

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 Travis Manfredi (State Bar No. 281779)
JOSEPH SAVERI LAW FIRM, LLP
 601 California Street, Suite 1000
 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
 tmanfredi@saverilawfirm.com

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

12 [Additional Counsel Listed on Signature Page]

13 UNITED STATES DISTRICT COURT
 14 NORTHERN DISTRICT OF CALIFORNIA
 15 SAN FRANCISCO DIVISION

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

22 Individual and Representative Plaintiffs,

23 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;
 Runway AI, Inc., a Delaware corporation;

28 Defendants.

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

FIRST 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	
4	COPYRIGHTED IMAGES WITHOUT CONSENT, CREDIT, OR COMPENSATION	
5	AND VIOLATE 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	VII. ARTISTS AND THEIR WORKS.....	11
12	VIII. THE SOURCE OF THE TRAINING DATASETS: LAION	13
13	IX. HOW AI IMAGE PRODUCTS WORK: CLIP-GUIDED DIFFUSION	20
14	X. PROTECTED EXPRESSION FROM TRAINING IMAGES IS COPIED,	
15	COMPRESSED, STORED, AND INTERPOLATED BY DIFFUSION MODELS	27
16	XI. EXAMPLES OF TEXT PROMPTS USING PLAINTIFF NAMES IN AI IMAGE	
17	PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY	35
18	XII. EXAMPLES OF IMAGE PROMPTS USING PLAINTIFF IMAGES IN AI IMAGE	
19	PRODUCTS OFFERED BY STABILITY, RUNWAY, AND MIDJOURNEY	41
20	XIII. USER AND LICENSEE ACTIVITY	50
21	XIV. DEFINITIONS FOR THE CAUSES OF ACTION.....	52
22	XV. CAUSES OF ACTION AGAINST STABILITY	52
23	XVI. CAUSES OF ACTION AGAINST MIDJOURNEY.....	61
24	XVII. CAUSES OF ACTION AGAINST RUNWAY	76
25	XVIII. CAUSES OF ACTION AGAINST DEVIANTART	81
26	XIX. JURY TRIAL DEMANDED	93
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,
5 Inc. (collectively “Stability”); Midjourney, Inc. (“Midjourney”); and DeviantArt, Inc.
6 (“DeviantArt”) (all collectively “Defendants”). Plaintiffs allege various violations of the
7 Copyright Act (17 U.S.C. § 501); Digital Millennium Copyright Act (17 U.S.C. §§ 1202(a) & (b))
8 as to all Defendants; violations of the Lanham Act (15 U.S.C. § 1125(a)(1)) as to Defendant
9 Midjourney; and breach of contract as to Defendant DeviantArt.

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

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

22 2. Under the hood, AI image products are powered by one or more *machine-learning*
23 *models* (or simply *models*). Models are not programmed directly in the manner of traditional
24 software, but rather *trained*. Training a model first requires amassing a huge corpus of data, called
25 a *dataset*. The models at issue in this complaint were trained on datasets containing billions of
26 images paired with descriptive captions. In this complaint, each image–caption pair is called a
27 *training image*. During training of the model, the training images in the dataset are directly copied
28

1 in full and then completely ingested by the model, meaning that protected expression from every
2 training image enters the model. As it copies and ingests billions of training images, the model
3 progressively develops the ability to generate outputs that mimic the protected expression copied
4 from the dataset. The outputs of a model are derived entirely and exclusively from what it has
5 extracted from the dataset.

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

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

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

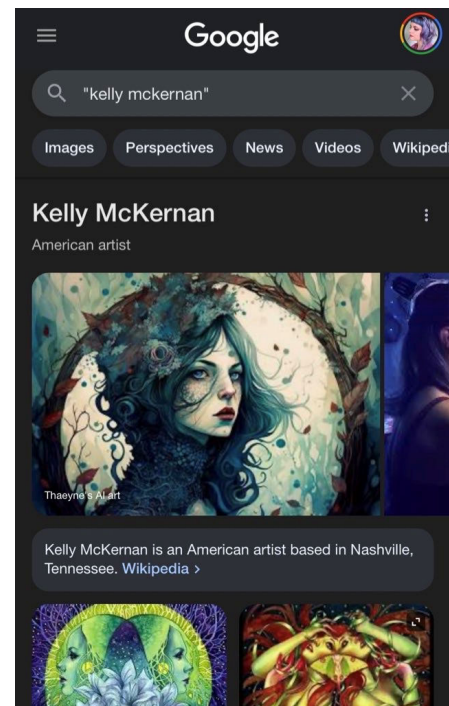
28 ² <https://discord.com/channels/729741769192767510/730095596861521970/1008530914525061190>

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

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

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

14 8. Worst of all, the Defendants hold out their AI image products as being able to
15 create substantially similar substitutes for the very works they
16 were trained on—either specific training images, or images
17 that imitate the trade dress of particular artists—including
18 Plaintiffs. This is already damaging the market for Plaintiffs'
19 artwork and labor, and the art market more broadly.
20 Midjourney, for instance, has repeatedly promoted the use of
21 artist names—including Plaintiffs' names—within text
22 prompts as a means of getting better results. Runway,
23 Stability, and Midjourney also encourage the use of images—
24 including images made by Plaintiffs—as a means of
25 prompting their AI image products. Recently, plaintiff Kelly
26 McKernan was astonished to find that the top internet search



1 result for their name is now an AI-generated image made with Midjourney, prompted with Mx.
2 McKernan's name.³ Without intervention, this is the grim future that awaits many other artists.

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

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

8 **II. JURISDICTION AND VENUE**
9

10 11. This Court has subject-matter jurisdiction under 28 U.S.C. § 1331 because this
11 case arises under the Copyright Act (17 U.S.C. § 501) and the Digital Millennium Copyright Act
12 (17 U.S.C. § 1202).

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

21 13. Under Civil Local Rule 3.2(c) and (e), assignment of this case to the San Francisco
22 Division is proper because defendant Midjourney is headquartered in San Francisco, and thus a
23 substantial part of the events giving rise to Plaintiffs' claims and the interstate trade and
24 commerce involved and affected by Defendants' illegal conduct occurred in this Division.
25
26
27

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

1 **III. PLAINTIFFS**

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

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

9 16. Karla Ortiz lives in California and owns the copyrights in training images shown
10 under her name in Exhibit A. Ms. Ortiz’s copyright registrations for certain works are included in
11 Exhibit C.⁴

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

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

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

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

24 21. Jingna Zhang lives in the state of Washington and owns the copyrights in training
25 images shown under her name in Exhibit A and Exhibit B. Ms. Zhang’s copyright registrations for
26

27 _____
28 ⁴ 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
4 under her name in Exhibit A. Ms. Kaye’s copyright registrations for these works are included in
5 Exhibit C.

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

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

23 IV. DEFENDANTS

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

1 the CompVis research group at Ludwig Maximilian University in Munich, where he was a
2 principal developer of 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
8 San 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
11 Angelo Sotira, Scott Jarkoff, and Matthew Stephens. In 2017, Wix.com, Inc. acquired
12 DeviantArt. Wix acquired all of DeviantArt's corporate stock for \$36 million. In April 2022, Moti
13 Levy became CEO 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
17 is 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
22 were authorized, ordered, or performed by the Defendants' respective officers, agents,
23 employees, representatives, or shareholders while actively engaged in the management, direction,
24 or control of the Defendants' businesses or affairs. The Defendants' agents operated under the
25 explicit and apparent authority of their principals. Each Defendant, and its subsidiaries, affiliates,
26 and agents operated as a single unified entity.

33. Various persons or firms not named as Defendants may have participated as co-conspirators in the violations alleged herein and may have performed acts and made statements in furtherance thereof. Each acted as the principal, agent, or joint venture of, or for other Defendants with respect to the acts, violations, and common course of conduct alleged herein.

VI. CLASS ALLEGATIONS

A. Class Definitions

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

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

All persons or entities nationalized or domiciled in the United States that own a copyright interest in any work 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.

