cause further infringement with Stable Diffusion 2.0 and Stable Diffusion XL 1.0. In an interview in September 2022, Stability CEO Emad Mostaque said: “So Stable Diffusion is the model itself. It’s a collaboration that we did with a whole bunch of people ... We took 100,000 gigabytes of images and compressed it to a two-gigabyte file **that can recreate any of those and iterations of those.**” (emphasis added). With this comment, Mostaque explicitly promoted the ability of Stable Diffusion to “recreate”—that is, infringe the copyright of—images in its training dataset, including the LAION-5B Registered Works.

244. The LAION-5B Registered Plaintiffs have been and continue to be injured by Stability’s inducement of copyright infringement. These Plaintiffs are entitled to statutory damages, actual damages, restitution of profits, and other remedies provided by law.

1			
2	[REDACTED]	[REDACTED]	[REDACTED]
3	[REDACTED]	[REDACTED]	[REDACTED]
4	[REDACTED]	[REDACTED]	[REDACTED]
5	[REDACTED]	[REDACTED]	[REDACTED]
6	[REDACTED]	[REDACTED]	[REDACTED]
7	[REDACTED]	[REDACTED]	[REDACTED]
8			
9	[REDACTED]	[REDACTED]	[REDACTED]
10	[REDACTED]		
11	[REDACTED]	[REDACTED]	[REDACTED]
12	[REDACTED]	[REDACTED]	[REDACTED]
13			

14 248. [REDACTED]

15 [REDACTED]

16 [REDACTED]

17 249. [REDACTED]

18 [REDACTED]

19 [REDACTED]

XVI. CAUSES OF ACTION AGAINST MIDJOURNEY

20 250. Midjourney promotes the Midjourney Image Product, which is accessed and run

21 through Discord. Midjourney maintains its own Discord server from which users can access the

22 Midjourney image generator. Midjourney’s Discord server also allows Midjourney executives and

23 other high-level employees to provide promotional communications to users and licensees of the

24 Midjourney Image Product.

25 251. In February 2022, near the release of the initial version of the Midjourney Image

26 Product, Midjourney CEO David Holz posted messages on the Midjourney Discord server

27

28

1 promoting the Midjourney Image Product’s ability to emulate existing artistic styles, in particular
2 the styles of certain artists.

3 252. Over a series of Discord messages, Holz said “i think you’re all gonna get [your]
4 mind blown by this style feature ... we were very liberal in building out the dictionary ... it has
5 cores and punks and artist names ... as much as we could dump in there ... i should be clear it’s not
6 just genres its also artist names ... it’s mostly artist names ... 4000 artist names.”⁴⁵

7 253. Holz then said, “here is our style list”⁴⁶ and posted a link to a spreadsheet on Google
8 Docs called “Midjourney Style List.”⁴⁷ One of the tabs on the spreadsheet was called “Artists” and
9 listed over 4,700 artist names. In other words, Holz published a list of artists who the Midjourney
10 Image Product recognizes with the express purpose of these names being used by users and
11 licensees of the Midjourney Image Product as terms in prompts. Holz’s comment, and the list, have
12 remained available ever since.

13 254. Below, this list is called the **Midjourney Name List**. A copy of this list appears in
14 **Exhibit J: Midjourney Name List**.

15 255. Plaintiffs Grzegorz Rutkowski, Sarah Andersen, Karla Ortiz, Gerald Brom, and Julia
16 Kaye appear in the Midjourney Name List. Below, this subset of Plaintiffs will be referred to as the
17 **Midjourney Named Plaintiffs**.

18 256. Midjourney also tracked the most popular artists used as prompts. At one point, the
19 Midjourney Bot maintained a count on Midjourney as recently as April 2023.

20 257. In the months before February 2022, Midjourney trained version 1 of the Midjourney
21 Model. In February 2022, on Midjourney’s Discord server Midjourney CEO David Holz described

23 ⁴⁵ Combining
24 <https://discord.com/channels/662267976984297473/938713143759216720/941972360171520001>,
25 <https://discord.com/channels/662267976984297473/938713143759216720/941972890520272906>,
26 and
<https://discord.com/channels/662267976984297473/938713143759216720/941976464704802836>
⁴⁶

27 <https://discord.com/channels/662267976984297473/938713143759216720/941987328828768256>
28 ⁴⁷ https://docs.google.com/spreadsheets/d/1MEglfejpqgVcaf-I-cgZ5ngV_MlaOTeGXAoBPJO69FM/edit#gid=1096178862

1 Midjourney’s training data: “we have some private data partners as well as some open ones like
 2 laion.”⁴⁸ In May 2022, Stability CEO Mostaque said “MidJourney is using a LAION 400m based
 3 model ... I just support MJ like many other labs/researchers in my nice way.”⁴⁹ In July 2022, a
 4 Midjourney-affiliated Discord moderator named Soar also confirmed that Midjourney was “using a
 5 modified version of the LAION 400m dataset.”⁵⁰ In other words, Midjourney trained its image
 6 generation product on at least the LAION-400M dataset.

7 258. The term **LAION-400M Works** denotes the works in **Exhibit B** (Plaintiff images in
 8 LAION-400M), all of which are part of the LAION-400M dataset.

9 259. The term **LAION-400M Plaintiffs** denotes the subset of Plaintiffs who are the
 10 authors of works included in the LAION-400M Works. The LAION-400M Plaintiffs hold the
 11 exclusive rights to their respective LAION-400M Works under 17 U.S.C. § 106, including the rights
 12 to make Statutory Copies, prepare Statutory Derivative Works, and distribute both Statutory Copies
 13 and Statutory Derivative Works.

14 260. The term **LAION-400M Registered Works** denotes the subset of works in the
 15 LAION-400M Works that are covered by registered copyrights.

16 261. The term **LAION-400M Registered Plaintiffs** denotes the subset of Plaintiffs who
 17 hold registered copyrights in these LAION-400M Registered Works, namely Sarah Andersen and
 18 Jingna Zhang.

19 262. Because LAION-400M is an openly accessible dataset, Midjourney knew that the
 20 LAION-400M dataset contained copyrighted works, including those of the LAION-400M
 21 Registered Plaintiffs. Further, because LAION also includes detection scores for watermarks for
 22
 23

24
 25 48

26 <https://discord.com/channels/662267976984297473/938713143759216720/943315577018126408>

27 49

28 <https://discord.com/channels/823813159592001537/912729332311556136/975894553225752626>

⁵⁰<https://discord.com/channels/662267976984297473/959962985655320616/100193813644575138>
 7

1 other datasets, and because of the ubiquity of artists affixing CMI to their works, Midjourney knew
2 that the LAION-400M dataset contained works with CMI affixed on them.

3 263. The LAION-400M Registered Works are included in the LAION-400M dataset.
4 Therefore, Midjourney used the LAION-400M Registered Works for training. Below, the term
5 **Midjourney 400M Models** refers to all models trained by Midjourney on the LAION-400M
6 Registered Works, including version 1 of the Midjourney Model.

