

1 Joseph R. Saveri (State Bar No. 130064)
2 Cadio Zirpoli (State Bar No. 179108)
3 Christopher K.L. Young (State Bar No. 318371)
4 Elissa A. Buchanan (State Bar No. 249996)
5 David Lerch (State Bar No. 229411)
6 Evan Creutz (State Bar No. 349728)

JOSEPH SAVERI LAW FIRM, LLP

601 California Street, Suite 1505
San Francisco, California 94108
Telephone: (415) 500-6800
Facsimile: (415) 395-9940
Email: jsaveri@saverilawfirm.com
czirpoli@saverilawfirm.com
cyoung@saverilawfirm.com
eabuchanan@saverilawfirm.com
dlerch@saverilawfirm.com
ecreutz@saverilawfirm.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**

17 Sarah Andersen, an individual; Kelly McKernan, an
18 individual; Karla Ortiz, an individual;
19 H. Southworth pka Hawke Southworth, an
20 individual; Grzegorz Rutkowski, an individual;
21 Gregory Manchess, an individual; Gerald Brom, an
22 individual; Jingna Zhang, an individual; Julia Kaye,
23 an individual; Adam Ellis, an individual;

Individual and Representative Plaintiffs,

v.

24 Stability AI Ltd., a UK corporation;
25 Stability AI, Inc., a Delaware corporation;
26 DeviantArt, Inc., a Delaware corporation;
27 Midjourney, Inc., a Delaware corporation;
28 Runway AI, Inc., a Delaware corporation;

Defendants.

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

SECOND AMENDED COMPLAINT

CLASS ACTION

DEMAND FOR JURY TRIAL

TABLE OF CONTENTS

Page(s)

1

2

3 I. AI IMAGE PRODUCTS ARE TRAINED ON VAST NUMBERS OF COPYRIGHTED

4 IMAGES WITHOUT CONSENT, CREDIT, OR COMPENSATION AND VIOLATE

5 THE RIGHTS OF MILLIONS OF ARTISTS 1

6 II. JURISDICTION AND VENUE 4

7 III. PLAINTIFFS..... 5

8 IV. DEFENDANTS 6

9 V. AGENTS AND CO-CONSPIRATORS 7

10 VI. CLASS ALLEGATIONS..... 8

11 A. CLASS DEFINITIONS 8

12 B. NUMEROSITY..... 9

13 C. TYPICALITY..... 9

14 D. COMMONALITY & PREDOMINANCE 10

15 E. ADEQUACY 10

16 F. OTHER CLASS CONSIDERATIONS 11

17 VII. ARTISTS AND THEIR WORKS..... 11

18 VIII. THE SOURCE OF THE TRAINING DATASETS: LAION 13

19 IX. HOW AI IMAGE PRODUCTS WORK: CLIP-GUIDED DIFFUSION 20

20 X. PROTECTED EXPRESSION FROM TRAINING IMAGES IS COPIED,

21 COMPRESSED, STORED, AND INTERPOLATED BY DIFFUSION MODELS 27

22 XI. EXAMPLES OF TEXT PROMPTS USING PLAINTIFF NAMES IN AI IMAGE

23 PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY 35

24 XII. EXAMPLES OF IMAGE PROMPTS USING PLAINTIFF IMAGES IN AI IMAGE

25 PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY 41

26 XIII. USER AND LICENSEE ACTIVITY 49

27

28

1 XIV. DEFINITIONS FOR THE CAUSES OF ACTION 51
2 XV. CAUSES OF ACTION AGAINST STABILITY 51
3 XVI. CAUSES OF ACTION AGAINST MIDJOURNEY 59
4 XVII. CAUSES OF ACTION AGAINST RUNWAY 72
5 XVIII. CAUSES OF ACTION AGAINST DEVIANTART 77
6 XIX. JURY TRIAL DEMANDED 84
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); Digital Millennium Copyright Act (17 U.S.C. §§ 1202(a) & (b)) as to all Defendants; and
8 violations of the Lanham Act (15 U.S.C. § 1125(a)(1)) as to Defendant Midjourney

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

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

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

1 The outputs of a model are derived entirely and exclusively from what it has extracted from the
2 dataset.

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

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

21 5. Defendant Midjourney has trained multiple models of its own and deployed them
22 commercially as part of its AI image product. Midjourney has downloaded copies of billions of
23 copyrighted images without permission to train its models—including those belonging to Plaintiffs.
24
25

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

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

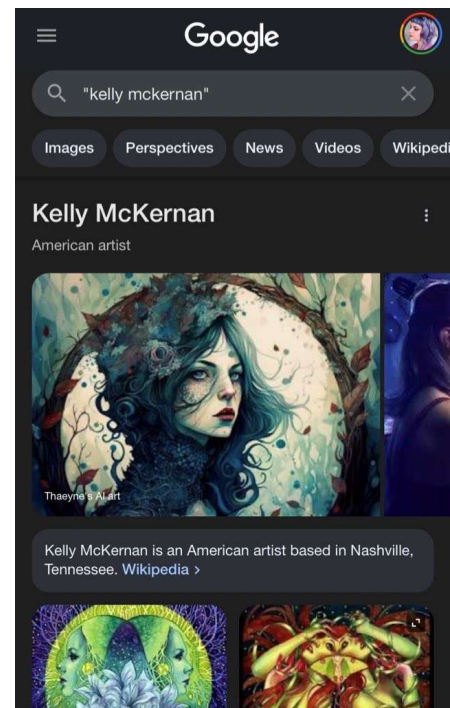
1 It has also copied the Stable Diffusion model and deployed it commercially as part of its AI image
2 product. It sells subscription access to its AI image product.

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

7 7. 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 8. 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://thayne.com/2022/12/11/image-in-the-style-of-kelly-mckernan-1/>

1 9. Though Defendants like to describe their AI image products in lofty terms, the reality
2 is grubbier and nastier: AI image products are primarily valued as copyright-laundering devices,
3 promising customers the benefits of art without the costs of artists.

4 10. Plaintiffs seek to end this unprecedented violation of their legal rights before their
5 jobs, their professions, and their creative communities are demolished by Defendants.

6 II. JURISDICTION AND VENUE

7 11. This Court has subject-matter jurisdiction under 28 U.S.C. § 1331 because this case
8 arises under the Copyright Act (17 U.S.C. § 501) and the Digital Millennium Copyright Act (17
9 U.S.C. § 1202).⁴

10 12. Jurisdiction and venue are proper in this judicial District under 27 U.S.C.
11 § 1391(c)(2) because Defendant Midjourney is headquartered in this District, and thus, a substantial
12 part of the events giving rise to Plaintiffs’ claims occurred in this District. Each Defendant has
13 transacted business, maintained substantial contacts, or committed overt acts in furtherance of the
14 illegal scheme and conspiracy throughout the United States, including in this District. Defendants’
15 conduct has had the intended and foreseeable effect of causing injury to persons residing in, located
16 in, or doing business throughout the United States, including in this District.

17 13. Under Civil Local Rule 3.2(c) and (e), assignment of this case to the San Francisco
18 Division is proper because Defendant Midjourney is headquartered in San Francisco, and thus, a
19 substantial part of the events giving rise to Plaintiffs’ claims and the interstate trade and commerce
20 involved and affected by Defendants’ illegal conduct occurred in this Division.

21
22
23 ⁴ In the order on the motion to dismiss, the Court stated that “plaintiffs may file a Second Amended
24 Complaint including the new plaintiffs [previously included in the First Amended Complaint] and
25 may attempt to plead unjust enrichment claims against any defendant based on theories (if any) that
26 are not preempted by the Copyright Act” (ECF 223 p.7 n. 6). In addition, as to the DMCA claims
27 which were included in the First Amended Complaint, the Ninth Circuit does not require that a
28 party replead claims to preserve them on appeal. *See Lacey v. Maricopa Cnty.*, 693 F.3d 896, 928 (9th
Cir. 2012) (en banc) (“For claims dismissed with prejudice and without leave to amend, we will not
require that they be replead in a subsequent amended complaint to preserve them for appeal.”). In
addition, the facts as alleged in this Second Amended Complaint are current as of the filing of the
First Amended Complaint.

III. PLAINTIFFS

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

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

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

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

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

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

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

21. Jingna Zhang lives in the State of Washington and owns the copyrights in training images shown under her name in **Exhibit A** and **Exhibit B**. Ms. Zhang's copyright registrations for

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

1 these works are included in **Exhibit C**. Ms. Zhang is a member of DeviantArt under the name
2 ‘Zemotion’, which is found in the text captions of some of her works.

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

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

9 24. The images shown in **Exhibit A** and **Exhibit B** are offered as a representative sample
10 of works by Plaintiffs that appear in the LAION datasets, not an exhaustive or complete list.
11 Plaintiffs confirmed that these particular images were in the LAION-5B and LAION-400M datasets
12 respectively by searching for their own names on two websites that allow searching of the LAION
13 datasets: <https://haveibeenentrained.com> and <https://rom1504.github.io/clip-retrieval/>. On
14 information and belief, all of Plaintiffs’ works that were registered as part of the collections in
15 **Exhibit C** and were online and were scraped into one or both of these datasets.

16 25. The registrations shown in **Exhibit C** are only a partial list of registered copyrights
17 owned by Plaintiffs.

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

23 IV. DEFENDANTS

24 27. Defendant Stability AI Ltd. is a UK corporation with its principal place of business at
25 88 Notting Hill Gate, London, England, W11 3HP. Stability was founded in 2020 by Mohammad
26 Emad Mostaque, a former hedge-fund manager. Mostaque is currently the Chief Executive Officer
27 of Stability AI. Stability AI also employs Robin Rombach, formerly a member of the CompVis
28

1 research group at Ludwig Maximilian University in Munich, where he was a principal developer of
2 the technology underlying Stable Diffusion.

3 28. Defendant Stability AI, Inc. is a Delaware corporation with its principal place of
4 business at 88 Notting Hill Gate, London, England, W11 3HP. Stability AI Ltd. is a wholly owned
5 subsidiary of Stability AI, Inc.

6 29. Defendant Midjourney, Inc. is a Delaware corporation with its principal place of
7 business at 333 Harrison Street, Apt. 605, San Francisco CA 94105. Midjourney was founded in San
8 Francisco in August 2021 by David Holz, who also serves as CEO.