“Damages Class” under Rule 23(b)(3):

All persons or entities nationalized or domiciled in the United States that own a copyright interest in any work 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.

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

All persons or entities nationalized or domiciled in the United States that own a registered copyright in any work in the LAION-5B dataset 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.

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

All persons or entities nationalized or domiciled in the United States 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 was offered directly or incorporated into another product by one or more Defendants during the Class Period.

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

All members of the Damages Class who (1) maintained an account on DeviantArt; (2) posted copyrighted work on DeviantArt; and (3) had that work used to train any version of an AI image product.

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

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

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

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

13 **B. Numerosity**

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

18 **C. Typicality**

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

1 **D. Commonality & Predominance**

2 37. 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 38. **Direct Copyright Infringement:** Whether Defendants violated the copyrights of
5 Plaintiffs and the Class when they downloaded and stored copies of the Works; Whether
6 Defendants violated the copyrights of Plaintiffs and the Class when they used copies of the Works
7 to train AI Image Products.

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

12 40. **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 41. **Lanham Act Violations:** Whether Defendants misappropriated Plaintiffs and
16 Class members' trade dress and distinctive look and feel in violation of the Lanham Act.

17 42. **Common Law Violations:** Whether the use of Plaintiffs and Class members'
18 works to train, develop, and promote Defendants AI Image Products constitute an unjust benefit
19 conferred upon Defendants to Plaintiffs' detriment.

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

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

24 **E. Adequacy**

25 45. Plaintiffs will fairly and adequately represent the interests of the Class because
26 they have experienced the same harms as the Class and have no conflicts with any other members
27 of the Class. Furthermore, Plaintiffs have retained sophisticated and competent counsel ("Class
28

1 Counsel”) who are experienced in prosecuting federal and state class actions throughout the
2 United States and other complex litigation and have extensive experience advising clients and
3 litigating intellectual property, competition, contract, and privacy matters.

4 **F. Other Class Considerations**

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

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

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

14 **VII. ARTISTS AND THEIR WORKS**

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

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

22 51. Copyright law protects artists’ works from infringement by creating *exclusive*
23 rights of artists to make copies of their works, to make derivative works of their copyrighted
24 works, and to distribute such copies and derivative works, which protects the interests of artists
25 and preserves the incentives for humans to producing art. As recognized by no higher source than
26 the United States Constitution itself, copyright law is intended to “promote the Progress of
27 Science and useful Arts.” By offering artists protection, they can be rewarded for their efforts. In
28

1 accordance with copyright law, many artists, including certain Plaintiffs, register copyrights in
2 their works.

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

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

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



9
10 Ex. A at p. 13 (Plaintiff Brom's personal website URL—lower left corner):



11
12 Ex. A at p. 62 (Plaintiff Zhang's personal website URL—lower-right corner):



13
14 Ex. A at p. 79 (Plaintiff Ortiz's signature—lower-right corner):



15
16 Ex. A at p. 86 (Plaintiff McKernan's mark—lower right corner):



17
18 Ex. A at p. 103 (Plaintiff Andersen's signature—lower right corner):



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

26 55. Machine-learning models depend on massive quantities of high-quality data that is
27 digitally copied without authorization to train the model. The quality of the dataset a model is
28

1 trained on determines the quality of the model itself. The models within AI image products are
2 no different. As a result, the images these models are trained on have an enormous impact on the
3 quality of the models and the outputs they can produce. There are many images that are in the
4 public domain, i.e., images that are not subject to copyright. The developers of many AI image
5 products, however, made the decision to greatly expand the datasets their models are trained on
6 by including billions of protected works. They made this choice because many desirable works are
7 not in the public domain, i.e., are not subject to unfettered use by anyone without consent, credit,
8 or compensation.

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

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

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

25 58. LAION's most well-known projects are the datasets of training images it has
26 released for training machine-learning models, which are now widely used in the AI industry.
27
28

1 59. In August 2021, LAION released LAION-400M, a dataset of 400 million training
2 images assembled from images accessible on the public internet. At the time, LAION-400M was
3 the largest freely available dataset of its kind. LAION distributes the LAION-400M dataset to the
4 public through its own website and elsewhere. Information about LAION-400M is available in an
5 accompanying paper by Schuhmann, Beaumont, and others titled “LAION-400M: Open Dataset
6 of CLIP-Filtered 400 Million Image-Text Pairs,” released in November 2021 (hereafter, the
7 “LAION-400M Paper”).⁵

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

13 61. The actual images referenced in the LAION-400M dataset records are not
14 included with the dataset. Anyone who wishes to use LAION-400M for training their own
15 machine-learning model must first acquire copies of the actual images from their URLs. To
16 facilitate the copying of these images, LAION provides a software tool called `img2dataset`⁶ that
17 takes the metadata records as input and makes copies of the referenced images from the URLs in
18 each metadata record, thereby creating local copies.

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

25 63. One of the entities that has made unauthorized copies of the LAION-400M
26 training images is LAION itself. According to the LAION-400M Paper, LAION made the dataset

27 _____
⁵ <https://arxiv.org/abs/2111.02114>

28 ⁶ <https://github.com/rom1504/img2dataset>

1 by starting with Common Crawl metadata records. Common Crawl is a corpus of 250 billion web
2 pages copied from the public web, including assets like Plaintiffs’ images
3 (<https://commoncrawl.org/>). The metadata records contain web URLs. According to the
4 LAION-400M Paper, LAION created training images by first “pars[ing] through [the metadata
5 records] from Common Crawl and pars[ing] out all HTML IMG tags containing an alt-text
6 attribute [that is, a text caption].” Then, LAION “download[ed] the raw images from the parsed
7 URLs”. LAION-400M Paper at 3. To ensure that the training images in the dataset had
8 reasonably accurate captions, LAION used a CLIP model to calculate the CLIP embeddings for
9 the image and text of each image–text pair. These two CLIP embeddings were compared to
10 measure how well the text described the image. Image–text pairs with low CLIP-similarity scores
11 were omitted from the dataset.

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

18 65. In October 2022, LAION released LAION-5B, a dataset of 5.85 billion training
19 images—more than 14 times bigger than LAION-400M. Information about LAION-5B is
20 available in an accompanying paper called “LAION-5B: An open large-scale dataset for training
21 next generation image-text models,” by Schuhmann, Beaumont, Crowson, and others (hereafter,
22 the “LAION-5B Paper”).¹⁰ According to the LAION-5B Paper, LAION-400M is a subset of
23 LAION-5B, meaning every image in LAION-400M is also in LAION-5B.

24 66. Much like the LAION-400M dataset, when one downloads the LAION-5B
25 dataset, one gets a list of metadata records, one for each training image. Each record includes the

26 ⁷ <https://discord.com/channels/662267976984297473/938713143759216720/954674533942591510>

27 ⁸ <https://twitter.com/EMostaque/status/1559332564787240962>

28 ⁹ <https://twitter.com/EMostaque/status/1629516125150011394>

¹⁰ <https://arxiv.org/abs/2210.08402>

1 URL of the image, the image caption, the similarity of the caption and image (as measured by
2 their respective CLIP embeddings), the width and height of the image, and other metadata fields.
3 Each record also includes a watermark detection score, which indicates the likelihood an image
4 possesses a distinctive mark of a copyright holder, e.g., the artist.

5 67. Just like the LAION-400M dataset, the actual images referenced in the LAION-
6 5B dataset records are not included with the dataset. Anyone who wishes to use LAION-5B for
7 training their own machine-learning model must first acquire copies of the actual images from
8 their URLs. To facilitate the copying of these images, LAION provides a software tool called
9 `img2dataset`¹¹ that takes the metadata records as input and makes copies of the referenced
10 images from the URLs in each record, thereby creating local copies.