7 264. Since October 2022, Midjourney has also incorporated a version of Stable Diffusion
8 into the Midjourney Image Product, which is accessible to users by adding the command ‘--test’ or
9 ‘--testp’ to a text prompt. According to Midjourney’s moderator Molang, “--test and --testp is a little
10 bit of SD [Stable Diffusion] mixed with a lot of Midjourney tweaks and magic.”⁵¹
11 Sometime after July 2022 and before March 2023, Midjourney adopted LAION-5B as its training
12 dataset. In July 2022, a Midjourney-affiliated Discord moderator named Danger Awesome said that
13 “the updated LAION 5B dataset”⁵² formed the basis for “the upcoming [Midjourney] dataset
14 update.” In March 2023, a Midjourney-affiliated Discord Moderator named Sunshineyday said that
15 “MJ is trained on a subset of Laion5b.”⁵³

16 265. Midjourney also created a spreadsheet entitled “datasets”, which lists [REDACTED] and
17 “datacomp” among other datasets, points to their storage locations on Midjourney’s servers, and
18 provides additional detail about them.

19
20 266. Version 5 of the Midjourney Model was released in March 2023. On information and
21 belief, version 5 of the Midjourney Model and subsequent versions were trained on LAION-5B.

22
23
24
25 _____
26 51

27 <https://discord.com/channels/662267976984297473/958069758211797092/1038899058636501132>

28 52

<https://discord.com/channels/662267976984297473/992207085146222713/998451098534817883>

53

<https://discord.com/channels/662267976984297473/992207085146222713/1082089794521268314>

1 exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs and continues to do
2 so today.

3 278. The LAION-5B dataset contains only URLs of training images, not the actual
4 training images. Therefore, anyone who wishes to use LAION-5B for training their own machine-
5 learning model must first acquire copies of the actual training images from their URLs by using the
6 img2dataset tool or another similar tool. Consistent with this, in preparation for training the
7 Midjourney 5B Models, Midjourney made one or more Statutory Copies of the LAION-5B
8 Registered Works so they could be fed to each Midjourney 5B Model as training data. The Statutory
9 Copies made of each registered work were substantially similar to that registered work.

10 279. During the training of Midjourney Model version 5.2 and the other Midjourney 5B
11 Models, Midjourney made a series of intermediate Statutory Copies of the LAION-5B Registered
12 Works. For instance, diffusion models are trained by creating “noised” copies of training images, as
13 described herein, all of which qualify as Statutory Copies. The intermediate Statutory Copies of
14 each registered work that Midjourney made during training of the Midjourney 5B Models were
15 substantially similar to that registered work.

16 280. By the end of training, Midjourney Model version 5.2 was capable of reproducing
17 protected expression from each of the LAION-5B Registered Works that was in each case
18 substantially similar to that registered work, as shown in **Exhibit F: Midjourney text prompts** and
19 **Exhibit I: Midjourney image prompts**. Therefore, Midjourney Model version 5.2 qualifies as an
20 infringing Statutory Copy of the LAION-5B Registered Works. Because Midjourney Model version
21 5.2 represents a transformation of the LAION-5B Registered Works into an alternative form,
22 Midjourney Model version 5.2 also qualifies as an infringing Statutory Derivative Work.

23 281. On information and belief, the other Midjourney 5B Models exhibit the same
24 properties, because they were trained on the same LAION-5B dataset.

25 282. The LAION-5B Registered Plaintiffs have been and continue to be injured by
26 Midjourney’s multiple acts of direct copyright infringement. These Plaintiffs are entitled to statutory
27 damages, actual damages, restitution of profits, and other remedies provided by law.
28

COUNT FIVE

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

[Redacted text block]

[Redacted text block]

[Redacted]	[Redacted]	[Redacted]

[Redacted text block]

1	[REDACTED]	[REDACTED]	[REDACTED]
2	[REDACTED]	[REDACTED]	[REDACTED]
3	[REDACTED]	[REDACTED]	[REDACTED]
4	[REDACTED]	[REDACTED]	[REDACTED]
5	[REDACTED]	[REDACTED]	[REDACTED]
6	[REDACTED]	[REDACTED]	[REDACTED]
7	[REDACTED]	[REDACTED]	[REDACTED]
8	[REDACTED]	[REDACTED]	[REDACTED]
9	[REDACTED]	[REDACTED]	[REDACTED]
10	[REDACTED]	[REDACTED]	[REDACTED]
11	[REDACTED]	[REDACTED]	[REDACTED]
12	[REDACTED]	[REDACTED]	[REDACTED]
13	[REDACTED]	[REDACTED]	[REDACTED]
14	[REDACTED]	[REDACTED]	[REDACTED]

15 286. [REDACTED]

16 [REDACTED]

17 [REDACTED]

18 287. [REDACTED]

19 [REDACTED]

20 [REDACTED]

COUNT SEVEN

Direct copyright infringement of the DataComp xLarge Registered Works against Midjourney on behalf of the DataComp xLarge Registered Plaintiffs and the Damages Subclass

25 288. The preceding factual allegations are incorporated by reference.

26 289. DataComp is a competition and benchmark focused on the design of multimodal
27 datasets for pre-training CLIP models. Unlike traditional machine learning benchmarks that focus
28

1 on iterating model architectures or hyperparameters, DataComp fixes the model architecture and
 2 hyperparameters, challenging participants to instead innovate on dataset curation to improve
 3 downstream accuracy. *See* <https://github.com/mlfoundations/datacomp>.

4 290. To facilitate this competition, the organizers provide a massive reservoir of
 5 uncurated data known as the “CommonPool.” This pool consists of billions of image-text pairs
 6 crawled from the public internet and was specifically designed and released to serve as the raw
 7 material for participants in the DataComp benchmark to test their filtering and curation strategies.
 8 *Id.*

9 291. The CommonPool is distributed in four distinct “scales” to accommodate participants
 10 with varying levels of computational resources: small (12.8 million examples), medium (128
 11 million examples), large (1.28 billion examples), and xLarge. The xLarge dataset is by far the most
 12 massive, containing a staggering 12.8 billion image-text pairs and requiring approximately 450
 13 terabytes of storage. *Id.*

14 292. Upon information and belief, Midjourney was not a participant in the DataComp
 15 competition, nor was it seeking to top the leaderboard in the noble pursuit of academic inquiry
 16 regarding dataset curation. Rather, Midjourney simply treated the competition’s “xLarge”
 17 CommonPool as a convenient, free-to-use reservoir of 12.8 billion training images to fuel its
 18 commercial products.