9 30. Defendant DeviantArt, Inc. is a Delaware corporation with its principal place of
10 business at 100 Gansevoort Street, New York NY 10014. DeviantArt was founded in 2000 by Angelo
11 Sotira, Scott Jarkoff, and Matthew Stephens. In 2017, Wix.com, Inc. acquired DeviantArt. Wix
12 acquired all of DeviantArt's corporate stock for \$36 million. In April 2022, Moti Levy became CEO
13 of DeviantArt.

14 31. Defendant Runway AI, Inc. is a Delaware corporation with its principal place of
15 business at 79 Walker Street, Floor 5, New York NY 10013. Runway was founded in New York in
16 2018 by Anastasis Germanidis, Alejandro Matamala-Ortiz and Cristóbal Valenzuela. Valenzuela is
17 currently the CEO of Runway. Runway also employs Patrick Esser, formerly a member of the
18 CompVis research group at Ludwig Maximilian University in Munich, where he was a principal
19 developer of the technology underlying Stable Diffusion.

20 **V. AGENTS AND CO-CONSPIRATORS**

21 32. The unlawful acts alleged against the Defendants in this class-action complaint were
22 authorized, ordered, or performed by the Defendants' respective officers, agents, employees,
23 representatives, or shareholders while actively engaged in the management, direction, or control of
24 the Defendants' businesses or affairs. The Defendants' agents operated under the explicit and
25 apparent authority of their principals. Each Defendant, and its subsidiaries, affiliates, and agents
26 operated as a single unified entity.

27 33. Various persons or firms not named as Defendants may have participated as co-
28

1 conspirators in the violations alleged herein and may have performed acts and made statements in
 2 furtherance thereof. Each acted as the principal, agent, or joint venture of, or for other Defendants
 3 with respect to the acts, violations, and common course of conduct alleged herein.

4 VI. CLASS ALLEGATIONS

5 A. Class Definitions

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

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

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

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

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

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

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

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

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

27 All persons or entities nationalized or domiciled in the United States
 28 that own a registered copyright in any work in the LAION-400M
 dataset that was used to train any version of an AI image product that

1 was offered directly or incorporated into another product by one or
2 more Defendants during the Class Period.

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

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

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

- 8 a. Any of the Defendants named herein;
- 9 b. Any of the Defendants’ co-conspirators;
- 10 c. Any of Defendants’ parent companies, subsidiaries, and affiliates;
- 11 d. Any of Defendants’ officers, directors, management, employees, subsidiaries,
12 affiliates, or agents;
- 13 e. All governmental entities; and
- 14 f. The judges and chambers staff in this case, as well as any members of their
15 immediate families.

16 **B. Numerosity**

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

21 **C. Typicality**

22 37. Plaintiffs’ claims are typical of the claims of their fellow Class members because
23 Plaintiffs’ claims arise out of the same course of conduct from which their injuries result. Plaintiffs
24 and all Class members own copyrights in the Works. Plaintiffs and the Class created or owned Works
25 that were published on the internet by themselves or others. The Works were used to train various
26 AI Image Products without permission. Plaintiffs and absent Class members were damaged by this
27 and other wrongful conduct of Defendants as alleged herein. Damages and the other relief sought
28 herein are common to all members of the Class.

1 **D. Commonality & Predominance**

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

4 39. **Direct Copyright Infringement:** Whether Defendants violated the copyrights of
5 Plaintiffs and the Class when they downloaded and stored copies of the Works; Whether Defendants
6 violated the copyrights of Plaintiffs and the Class when they used copies of the Works to train AI
7 Image Products.

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

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

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

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

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

21 **E. Adequacy**

22 45. Plaintiffs will fairly and adequately represent the interests of the Class because they
23 have experienced the same harms as the Class and have no conflicts with any other members of the
24 Class. Furthermore, Plaintiffs have retained sophisticated and competent counsel ("Class Counsel")
25 who are experienced in prosecuting federal and state class actions throughout the United States and
26 other complex litigation and have extensive experience advising clients and litigating intellectual
27 property, competition, contract, and privacy matters.
28

1 **F. Other Class Considerations**

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

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

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

11 **VII. ARTISTS AND THEIR WORKS**

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

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

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

26 52. Artists also protect their works in other ways. One of the most common ways is to
27 affix a distinctive mark, watermark, signature, website URL or other identifying mark to their work.
28

1 These marks ensure that artists receive credit and recognition for the artwork they have created and
2 serves as a means to manage the use of their work.

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

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



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



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



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



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



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



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

21 55. Machine-learning models depend on massive quantities of high-quality data that is
22 digitally copied without authorization to train the model. The quality of the dataset a model is
23 trained on determines the quality of the model itself. The models within AI image products are no
24 different. As a result, the images these models are trained on have an enormous impact on the
25 quality of the models and the outputs they can produce. There are many images that are in the
26 public domain, *i.e.*, images that are not subject to copyright. The developers of many AI image
27 products, however, made the decision to greatly expand the datasets their models are trained on by
28

1 including billions of protected works. They made this choice because many desirable works are not
2 in the public domain, *i.e.*, are not subject to unfettered use by anyone without consent, credit, or
3 compensation.

4 56. These AI image products compete with the artists whose very works comprise the
5 raw material for the models within the AI image products. Rather than pay or commission an artist
6 to create a particular work, pay an artist for a print or copy of a particular work, or pay for the artist's
7 permission to use the work, now users and licensees of these AI image products merely have to
8 prompt the AI image product using terms—such as an artist's name, or titles of their works—to
9 generate an image indistinguishable from one the artist might've created themselves. Worse, certain
10 users of AI image products have at times abused this “feature” to harass and annoy the artists
11 themselves by making knockoff versions of their artwork and publicizing it.

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

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

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

21 59. In August 2021, LAION released LAION-400M, a dataset of 400 million training
22 images assembled from images accessible on the public internet. At the time, LAION-400M was the
23 largest freely available dataset of its kind. LAION distributes the LAION-400M dataset to the public
24 through its own website and elsewhere. Information about LAION-400M is available in an
25 accompanying paper by Schuhmann, Beaumont, and others titled “LAION-400M: Open Dataset of
26 CLIP-Filtered 400 Million Image-Text Pairs,” released in November 2021 (hereafter, the “LAION-
27
28

1 400M Paper”).⁶

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

7 61. The actual images referenced in the LAION-400M dataset records are not included
8 with the dataset. Anyone who wishes to use LAION-400M for training their own machine-learning
9 model must first acquire copies of the actual images from their URLs. To facilitate the copying of
10 these images, LAION provides a software tool called ‘img2dataset’⁷ that takes the metadata records
11 as input and makes copies of the referenced images from the URLs in each metadata record, thereby
12 creating local copies.

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

19 63. One of the entities that has made unauthorized copies of the LAION-400M training
20 images is LAION itself. According to the LAION-400M Paper, LAION made the dataset by starting
21 with Common Crawl metadata records. Common Crawl is a corpus of 250 billion web pages copied
22 from the public web, including assets like Plaintiffs’ images (<https://commoncrawl.org/>). The
23 metadata records contain web URLs. According to the LAION-400M Paper, LAION created
24 training images by first “pars[ing] through [the metadata records] from Common Crawl and
25 pars[ing] out all HTML IMG tags containing an alt-text attribute [that is, a text caption].” Then,
26

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

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

1 LAION “download[ed] the raw images from the parsed URLs”. *See* LAION-400M Paper at 3. To
2 ensure that the training images in the dataset had reasonably accurate captions, LAION used a CLIP
3 model to calculate the CLIP embeddings for the image and text of each image–text pair. These two
4 CLIP embeddings were compared to measure how well the text described the image. Image–text
5 pairs with low CLIP-similarity scores were omitted from the dataset.

6 64. Sometime after the release of LAION-400M in August 2021, Stability funded
7 LAION’s creation of a similar dataset, but much larger. In March 2022, Stability CEO Mostaque
8 called himself “the biggest backer of LAION.”⁸ In August 2022, Stability CEO Mostaque said, “I
9 funded LAION, underlying dataset for ... stable diffusion.”⁹ (After the initial complaint in this
10 action was filed, Mostaque changed his story, saying “We actually have/had no influence on
11 CompVis or LAION, did not funded [sic] either.”¹⁰)

12 65. In October 2022, LAION released LAION-5B, a dataset of 5.85 billion training
13 images—more than 14 times bigger than LAION-400M. Information about LAION-5B is available
14 in an accompanying paper called “LAION-5B: An open large-scale dataset for training next
15 generation image-text models,” by Schuhmann, Beaumont, Crowson, and others (hereafter, the
16 “LAION-5B Paper”).¹¹ According to the LAION-5B Paper, LAION-400M is a subset of LAION-
17 5B, meaning every image in LAION-400M is also in LAION-5B.

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

24
25
26 ⁸ <https://discord.com/channels/662267976984297473/938713143759216720/954674533942591510>

27 ⁹ <https://twitter.com/EMostaque/status/1559332564787240962>

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

¹¹ <https://arxiv.org/abs/2210.08402>

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

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

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

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

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

25 71. Despite this warning that LAION-5B “should only be used for academic research
26 purposes,” all the Defendants—Runway, Stability, Midjourney, and DeviantArt—have trained,
27

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

1 distributed, promoted, or deployed commercial systems that rely on LAION-5B, directly in
2 contravention of this warning.

3 72. Though LAION-400M and LAION-5B are often used to train diffusion models, they
4 are used to train other models as well. For instance, Stability sponsored LAION's creation of
5 OpenCLIP, a CLIP model trained on a subset of images from LAION-5B called LAION-2B
6 consisting of training images with English-language captions.¹³ On information and belief, because
7 all the Plaintiff images in **Exhibit A** (Plaintiff images in LAION-5B) have English-language captions,
8 they are also part of LAION-2B and were thus used to train OpenCLIP. To create OpenCLIP,
9 LAION necessarily had to create one or more copies of these images. Plaintiffs never gave their
10 permission to LAION to copy their images or use them to train OpenCLIP.