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

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

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

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

¹¹ <https://github.com/rom1504/img2dataset>

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

5 72. Though LAION-400M and LAION-5B are often used to train diffusion models,
6 they are used to train other models as well. For instance, Stability sponsored LAION’s creation of
7 OpenCLIP, a CLIP model trained on a subset of images from LAION-5B called LAION-2B
8 consisting of training images with English-language captions.¹² On information and belief,
9 because all the Plaintiff images in Exhibit A (Plaintiff images in LAION-5B) have English-
10 language captions, they are also part of LAION-2B and were thus used to train OpenCLIP. To
11 create OpenCLIP, LAION necessarily had to create one or more copies of these images. Plaintiffs
12 never gave their permission to LAION to copy their images or use them to train OpenCLIP.

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

19
20
21
22
23
24
25
26
27
28

¹² <https://huggingface.co/laion/CLIP-ViT-H-14-laion2B-s32B-b79K>

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



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

Backend url: <https://rom1504.github.io/clip-retrieval/>
 Index: laion5B-H-14

0.8704 Serf and Greyhounds 8.5 x10 Graphite Moleskin Show II, Spoke Art Gallery, 2012

0.8637 Karla_Ortiz_Concept_Art_1

0.7958 Serf and Greyhounds id 8.5 x10 Graphite Moleskin Show II, Spoke Art Gallery, 2012

0.6608 The Art of Karla Ortiz: Drawing process.

0.6331 ilustración de Jennifer A. Hell

0.6316 E: lynx ladyandbeast

0.6257 ""Vania Zouravliov / Personal Work / Kitsune: slide numbersslide_number1</5>

0.6220

0.6200

0.6196 Theodor Severin

0.6186 ""Czerwony Kapturek"" ilustracja dla dzieci Joanny Concejo

0.6185 shockyou-jma-brice-sketchbooks?

0.6166

0.6156 Borzoi Oil based graphite on Italian laid paper 2015 A meeting of two beings in the woods

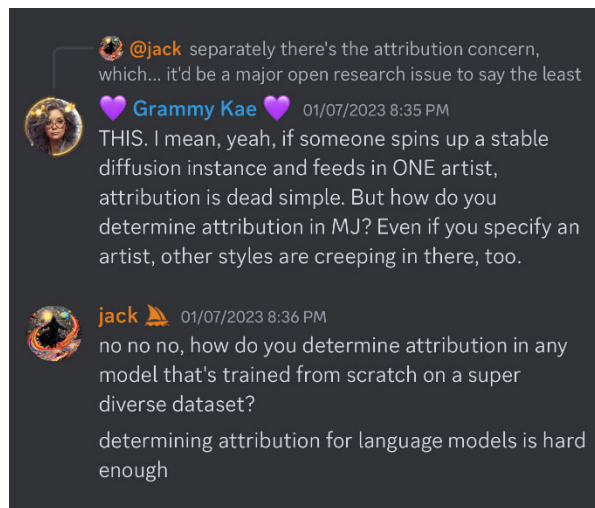
76. The caption shown for the first image in the search results of the LAION dataset is exactly the same as the caption shown on Plaintiff Ortiz’s website: “Serf and Greyhounds / 8.5 x 10 Graphite / Moleskin Show II, Spoke Art Gallery, 2012”.

77. Further, inspection of the metadata given by <https://rom1504.github.io/clip-retrieval/> indicates that the image was scraped from the following URL: <https://images.squarespace-cdn.com/content/v1/510a0982e4b08fd84ce45a43/1359614369317->

1 Q1QXIFKXQCVCO4I62D06/ke17ZwdGBToddI8pDm48kNFwqbaKfT7OPsXFUnn0nBkUqsx
 2 RUqqbr1mOJYKfIPR7LoDQ9mXPOjoJoqy81S2I8N_N4V1vUb5AoIIIbLZhVYxCRW4BPu10St
 3 3TBAUQYVKczZ8BZeDbXUHhGUs_1S_OvE6uym2C-
 4 ge4vqvF4L8FpMvaIRyuEhmGLRxo5gMFxPRT/drawing_serfandhound.jpg, which indicates the
 5 source of the image was Ms. Ortiz's own website. In other words, the metadata indicates that this
 6 image was scraped from Ms. Ortiz's personal website for inclusion into the LAION datasets.

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

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



24 80. Stability CEO Emad Mostaque has publicly acknowledged the importance of using
 25 licensed training images, saying that future versions of Stable Diffusion would be based on “fully
 26 licensed” training images.¹³ But so far, Stability has taken no steps to negotiate suitable licenses.

27

28 ¹³@EMostaque, Twitter (Dec. 15, 2022, 8:03 AM), <https://twitter.com/EMostaque/status/1603390169192833027>.

1 Neither has Runway. Neither has Midjourney. They all just use LAION datasets—with no
2 consent, no credit, and no compensation to the artists.

3 81. In July 2023, the topic of AI training reached the U.S. Senate.¹⁴ During a hearing
4 convened by a subcommittee of the Senate Judiciary Committee, Sen. Mazie Hirono quizzed Ben
5 Brooks, a representative from Stability, about Stability’s position on licensing training data. Sen.
6 Hirono asked directly, “So basically you don’t pay for the data that you put into your—to train
7 your models?” Mr. Brooks replied, “There is no arrangement in place.” Sen. Hirono then turned
8 to plaintiff Karla Ortiz, who was testifying on the same panel. “So you have Ms. Ortiz, who says
9 that that is wrong. Is that correct, Ms. Ortiz?” Ms. Ortiz replied, “A hundred percent, Senator.”

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

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

15 83. Stable Diffusion is a model that generates images using CLIP-guided diffusion.
16 Therefore, AI image products based on Stable Diffusion—including those offered by Runway,
17 Stability, and DeviantArt—rely on CLIP-guided diffusion as well. In June 2022, a Midjourney
18 developer confirmed on their public message board that “we use clip guided diffusion” too.¹⁵
19 Thus, the description below applies to all the AI image products at issue in this complaint.

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

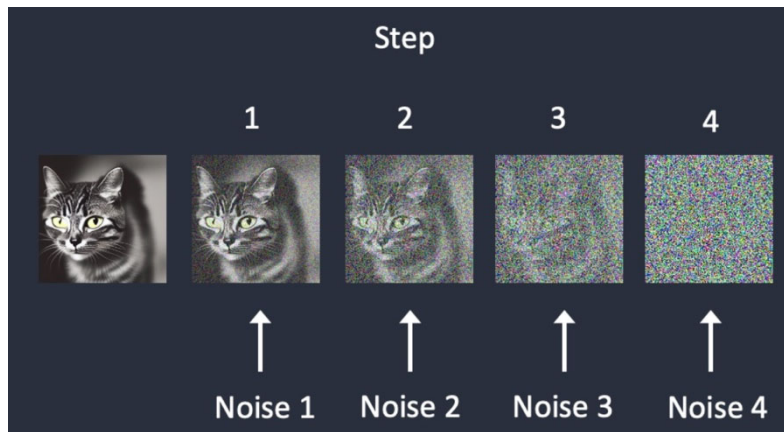
25
26
27 ¹⁴ <https://youtu.be/uoCJun7gkbA?t=3578>

¹⁵ <https://discord.com/channels/662267976984297473/938713143759216720/982136076888068156>

28 ¹⁶ <https://arxiv.org/abs/1503.03585>

85. Diffusion proceeds in two phases. To begin the training phase, initial copies are made of many training images. For each training image, progressively more noise is added over a series of steps. At each step, the model records how the addition of noise changes the image. By the last step, the image has been “diffused” into essentially random noise. A simplified version of this process is shown in the diagram below:¹⁷


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 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.



¹⁷ Noising & denoising illustrations from https://stable-diffusion-art.com/how-stable-diffusion-work/#How_training_is_done

1 90. In sum, diffusion is a way for a machine-learning model to calculate how to
2 reconstruct a copy of its training images. For each training image, a diffusion model finds the
3 sequence of denoising steps to reconstruct that specific image. Then it stores this sequence of
4 steps. In practice, this training would be repeated for many images—likely millions or billions. A
5 diffusion model is then able to reconstruct copies of each training image based on this denoising
6 process. Furthermore, being able to reconstruct copies of the training images is not an incidental
7 side effect. The primary objective of a diffusion model is to reconstruct copies of its training
8 images with maximum accuracy and fidelity. Or in the words of prominent machine-learning
9 researcher Nicholas Carlini, who has studied the behavior of diffusion models: “diffusion models
10 are explicitly trained to reconstruct the training set.”¹⁸

11 91. In December 2020, the diffusion technique was improved by a team of researchers
12 at UC Berkeley led by Jonathan Ho. These ideas were introduced in their paper “Denoising
13 Diffusion Probabilistic Models.”¹⁹


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

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

27 ¹⁸ “Extracting Training Data from Diffusion Models,” p. 12, available at <https://arxiv.org/abs/2301.13188>

28 ¹⁹ 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 detail of the two interpolated images is shown below:

1 99.