19 293. Based on information learned through discovery, in developing its generative AI
 20 models, Midjourney downloaded the xLarge CommonPool dataset hosted by DataComp, including
 21 infringing copies of the following registered copyrighted works by Plaintiffs (each previously
 22 asserted in Plaintiffs’ Second Amended Complaint):

Asserted Artwork, as Referenced in Second Amended Complaint	Asserted Artwork, as Referenced in Exhibits to Third Amended Complaint	Present in DataComp xLarge Dataset?
Adam Ellis: training image 1 (ECF No. 238-1 at 2).	TAC, Exh. E, pg. 1.	Yes.
Adam Ellis: training image 2 (ECF No. 238-1 at 3).	TAC, Exh. E, pg. 2.	Yes.

1	Gerald Brom: training image 1 (ECF No. 238-1 at 4).	TAC, Exh. E, pg. 3.	Yes.
2			
3	Gerald Brom: training image 2 (ECF No. 238-1 at 5).	TAC, Exh. E, pg. 4.	Yes.
4			
5	Gerald Brom: training image 3 (ECF No. 238-1 at 6).	TAC, Exh. E, pg. 5.	Yes.
6			
7	Gerald Brom: training image 4 (ECF No. 238-1 at 7).	TAC, Exh. E, pg. 6.	Yes.
8			
9	Gerald Brom: training image 5 (ECF No. 238-1 at 8).	TAC, Exh. E, pg. 7.	Yes.
10			
11	Gerald Brom: training image 7 (ECF No. 238-1 at 10).	TAC, Exh. E, pg. 8.	Yes.
12			
13	Gerald Brom: training image 9 (ECF No. 238-1 at 12).	TAC, Exh. E, pg. 9.	Yes.
14			
15	Gerald Brom: training image 10 (ECF No. 238-1 at 13).	TAC, Exh. E, pg. 10.	Yes.
16			
17	Gerald Brom: training image 11 (ECF No. 238-1 at 14).	TAC, Exh. E, pg. 11.	Yes.
18			
19	Gregory Manchess: training image 1 (ECF No. 238-1 at 17).	TAC, Exh. E, pg. 12.	Yes.
20			
21	Gregory Manchess: training image 2 (ECF No. 238-1 at 18).	TAC, Exh. E, pg. 13.	Yes.
22			
23	Gregory Manchess: training image 6 (ECF No. 238-1 at 23).	TAC, Exh. E, pg. 14.	Yes.
24			
25	Gregory Manchess: training image 7 (ECF No. 238-1 at 24).	TAC, Exh. E, pg. 15.	Yes.
26			
27	Grzegorz Rutkowski: training image 1 (ECF No. 238-1 at 25).	TAC, Exh. E, pg. 16.	Yes.
28			
	Grzegorz Rutkowski: training image 2 (ECF No. 238-1 at 26).	TAC, Exh. E, pg. 17.	Yes.
	Grzegorz Rutkowski: training image 7 (ECF No. 238-1 at 31).	TAC, Exh. E, pg. 18.	Yes.

1			
2	Grzegorz Rutkowski: training image 8 (ECF No. 238-1 at 32).	TAC, Exh. E, pg. 19.	Yes.
3	Jingna Zhang: training image 6 (ECF No. 238-1 at 45).	TAC, Exh. E, pg. 20.	Yes.
4	Jingna Zhang: training image 8 (ECF No. 238-1 at 47).	TAC, Exh. E, pg. 21.	Yes.
5	Jingna Zhang: training image 10 (ECF No. 238-1 at 49).	TAC, Exh. E, pg. 22.	Yes.
6	Jingna Zhang: training image 11 (ECF No. 238-1 at 50).	TAC, Exh. E, pg. 23.	Yes.
7	Jingna Zhang: training image 12 (ECF No. 238-1 at 51).	TAC, Exh. E, pg. 24.	Yes.
8	Jingna Zhang: training image 13 (ECF No. 238-1 at 52).	TAC, Exh. E, pg. 25.	Yes.
9	Jingna Zhang: training image 17 (ECF No. 238-1 at 56).	TAC, Exh. E, pg. 26.	Yes.
10	Jingna Zhang: training image 20 (ECF No. 238-1 at 59).	TAC, Exh. E, pg. 27.	Yes.
11	Jingna Zhang: training image 21 (ECF No. 238-1 at 60).	TAC, Exh. E, pg. 28.	Yes.
12	Jingna Zhang: training image 23 (ECF No. 238-1 at 62).	TAC, Exh. E, pg. 29.	Yes.
13	Julia Kaye: training image 1 (ECF No. 238-1 at 64).	TAC, Exh. E, pg. 30.	Yes.
14	Julia Kaye: training image 2 (ECF No. 238-1 at 65).	TAC, Exh. E, pg. 31.	Yes.
15	Julia Kaye: training image 3 (ECF No. 238-1 at 66).	TAC, Exh. E, pg. 32.	Yes.
16	Karla Ortiz: training image 1 (ECF No. 238-1 at 70).	TAC, Exh. E, pg. 33.	Yes.
17			
18			
19			
20			
21			
22			
23			
24			
25			
26			
27			
28			

1	Karla Ortiz: training image 2 (ECF No. 238-1 at 71).	TAC, Exh. E, pg. 34.	Yes.
2			
3	Karla Ortiz: training image 3 (ECF No. 238-1 at 72).	TAC, Exh. E, pg. 35.	Yes.
4			
5	Karla Ortiz: training image 5 (ECF No. 238-1 at 74).	TAC, Exh. E, pg. 36.	Yes.
6			
7	Karla Ortiz: training image 6 (ECF No. 238-1 at 75).	TAC, Exh. E, pg. 37.	Yes.
8			
9	Karla Ortiz: training image 8 (ECF No. 238-1 at 77).	TAC, Exh. E, pg. 38.	Yes.
10			
11	Karla Ortiz: training image 9 (ECF No. 238-1 at 78).	TAC, Exh. E, pg. 39.	Yes.
12			
13	Karla Ortiz: training image 11 (ECF No. 238-1 at 80).	TAC, Exh. E, pg. 40.	Yes.
14			
15	Karla Ortiz: training image 12 (ECF No. 238-1 at 81).	TAC, Exh. E, pg. 41.	Yes.
16			
17	Karla Ortiz: training image 13 (ECF No. 238-1 at 82).	TAC, Exh. E, pg. 42.	
18			
19	Karla Ortiz: training image 14 (ECF No. 238-1 at 83).	TAC, Exh. E, pg. 43.	Yes.
20			
21	Karla Ortiz: training image 15 (ECF No. 238-1 at 84).	TAC, Exh. E, pg. 44.	Yes.
22			
23	Sarah Andersen: training image 3 (ECF No. 238-1 at 96).	TAC, Exh. E, pg. 45.	Yes.
24			
25	Sarah Andersen: training image 6 (ECF No. 238-1 at 99).	TAC, Exh. E, pg. 46.	Yes.
26			
27	Sarah Andersen: training image 8 (ECF No. 238-1 at 101).	TAC, Exh. E, pg. 47.	Yes.
28			
	Sarah Andersen: training image 9 (ECF No. 238-1 at 102).	TAC, Exh. E, pg. 48.	Yes.
	Sarah Andersen: training image 12 (ECF No. 238-1 at 105).	TAC, Exh. E, pg. 49.	Yes.