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

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

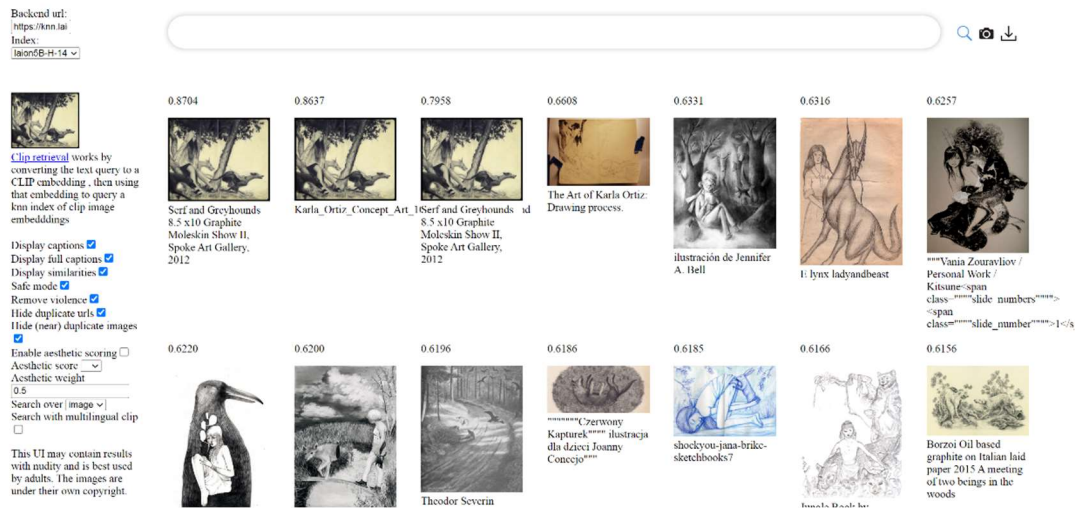


27
28 ¹³ <https://huggingface.co/laion/CLIP-ViT-H-14-laion2B-s32B-b79K>

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

3

4 Backend url:
https://knn.lai
Index:
laion5B_H-14

5 

6 [Clip retrieval](#) works by converting the text query to a CLIP embedding, then using that embedding to query a knn index of clip image embeddings

7 Display captions
 Display full captions
 Display similarities
 Safe mode
 Remove violence
 Hide duplicate urls
 Hide (near) duplicate images
 Enable aesthetic scoring
Aesthetic score:
Aesthetic weight:
Search over image
Search with multilingual clip

8 This UI may contain results with nudity and is best used by adults. The images are under their own copyright.

9

10

11

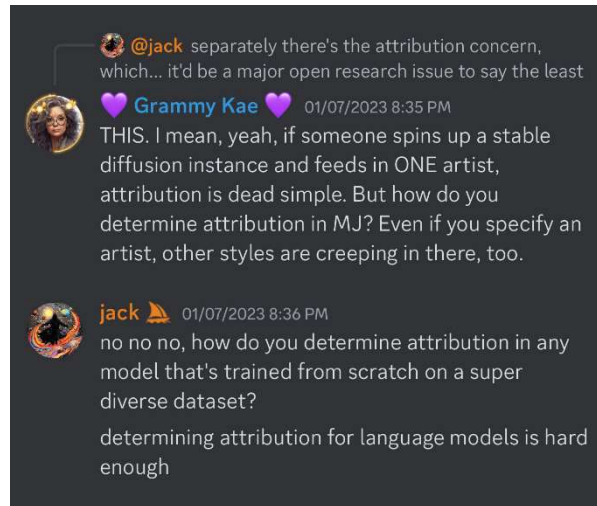
12 76. The caption shown for the first image in the search results of the LAION dataset is
13 exactly the same as the caption shown on Plaintiff Ortiz’s website: “Serf and Greyhounds / 8.5 x 10
14 Graphite / Moleskin Show II, Spoke Art Gallery, 2012”.

15 77. Further, inspection of the metadata given by [https://rom1504.github.io/clip-](https://rom1504.github.io/clip-retrieval/)
16 [retrieval/](https://rom1504.github.io/clip-retrieval/) indicates that the image was scraped from the following URL: [https://images.squarespace-](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg)
17 [cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg)
18 [Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRU](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg)
19 [qqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg)
20 [UQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg)
21 [ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg](https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317-Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsxRUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St3TBAAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg), which indicates the
22 source of the image was Ms. Ortiz’s own website. In other words, the metadata indicates that this
23 image was scraped from Ms. Ortiz’s personal website for inclusion into the LAION datasets.

24 78. In this way, the captions included in the training images also function as copyright-
25 management information. Much as a music publisher may search on a streaming platform for
26 unlawful uses of their work in order to conduct a DMCA strike, artists utilize keywords (for
27 example, their name) as search terms on <https://haveibeen trained.com> to identify whether their
28

1 works have been scraped and used as training material for AI image products.

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



7

8

9

10

11

12

13

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

19 81. In July 2023, the topic of AI training reached the U.S. Senate.¹⁵ During a hearing
20 convened by a subcommittee of the Senate Judiciary Committee, Sen. Mazie Hirono quizzed Ben
21 Brooks, a representative from Stability, about Stability’s position on licensing training data. Sen.
22 Hirono asked directly, “So basically you don’t pay for the data that you put into your—to train your
23 models?” Mr. Brooks replied, “There is no arrangement in place.” Sen. Hirono then turned to
24

25

26

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

28 ¹⁵ <https://youtu.be/uoCJun7gkbA?t=3578>

1 Plaintiff Karla Ortiz, who was testifying on the same panel. “So you have Ms. Ortiz, who says that
2 that is wrong. Is that correct, Ms. Ortiz?” Ms. Ortiz replied, “A hundred percent, Senator.”

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

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

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

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

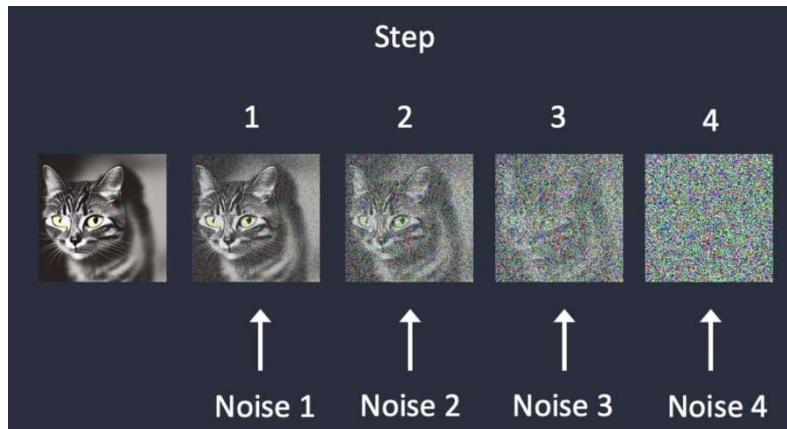
18 85. Diffusion proceeds in two phases. To begin the training phase, initial copies are made
19 of many training images. For each training image, progressively more noise is added over a series of
20 steps. At each step, the model records how the addition of noise changes the image. By the last step,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

26 ¹⁶ <https://discord.com/channels/662267976984297473/938713143759216720/982136076888068156>

27 ¹⁷ <https://arxiv.org/abs/1503.03585>

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

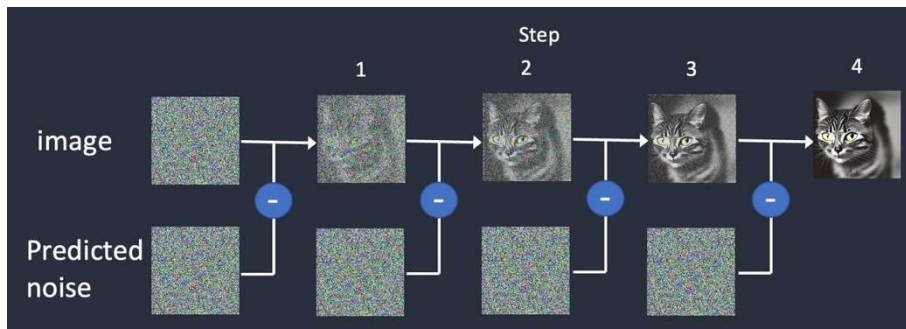
86.



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 Nicholas
5 Carlini, who has studied the behavior of diffusion models: “diffusion models are explicitly trained to
6 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 images
13 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 diffusion
16 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 in
14 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 by
19 a user that might describe elements of the image, *e.g.*— “a dog wearing a baseball cap while eating
20 ice cream.” As the denoising process proceeds, the conditioning data is used to nudge the denoising
21 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
12 prompt, first the CLIP model converts the prompt into a CLIP embedding. Then, as the diffusion
13 model 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,**
14 **COMPRESSED, 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 takes
23 data points and turns them into a query-able structure that **enables retrieval and interpolation**
24 **between the points**. You could think of it as a continuous generalization of database technology.
25 ... Because it is analogous to a database, the usefulness of a deep learning system depends entirely
26 on the data points it was constructed with. **You get back what you put in (or interpolations of the**
27
28

1 **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 other
 22 natural phenomena or technical innovations that were once “black
 23 boxes” to people, significant confusion and anxiety is arising in
 24 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 compresses**
21 **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 ²⁵ <https://stability.ai/blog/stable-diffusion-public-release>

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

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

28 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 Carlini
17 notes that “diffusion models are explicitly trained to reconstruct the training set.” *See* Carlini Paper
18 at 12. Based on further experiments, Carlini concludes “that state-of-the-art diffusion models *do*
19 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 directly
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 below in the “Generated” line is an image output from Stable Diffusion. As Carlini notes, the
 2 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] succeed,
 17 this question remains open.” *Id.* François Chollet has taken an even stronger position, saying that
 18 “It’s accurate that generative art models create new content by recombining images from their
 19 training data.”³¹ Yaodong Yu concurs, stating “large image generation models ... are unlikely to do
 20 anything beyond purely mechanical data compression (encoding) and interpolation (decoding).” *See*
 21 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 Webster
9 found that by using a less strict technique for detecting stored copies of training images, more
10 instances of stored copies could be discovered. Webster tested several diffusion models, including
11 Stable Diffusion 2.0 and version 4 of the Midjourney Model, and found stored copies of training
12 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 models
15 with more parameters are considered “larger” than those with fewer.) But in July 2023, Stability
16 released Stable Diffusion XL 1.0, which has 3.5 billion parameters. On information and belief, based
17 on Carlini’s theory that larger models are more likely to “memorize and regenerate individual
18 training images,” a model like Stable Diffusion XL 1.0 is even more likely to exhibit this behavior
19 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 works
6 made by artists whose names were used to generated them ... Overall, these results suggest that
7 Stable Diffusion has a broad-ranging ability to imitate the style of individual artists.” *See* Casper
8 Paper at 3.