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

11 101. In April 2022, the diffusion technique was further improved by a team of
12 researchers led by Robin Rombach at Ludwig Maximilian University of Munich. These ideas
13 were introduced in his paper “High-Resolution Image Synthesis with Latent Diffusion
14 Models.”²⁰ (Rombach is currently employed by Stability as a developer of Stable Diffusion.)

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

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

23 104. In 2021, researchers from OpenAI introduced the idea of a CLIP model in a paper
24 called “Learning Transferable Visual Models From Natural Language Supervision.”²¹ A CLIP
25 model quantifies the semantic correlation between images and captions.

26

27

²⁰ <https://arxiv.org/abs/2112.10752>

²¹ <https://arxiv.org/abs/2103.00020>

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

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

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

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

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

1 110. The CLIP model is able to do this by converting both images and text captions into
2 a 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
14 the image-denoising process toward a more dog-like result. Similarly, if a user includes the name
15 of plaintiff "Grzegorz Rutkowski" in the text prompt, then as the diffusion model iteratively
16 denoises the image, the CLIP embedding for the phrase "Grzegorz Rutkowski" nudges the
17 image-denoising 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
23 and diffusion model is called *CLIP-guided diffusion*, and is used by all the AI image products at
24 issue in 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,
27 since either can be converted into a CLIP embedding. Image prompts, however, tend to produce
28

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

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

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

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

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

1 system depends entirely on the data points it was constructed with. **You get back what you put**
2 **in (or interpolations of the same).**”²² (Emphases added.)

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

13 119. Though the framing of machine learning as a form of data compression has been
14 resisted by some, research shows an ever-stronger connection between the two, and between
15 diffusion models and compression in particular. In November 2023, a team of machine-learning
16 researchers led by Yaodong Yu at UC Berkeley published a paper called “White-Box
17 Transformers via Sparse Rate Reduction: Compression Is All There Is?”²³ (Below, the “Yu
18 Paper”.) In their paper, the authors describe in detail a strong mathematical and experimental
19 correspondence between diffusion models and data compression, and conclude by saying (italics
20 in original, bold emphasis added):

21
22
23
24
25
26
27
28

²² <https://twitter.com/fchollet/status/1563153087514419206>

²³ <https://arxiv.org/abs/2311.13110>

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

13 Yu Paper at 53.

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

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

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

25 123. In August 2022, Stability said in its launch announcement for Stable Diffusion that
 26 it “is the culmination of many hours of collective effort to create a single file that **compresses the
 27 visual information** of humanity into a few gigabytes.”²⁴

28 124. In September 2022, Mostaque said in a podcast interview: “Stable Diffusion is the
 model itself. ... 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.”²⁵

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

²⁵ <https://narrativespodcast.com/2022/09/19/112-emad-mostaque-ai-alignment-and-stable-diffusion/>

1 125. In January 2023, Mostaque said in a podcast interview: “We took 100,000
2 gigabytes of image-label pairs—2 billion images—and created a 1.6 gigabyte file ... that **basically**
3 **compresses the visual information** of a snapshot of the internet.”²⁶

4 126. In February 2023, Mostaque said in a podcast interview: “We’ve created the **most**
5 **efficient compression** in the world.”²⁷

6 127. In May 2023, Mostaque said to a tech journalist that Stable Diffusion is “a
7 hundred thousand gigabytes of **images compressed** to a two-gigabyte file.”²⁸

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

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

15 130. This topic was explored in a January 2023 paper called “Extracting Training Data
16 from Diffusion Models” by Nicholas Carlini of Google and others.²⁹ (Below, the “Carlini
17 Paper”.) Carlini is one of the world’s leading AI researchers. He often studies the security of
18 machine-learning models, in particular issues related to the privacy and security of training data
19 after it has been ingested into the model.

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

25
26 ²⁶ <https://www.youtube.com/watch?v=jgTv2W0mUP0>

27 ²⁷ <https://sarahguo.com/blog/emadmostaque>

28 ²⁸ <https://www.zdnet.com/article/why-open-source-is-essential-to-allaying-ai-fears-according-to-stability-ai-founder/>

²⁹ <https://arxiv.org/abs/2301.13188>

132. Carlini’s experiment involved supplying text prompts to Stable Diffusion 1.4 to see if the prompts could elicit images essentially identical to those found in the training dataset. In many instances Carlini was able to coax Stable Diffusion 1.4 to output copies of its training images. An example from the paper is shown below, comparing certain training images with images output by Stable Diffusion. In each case, the image in the “Original” line is a training image; the image directly below in the “Generated” line is an image output from Stable Diffusion. As Carlini notes, the generated images are nearly identical to the training images:

133.



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

135. Carlini also poses a question: “[d]o large-scale models work by generating novel output, or do they just copy and interpolate between individual training examples?” Carlini Paper at 15. He concludes that “because our attacks [i.e., attempts to elicit stored copies of training images] succeed, this question remains open.” Carlini Paper at 15. François Chollet has taken an even stronger position, saying that “It’s accurate that generative art models create new content by recombining images from their training data.”³⁰ Yaodong Yu concurs, stating “large image generation models ... are unlikely to do anything beyond purely mechanical data compression (encoding) and interpolation (decoding).” Yu Paper at 53.

³⁰ <https://twitter.com/fchollet/status/1600230516934209536>

1 136. Carlini notes that a limitation of his experiment is that it relied on a very strict
2 “definition of ‘memorization’: whether diffusion models can be induced to generate” essentially
3 identical copies of certain training images “when prompted with appropriate instructions.”
4 Carlini Paper at 4. Carlini says later: “[o]ur work highlights the difficulty in defining *memorization*
5 ... a more comprehensive analysis will be necessary to accurately capture more nuanced
6 definitions of memorization that allow for more human-aligned notions of data copying.” Carlini
7 Paper at 15.

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

11 138. In May 2023, researcher Ryan Webster extended Carlini’s research in this
12 direction in a paper called “A Reproducible Extraction of Training Images from Diffusion
13 Models.”³¹ Webster found that by using a less strict technique for detecting stored copies of
14 training images, more instances of stored copies could be discovered. Webster tested several
15 diffusion models, including Stable Diffusion 2.0 and version 4 of the Midjourney Model, and
16 found stored copies of training images within all of them.

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

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

27
28

³¹ <https://arxiv.org/abs/2305.08694>

1 141. A related topic was explored in a July 2023 paper called “Measuring the Success
2 of Diffusion Models at Imitating Human Artists”³² by Stephen Casper of MIT and others.
3 (Below, the “Casper Paper”.)

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

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

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

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

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

24 147. Stable Diffusion is held out as an “open-source” program. But in the hands of
25 Runway and Stability, the term “open source” is more of a marketing and competitive gimmick
26 than a substantive virtue, intended mostly to ensure the widest distribution of Stable Diffusion,
27

28 ³² <https://arxiv.org/pdf/2307.04028.pdf>.