1			
2	Sarah Andersen: training image 13 (ECF No. 238-1 at 106).	TAC, Exh. E, pg. 50.	Yes.
3	Sarah Andersen: training image 15 (ECF No. 238-1 at 108).	TAC, Exh. E, pg. 51.	Yes.
4			
5	Sarah Andersen: training image 16 (ECF No. 238-1 at 109).	TAC, Exh. E, pg. 52.	Yes.
6			
7	Sarah Andersen: training image 17 (ECF No. 238-1 at 110).	TAC, Exh. E, pg. 53.	Yes.
8			
9	Sarah Andersen: training image 18 (ECF No. 238-1 at 111).	TAC, Exh. E, pg. 54.	Yes.
10			
11	Sarah Andersen: training image 19 (ECF No. 238-1 at 112).	TAC, Exh. E, pg. 55.	Yes.
12			
13	Sarah Andersen: training image 20 (ECF No. 238-1 at 113).	TAC, Exh. E, pg. 56.	Yes.
14			
15	Sarah Andersen: training image 21 (ECF No. 238-1 at 114).	TAC, Exh. E, pg. 57.	Yes.
16			
17	Sarah Andersen: training image 22 (ECF No. 238-1 at 115).	TAC, Exh. E, pg. 58.	Yes.
18			
19	Sarah Andersen: training image 23 (ECF No. 238-1 at 116).	TAC, Exh. E, pg. 59.	Yes.
20			
21	Sarah Andersen: training image 25 (ECF No. 238-1 at 118).	TAC, Exh. E, pg. 60.	Yes.
22			
23	Sarah Andersen: training image 27 (ECF No. 238-1 at 120).	TAC, Exh. E, pg. 61.	Yes.
24			
25	Sarah Andersen: training image 29 (ECF No. 238-1 at 122).	TAC, Exh. E, pg. 62.	Yes.
26			
27	Jingna Zhang: training image 5 (ECF No. 238-2 at 6).	TAC, Exh. E, pg. 63.	Yes.
28	Jingna Zhang: training image 7 (ECF No. 238-2 at 8)	TAC, Exh. E, pg. 64.	Yes.

1 capabilities of Midjourney’s image generator to emulate and create work that is indistinguishable
2 from that of the artists whose names were published.

3 298. Midjourney’s use of the Midjourney Named Plaintiffs’ names was unauthorized and
4 without their consent.

5 299. Midjourney’s commercial speech created a likelihood of confusion over whether the
6 Midjourney Named Plaintiffs actually endorsed the Midjourney Image Product, and over the
7 affiliation, connection, or association that the Midjourney Named Plaintiffs might have with
8 Midjourney.

9 300. Midjourney’s commercial speech and use of the names of the Midjourney Named
10 Plaintiffs have deceived consumers as to their affiliation, connection, or association with
11 Midjourney.

12 301. A reasonably prudent consumer in the marketplace for art products likely would be
13 confused as to whether the Midjourney Named Plaintiffs included in the Midjourney Name List
14 sponsored or approved of Midjourney’s image generator.

15 302. The Midjourney Named Plaintiffs have a high level of recognition among
16 Midjourney’s users and consumers. In fact, Midjourney relies on this high level of recognition to
17 advertise the capabilities of its image product by publishing artists’ names. Midjourney strives to
18 capitalize off Midjourney Named Plaintiffs’ reputation as artists to induce users to use its image
19 generator.

20 303. The Midjourney Named Plaintiffs’ actual names were used by Midjourney.

21 304. Midjourney marketed its Midjourney Image Product on channels heavily trafficked
22 by its users and consumers such as on message boards. A link to the Midjourney Name List was
23 published on Discord by Midjourney CEO David Holz, which is frequented by Midjourney’s users
24 and potential consumers.

25 305. There is consumer appetite both for the Midjourney Named Plaintiffs’ art products as
26 well as potentially cheaper, or even free, imitations of such art.

- d. Grzegorz Rutkowski is known for lavish fantasy scenes rendered in a classical painting style.
- e. Julia Kaye is known for three-panel black-and-white comics, loosely inked with a thin fixed-width pen, wherein each individual comic is a micro-vignette in the artist's life.

312. Midjourney put the names of the Midjourney Named Plaintiffs on the Midjourney Name List because Midjourney makes use of a CLIP model that has been trained on the work of the Midjourney Named Plaintiffs. For artists like the Midjourney Named Plaintiffs and others on the Midjourney Name List, the CLIP model essentially acts as a trade-dress database.

313. The trade dress of each of the Midjourney Named Plaintiffs is inherently distinctive in look and feel as used in connection with their artwork and art products. On information and belief, a significant portion of consumers readily identify each of the Midjourney Named Plaintiffs' trade dress with the individual Midjourney Named Plaintiff.

314. On information and belief, Midjourney ensured that its CLIP model was trained to successfully and convincingly imitate the trade dress of the Midjourney Named Plaintiffs and the other artists on the Midjourney Name List. In other words, Midjourney ensured its CLIP model could appropriate the distinctive look and feel of each Midjourney Named Plaintiffs' trade dress.

315. As a result, the Midjourney Image Product can and frequently does generate images featuring protectable trade dress that are likely to cause confusion in consumers. The Midjourney Named Plaintiffs never authorized Midjourney to copy, emulate, or otherwise recreate their trade dress; nor did the Midjourney Named Plaintiffs authorize Midjourney to use, in conjunction with the advertisement and sale of its services, images featuring their trade dress.

316. Midjourney acknowledges and, in fact, relies on the inherent distinctiveness of the Midjourney Named Plaintiffs' respective trade dress to market its image generator by advertising that users can generate images in the style of particular artists simply by typing in their name. In this way, users do not have to describe specific design or artistic elements in the prompt to generate

1 an image in the artist’s style—they merely need to type in that artist’s name. Examples of
2 Midjourney text prompts featuring Sarah Andersen and Gerald Brom are shown in **Exhibit F**.

3 317. Midjourney vicariously infringes on the Midjourney Named Plaintiffs’ trade-dress
4 rights by encouraging and inducing the users of the Midjourney Image Product to enter artist-name
5 prompts and generate images featuring the Midjourney Named Plaintiffs’ protectable trade dress.
6 For example—

- 7 a. In its original online documentation offering “tips for text-prompts,”
8 Midjourney recommended that users should “try invoking unique artists to
9 get a unique style,” and offered a list that included “Greg Rutkowski,” who is a
10 one of the Midjourney Named Plaintiffs. The documentation also
11 recommended that users should “Combine names for new styles: ‘A temple
12 by Greg Rutkowski and Ross Tran.’” (Midjourney deleted these pages from
13 its public website two weeks after the initial complaint was filed in this
14 action.)
- 15 b. Midjourney currently promotes images made with artist-name prompts in an
16 online marketing gallery accessible to subscribers called “Showcase” (at
17 <https://midjourney.com/showcase>). **Exhibit K: Midjourney Showcase**
18 contains examples of images from the Showcase gallery featuring Plaintiff
19 names, including two of the Midjourney Named Plaintiffs: Sarah Andersen
20 and Gerald Brom.