9 144. In sum—based on work by leading AI researchers, AI image products are only getting
10 better and better at storing copies of training images and can even produce images indistinguishable
11 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 artist’s
20 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
3 As a rule, ‘open’ refers to systems that offer transparency,
4 reusability, and extensibility—they can be scrutinized, reused, and
5 built on. ... we also find that marketing around openness and
6 investment in (somewhat) open AI systems is being leveraged by
7 powerful companies to bolster their positions in the face of growing
8 interest in AI regulation. And that some companies have moved to
9 embrace ‘open’ AI as a mechanism to entrench dominance, using the
10 rhetoric of ‘open’ AI to expand market power, and investing in
11 ‘open’ AI efforts in ways that allow them to set standards of
12 development while benefiting from the free labor of open source
13 contributors ... **Companies like ... Stability AI offer open source
14 AI models to their customers and the public. Their business
15 models rely not on licensing proprietary models themselves, but
16 on charging for extra features and services on top of open
17 models,** features such as API access, model training on custom data,
18 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* file
16 is a binary file that encodes all the information that the model captured during training by copying
17 protected expression from training images. An example of a Stable Diffusion weights file is available
18 at <https://huggingface.co/stabilityai/stable-diffusion-2/blob/main/768-v-ema.ckpt>. This is a mass
19 of binary data that is meaningful when accessed via the Stable Diffusion scripts, but otherwise not
20 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 an
22 online AI image product called DreamStudio (<https://beta.dreamstudio.ai>). DreamStudio was used
23 to make the text-prompt examples below. The version of the Stable Diffusion model made available
24 in DreamStudio is called Stable Diffusion XL 1.0. This version of Stable Diffusion was trained on
25 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 their
19 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 "*teacher*"
17 prompt shows what Stable Diffusion produces by default, without an artist name in the prompt. The
18 Manchess-inspired teachers are rendered in calligraphic brushwork with realistic settings, as
19 frequently found in the artwork of Mr. Manchess. The Brom-inspired teachers are demonic, and
20 feature images of weapons and skulls, like a characteristic artwork by Mr. Brom. The Zhang-inspired
21 teachers are Asian and rendered photographically, like many of Ms. Zhang's images. As with "chef,"
22 the addition of the artist's name causes the generic term "teacher" to be rendered in a manner
23 characteristic of the artist. Furthermore, the changes provoked by the insertion of the artist's name
24 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 Plaintiff
2 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*.” **Exhibit E**,
17 pp. 1–4. The initial “*chef*” prompt shows what the Runway image product produces by default,
18 without an artist name in the prompt. Much like the Stability results, the Manchess-inspired chefs
19 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 andersen*
2 *teacher.*” **Exhibit E**, pp. 5–8. The initial “*teacher*” prompt shows what the Runway image product
3 produces by default, without an artist name in the prompt. Much like the Stability results, the
4 Manchess-inspired teachers are rendered in calligraphic brushwork with realistic settings, common
5 in the work of Mr. Manchess. The Brom-inspired teachers are fantastic, otherworldly creatures,
6 resembling those often found in Mr. Brom’s work. The Zhang-inspired teachers are Asian and
7 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 available
13 through an online discussion service called Discord (<https://discord.com>). Midjourney’s AI image
14 product is also called “Midjourney.” Thus, for clarity below, the Midjourney AI image product as
15 presented via Discord will be called the **Midjourney Image Product**. The text-prompt examples
16 below were made using the Midjourney Image Product. Because the Midjourney Image Product
17 incorporates both Stable Diffusion as an underlying model as well as a custom model trained by
18 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*.” **Exhibit F**,
7 pp. 1–6. The initial “*chef*” prompt shows what the Midjourney Model produces by default, without
8 an artist name in the prompt. Much like the Stability results, the Manchess-inspired chefs are
9 rendered in the calligraphic brushwork characteristic of Mr. Manchess’s oil paintings. The Brom-
10 inspired chefs have gothic and countercultural air, including several fantastic creatures, typical of
11 artwork by Mr. Brom. The Zhang-inspired chefs are Asian and rendered photographically, like many
12 of Ms. Zhang’s works. The McKernan-inspired chefs feature colorful female-presenting faces with
13 elaborate hair and decorations, commonly found in Mx. McKernan’s work. The Andersen-inspired
14 chefs are all obvious variations on the main character of Ms. Andersen’s celebrated comic “Sarah’s
15 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 andersen*
18 *teacher*.” **Exhibit F**, pp. 7–12. The initial “*teacher*” prompt shows what the Midjourney Model
19 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 skulls,
22 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 by
12 Stability, Runway, and Midjourney (as of November 2023) showing the results of prompting these
13 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 a
16 CLIP model to convert each text prompt to a numerical descriptor called a CLIP embedding, which
17 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!

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

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

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

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

25 185. **Runway:** The next set of image-prompt examples were made with Runway’s AI
26

27 _____
28 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 merges
20 them together into novel interpretations.”³⁸ According to Midjourney’s documentation, using an
21 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 Midjourney
23 with training images from **Exhibit A** (Plaintiff images in LAION-5B). These training images were
24 made by Plaintiffs Gerald Brom, Gregory Manchess, Grzegorz Rutkowski, Hawke Southworth,
25 Jingna Zhang, Karla Ortiz, and Kelly McKernan. Midjourney requires that an image prompt be
26

27 ³⁸ <https://discord.com/channels/662267976984297473/952771221915840552/1064031587735445546>

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

1 accompanied by a text prompt, so in these cases, each image prompt was supplemented by the
2 artist's name. Each prompt was also supplemented with the command '--iw 2' to ensure that the
3 image portion of the prompt was treated as the primary part (where "iw" means "image weight"),
4 thereby maximizing "the 'concepts' and 'vibes'" derived from the image.

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

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

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

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

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

25 195. In this example (**Ex. G** at p. 2), the original image is in the top-left quadrant. Plaintiff
26 Brom's CMI in the form of the URL for his personal website is clearly visible. None of the visually
27 similar copies of the original image generated by the Stable Diffusion XL model contained the
28

original CMI.



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

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



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 CMI
2 in the form of the URL for his personal website is clearly visible. The Runway model generated
3 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



1
2
3
4
5
6
7
8
9
10
11
12 200. In each of these examples, the copies generated by the AI image product could not
13 have been generated but for copying the original image which included CMI by operation of the
14 processes described herein.

15 **XIII. USER AND LICENSEE ACTIVITY**

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

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

26 203. An AI image product user posted onto the internet a purported study representing
27 “[a] collection of images from Midjourney that demonstrate the recognized artist styles and their
28

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
7 as 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 which
8 the work can be perceived, reproduced, or otherwise communicated, either directly or with the aid
9 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 XV. CAUSES OF ACTION AGAINST STABILITY

26 214. Between April and November 2022, Stability trained an image model called Stable
27 Diffusion 2.0. According to Stability, “The model developers used the following dataset for training
28

1 the model: LAION-5B and subsets 5B” (See [https://huggingface.co/stabilityai/stable-diffusion-](https://huggingface.co/stabilityai/stable-diffusion-2#training)
2 [2#training](https://huggingface.co/stabilityai/stable-diffusion-2#training)).

3 215. Between November 2022 and July 2023, Stability trained an image model called
4 Stable Diffusion XL 1.0. On information and belief, Stable Diffusion XL is also trained on LAION-
5 5B, because Stability has funded LAION and used LAION datasets for all its previous models. AI
6 chipmaker NVIDIA features Stable Diffusion XL as one of its “AI Foundation Models.” On its
7 information page for Stable Diffusion XL, the training dataset is listed as LAION-5B.⁴³

8 216. Because LAION-5B is an openly accessible dataset, Stability knew the LAION-5B
9 dataset contained copyrighted works, including those of the LAION-5B Registered Plaintiffs.
10 Additionally, because the LAION-5B dataset contains detection scores for watermarks and because
11 CMI is ubiquitous in art, Stability also knew that the LAION-5B dataset contained copyrighted
12 works with CMI affixed to them.

13 217. The LAION-5B Registered Works are included in the LAION-5B dataset. Therefore,
14 Stability used the LAION-5B Registered Works for training. Below, the term **Stability Models**
15 refers to all models trained by Stability on the LAION-5B Registered Works, including Stable
16 Diffusion 2.0 and Stable Diffusion XL 1.0.

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

22
23
24
25
26
27
28 ⁴³ See <https://catalog.ngc.nvidia.com/orgs/nvidia/teams/ai-foundation/models/sdxl/overview>

COUNT ONE

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

219. The preceding factual allegations are incorporated by reference.

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

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

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


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

1 infringing Statutory Derivative Work.

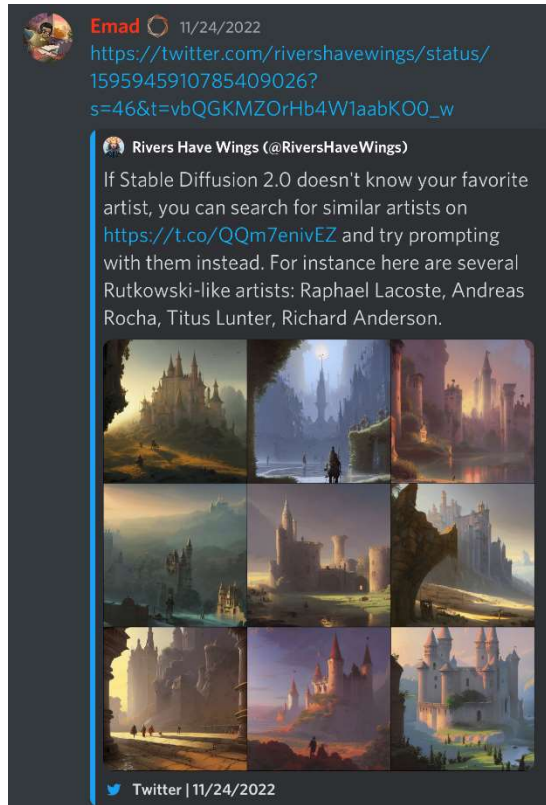
2 224. Executives and high-level employees of Stability know that one of the most attractive
3 features of the Stability models is its ability to mimic and copy artists' works, including Plaintiffs. As
4 such, they routinely advertise the Stability Models' ability to mimic artwork.