1 and the economic opportunities that result. David Widder and others strongly criticized this
 2 corruption of the traditional meaning of “open source” in an August 2023 paper called “Open
 3 (for Business): Big Tech, Concentrated Power, and the Political Economy of Open API”³³
 4 (emphasis added below):

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

21 Widder at 11.

22 148. If one downloads the Stable Diffusion 2.0 model from Stability via GitHub,³⁴ one
 23 does not get everything one needs to operate Stable Diffusion 2.0, let alone train a comparable
 24 model from scratch. Rather, one gets a set of scripts (mostly written in Python) and configuration
 25 files for generating images using a weights file (not included, some assembly required). A *weights*
 26 file is a binary file that encodes all the information that the model captured during training by
 27 copying protected expression from training images. An example of a Stable Diffusion weights file
 28 is available at <https://huggingface.co/stabilityai/stable-diffusion-2/blob/main/768-v-ema.ckpt>.
 This is a mass of binary data that is meaningful when accessed via the Stable Diffusion scripts,
 but otherwise not intelligible to humans. In that way, the weights file has a status similar to that of
 a videogame cartridge or DVD—it contains protected expression that can be perceived only with
 the aid of a device (in this case, the accompanying software scripts), which can only be seen when
 interacted with using the appropriate device, for example, a videogame console or a DVD player.

³³ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543807

³⁴ <https://github.com/Stability-AI/stablediffusion>

1 In other words, simply because certain code for a particular model is labeled “open source” does
2 not mean one can meaningfully interrogate the model or experiment with it. All the value of the
3 model is encapsulated in the weights file, and a weights file is unintelligible to human beings.

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

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

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

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

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

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

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

1 155. Plaintiffs Sarah Andersen, Kelly McKernan, and Karla Ortiz cannot show samples
2 of DreamStudio images with their names in the text prompts because their names have been
3 blocked since the initial complaint in this action was filed. Using one of these three names in
4 DreamStudio produces the error “Something isn’t quite right with your prompts.” Before
5 Plaintiffs Andersen, McKernan and Ortiz filed their initial complaint, however, their names could
6 be used as prompts to generate images.

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

10 157. Plaintiff Grzegorz Rutkowski cannot show samples of Stable Diffusion images with
11 his name in prompts. Due to the massive popularity of his name in Stable Diffusion prompts—
12 one report estimated Mr. Rutkowski’s name had been invoked “over 400,000 times”³⁵—Mr.
13 Rutkowski was also removed as a possible prompt.

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

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

- 20 a. Gregory Manchess is known for his classic oil paintings distinguished by
21 their handcrafted brushwork, calligraphic style, and realistic themes.
- 22 b. Gerald Brom is known for his gritty, dark, fantasy images that combine
23 classical realism, gothic, and countercultural aesthetics.
- 24 c. Jingna Zhang is known for her painterly and romantic photography, with
25 special attention to color, movement, and props. (Many of the subjects of
26

27
28 ³⁵ See “Greg Rutkowski Was Removed From Stable Diffusion, But AI Artists Brought Him Back,”
<https://decrypt.co/150575/greg-rutkowski-removed-from-stable-diffusion-but-brought-back-by-ai-artists>

1 Ms. Zhang’s photographs happen to be Asian, a fact that is incidental to
2 her work, but affects diffusion models in a peculiar way.)

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

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

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

26 163. **Runway:** Runway makes a text-to-image generator available via its online AI image
27 product called AI Magic Tools (<https://app.runwayml.com/>). This online app was used to make
28

1 the text-prompt examples below. On information and belief, Runway’s AI Magic Tools app uses
2 Stable Diffusion 1.5, because Runway trained that version of Stable Diffusion, and trained it on
3 the Plaintiff works in Exhibit A.

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

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

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

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

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

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

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

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

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

26
27
28

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

4 173. By default, the Midjourney Model layers onto every output image what it calls the
5 “Midjourney default aesthetic.” Because this “default aesthetic” is an overbaked visual style
6 supplied by designers at Midjourney, not the training images, it has been turned off in the
7 examples below by appending “--style raw” to each prompt listed, thereby more accurately
8 revealing the style changes that arise from changing the artist names.

9 174. The first set of Midjourney text prompts consists of “*chef*,” “*gregory manchess*
10 *chef*,” “*gerald brom chef*,” “*jingna zhang chef*,” “*kelly mckernan chef*,” and “*sarah andersen chef*.”
11 Exhibit F, pp. 1–6. The initial “*chef*” prompt shows what the Midjourney Model produces by
12 default, without an artist name in the prompt. Much like the Stability results, the Manchess-
13 inspired chefs are rendered in the calligraphic brushwork characteristic of Mr. Manchess’s oil
14 paintings. The Brom-inspired chefs have a gothic and countercultural air, including several
15 fantastic creatures, typical of artworks by Mr. Brom. The Zhang-inspired chefs are Asian and
16 rendered photographically, like many of Ms. Zhang’s works. The McKernan-inspired chefs
17 feature colorful female-presenting faces with elaborate hair and decorations, commonly found in
18 Mx. McKernan’s work. The Andersen-inspired chefs are all obvious variations on the main
19 character of Ms. Andersen’s celebrated comic “Sarah’s Scribbles.”

20 175. The next set of Midjourney text prompts consists of “*teacher*,” “*gregory manchess*
21 *teacher*,” “*gerald brom teacher*,” “*jingna zhang teacher*,” “*kelly mckernan teacher*,” and “*sarah*
22 *andersen teacher*.” Exhibit F, pp. 7–12. The initial “*teacher*” prompt shows what the Midjourney
23 Model produces by default, without an artist name in the prompt. Much like the Stability results,
24 the Manchess-inspired teachers are rendered in calligraphic brushwork with realistic settings,
25 common in the work of Mr. Manchess. The Brom-inspired teachers are demonic, and feature
26 images of skulls, common motifs in Mr. Brom’s work. The Zhang-inspired teachers are Asian and
27 rendered photographically, like many of Ms. Zhang’s images. The McKernan-inspired teachers
28

1 feature colorful female-presenting faces with elaborate hair and decorations, as often seen in Mx.
2 McKernan's work. The Andersen-inspired teachers are all obvious variations on the main
3 character of Ms. Andersen's comic "Sarah's Scribbles."

4 176. In sum, Stability's diffusion model (Stable Diffusion XL 1.0), Runway's diffusion
5 model (inferred to be Stable Diffusion 1.5), and the Midjourney Model version 5.2 demonstrate
6 behavior similar to that described in the Casper research paper: by adding a certain artist name to
7 a prompt, one can consistently elicit characteristic elements of that artist's body of work in the
8 output images, allowing the creation of unlimited output images that are substantially similar to,
9 and could be mistaken for those of the original artist. These results are consistent between
10 prompts and between models. This strongly suggests that the Stable Diffusion XL, the Runway
11 model, and the Midjourney Model store copies of protected expression after copying and
12 ingesting training images.

13 **XII. EXAMPLES OF IMAGE PROMPTS USING PLAINTIFF IMAGES**
14 **IN AI IMAGE PRODUCTS OFFERED BY STABILITY, RUNWAY,**
15 **AND MIDJOURNEY**

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

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

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

1 180. **Stability:** The first set of image-prompt examples were made with a Stability AI
2 image product called Reimagine XL, released in May 2023. Reimagine XL is built atop the Stable
3 Diffusion XL model that was also used for the text-prompt examples in the previous section. The
4 difference is that Reimagine XL accepts image prompts rather than text prompts. As Stability
5 explains³⁶ (emphasis added below)—

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

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

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

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

25 184. On each page of this exhibit, the original training image is positioned in the upper
26 left; the other three images are output images. In every case, the output images are not merely
27 similar to the training image, but substantially similar—in some cases startlingly so. On
28 information and belief, because Stability says “not ... a single pixel” from the input image is being
passed into Stable Diffusion XL (via the Reimagine XL image product), it would not be possible

³⁶ <https://clipdrop.co/stable-diffusion-reimagine>

1 for Stable Diffusion XL to produce output images substantially similar to the training images
2 unless it had stored copies of protected expression from those training images, and the CLIP
3 embedding generated from the image prompt was eliciting the output of this copied expression.

4 185. **Runway:** The next set of image-prompt examples were made with Runway’s AI
5 Magic Tools using its “Image Variation” feature. On information and belief, based on the output
6 of this tool, it works in a manner similar to Stability’s Reimagine XL: none of the pixels of the
7 input image are retained, but rather a CLIP embedding is generated from the input image, which
8 guides the subsequent diffusion process.