21 318. Midjourney exercises control over the infringing images by including the CLIP
22 model in its image pipeline, and by marketing artist-name prompts as a key feature of its image
23 generator via the Midjourney Name List. Without the CLIP model, Midjourney’s users would not be
24 able to infringe on the Midjourney Named Plaintiffs’ trade-dress rights or those of the other artists
25 on the Midjourney Name List.

26 319. Each of the Midjourney Named Plaintiffs’ respective trade dress has no intrinsic
27 functional value. The unique combination of particular artistic elements does not confer any
28

1 utilitarian advantages on their art products and are purely ornamental and aesthetic. There remains
2 an unlimited number of alternative artistic styles available beyond the trade dress owned by the
3 Midjourney Named Plaintiffs.

4 320. Each of the Midjourney Named Plaintiffs' trade dress possesses secondary meaning
5 because the trade dress of their art products invoke a mental association by a substantial segment of
6 potential consumers between the trade dress and the creator of the art product.

7 321. Midjourney's vicarious infringement of the Midjourney Named Plaintiffs' trade-dress
8 rights are committed with actual and constructive knowledge of their trade dress, and with the intent
9 to cause confusion, mistake, or deception.

10 322. As a direct and proximate cause of Midjourney's conduct, the Midjourney Named
11 Plaintiffs have suffered, and will continue to suffer, significant damage in the form of loss of
12 revenue, income, profits, and goodwill, which will increase if not enjoined. Midjourney has, and
13 will unfairly, acquire revenue, income, profits, and goodwill at the expense of the Midjourney
14 Named Plaintiffs.

15 323. Midjourney's trade-dress infringement will also continue to cause irreparable harm if
16 Midjourney is not restrained by this Court from further violation of the rights of the Midjourney
17 Named Plaintiffs. The Midjourney Named Plaintiffs have no adequate remedy at law for the harm
18 being caused by Midjourney, particularly in regard to the loss of their goodwill and market share
19 due to Midjourney's infringing conduct. The Midjourney Named Plaintiffs are, therefore, entitled to
20 and seek temporary and permanent injunctive relief.

21 324. Midjourney has, and continues to, vicariously infringe on the trade-dress rights of
22 the Midjourney Named Plaintiffs in violation of section 43(a) of the Lanham Act, 15 U.S.C.
23 § 1125(a).

24 325. Midjourney's past and continuing infringement of the Midjourney Named Plaintiffs'
25 trade dress is an exceptional case and was willful and intentional, as evidenced by a) Midjourney's
26 intentional inclusion of the CLIP model in the design of the Midjourney Image Product and b) its
27 open advertisement of the Midjourney Image Product's ability to replicate an artist's trade dress via
28

1 the Midjourney Name List. Thus, the Midjourney Named Plaintiffs are entitled to treble their actual
2 damages and to an award of attorneys' fees under 15 U.S.C. § 1117(a), and all other available
3 remedies.

4 **XVII. CAUSES OF ACTION AGAINST RUNWAY**

5 326. Between April and October 2022, Runway trained an image model called Stable
6 Diffusion 1.5. According to Runway, Stable Diffusion 1.5 “was trained on a large-scale dataset
7 [called] LAION-5B” (See <https://huggingface.co/runwayml/stable-diffusion-v1-5#limitations>).

8 327. Stable Diffusion 1.5 is still sought out by many users of AI image products for,
9 among other things, its ability to mimic artists.

10 328. Because LAION-5B is an openly accessible dataset, Runway knew that the LAION-
11 5B dataset contained copyrighted works, including those of the LAION-5B Registered Plaintiffs
12 and Karla Ortiz.

13 329. The LAION-5B Registered Works are included in the LAION-5B dataset. Because
14 Runway admits to using the LAION-5B dataset for training, it must've also used the LAION-5B
15 Registered Works for training. Below, the term **Runway Models** refers to all models trained by
16 Runway on the LAION-5B Registered Works, including Stable Diffusion 1.5.

17 **COUNT TEN**

18 **Direct copyright infringement of the LAION-5B Registered Works by training the Runway** 19 **Models, including Stable Diffusion 1.5 against Runway on behalf of the LAION-5B Registered** 20 **Plaintiffs, LAION-5B Subclass, and Karla Ortiz Individually**

21 330. The preceding factual allegations are incorporated by reference.

22 331. The LAION-5B Registered Plaintiffs and Karla Ortiz never authorized Runway to
23 use their respective LAION-5B Registered Works in any way. Nevertheless, Runway repeatedly
24 violated the exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs and
25 Karla Ortiz and continues to do so today.

26 332. The LAION-5B dataset contains only URLs of training images, not the actual
27 training images. Therefore, anyone who wishes to use LAION-5B for training their own machine-
28

1 learning model must first acquire copies of the actual training images from their URLs by using the
2 ‘img2dataset’ tool or another similar tool. Consistent with this, in preparation for training the
3 Runway Models, Runway made one or more Statutory Copies of the LAION-5B Registered Works
4 so they could be fed to each Runway Model as training data. The Statutory Copies made of each
5 registered work were substantially similar to that registered work.

6 333. During the training of each Runway Model, Runway made a series of intermediate
7 Statutory Copies of the LAION-5B Registered Works. For instance, diffusion models are trained by
8 creating “noised” copies of training images, as described herein, all of which qualify as Statutory
9 Copies. The intermediate Statutory Copies of each registered work that Runway made during
10 training of the Runway Models were substantially similar to that registered work.

11 334. By the end of training, Stable Diffusion 1.5 was capable of reproducing protected
12 expression from each of the LAION-5B Registered Works that was in each case substantially
13 similar to that registered work, as shown in **Exhibit E: Runway text prompts** and **Exhibit H:**
14 **Runway image prompts**. Therefore, Stable Diffusion 1.5 qualifies as an infringing Statutory Copy
15 of the LAION-5B Registered Works. Because Stable Diffusion 1.5 represents a transformation of
16 the LAION-5B Registered Works into an alternative form, Stable Diffusion 1.5 also qualifies as an
17 infringing Statutory Derivative Work.

18 335. On information and belief, the other Runway Models exhibit the same properties,
19 because they were trained on the same LAION-5B dataset.

20 336. Since October 2022, Runway has distributed Stable Diffusion 1.5 to the public, for
21 instance via websites like GitHub (*See* <https://github.com/runwayml/stable-diffusion>) and Hugging
22 Face (*See* <https://huggingface.co/runwayml/stable-diffusion-v1-5>). In so doing, Runway infringed
23 the exclusive distribution rights of the LAION-5B Registered Plaintiffs and Karla Ortiz.