5 225. For example, once Stability prohibited the use of Plaintiff Rutkowski's name as a
6 prompt, Stability's employees and executives encouraged the use of similar artist names in lieu of
7 Plaintiff Rutkowski's in order to achieve similar results.

8 226. For example, Katherine Crowson, a principal researcher at Stability AI tweeted the
9 following on November 24, 2022:



1 227. Emad Mostaque, Stability's CEO retweeted Crowson's advice:



15

16 228. 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

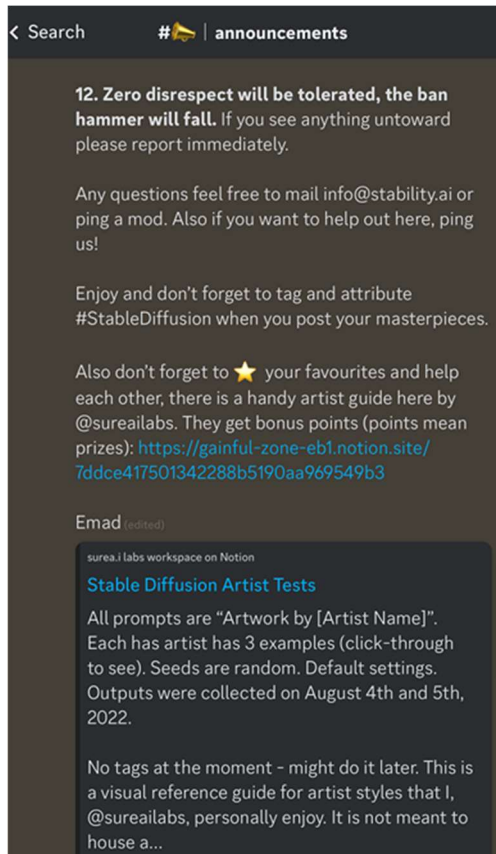
24

25

26

27

28



26 229. On information and belief, the other Stability Models exhibit the same properties,

27 because they were trained on the same LAION-5B dataset.

28 230. 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.

231. 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

232. The preceding factual allegations are incorporated by reference.

233. 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>).

234. 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.

235. 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.

236. 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.

237. 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.

COUNT THREE

**DMCA violations by removing and altering CMI of training images
against Stability on behalf of all Plaintiffs, the Damages and the Injunctive Classes**

238. The preceding factual allegations are incorporated by reference.

239. The LAION-5B Plaintiffs included one or more forms of CMI (as defined in Section 1202(c) of the DMCA) in each of their respective works in the LAION-5B Works, including captions, and distinctive marks such as URLs to personal websites, signatures, and watermarks.

240. Stability did not contact Plaintiffs and the Class to obtain authority to remove or alter CMI from their works within the meaning of the DMCA.

241. Stability knew the LAION-5B dataset contained CMI. The LAION-5B dataset includes a detection score for watermarks which indicates the likelihood of a particular image in the dataset contains a watermark or other distinctive mark signaling the presence of CMI. Stability therefore could have trained the Stability Models on images free of CMI, but instead chose not to because images with CMI tend to be high-quality.

242. Stability had access to but were not licensed by Plaintiffs or the Class to incorporate their works in the Stability Models.

243. Stability had access to but were not licensed by Plaintiffs or the Class to create copies based on their works into the Stability Models.

244. Stability had access to but were not licensed by Plaintiffs or the Class to distribute their works as Stability does through the Stability Models.

245. Without the authority of the Plaintiffs, Stability directly copied the LAION-5B Works and used these Statutory Copies as training data for the Stability Models. The works copied by Stability included CMI, including in the form of distinctive marks such as watermarks or signatures, and as the captions in the image-text pairs. The training process is designed to remove or alter CMI from the training images. Therefore, Stability intentionally removed or altered CMI from the Plaintiffs' works in violation of 17 U.S.C. § 1202(b)(1).

246. Stability also knew that the distribution of works without CMI would lead to further

1 infringement. Stability encourages the use of artist names as prompts, *i.e.*, encourages the Stability
2 Models' users and licensees to infringe on an artists' work. Because the Stability Models do not
3 preserve CMI, users and licensees also create infringing works without CMI which can reasonably
4 lead to further infringement.

5 247. Without the authority of the LAION-5B Plaintiffs, Stability directly copied the
6 LAION-5B Works and used these Statutory Copies as training data for the Stability Models. The
7 training process is designed to remove or alter CMI from the training images. As demonstrated
8 herein, the Stability Models generate copies based on original images with the CMI removed and/or
9 altered as output. Therefore, Stability intentionally removed or altered CMI from the LAION-5B
10 Works in violation of 17 U.S.C. § 1202(b)(1).

11 248. Stability distributes the Stability Models under the MIT License (*See, e.g.* —
12 <https://github.com/Stability-AI/stablediffusion/blob/main/LICENSE>). Within this license,
13 Stability asserts copyright in the Stability Models. By asserting copyright in the Stability Models,
14 which infringe the copyrights of the LAION-5B Plaintiffs, Stability is providing and distributing false
15 CMI in violation of 17 U.S.C. § 1202(a).

16 249. Stability knew or had reasonable grounds to know that this removal of the LAION-5B
17 Plaintiffs' CMI would facilitate copyright infringement by concealing the facts that a) the Stability
18 Models are infringing Statutory Copies of the LAION-5B Works, and b) the Stability Models are
19 infringing Statutory Derivative Works based on the LAION-5B Works.

20 250. The LAION-5B Plaintiffs have been injured by Stability's removal or alteration of
21 CMI. The LAION-5B Plaintiffs have been injured by Stability's falsification of CMI by claiming
22 false copyright in the Stability Models. These Plaintiffs are entitled to statutory damages, actual
23 damages, restitution of profits, and other remedies provided by law.

24 **XVI. CAUSES OF ACTION AGAINST MIDJOURNEY**

25 251. Midjourney promotes the Midjourney Image Product, which is accessed and run
26 through Discord. Midjourney maintains its own Discord server from which users can access the
27 Midjourney image generator. Midjourney's Discord server also allows Midjourney executives and
28

1 other high-level employees to provide promotional communications to users and licensees of the
2 Midjourney Image Product.

3 252. In February 2022, near the release of the initial version of the Midjourney Image
4 Product, Midjourney CEO David Holz posted messages on the Midjourney Discord server
5 promoting the Midjourney Image Product's ability to emulate existing artistic styles, in particular
6 the styles of certain artists.

7 253. Over a series of Discord messages, Holz said "i think you're all gonna get [your]
8 mind blown by this style feature ... we were very liberal in building out the dictionary ... it has cores
9 and punks and artist names ... as much as we could dump in there ... i should be clear it's not just
10 genres its also artist names ... it's mostly artist names ... 4000 artist names."⁴⁴

11 254. Holz then said, "here is our style list"⁴⁵ and posted a link to a spreadsheet on Google
12 Docs called "Midjourney Style List."⁴⁶ One of the tabs on the spreadsheet was called "Artists" and
13 listed over 4,700 artist names. In other words, Holz published a list of artists who the Midjourney
14 Image Product recognizes with the express purpose of these names being used by users and licensees
15 of the Midjourney Image Product as terms in prompts. Holz's comment, and the list, have remained
16 available ever since.

17 255. Below, this list is called the **Midjourney Name List**. A copy of this list appears in
18 **Exhibit J: Midjourney Name List**.

19 256. Plaintiffs Grzegorz Rutkowski, Sarah Andersen, Karla Ortiz, Gerald Brom, and Julia
20 Kaye appear in the Midjourney Name List. Below, this subset of Plaintiffs will be referred to as the
21 **Midjourney Named Plaintiffs**.

22
23
24 ⁴⁴ Combining
25 <https://discord.com/channels/662267976984297473/938713143759216720/941972360171520001>,
26 <https://discord.com/channels/662267976984297473/938713143759216720/941972890520272906>,
27 and
28 <https://discord.com/channels/662267976984297473/938713143759216720/941976464704802836>

⁴⁵ <https://discord.com/channels/662267976984297473/938713143759216720/941987328828768256>

⁴⁶ https://docs.google.com/spreadsheets/d/1MEglfejpqgVcaf-I-cgZ5ngV_MlaOTeGXAoBPJO69FM/edit#gid=1096178862

1 257. Midjourney also tracked the most popular artists used as prompts. At one point, the
2 Midjourney Bot maintained a count on Midjourney as recently as April 2023.

3 258. In the months before February 2022, Midjourney trained version 1 of the Midjourney
4 Model. In February 2022, on Midjourney’s Discord server Midjourney CEO David Holz described
5 Midjourney’s training data: “we have some private data partners as well as some open ones like
6 laion.”⁴⁷ In May 2022, Stability CEO Mostaque said “MidJourney is using a LAION 400m based
7 model ... I just support MJ like many other labs/researchers in my nice way.”⁴⁸ In July 2022, a
8 Midjourney-affiliated Discord moderator named Soar also confirmed that Midjourney was “using a
9 modified version of the LAION 400m dataset.”⁴⁹ In other words, Midjourney trained its image
10 generation product on at least the LAION-400M dataset.

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

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

18 261. The term **LAION-400M Registered Works** denotes the subset of works in the
19 LAION-400M Works that are covered by registered copyrights.

20 262. The term **LAION-400M Registered Plaintiffs** denotes the subset of Plaintiffs who
21 hold registered copyrights in these LAION-400M Registered Works, namely Sarah Andersen and
22 Jingna Zhang.

23 263. Because LAION-400M is an openly accessible dataset, Midjourney knew that the
24 LAION-400M dataset contained copyrighted works, including those of the LAION-400M
25

26
27 ⁴⁷ <https://discord.com/channels/662267976984297473/938713143759216720/943315577018126408>

⁴⁸ <https://discord.com/channels/823813159592001537/912729332311556136/975894553225752626>

28 ⁴⁹ <https://discord.com/channels/662267976984297473/959962985655320616/1001938136445751387>

1 Registered Plaintiffs. Further, because LAION also includes detection scores for watermarks for
 2 other datasets, and because of the ubiquity of artists affixing CMI to their works, Midjourney knew
 3 that the LAION-400M dataset contained works with CMI affixed on them.