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

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

20 188. **Midjourney:** The next set of image-prompt examples were made with Midjourney
21 using its image-prompting feature. This feature was released by Midjourney one day after the
22 initial complaint in this action was filed. According to Midjourney CEO David Holz, this feature
23 does not copy pixels from the input, but rather “looks at the ‘concepts’ and ‘vibes’ of your images
24 and merges them together into novel interpretations.”³⁷ According to Midjourney’s
25 documentation, using an image as a prompt merely “influence[s] a Job’s composition, style, and
26 colors.”³⁸

27 _____
³⁷ <https://discord.com/channels/662267976984297473/952771221915840552/1064031587735445546>

28 ³⁸ <https://docs.midjourney.com/docs/image-prompts>

1 189. **Exhibit I: Midjourney image prompts** contains examples of prompting
2 Midjourney with training images from Exhibit A (Plaintiff images in LAION-5B). These training
3 images were made by plaintiffs Gerald Brom, Gregory Manchess, Grzegorz Rutkowski, Hawke
4 Southworth, Jingna Zhang, Karla Ortiz, and Kelly McKernan. Midjourney requires that an image
5 prompt be accompanied by a text prompt, so in these cases, each image prompt was
6 supplemented by the artist name. Each prompt was also supplemented with the command `--iw
7 2` to ensure that the image portion of the prompt was treated as the primary part (where “iw”
8 means “image weight”), thereby maximizing “the ‘concepts’ and ‘vibes’” derived from the
9 image.

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

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

23 192. Taken together, these examples of text prompting and image prompting strongly
24 imply that diffusion models like the ones shown above store copies of protected expression from
25 copying and ingesting training images.
26
27
28

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

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

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

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



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



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



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



16
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
20 right, 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
17 AI image products to create images that mimic and imitate Plaintiffs and Class members’ work.
18 Further, users and licensees, with assistance from Defendants, track and update the specific
19 artists (including Plaintiffs) which the AI image products are able to mimic or imitate.

20 202. Midjourney Image Product users and licensees maintain a spreadsheet that
21 features community-created images and tracks the specific artists that the Midjourney Image
22 Product can successfully mimic or imitate. The list includes thousands of artists and tracks user-
23 and licensee-generated works using artist names as prompts for various versions of the
24 Midjourney Model. Plaintiffs listed in the Midjourney community-generated spreadsheet include
25 Sarah Andersen, 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).”³⁹ The purpose of the post was, in the words of the user, to
2 “help decide the style to invoke when prompting the AI to create your desired images. The user
3 used the following prompt into Midjourney “Art by <artist name>.” Included in the study were
4 several Plaintiffs including Sarah Andersen, Kelly McKernan and Jingna Zhang.

5 204. That same individual later posted an “artist study” wherein the individual used
6 the prompt “art by<artist>” with a “negative prompt” of “blurry, soft, low quality”; the artist’s
7 name as the only change in the prompt into Stable Diffusion XL.⁴⁰ As demonstrated by the user,
8 “[t]he prompt was straight forward ‘art by <artist>’ which would get the SDXL mode [sic] to
9 emulate the style and creations of that artists [sic].” Included in this artist study were several of
10 the Plaintiffs including Sarah Andersen, Gerald Brom, Kelly McKernan, Karla Ortiz and Jingna
11 Zhang.

12 205. Another example involves users conducting “Artist Style Studies” using “Stable
13 Diffusion V1”.⁴¹ The user input a series of six prompts, where in prompt 1-3 included “a portrait
14 of a character in a scenic environment by [artist]” and prompts 4-6 included “a building in a
15 stunning landscape by [artist].” This particular study includes over 1781 artists to date, and
16 includes a form to recommend other artist names to input. Plaintiffs who have been included in
17 this study include Plaintiffs Andersen, Brom, McKernan, Ortiz, and Zhang.

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

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

26 ³⁹ <https://weirdwonderfulai.art/resources/artist-styles-on-midjourney-v4/>

27 ⁴⁰ <https://weirdwonderfulai.art/resources/stable-diffusion-xl-sd-xl-artist-study/>

28 ⁴¹

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

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

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

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

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

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

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

21 213. The term **LAION-5B Registered Plaintiffs** denotes the subset of plaintiffs who
22 hold copyrights in these LAION-5B Registered Works that were registered before the filing of the
23 initial complaint in this action, namely Sarah Andersen, Jingna Zhang, Gerald Brom, Gregory
24 Manchess, 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
28

1 training the model: LAION-5B and subsets 5B” (see [https://huggingface.co/stabilityai/stable-](https://huggingface.co/stabilityai/stable-diffusion-2#training)
2 [diffusion-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
5 LAION-5B, because Stability has funded LAION and used LAION datasets for all its previous
6 models. AI chipmaker NVIDIA features Stable Diffusion XL as one of its “AI Foundation
7 Models.” On its information page for Stable Diffusion XL, the training dataset is listed as
8 LAION-5B.⁴²

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

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

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

23
24
25
26
27
28

⁴² 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

1 represented a transformation of the LAION-5B Registered Works into an alternative form, Stable
2 Diffusion XL 1.0 also qualifies as an infringing Statutory Derivative Work.

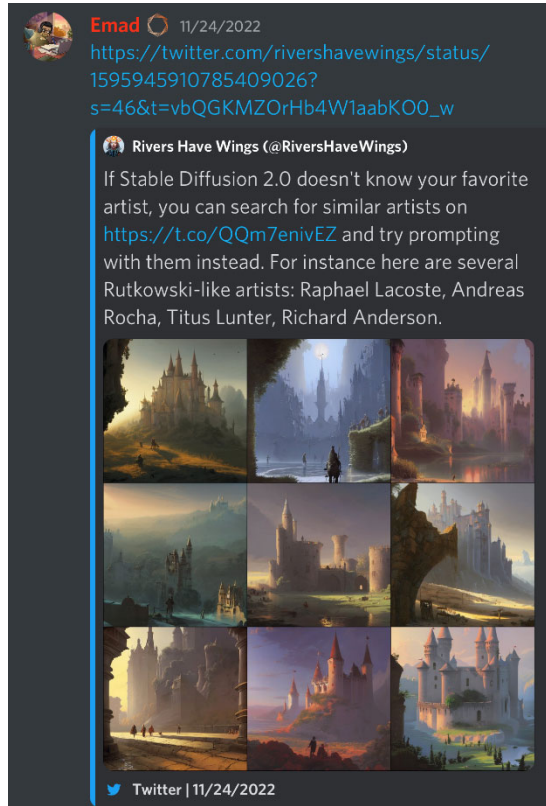
3 224. Executives and high-level employees of Stability know that one of the most
4 attractive features of the Stability models is its ability to mimic and copy artists' works, including
5 Plaintiffs. As such, they routinely advertise the Stability Models' ability to mimic artwork.

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


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



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



15 228. Stability also maintains a Discord channel where executives routinely offered
16 resources to users including encouragement to use artist names as prompts:



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

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 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

1 remove or alter CMI from the training images. Therefore, Stability intentionally removed or
2 altered CMI from the Plaintiffs' works in violation of 17 U.S.C. § 1202(b)(1).

3 246. Stability also knew that the distribution of works without CMI would lead to
4 further infringement. Stability encourages the use of artist names as prompts, i.e., encourages the
5 Stability Models' users and licensees to infringe on an artists' work. Because the Stability Models
6 do not preserve CMI, users and licensees also create infringing works without CMI which can
7 reasonably lead to further infringement.

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

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

19 249. Stability knew or had reasonable grounds to know that this removal of the LAION-
20 5B Plaintiffs' CMI would facilitate copyright infringement by concealing the facts that a) the
21 Stability Models are infringing Statutory Copies of the LAION-5B Works, and b) the Stability
22 Models are infringing Statutory Derivative Works based on the LAION-5B Works.

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

1 **COUNT FOUR**

2 **Unjust enrichment**
3 **under Cal. Bus. & Prof. Code § 17200 and California Common Law**
4 **against Stability**
5 **on behalf of all Plaintiffs, the Damages and the Injunctive Classes**

6 251. The preceding factual allegations are incorporated by reference.