24 337. The LAION-5B Registered Plaintiffs and Karla Ortiz have been and continue to be
25 injured by Runway’s multiple acts of direct copyright infringement. These Plaintiffs are entitled to
26 statutory damages, actual damages, restitution of profits, and other remedies provided by law.
27
28

COUNT ELEVEN

Inducement of copyright infringement by distributing Stable Diffusion 1.5 for free against Runway on behalf of the LAION-5B Registered Plaintiffs and Subclass

338. The preceding factual allegations are incorporated by reference.

339. Runway distributes Stable Diffusion 1.5 under the “CreativeML Open RAIL-M” license, which allows anyone to download, use, and deploy Stable Diffusion 1.5 for free. For instance, via websites like GitHub (*See* <https://github.com/runwayml/stable-diffusion>) and Hugging Face (*See* <https://huggingface.co/runwayml/stable-diffusion-v1-5>).

340. Stable Diffusion 1.5 violates the exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs. Therefore, anyone who in fact downloads, uses, or deploys Stable Diffusion 1.5 is engaged in infringing activity.

341. Runway has made a material contribution to this infringing activity by training Stable Diffusion 1.5 and then distributing it for free.

342. Runway intends to cause further infringement with Stable Diffusion 1.5. In February 2023, Stability CEO Mostaque said that Stable Diffusion 1.5 was “the most popular model by far by [a] for profit company.”⁵⁵

343. The LAION-5B Registered Plaintiffs have been and continue to be injured by Runway’s inducement of copyright infringement. These Plaintiffs are entitled to statutory damages, actual damages, restitution of profits, and other remedies provided by law.

XVIII. CAUSES OF ACTION AGAINST DEVIANTART

344. Since its founding in 2000, DeviantArt has held itself out as an online community friendly to artists, colloquially known on the site as “deviants.” A primary activity of artists on DeviantArt is sharing digital images of their artwork, colloquially called “deviations.” Today, DeviantArt bills itself as “the world’s largest art community,” hosting millions of such images.

⁵⁵ <https://twitter.com/EMostaque/status/1629514395825983489>

1 345. Plaintiffs Kelly McKernan, Hawke Southworth, Jingna Zhang, and Grzegorz
2 Rutkowski are DeviantArt users. Below, they are called the **DeviantArt Plaintiffs**.

3 346. On November 9, 2022, DeviantArt released DreamUp, an AI image product.
4 DeviantArt claims that DreamUp “lets you create AI art knowing that creators and their work are
5 treated fairly.” DreamUp is only available to paying customers of DeviantArt. DeviantArt offers
6 paid subscriptions to its members called “Core Plans.” Custom Core Plans typically range in price
7 from \$3.95 to \$14.95 per month. To use DreamUp, a member must first subscribe to a Core Plan. A
8 Core Plan subscriber is allowed to use DreamUp for a certain number of Text Prompts per month.
9 For instance, the \$9.95 “Pro” level permits 200 DreamUp Text Prompts per month. Core Plan
10 members can purchase additional Text Prompts by purchasing packages of “points.” DeviantArt
11 charges \$1 for 80 points, with a minimum purchase of 400 points for \$5.

12 347. Because DeviantArt holds itself out as an art community, DeviantArt chooses to
13 provide many features that artists may prefer. For example, given the ubiquity of affixing CMI such
14 as distinctive marks onto deviations that are being uploaded on to DeviantArt.

15 348. DeviantArt is the source of millions of images in the LAION-5B dataset. Users of
16 the LAION-5B dataset have copied these millions of images many times over by downloading them
17 from DeviantArt.

18 349. On information and belief, DeviantArt was aware that LAION-5B contained
19 references to millions of DeviantArt images, and that Stability downloaded these millions of images
20 from the DeviantArt website as a necessary preliminary step in the training of the Stability Models.

21 350. Each of the DeviantArt Plaintiffs has stored images on DeviantArt that were later
22 incorporated into the LAION-5B dataset. These images were therefore copied by Runway and
23 Stability to train versions of Stable Diffusion. **Exhibit A** contains a sampling—but not an
24 exhaustive listing—of images created by the DeviantArt Plaintiffs that are contained in LAION-5B
25 and were copied from DeviantArt. They can be identified through their LAION-5B URL, which
26 comes from the “wixmp.com” domain. This domain is used by DeviantArt to store member images.
27 (“Wix” in the domain name refers to the parent company of DeviantArt.)
28

1 351. DreamUp relies on Stable Diffusion to produce images. The DreamUp app
2 incorporates a copy of Stable Diffusion. The terms of service for DreamUp do not disclose the
3 specific version of Stable Diffusion that is incorporated within the app.

4 352. But the DreamUp terms require users to also accept the terms of the CreativeML
5 Open RAIL-M License linked at <https://huggingface.co/spaces/CompVis/stable-diffusion-license>.
6 Because this URL refers to “CompVis” and the license itself is dated August 22, 2022, DreamUp
7 must be based on Stable Diffusion version 1.4, which was trained by CompVis and released on
8 August 22, 2022. Below, the model inside DreamUp will be called the **DreamUp–CompVis**
9 **Model**.

10 353. CompVis is the shorthand name of the Computer Vision and Learning Group at
11 Ludwig Maximilian University in Munich, where the original research underlying Stable Diffusion
12 was first conducted. According to the GitHub page for Stable Diffusion 1.4, “Stable Diffusion was
13 made possible thanks to a collaboration with Stability AI and Runway.”⁵⁶

14 354. According to CompVis, Stable Diffusion 1.4 “was trained on a large-scale dataset
15 LAION-5B.”⁵⁷

16 355. The LAION-5B dataset contains only URLs of training images, not the actual
17 training images. Therefore, anyone who wishes to use LAION-5B for training their own machine-
18 learning model must first acquire copies of the actual training images from their URLs by using the
19 ‘img2dataset’ tool or another similar tool. Consistent with this, in preparation for training Stable
20 Diffusion 1.4, CompVis made one or more Statutory Copies of the LAION-5B Registered Works so
21 they could be fed to Stable Diffusion 1.4 as training data. The Statutory Copies made of each
22 registered work were substantially similar to that registered work.

23 356. During the training of Stable Diffusion 1.4, CompVis made a series of intermediate
24 Statutory Copies of the LAION-5B Registered Works. For instance, diffusion models are trained by
25 creating “noised” copies of training images, as described herein, all of which qualify as Statutory
26

27 ⁵⁶ See <https://github.com/CompVis/stable-diffusion>

28 ⁵⁷ See <https://huggingface.co/CompVis/stable-diffusion-v1-4>

1 Copies. The intermediate Statutory Copies of each registered work that CompVis made during
2 training of Stable Diffusion 1.4 were substantially similar to that registered work.