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

8 265. Since October 2022, Midjourney has also incorporated a version of Stable Diffusion
 9 into the Midjourney Image Product, which is accessible to users by adding the command ‘--test’ or ‘-
 10 -testp’ to a text prompt. According to Midjourney’s moderator Molang, “--test and -testp is a little
 11 bit of SD [Stable Diffusion] mixed with a lot of Midjourney tweaks and magic.”⁵⁰

12 266. Sometime after July 2022 and before March 2023, Midjourney adopted LAION-5B
 13 as its training dataset. In July 2022, a Midjourney-affiliated Discord moderator named Danger
 14 Awesome said that “the updated LAION 5B dataset”⁵¹ formed the basis for “the upcoming
 15 [Midjourney] dataset update.” In March 2023, a Midjourney-affiliated Discord Moderator named
 16 Sunshineyday said that “MJ is trained on a subset of Laion5b.”⁵²

17 267. Version 5 of the Midjourney Model was released in March 2023. On information and
 18 belief, version 5 of the Midjourney Model and subsequent versions were trained on LAION-5B.

19 268. Because LAION-5B is an openly accessible dataset, Midjourney knew that the
 20 LAION-5B dataset contained copyrighted works, including those of the LAION-5B Registered
 21 Plaintiffs. Because LAION-5B also includes detection scores for watermarks, Midjourney knew that
 22 the LAION-5B dataset contained works with CMI affixed to them.

24 50

25 <https://discord.com/channels/662267976984297473/958069758211797092/1038899058636501132>
 26 51

27 <https://discord.com/channels/662267976984297473/992207085146222713/998451098534817883>
 28 52

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

1 269. The LAION-5B Registered Works are included in the LAION-5B dataset. Therefore,
2 Midjourney used the LAION-5B Registered Works for training. Below, the term **Midjourney 5B**
3 **Models** refers to all models trained by Midjourney on the LAION-5B Registered Works, including
4 version 5 of the Midjourney Model and the current version 5.2 of the Midjourney Model.

5 270. Midjourney also publishes the Midjourney Magazine, which is analog only and is
6 unavailable in digital form. Those who wish to subscribe to the Midjourney Magazine need to
7 subscribe for a monthly fee. Each issue of the Midjourney Magazine features a selection of output
8 generated by the Midjourney Model and along with the prompts that created them. Each image is
9 selected by Midjourney for inclusion in the Midjourney Magazine. Numerous exemplar prompts
10 feature the use of artist names as a keyword, including those of Plaintiffs Brom and Rutkowski.

11 **COUNT FOUR**

12 **Direct copyright infringement of the LAION-400M Registered Works by training the** 13 **Midjourney 400M Models, including Midjourney Model version 1 against Midjourney on** 14 **behalf of the LAION-400M Registered Plaintiffs and Damages Subclass**

15 271. The preceding factual allegations are incorporated by reference.

16 272. The LAION-400M Registered Plaintiffs never authorized Midjourney to use their
17 respective LAION-400M Registered Works in any way. Nevertheless, Midjourney repeatedly
18 violated the exclusive rights (under 17 U.S.C. § 106) of the LAION-400M Registered Plaintiffs and
19 continues to do so today.

20 273. The LAION-400M dataset contains only URLs of training images, not the actual
21 training images. Therefore, anyone who wishes to use LAION-400M for training their own
22 machine-learning model must first acquire copies of the actual training images from their URLs.
23 Consistent with this, in preparation for training the Midjourney 400M Models, Midjourney made
24 one or more Statutory Copies of the LAION-400M Registered Works so they could be fed to each
25 Midjourney 400M Model as training data. The Statutory Copies made of each registered work were
26 substantially similar to that registered work.

27 274. During the training of each Midjourney 400M Model, Midjourney made a series of
28

1 intermediate Statutory Copies of the LAION-400M Registered Works. For instance, diffusion
2 models are trained by creating “noised” copies of training images, as described herein, all of which
3 qualify as Statutory Copies. The intermediate Statutory Copies of each registered work that
4 Midjourney made during training of the Midjourney 400M Models were substantially similar to that
5 registered work.

6 275. The LAION-400M Registered Plaintiffs have been and continue to be injured by
7 Midjourney’s multiple acts of direct copyright infringement. These Plaintiffs are entitled to
8 statutory damages, actual damages, restitution of profits, and other remedies provided by law.

9 COUNT FIVE

10 **Direct copyright infringement of the LAION-5B Registered Works** 11 **by training the Midjourney 5B Models, including Midjourney Model version 5.2** 12 **against Midjourney on behalf of the LAION-5B Registered Plaintiffs and Damages Subclass**

13 276. The preceding factual allegations are incorporated by reference.

14 277. The LAION-5B Registered Plaintiffs never authorized Stability to use their
15 respective LAION-5B Registered Works in any way. Nevertheless, Stability repeatedly violated the
16 exclusive rights (under 17 U.S.C. § 106) of the LAION-5B Registered Plaintiffs and continues to do
17 so today.

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

25 279. During the training of Midjourney Model version 5.2 and the other Midjourney 5B
26 Models, Midjourney made a series of intermediate Statutory Copies of the LAION-5B Registered
27 Works. For instance, diffusion models are trained by creating “noised” copies of training images, as
28

1 described herein, all of which qualify as Statutory Copies. The intermediate Statutory Copies of
2 each registered work that Midjourney made during training of the Midjourney 5B Models were
3 substantially similar to that registered work.

4 280. By the end of training, Midjourney Model version 5.2 was capable of reproducing
5 protected expression from each of the LAION-5B Registered Works that was in each case
6 substantially similar to that registered work, as shown in **Exhibit F: Midjourney text prompts** and
7 **Exhibit I: Midjourney image prompts**. Therefore, Midjourney Model version 5.2 qualifies as an
8 infringing Statutory Copy of the LAION-5B Registered Works. Because Midjourney Model version
9 5.2 represents a transformation of the LAION-5B Registered Works into an alternative form,
10 Midjourney Model version 5.2 also qualifies as an infringing Statutory Derivative Work.

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

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

16 COUNT SIX

17 **DMCA violations by removing and altering CMI of training images** 18 **against Midjourney on behalf of All Plaintiffs, the Damages and Injunctive Class**

19 283. The preceding factual allegations are incorporated by reference.

20 284. The LAION-400M Plaintiffs included one or more forms of CMI (as defined in
21 Section 1202(c) of the DMCA) in each of their respective works in the LAION-400M Works,
22 including captions, and distinctive marks such as URLs to personal websites, signatures, and
23 watermarks.

24 285. The LAION-5B Plaintiffs included one or more forms of copyright-management
25 information (as defined in Section 1202(c) of the DMCA) in each of their respective works in the
26 LAION-5B Works, including captions, URLs, signatures, and watermarks.

27 286. Midjourney did not contact Plaintiffs and the Class to obtain authority to remove or
28

1 alter CMI from their works within the meaning of the DMCA.

2 287. Midjourney knew the LAION-5B dataset contained CMI. The LAION-5B dataset
3 includes a detection score for watermarks which indicates the likelihood of a particular image in the
4 dataset contains a watermark or other distinctive mark signaling the presence of CMI. Midjourney
5 therefore could have trained the Midjourney Model on images free of CMI but chose not to because
6 images with CMI tend to be high-quality.

7 288. Midjourney also knew the LAION-400M dataset contained CMI. Given that CMI is
8 ubiquitous and the LAION-400M dataset contained copyrighted works, it is certain that many works
9 within the dataset contained CMI.

10 289. Midjourney had access to but were not licensed by Plaintiffs or the Class to
11 incorporate their works in Midjourney Image Product.

12 290. Midjourney had access to but were not licensed by Plaintiffs or the Class to create
13 copies based on their works into the Midjourney Image Product.

14 291. Midjourney had access to but were not licensed by Plaintiffs or the Class to distribute
15 their works as Midjourney does through the Midjourney Image Product.

16 292. Without the authority of the Plaintiffs, Midjourney directly copied the LAION-5B
17 Works and used these Statutory Copies as training data for the Midjourney Image Product. The
18 works copied by Midjourney included CMI, including in the form of distinctive marks such as
19 watermarks or signatures, and as the captions in the image-text pairs. The training process is
20 designed to remove or alter CMI from the training images. Therefore, Midjourney intentionally
21 removed or altered CMI from the Plaintiffs' works in violation of 17 U.S.C. § 1202(b)(1).

22 293. Without the authority of the Plaintiffs, Midjourney copied the LAION-400M and
23 LAION-5B Works and used these copies as training data for the Midjourney Models. The training
24 process is designed to remove or alter CMI from the training images. Therefore, Midjourney
25 intentionally removed or altered CMI from the Plaintiffs' Works in violation of 17 U.S.C.
26 § 1202(b)(1).

27 294. Midjourney also knew that the distribution of works without CMI would lead to
28

1 further infringement. Midjourney encourages the use of artist names as prompts, *i.e.*, encourages
2 the Midjourney Image Product's users and licensees to infringe on an artist's work. As
3 demonstrated herein, the Midjourney Model generates copies of original works with their CMI
4 removed and/or altered. Because the Midjourney Image Product does not preserve CMI, users and
5 licensees also create infringing works without CMI which can reasonably lead to further
6 infringement.

7 295. The LAION-400M Plaintiffs have been injured by Midjourney's removal or
8 alteration of CMI. The LAION-400M Plaintiffs have also been injured by Midjourney's falsification
9 of CMI. These Plaintiffs are entitled to statutory damages, actual damages, restitution of profits, and
10 other remedies provided by law.

11 COUNT SEVEN

12 **Lanham Act — false endorsement by unauthorized commercial use of artists' names** 13 **against Midjourney on behalf of the Midjourney Named Plaintiffs and Class**

14 296. The preceding factual allegations are incorporated by reference.

15 297. Midjourney engaged in commercial speech that sought to capitalize upon the
16 Midjourney Named Plaintiffs' popularity, recognition, and appeal among consumers of art products.
17 Midjourney's use of the Midjourney Named Plaintiffs' names was purely to advertise its image
18 generator. This use does not contribute significantly to a matter of public interest. The purpose of
19 publishing over 4,700 names in the Midjourney Name List was to promote and highlight the
20 capabilities of Midjourney's image generator to emulate and create work that is indistinguishable
21 from that of the artists whose names were published.