7 252. Stability has unjustly misappropriated the LAION-5B Works in order to develop,
8 train and promote the Stability Models, enabling it to receive profit and other benefits. It would
9 be unjust for Stability to retain these benefits.

10 253. Plaintiffs and the Class have invested substantial time and energy in creating their
11 works, including those included as LAION-5B Works.

12 254. By using Plaintiffs' works to train, develop and promote the Stability Models,
13 Plaintiffs and the Class were deprived of the benefit of the value of their works, including
14 monetary damages.

15 255. Plaintiffs did not consent to the unauthorized use of their works to train, develop
16 and promote the Stability Models.

17 256. Stability derived profit and/or other benefits from the use of Plaintiffs' works to
18 train, develop and promote the Stability Models.

19 257. It would be unjust for Stability to retain those benefits.

20 258. Plaintiffs are entitled to restitution, including disgorgement of profits and a
21 constructive trust over all assets created with the Stability Models.

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

23 259. Midjourney promotes the Midjourney Image Product, which is accessed and run
24 through Discord. Midjourney maintains its own Discord server from which users can access the
25 Midjourney image generator. Midjourney's Discord server also allows Midjourney executives and
26 other high-level employees to provide promotional communications to users and licensees of the
27 Midjourney Image Product.

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

5 261. Over a series of Discord messages, Holz said “i think you’re all gonna get [your]
6 mind blown by this style feature ... we were very liberal in building out the dictionary ... it has
7 cores and punks and artist names ... as much as we could dump in there ... i should be clear it’s
8 not just genres its also artist names ... it’s mostly artist names ... 4000 artist names.”⁴³

9 262. Holz then said, “here is our style list”⁴⁴ and posted a link to a spreadsheet on
10 Google Docs called “Midjourney Style List.”⁴⁵ One of the tabs on the spreadsheet was called
11 “Artists” and listed over 4700 artist names. In other words, Holz published a list of artists who
12 the Midjourney Image Product recognizes with the express purpose of these names being used by
13 users and licensees of the Midjourney Image Product as terms in prompts. Holz’s comment, and
14 the list, have remained available ever since.

15 263. Below, this list is called the **Midjourney Name List**. A copy of this list appears in
16 **Exhibit J: Midjourney Name List**.

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

20 265. Midjourney also tracked the most popular artists used as prompts. At one point,
21 the Midjourney Bot maintained a count on Midjourney as recently as April 2023.

22 266. In the months before February 2022, Midjourney trained version 1 of the
23 Midjourney Model. In February 2022, on Midjourney’s Discord server Midjourney CEO David
24 Holz described Midjourney’s training data: “we have some private data partners as well as some

25 _____
26 ⁴³ Combining <https://discord.com/channels/662267976984297473/938713143759216720/941972360171520001>,
<https://discord.com/channels/662267976984297473/938713143759216720/941972890520272906>, and
<https://discord.com/channels/662267976984297473/938713143759216720/941976464704802836>

27 ⁴⁴ <https://discord.com/channels/662267976984297473/938713143759216720/941987328828768256>

28 ⁴⁵ https://docs.google.com/spreadsheets/d/1MEglfejppqgVcaf-I-cgZ5ngV_MlaOTeGXAoBPJO69FM/edit#gid=1096178862

1 open ones like laion.”⁴⁶ In May 2022, Stability CEO Mostaque said “MidJourney is using a
2 LAION 400m based model ... I just support MJ like many other labs/researchers in my nice
3 way.”⁴⁷ In July 2022, a Midjourney-affiliated Discord moderator named Soar also confirmed that
4 Midjourney was “using a modified version of the LAION 400m dataset.”⁴⁸ In other words,
5 Midjourney trained its image generation product on at least the LAION-400M dataset.

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

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

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

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

18 271. Because LAION-400M is an openly accessible dataset, Midjourney knew that the
19 LAION-400M dataset contained copyrighted works, including those of the LAION-400M
20 Registered Plaintiffs. Further, because LAION also includes detection scores for watermarks for
21 other datasets, and because of the ubiquity of artists affixing CMI to their works, Midjourney
22 knew that the LAION-400M dataset contained works with CMI affixed on them.

23 272. The LAION-400M Registered Works are included in the LAION-400M dataset.
24 Therefore, Midjourney used the LAION-400M Registered Works for training. Below, the term
25
26

27 ⁴⁶ <https://discord.com/channels/662267976984297473/938713143759216720/943315577018126408>

⁴⁷ <https://discord.com/channels/823813159592001537/912729332311556136/975894553225752626>

28 ⁴⁸ <https://discord.com/channels/662267976984297473/959962985655320616/1001938136445751387>

1 **Midjourney 400M Models** refers to all models trained by Midjourney on the LAION-400M
2 Registered Works, including version 1 of the Midjourney Model.

3 273. Since October 2022, Midjourney has also incorporated a version of Stable
4 Diffusion into the Midjourney Image Product, which is accessible to users by adding the
5 command `--test` or `--testp` to a text prompt. According to Midjourney's moderator Molang,
6 "--test and --testp is a little bit of SD [Stable Diffusion] mixed with a lot of Midjourney tweaks and
7 magic."⁴⁹

8 274. Sometime after July 2022 and before March 2023, Midjourney adopted LAION-
9 5B as its training dataset. In July 2022, a Midjourney-affiliated Discord moderator named Danger
10 Awesome said that "the updated LAION 5B dataset"⁵⁰ formed the basis for "the upcoming
11 [Midjourney] dataset update." In March 2023, a Midjourney-affiliated Discord Moderator named
12 Sunshineyday said that "MJ is trained on a subset of Laion5b."⁵¹

13 275. Version 5 of the Midjourney Model was released in March 2023. On information
14 and belief, version 5 of the Midjourney Model and subsequent versions were trained on LAION-
15 5B.

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

20 277. The LAION-5B Registered Works are included in the LAION-5B dataset.
21 Therefore, Midjourney used the LAION-5B Registered Works for training. Below, the term
22 **Midjourney 5B Models** refers to all models trained by Midjourney on the LAION-5B Registered
23 Works, including version 5 of the Midjourney Model and the current version 5.2 of the
24 Midjourney Model.

25
26
27 ⁴⁹ <https://discord.com/channels/662267976984297473/958069758211797092/1038899058636501132>

⁵⁰ <https://discord.com/channels/662267976984297473/992207085146222713/998451098534817883>

28 ⁵¹ <https://discord.com/channels/662267976984297473/992207085146222713/1082089794521268314>

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

7
8 **COUNT FIVE**

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

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

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

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

25 282. During the training of each Midjourney 400M Model, Midjourney made a series
26 of intermediate Statutory Copies of the LAION-400M Registered Works. For instance, diffusion
27 models are trained by creating “noised” copies of training images, as described herein, all of
28

1 which qualify as Statutory Copies. The intermediate Statutory Copies of each registered work
2 that Midjourney made during training of the Midjourney 400M Models were substantially similar
3 to that registered work.

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

7
8 **COUNT SIX**

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

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

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

18 286. 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
21 the 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
24 Statutory Copies made of each registered work were substantially similar to that registered work.

25 287. 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,
28

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

4 288. 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**
7 and **Exhibit I: Midjourney image prompts**. Therefore, Midjourney Model version 5.2 qualifies
8 as an infringing Statutory Copy of the LAION-5B Registered Works. Because Midjourney Model
9 version 5.2 represents a transformation of the LAION-5B Registered Works into an alternative
10 form, Midjourney Model version 5.2 also qualifies as an infringing Statutory Derivative Work.

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

13 290. 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

17 **COUNT SEVEN**

18 **DMCA violations**

19 **by removing and altering CMI of training images**

20 **against Midjourney**

21 **on behalf of All Plaintiffs, the Damages and Injunctive Class**

22 291. The preceding factual allegations are incorporated by reference.