3 357. On information and belief, by the end of training, Stable Diffusion 1.4 was capable
4 of reproducing protected expression from each of the LAION-5B Registered Works that was in each
5 case substantially similar to that registered work, because—

- 6 a. In the Carlini Paper, Nicholas Carlini tested Stable Diffusion 1.4 and found
7 that it could emit stored copies of its training images;
- 8 b. The training procedure for Stable Diffusion 1.4 was very similar to that of
9 Stable Diffusion 1.5, which was shown in **Exhibit E: Runway text prompts**
10 and **Exhibit H: Runway image prompts** to be capable of emitting stored
11 copies of protected expression.

12 358. Therefore, like Stable Diffusion 1.5, Stable Diffusion 1.4 also qualifies as an
13 infringing Statutory Copy of the LAION-5B Registered Works. Because Stable Diffusion 1.4
14 represents a transformation of the LAION-5B Registered Works into an alternative form, Stable
15 Diffusion 1.4 also qualifies as an infringing Statutory Derivative Work.

16 359. DeviantArt continues to obfuscate the source of DreamUp’s training data. One of the
17 questions in DeviantArt’s frequently asked questions (“FAQ”) section for DreamUp on its website
18 is “Does DreamUp use art submitted on the DeviantArt platform to train the AI models.”

19 DeviantArt responds that:

20 DreamUp is based on 3rd-party technologies (like Stable
21 Diffusion) which train their models based on the open web.
22 DreamUp uses semantic interpretation of a textual prompt and then
23 translates it to input for these models.

24 DeviantArt does NOT add images from DeviantArt to the training
25 sets of 3rd-party technologies, and DeviantArt does NOT provide
26 data to expand distribution of images that 3rd-party technologies
27 can generate.

28 DeviantArt lets you declare whether or not external AI models and
platforms can train based on your deviations. When submitting a
deviation, you’ll be able to check a box that informs third parties

1 whether or not you authorize that submission being included in
2 datasets used to train AI models like AI image generators.⁵⁸

3 360. DeviantArt’s answer is misleading. As confirmed by the FAQ, while DeviantArt did
4 not “add” images to the training sets of DreamUp, it made no mention of any images already in the
5 training set for DreamUp’s underlying models. DeviantArt knew that Stable Diffusion had already
6 been trained on images scraped from DeviantArt itself. DeviantArt thus misled its community
7 because art from DeviantArt was already in DreamUp because Stable Diffusion had already been
8 trained on them.

9 361. This has been further confirmed by DeviantArt CTO Chris Nell. In November 2022,
10 on the public LAION Discord server, Nell described himself as “one of the people at DeviantArt
11 working on improving acceptance of AI generated/augmented art in the broader online arts
12 community” and added “I think our goals at DA [DeviantArt] are very aligned with LAION’s
13 ... and want to collaborate as much as possible.”⁵⁹ Nell said of DreamUp: “we did not fine tune
14 [meaning, perform additional training on] SD [= Stable Diffusion] at all, so there aren’t novel
15 weights to share. [W]e do perform additional guidance at generation time ... so it’s not exactly
16 unmodified SD [= Stable Diffusion] output, but that is more akin to prompt tuning.”⁶⁰ As confirmed
17 by Nell, DeviantArt was well aware of how Stable Diffusion was developed and did not do any fine-
18 tuning of the weights included in the Stable Diffusion model DreamUp was based on. In other
19 words, this implies all of the images copied in training Stable Diffusion were included in the
20 DreamUp model.

21
22
23
24 ⁵⁸ <https://www.deviantartsupport.com/en/dreamup>

25 ⁵⁹

26 <https://discord.com/channels/823813159592001537/1006139459860975716/104253965639641100>

27 4

28 ⁶⁰

<https://discord.com/channels/823813159592001537/1006139459860975716/104254383742543880>
4

1 362. DeviantArt is also aware that DreamUp can be used by DreamUp’s users and
2 licensees to create potentially infringing works based on artists’ underlying work. This is evidenced
3 by another provision of DeviantArt’s DreamUp FAQ which provides:

4 DreamUp is an AI-based image-generation tool used to create art
5 using free-form text prompts. **Certain art styles can sometimes be**
6 **achieved by referencing names of real artists** such as Thomas
7 **Kinkade, Picasso, and Gustave Doré in text prompts. Referencing**
8 **artists when having the AI create your work can give the**
9 **resulting piece a unique “look,” inspired by the style of that**
10 **particular artist.**

11 If you refer to an artist in a DreamUp prompt, you must also tag that
12 artist when submitting the resulting image to DeviantArt. Failure to
13 do so is a violation of our DreamUp Policy and can result in your
14 deviation’s deletion or an account suspension.

15 363. Again, DeviantArt’s FAQ misleads by omission. DeviantArt tellingly is only
16 concerned with images posted on DeviantArt itself, even though the infringing art would have been
17 created with DeviantArt’s product. Further, because DeviantArt knew Stable Diffusion contained
18 copies of training images (including those scraped from DeviantArt), and thus, so did DreamUp, it
19 knew that there was a real possibility that DreamUp could regenerate images in the training set,
20 requiring it to include a provision in its FAQ addressing the possibility. Furthermore, even with the
21 risk that DreamUp could generate images based on protected images, whenever a user uses
22 DreamUp, it asks users to resubmit their generated outputs to use as image prompts with other text
23 in order to generate more images.

24 364. DeviantArt’s embrace of generative AI art was seen as a betrayal by its art
25 community.

26 365. The scope of DeviantArt’s betrayal of its artist community by embracing Stable
27 Diffusion was evident in a group audio session held by DeviantArt management on November 11,
28 2022 from approximately 1:00–2:30 pm Pacific Time. DeviantArt scheduled the discussion
specifically to allay the well-founded concerns of DeviantArt members that DeviantArt’s embrace
of AI art was a complete repudiation of its longstanding community principles, as well as
economically and legally unfair.

1 366. At one point in the audio session, CEO Moti Levy explicitly took ownership of the
2 decision to bring Stable Diffusion (the basis of the DreamUp–CompVis Model) onto DeviantArt via
3 the DreamUp app: “The reason why we’re using Stable Diffusion because it’s the only option for us
4 to take an open source [software engine] and modify it The other platforms or the other
5 companies do not allow it. . . . [A]nd by the way, that was my decision. **That’s our decision by me
6 as the CEO. That’s my decision to take Stable Diffusion.**” (emphasis added.)

7 367. Levy also said, “DeviantArt expects all users accessing our service or the DeviantArt
8 site to respect creators’ choices about the acceptable use of their content, including for AI purposes.
9 When a DeviantArt user doesn’t consent to third party use of their content for AI purposes, other
10 users of the service and third parties accessing the DeviantArt site are prohibited from using such
11 content to train an AI system, as input into any previously trained AI system or to make available
12 any derivative copy unless usage of that copy is subject to conditions at least as restrictive as those
13 set out in the DeviantArt terms of service.”