22 298. Midjourney's use of the Midjourney Named Plaintiffs' names was unauthorized and
23 without their consent.

24 299. Midjourney's commercial speech created a likelihood of confusion over whether the
25 Midjourney Named Plaintiffs actually endorsed the Midjourney Image Product, and over the
26 affiliation, connection, or association that the Midjourney Named Plaintiffs might have with
27 Midjourney.
28

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

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

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

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

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

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

19 306. Midjourney’s use of the Midjourney Named Plaintiffs’ names was intentional, in
20 order to capitalize on their fame and goodwill as popular artists.

21 307. If its conduct is left unchecked, Midjourney likely will continue to use artist names to
22 advertise the capabilities of its image generator, which undergoes frequent updates.

23 308. The Midjourney Named Plaintiffs have been, and likely will continue to be, harmed
24 by Midjourney’s misrepresentation of fact in terms of their reputation and goodwill.

25 309. The Midjourney Named Plaintiffs are entitled to statutory damages, actual damages,
26 restitution of profits, and other remedies provided by law.

COUNT EIGHT

Lanham Act — vicarious trade-dress violation by profiting from imitations of protectable trade dress against Midjourney on behalf of the Midjourney Named Plaintiffs and Class

310. The preceding factual allegations are incorporated by reference.

311. The Midjourney Named Plaintiffs each sell original art, art reproductions, and art products, all of which feature respective protectable and distinctive trade dress. This trade dress consists of a set of recurring visual elements and artistic techniques, the particular combination of which are distinctive to each of the Midjourney Named Plaintiffs, associated with them and their work, and desirable to customers. For instance—

- a. Sarah Andersen is known for work that is simple, cartoony, and often strictly in black and white. In particular, she is known for “Sarah’s Scribbles,” a comic featuring a young woman with dark hair, big eyes, and a striped shirt.
- b. Karla Ortiz is known for a mixture of classical realism and impressionism, often delving into fantastical, macabre and surrealist themes, and inspired by the technical prowess of American Renaissance movements with a strong influence of contemporary media.
- c. Gerald Brom is known for gritty, dark, fantasy images, painted in traditional media, combining classical realism, gothic and counterculture aesthetics.
- 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.

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

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

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

14 316. Midjourney acknowledges and, in fact, relies on the inherent distinctiveness of the
15 Midjourney Named Plaintiffs' respective trade dress to market its image generator by advertising
16 that users can generate images in the style of particular artists simply by typing in their name. In this
17 way, users do not have to describe specific design or artistic elements in the prompt to generate an
18 image in the artist's style—they merely need to type in that artist's name. Examples of Midjourney
19 text prompts featuring Sarah Andersen and Gerald Brom are shown in **Exhibit F**.

20 317. Midjourney vicariously infringes on the Midjourney Named Plaintiffs' trade-dress
21 rights by encouraging and inducing the users of the Midjourney Image Product to enter artist-name
22 prompts and generate images featuring the Midjourney Named Plaintiffs' protectable trade dress.
23 For example—

- 24 a. In its original online documentation offering “tips for text-prompts,”
25 Midjourney recommended that users should “try invoking unique artists to
26 get a unique style,” and offered a list that included “Greg Rutkowski,” who is a
27 one of the Midjourney Named Plaintiffs. The documentation also
28

1 recommended that users should “Combine names for new styles: ‘A temple
2 by Greg Rutkowski and Ross Tran.’” (Midjourney deleted these pages from
3 its public website two weeks after the initial complaint was filed in this
4 action.)

- 5 b. Midjourney currently promotes images made with artist-name prompts in an
6 online marketing gallery accessible to subscribers called “Showcase” (at
7 <https://midjourney.com/showcase>). **Exhibit K: Midjourney Showcase**
8 contains examples of images from the Showcase gallery featuring Plaintiff
9 names, including two of the Midjourney Named Plaintiffs: Sarah Andersen
10 and Gerald Brom.

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

16 319. Each of the Midjourney Named Plaintiffs’ respective trade dress has no intrinsic
17 functional value. The unique combination of particular artistic elements does not confer any
18 utilitarian advantages on their art products and are purely ornamental and aesthetic. There remains
19 an unlimited number of alternative artistic styles available beyond the trade dress owned by the
20 Midjourney Named Plaintiffs.

21 320. Each of the Midjourney Named Plaintiffs’ trade dress possesses secondary meaning
22 because the trade dress of their art products invoke a mental association by a substantial segment of
23 potential consumers between the trade dress and the creator of the art product.

24 321. Midjourney’s vicarious infringement of the Midjourney Named Plaintiffs’ trade-
25 dress rights are committed with actual and constructive knowledge of their trade dress, and with the
26 intent to cause confusion, mistake, or deception.

27 322. As a direct and proximate cause of Midjourney’s conduct, the Midjourney Named
28

1 Plaintiffs have suffered, and will continue to suffer, significant damage in the form of loss of revenue,
2 income, profits, and goodwill, which will increase if not enjoined. Midjourney has, and will unfairly,
3 acquire revenue, income, profits, and goodwill at the expense of the Midjourney Named Plaintiffs.

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

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

12 325. Midjourney's past and continuing infringement of the Midjourney Named Plaintiffs'
13 trade dress is an exceptional case and was willful and intentional, as evidenced by a) Midjourney's
14 intentional inclusion of the CLIP model in the design of the Midjourney Image Product and b) its
15 open advertisement of the Midjourney Image Product's ability to replicate an artist's trade dress via
16 the Midjourney Name List. Thus, the Midjourney Named Plaintiffs are entitled to treble their actual
17 damages and to an award of attorneys' fees under 15 U.S.C. § 1117(a), and all other available
18 remedies.

19 XVII. CAUSES OF ACTION AGAINST RUNWAY

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

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

25 328. Because LAION-5B is an openly accessible dataset, Runway knew that the LAION-
26 5B dataset contained copyrighted works, including those of the LAION-5B Registered Plaintiffs and
27 Karla Ortiz.

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

5 **COUNT NINE**

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

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

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

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

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

26 334. By the end of training, Stable Diffusion 1.5 was capable of reproducing protected
 27 expression from each of the LAION-5B Registered Works that was in each case substantially similar
 28

1 to that registered work, as shown in **Exhibit E: Runway text prompts** and **Exhibit H: Runway**
2 **image prompts**. Therefore, Stable Diffusion 1.5 qualifies as an infringing Statutory Copy of the
3 LAION-5B Registered Works. Because Stable Diffusion 1.5 represents a transformation of the
4 LAION-5B Registered Works into an alternative form, Stable Diffusion 1.5 also qualifies as an
5 infringing Statutory Derivative Work.

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

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

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

15 COUNT TEN

16 **Inducement of copyright infringement by distributing Stable Diffusion 1.5 for free** 17 **against Runway on behalf of the LAION-5B Registered Plaintiffs and Subclass**

18 338. The preceding factual allegations are incorporated by reference.

19 339. Runway distributes Stable Diffusion 1.5 under the “CreativeML Open RAIL-M”
20 license, which allows anyone to download, use, and deploy Stable Diffusion 1.5 for free. 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>).

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

26 341. Runway has made a material contribution to this infringing activity by training Stable
27 Diffusion 1.5 and then distributing it for free.
28

1 342. Runway intends to cause further infringement with Stable Diffusion 1.5. In February
2 2023, Stability CEO Mostaque said that Stable Diffusion 1.5 was “the most popular model by far by
3 [a] for profit company.”⁵³

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

7 COUNT ELEVEN

8 **DMCA violations by removing and altering CMI of training images** 9 **against Runway on behalf of all Plaintiffs, the Damages and Injunctive Classes**

10 344. The preceding factual allegations are incorporated by reference.

11 345. The LAION-5B Plaintiffs included one or more forms of copyright-management
12 information (“CMI”) (as defined in Section 1202(c) of the DMCA) in each of their respective works
13 in the LAION-5B Works, including captions in image-text pairs, and distinctive marks such as URLs
14 to personal webpages, signatures, and watermarks.

15 346. Runway did not contact Plaintiffs and the Class to obtain authority to remove or alter
16 CMI from their works within the meaning of the DMCA.

17 347. Runway knew the LAION-5B dataset contained CMI. The LAION-5B dataset
18 includes a detection score for watermarks which indicates the likelihood of a particular image in the
19 dataset contains a watermark or other distinctive mark signaling the presence of CMI. Runway thus
20 could have trained the Runway Models on images free of CMI but chose not to because images with
21 CMI tend to be high-quality.

22 348. Runway had access to but were not licensed by Plaintiffs or the Class to incorporate
23 their works in the LAION-5B dataset into the Runway Models.

24 349. Runway had access to but were not licensed by Plaintiffs or the Class to create copies
25 based on their works in the LAION-5B dataset into the Runway Models.

26
27
28 ⁵³ <https://twitter.com/EMostaque/status/1629514395825983489>

1 350. Runway had access to but were not licensed by Plaintiffs or the Class to distribute
2 their works in the LAION-5B dataset as Runway does through the Runway Models.

3 351. Without the authority of the LAION-5B Plaintiffs, Runway directly copied the
4 LAION-5B Works and used these Statutory Copies as training data for the Runway Models. The
5 works copied by Runway included CMI, including in the form of distinctive marks such as
6 watermarks or signatures. The training process is designed to remove or alter CMI from the training
7 images. Therefore, Runway intentionally removed or altered CMI from the LAION-5B Works in
8 violation of 17 U.S.C. § 1202(b)(1).

9 352. Runway distributes Stable Diffusion 1.5 subject to the “CreativeML Open RAIL M
10 License” (*See, e.g.* — <https://github.com/runwayml/stable-diffusion/blob/main/LICENSE>). As
11 alleged above, Stable Diffusion 1.5 and the other Runway Models themselves constitute Statutory
12 Copies of the LAION-5B works or Statutory Derivative Works. The license distributed by Runway
13 asserts that copyright in the model belongs to “Robin Rombach and Patrick Esser and contributors.”
14 By asserting that these third parties have copyright in the Stable Diffusion 1.5, which infringe the
15 copyrights of the LAION-5B Plaintiffs, Runway is providing and distributing false CMI in violation
16 of 17 U.S.C. § 1202(a).

17 353. As demonstrated herein, the Runway Models generate output that are copies of
18 original images with CMI with the CMI removed and/or altered.