14 368. Shortly after the end of this audio session, DeviantArt updated its terms of service.
15 DeviantArt added a new paragraph about “Data Scraping & Machine Learning Activities” that
16 explicitly *permits* this kind of usage under certain circumstances, so that Stable Diffusion and future
17 generative AI services can continue to scrape DeviantArt for images. In so doing, DeviantArt has
18 reneged on its promises. It plainly switched its loyalties from its artist members to the AI
19 companies, like Stability, infringing Plaintiffs’ and the Class’s intellectual property rights in the
20 work of those members. (According to the Internet Archive, this new data-scraping provision was
21 added to the DeviantArt terms of service on November 11, 2022, sometime between 1:41pm and
22 4:22pm Pacific Time.)

23 369. Furthermore, although the new “Data Scraping” provision acknowledges that certain
24 kinds of data scraping will continue to be an “unauthorized use” of the DeviantArt website, that
25 “owners of the works are responsible for policing their own works.” In other words, despite its
26 professed interest in using its terms of service to protect artists, DeviantArt is washing its hands of
27 the matter. Instead of standing up for artists and using its resources to combat illegal AI data
28

1 scraping, it is forcing artists to take matters into their own hands. What is more, while DeviantArt
2 purported to spearhead a system for artists to opt-out of having their works trained upon, these
3 promises are mostly hollow.

4 370. DeviantArt’s proposal for artists to opt out was to utilize a system of HTML tags.
5 Artists who do not wish to have their content used for AI training can append the “noai” and
6 “noaimageai” hashtags to the HTML page associated with their art.

7 371. This promise is misleading.

8 372. Even if an artist indicates they do not want their artwork used by affixing the “noai”
9 and “noimageai” directives to their HTML pages, it does not apply retroactively to AI image
10 products that have already been trained on their works, such as all the models at issue in this
11 Complaint.

12 373. Further, even if an artist appends “noai” or “noimageai” directives, however, that is
13 still not a guarantee that their work will not be used to train AI models. As indicated in DeviantArt’s
14 own TOS, “DeviantArt provides no guarantees that ‘noai’ or ‘noimageai’ directives will be present
15 each time Content is accessed, even if the creator does not consent to use of that Content for
16 Artificial Intelligence Purposes; and absence of such directives does not imply creator consent has
17 been granted. . . Users acknowledge that by uploading Content to DeviantArt, third parties may
18 scrape or otherwise use their works without permission. DeviantArt provides no guarantees that
19 third parties will not include certain Content in external data sources, or otherwise use a creator's
20 work for Artificial Intelligence Purposes, even when such directives are present. By prohibiting
21 such conduct, DeviantArt makes no guarantees that it will pursue each unauthorized use of the
22 Service, and the owners of the works are responsible for policing their own works to the extent
23 permitted by law.”

24 COUNT TWELVE

25 **Direct copyright infringement by copying the DreamUp–CompVis Model and incorporating it** 26 **into DreamUp against DeviantArt on behalf of the LAION-5B Registered Plaintiffs**

27 374. The preceding factual allegations are incorporated by reference.
28

1 375. Because Stable Diffusion 1.4 is an infringing Statutory Copy of the LAION-5B
2 Registered Works, the DreamUp–CompVis Model is too.

3 376. Because Stable Diffusion 1.4 is an infringing Statutory Derivative Work based on the
4 LAION-5B Registered Works, the DreamUp–CompVis Model is too.

5 377. The DreamUp–CompVis Model infringes the exclusive rights (under 17 U.S.C.
6 § 106) of the LAION-5B Registered Plaintiffs.

7 378. Because the DreamUp app contains a copy of DreamUp–CompVis Model, the
8 DreamUp app infringes copyrights owned by the LAION-5B Registered Plaintiffs.

9 379. The LAION-5B Registered Plaintiffs have been and continue to be injured by
10 DeviantArt’s multiple acts of direct copyright infringement. These Plaintiffs are entitled to statutory
11 damages, actual damages, restitution of profits, and other remedies provided by law.

12 COUNT THIRTEEN

13 **Direct copyright infringement of the DeviantArt DreamUp Registered Works** 14 **by training the DeviantArt Clip Model against DeviantArt on behalf of Jingna Zhang,** 15 **individually, the DreamUp Registered Plaintiffs, and the DeviantArt DreamUp Damages** 16 **Subclass**

16 380. The preceding factual allegations are incorporated by reference.

17 381. Based on information learned through discovery, in developing its DreamUp
18 product, DeviantArt trained its CLIP Model on millions of artworks uploaded by DeviantArt users
19 to the DeviantArt website, including infringing copies of the following registered copyrighted
20 works by Jingna Zhang (each previously asserted in Plaintiffs’ Second Amended Complaint):

21 Asserted Artwork, as 22 Referenced in Second 23 Amended Complaint	24 Asserted Artwork, as Referenced 25 in Exhibits to Third Amended 26 Complaint	27 Present in DreamUp Data?
28 ECF No. 238-1 at 45.	TAC, Exh. F, pg. 1.	Yes.
ECF No. 238-1 at 48.	TAC, Exh. F, pg. 2.	Yes.
ECF No. 238-1 at 52.	TAC, Exh. F, pg. 3.	Yes.
ECF No. 238-1 at 53.	TAC, Exh. F, pg. 4.	Yes.
ECF No. 238-1 at 56.	TAC, Exh. F, pg. 5.	Yes.

1 Dated: February 27, 2026

By: /s/ Joseph R. Saveri
Joseph R. Saveri

2 Joseph R. Saveri (SBN 130064)
3 Cadio Zirpoli (SBN 179108)
4 Christopher K.L. Young (SBN 318371)
5 Elissa A. Buchanan (SBN 249996)
6 Evan Creutz (SBN 349728)
7 Alexandra Fernandez (SBN 330518)
8 Holden Benon (SBN 325847)
9 Aaron Cera (SBN 351163)
10 Alexander Zeng (SBN 360220)
11 **JOSEPH SAVERI LAW FIRM, LLP**
12 601 California Street, Suite 1505
13 San Francisco, CA 94108
14 Telephone: (415) 500-6800
15 Facsimile: (415) 395-9940
16 Email: jsaveri@saverilawfirm.com
17 Email: czirpoli@saverilawfirm.com
18 Email: cyoung@saverilawfirm.com
19 Email: ebuchanan@saverilawfirm.com
20 Email: ecreutz@saverilawfirm.com
21 Email: afernandez@saverilawfirm.com
22 Email: hbenon@saverilawfirm.com
23 Email: acera@saverilawfirm.com
24 Email: azeng@saverilawfirm.com

25 Matthew Butterick (State Bar No. 250953)
26 1920 Hillhurst Avenue, #406
27 Los Angeles, CA 90027
28 Telephone: (323) 968-2632
Facsimile: (415) 395-9940
Email: mb@buttericklaw.com