19 354. Runway knows that the Runway Models are being used by users and/or licenses of
20 the Runway Models to create infringing copies of Plaintiffs and Class Members’ works. Indeed, one
21 of the reasons for the Runway Models’ popularity is because of the models’ ability to mimic or
22 imitate artists whose works are in the LAION-5B dataset. Thus, Runway knew or reasonably should
23 have known that the Runway Models removal and alteration of CMI would induce, enable, facilitate,
24 or conceal further infringement.

25 355. The LAION-5B Plaintiffs have been injured by Runway’s removal or alteration of
26 CMI. The LAION-5B Plaintiffs have been injured by Runway’s falsification of CMI by claiming false
27 copyright in Stable Diffusion 1.5. These Plaintiffs are entitled to statutory damages, actual damages,
28

1 restitution of profits, and other remedies provided by law.

2 XVIII. CAUSES OF ACTION AGAINST DEVIANTART

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

7 357. Plaintiffs Kelly McKernan, Hawke Southworth, Jingna Zhang, and Grzegorz
8 Rutkowski are DeviantArt users. Below, they are called the **DeviantArt Plaintiffs**.

9 358. On November 9, 2022, DeviantArt released DreamUp, an AI image product.
10 DeviantArt claims that DreamUp “lets you create AI art knowing that creators and their work are
11 treated fairly.” DreamUp is only available to paying customers of DeviantArt. DeviantArt offers paid
12 subscriptions to its members called “Core Plans.” Custom Core Plans typically range in price from
13 \$3.95 to \$14.95 per month. To use DreamUp, a member must first subscribe to a Core Plan. A Core
14 Plan subscriber is allowed to use DreamUp for a certain number of Text Prompts per month. For
15 instance, the \$9.95 “Pro” level permits 200 DreamUp Text Prompts per month. Core Plan
16 members can purchase additional Text Prompts by purchasing packages of “points.” DeviantArt
17 charges \$1 for 80 points, with a minimum purchase of 400 points for \$5.

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

21 360. DeviantArt is the source of millions of images in the LAION-5B dataset. Users of the
22 LAION-5B dataset have copied these millions of images many times over by downloading them from
23 DeviantArt.

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

27 362. Each of the DeviantArt Plaintiffs has stored images on DeviantArt that were later
28

1 incorporated into the LAION-5B dataset. These images were therefore copied by Runway and
2 Stability to train versions of Stable Diffusion. **Exhibit A** contains a sampling—but not an exhaustive
3 listing—of images created by the DeviantArt Plaintiffs that are contained in LAION-5B and were
4 copied from DeviantArt. They can be identified through their LAION-5B URL, which comes from
5 the “wixmp.com” domain. This domain is used by DeviantArt to store member images. (“Wix” in
6 the domain name refers to the parent company of DeviantArt.)

7 363. DreamUp relies on Stable Diffusion to produce images. The DreamUp app
8 incorporates a copy of Stable Diffusion. The terms of service for DreamUp do not disclose the
9 specific version of Stable Diffusion that is incorporated within the app.

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

15 365. CompVis is the shorthand name of the Computer Vision and Learning Group at
16 Ludwig Maximilian University in Munich, where the original research underlying Stable Diffusion
17 was first conducted. According to the GitHub page for Stable Diffusion 1.4, “Stable Diffusion was
18 made possible thanks to a collaboration with Stability AI and Runway.”⁵⁴

19 366. According to CompVis, Stable Diffusion 1.4 “was trained on a large-scale dataset
20 LAION-5B.”⁵⁵

21 367. The LAION-5B dataset contains only URLs of training images, not the actual
22 training images. Therefore, anyone who wishes to use LAION-5B for training their own machine-
23 learning model must first acquire copies of the actual training images from their URLs by using the
24 ‘img2dataset’ tool or another similar tool. Consistent with this, in preparation for training Stable
25 Diffusion 1.4, CompVis made one or more Statutory Copies of the LAION-5B Registered Works so
26

27 ⁵⁴ See <https://github.com/CompVis/stable-diffusion>

28 ⁵⁵ See <https://huggingface.co/CompVis/stable-diffusion-v1-4>

1 they could be fed to Stable Diffusion 1.4 as training data. The Statutory Copies made of each
2 registered work were substantially similar to that registered work.

3 368. During the training of Stable Diffusion 1.4, CompVis made a series of intermediate
4 Statutory Copies of the LAION-5B Registered Works. For instance, diffusion models are trained by
5 creating “noised” copies of training images, as described herein, all of which qualify as Statutory
6 Copies. The intermediate Statutory Copies of each registered work that CompVis made during
7 training of Stable Diffusion 1.4 were substantially similar to that registered work.

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

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

17 370. Therefore, like Stable Diffusion 1.5, Stable Diffusion 1.4 also qualifies as an infringing
18 Statutory Copy of the LAION-5B Registered Works. Because Stable Diffusion 1.4 represents a
19 transformation of the LAION-5B Registered Works into an alternative form, Stable Diffusion 1.4
20 also qualifies as an infringing Statutory Derivative Work.

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

24 DeviantArt responds that:

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

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

4 DeviantArt lets you declare whether or not external AI models and
 5 platforms can train based on your deviations. When submitting a
 6 deviation, you'll be able to check a box that informs third parties
 7 whether or not you authorize that submission being included in
 datasets used to train AI models like AI image generators.⁵⁶

8 372. DeviantArt's answer is misleading. As confirmed by the FAQ, while DeviantArt did
 9 not "add" images to the training sets of DreamUp, it made no mention of any images already in the
 10 training set for DreamUp's underlying models. DeviantArt knew that Stable Diffusion had already
 11 been trained on images scraped from DeviantArt itself. DeviantArt thus misled its community
 12 because art from DeviantArt was already in DreamUp because Stable Diffusion had already been
 13 trained on them.

14 373. This has been further confirmed by DeviantArt CTO Chris Nell. In November 2022,
 15 on the public LAION Discord server, Nell described himself as "one of the people at DeviantArt
 16 working on improving acceptance of AI generated/augmented art in the broader online arts
 17 community" and added "I think our goals at DA [DeviantArt] are very aligned with LAION's
 18 ... and want to collaborate as much as possible."⁵⁷ Nell said of DreamUp: "we did not fine tune
 19 [meaning, perform additional training on] SD [= Stable Diffusion] at all, so there aren't novel
 20 weights to share. [W]e do perform additional guidance at generation time ... so it's not exactly
 21 unmodified SD [= Stable Diffusion] output, but that is more akin to prompt tuning."⁵⁸ As confirmed
 22 by Nell, DeviantArt was well aware of how Stable Diffusion was developed and did not do any fine-
 23 tuning of the weights included in the Stable Diffusion model DreamUp was based on. In other
 24

25 ⁵⁶ <https://www.deviantartsupport.com/en/dreamup>

26 ⁵⁷

27 <https://discord.com/channels/823813159592001537/1006139459860975716/1042539656396411004>

28 <https://discord.com/channels/823813159592001537/1006139459860975716/1042543837425438804>

1 words, this implies all of the images copied in training Stable Diffusion were included in the
2 DreamUp model.

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

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

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

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

27 376. DeviantArt's embrace of generative AI art was seen as a betrayal by its art
28 community.

377. The scope of DeviantArt's betrayal of its artist community by embracing Stable
Diffusion was evident in a group audio session held by DeviantArt management on November 11,
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

1 AI art was a complete repudiation of its longstanding community principles, as well as economically
2 and legally unfair.

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

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

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

25 381. Furthermore, although the new “Data Scraping” provision acknowledges that certain
26 kinds of data scraping will continue to be an “unauthorized use” of the DeviantArt website, that
27 “owners of the works are responsible for policing their own works.” In other words, despite its
28

1 professed interest in using its terms of service to protect artists, DeviantArt is washing its hands of
2 the matter. Instead of standing up for artists and using its resources to combat illegal AI data
3 scraping, it is forcing artists to take matters into their own hands.

4 382. What is more, while DeviantArt purported to spearhead a system for artists to opt-
5 out of having their works trained upon, these promises are mostly hollow.

6 383. DeviantArt's proposal for artists to opt out was to utilize a system of HTML tags.
7 Artists who do not wish to have their content used for AI training can append the "noai" and
8 "noimageai" hashtags to the HTML page associated with their art.

9 384. This promise is misleading.

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

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

COUNT TWELVE

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

387. The preceding factual allegations are incorporated by reference.

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

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

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

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

392. The LAION-5B Registered Plaintiffs have been and continue to be injured by DeviantArt’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.

XIX. JURY TRIAL DEMANDED

Pursuant to Federal Rule of Civil Procedure 38(b), Plaintiffs demand a trial by jury of all the claims asserted in this Complaint so triable.

1 Dated: October 31, 2024

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

4 Joseph R. Saveri (State Bar No. 130064)
5 Cadio Zirpoli (State Bar No. 179108)
6 Christopher K.L. Young (State Bar No. 318371)
7 Elissa A. Buchanan (State Bar No. 249996)
8 David Lerch (State Bar No. 229411)
9 Evan Creutz (State Bar No. 349728)
10 **JOSEPH SAVERI LAW FIRM, LLP**
11 601 California Street, Suite 1505
12 San Francisco, California 94108
13 Telephone: (415) 500-6800
14 Facsimile: (415) 395-9940
15 Email: jsaveri@saverilawfirm.com
16 czirpoli@saverilawfirm.com
17 cyoung@saverilawfirm.com
18 eabuchanan@saverilawfirm.com
19 dlerch@saverilawfirm.com
20 ecreutz@saverilawfirm.com

21 Matthew Butterick (State Bar No. 250953)
22 1920 Hillhurst Avenue, #406
23 Los Angeles, CA 90027
24 Telephone: (323) 968-2632
25 Facsimile: (415) 395-9940
26 Email: mb@buttericklaw.com

27 Brian D. Clark (*pro hac vice*)
28 Laura M. Matson (*pro hac vice*)
Arielle S. Wagner (*pro hac vice*)
Eura Chang (*pro hac vice*)
LOCKRIDGE GRINDAL NAUEN P.L.L.P.
100 Washington Avenue South, Suite 2200
Minneapolis, MN 55401
Telephone: (612)339-6900
Facsimile: (612)339-0981
Email: bdclark@locklaw.com
lmmatson@locklaw.com
aswagner@locklaw.com
echang@locklaw.com

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