23 292. The LAION-400M Plaintiffs included one or more forms of CMI (as defined in
24 Section 1202(c) of the DMCA) in each of their respective works in the LAION-400M Works,
25 including captions, and distinctive marks such as URLs to personal websites, signatures, and
26 watermarks.
27
28

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

4 294. Midjourney did not contact Plaintiffs and the Class to obtain authority to remove
5 or alter CMI from their works within the meaning of the DMCA.

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

11 296. Midjourney also knew the LAION-400M dataset contained CMI. Given that CMI
12 is ubiquitous and the LAION-400M dataset contained copyrighted works, it is a certainty that
13 many works within the dataset contained CMI.

14 297. Midjourney had access to but were not licensed by Plaintiffs or the Class to
15 incorporate their works in Midjourney Image Product.

16 298. Midjourney had access to but were not licensed by Plaintiffs or the Class to create
17 copies based on their works into the Midjourney Image Product.

18 299. Midjourney had access to but were not licensed by Plaintiffs or the Class to
19 distribute their works as Midjourney does through the Midjourney Image Product.

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

26 301. Without the authority of the Plaintiffs, Midjourney copied the LAION-400M and
27 LAION-5B Works and used these copies as training data for the Midjourney Models. The
28

1 training process is designed to remove or alter CMI from the training images. Therefore,
2 Midjourney intentionally removed or altered CMI from the Plaintiffs' Works in violation of 17
3 U.S.C. § 1202(b)(1).

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

11 303. The LAION-400M Plaintiffs have been injured by Midjourney's removal or
12 alteration of CMI. The LAION-400M Plaintiffs have also been injured by Midjourney's
13 falsification of CMI. These plaintiffs are entitled to statutory damages, actual damages,
14 restitution of profits, and other remedies provided by law.

15 16 **COUNT EIGHT**

17 **Lanham Act — false endorsement** 18 **by unauthorized commercial use of artists' names** 19 **against Midjourney** 20 **on behalf of the Midjourney Named Plaintiffs and Class**

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

22 305. Midjourney engaged in commercial speech that sought to capitalize upon the
23 Midjourney Named Plaintiffs' popularity, recognition, and appeal among consumers of art
24 products. Midjourney's use of the Midjourney Named Plaintiffs' names was purely to advertise
25 its image generator. This use does not contribute significantly to a matter of public interest. The
26 purpose of publishing over 4700 names in the Midjourney Name List was to promote and
27
28

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

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

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

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

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

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

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

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

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

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

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

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

7 317. The Midjourney Named Plaintiffs are entitled to statutory damages, actual
8 damages, restitution of profits, and other remedies provided by law.

9
10 **COUNT NINE**

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

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

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

- 21 a. Sarah Andersen is known for work that is simple, cartoony, and often
22 strictly in black and white. In particular, she is known for “Sarah’s
23 Scribbles,” a comic featuring a young woman with dark hair, big eyes, and a
24 striped shirt.
- 25 b. Karla Ortiz is known for a mixture of classical realism and impressionism,
26 often delving into fantastical, macabre and surrealist themes, and inspired
27
28

1 by the technical prowess of American Renaissance movements with a
2 strong influence of contemporary media.

3 c. Gerald Brom is known for gritty, dark, fantasy images, painted in
4 traditional media, combining classical realism, gothic and counterculture
5 aesthetics.

6 d. Grzegorz Rutkowski is known for lavish fantasy scenes rendered in a
7 classical painting style.

8 e. Julia Kaye is known for three-panel black-and-white comics, loosely inked
9 with a thin fixed-width pen, wherein each individual comic is a micro-
10 vignette in the artist's life.

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

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

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

23 323. As a result, the Midjourney Image Product can and frequently does generate
24 images featuring protectable trade dress that are likely to cause confusion in consumers. The
25 Midjourney Named Plaintiffs never authorized Midjourney to copy, emulate, or otherwise
26 recreate their trade dress; nor did the Midjourney Named Plaintiffs authorize Midjourney to use,
27 in conjunction with the advertisement and sale of its services, images featuring their trade dress.
28

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

7 325. Midjourney vicariously infringes on the Midjourney Named Plaintiffs’ trade-dress
8 rights by encouraging and inducing the users of the Midjourney Image Product to enter artist-
9 name prompts and generate images featuring the Midjourney Named Plaintiffs’ protectable trade
10 dress. For example—

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

25 326. Midjourney exercises control over the infringing images by including the CLIP
26 model in its image pipeline, and by marketing artist-name prompts as a key feature of its image
27 generator via the Midjourney Name List. Without the CLIP model, Midjourney’s users would
28

1 not be able to infringe on the Midjourney Named Plaintiffs' trade-dress rights or those of the
2 other artists on the Midjourney Name List.

3 327. Each of the Midjourney Named Plaintiffs' respective trade dress has no intrinsic
4 functional value. The unique combination of particular artistic elements does not confer any
5 utilitarian advantages on their art products and are purely ornamental and aesthetic. There
6 remains an unlimited number of alternative artistic styles available beyond the trade dress owned
7 the Midjourney Named Plaintiffs.

8 328. Each of the Midjourney Named Plaintiffs' trade dress possesses secondary
9 meaning because the trade dress of their art products invoke a mental association by a substantial
10 segment of potential consumers between the trade dress and the creator of the art product.

11 329. Midjourney's vicarious infringement of the Midjourney Named Plaintiffs' trade-
12 dress rights are committed with actual and constructive knowledge of their trade dress, and with
13 the intent to cause confusion, mistake, or deception.

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

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

25 332. Midjourney has, and continues to, vicariously infringe on the trade-dress rights of
26 the Midjourney Named Plaintiffs in violation of section 43(a) of the Lanham Act, 15 U.S.C.
27 § 1125(a).

1 333. Midjourney’s past and continuing infringement of the Midjourney Named
2 Plaintiffs’ trade dress is an exceptional case and was willful and intentional, as evidenced by
3 a) Midjourney’s intentional inclusion of the CLIP model in the design of the Midjourney Image
4 Product and b) its open advertisement of the Midjourney Image Product’s ability to replicate an
5 artist’s trade dress via the Midjourney Name List. Thus, the Midjourney Named Plaintiffs are
6 entitled to treble their actual damages and to an award of attorneys’ fees under 15 U.S.C.
7 § 1117(a), and all other available remedies.
8

9 **COUNT TEN**

10 **Unjust enrichment**
11 **under Cal. Bus. & Prof. Code § 17200 and California Common Law**
12 **against Midjourney**
13 **on behalf of all Plaintiffs, and the Damages and Injunctive Class**

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

15 335. Midjourney has unjustly misappropriated the LAION-400M Works and the
16 LAION-5B Works, enabling it to receive profit and other benefits in order to train, develop and
17 promote the Midjourney Image Product. It would be unjust for Midjourney to retain these
18 benefits.

19 336. Plaintiffs and the Class have invested substantial time and energy in creating their
20 works, including those included as LAION-400M Works and LAION-5B Works.

21 337. By using Plaintiffs’ works to train, develop and promote the Midjourney Image
22 Product, Plaintiffs and the Class were deprived of the benefit of the value of their works,
23 including monetary damages.

24 338. Plaintiffs did not consent to the unauthorized use of their works to train, develop
25 and promote the Midjourney Image Product.

26 339. Midjourney derived profit and/or other benefits from the use of Plaintiffs’ works
27 to train, develop and promote the Midjourney Image Product.
28

1 340. It would be unjust for Midjourney to retain those benefits.

2 341. Plaintiffs are entitled to restitution, including disgorgement of profits and a
3 constructive trust over all assets created with the Midjourney 400M Models and Midjourney 5B
4 Models.

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

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

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

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

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

18
19 **COUNT ELEVEN**

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

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

26 347. The LAION-5B Registered Plaintiffs and Karla Ortiz never authorized Runway to
27 use their respective LAION-5B Registered Works in any way. Nevertheless, Runway repeatedly
28

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

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

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

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

22 351. On information and belief, the other Runway Models exhibit the same properties,
23 because they were trained on the same LAION-5B dataset.

24 352. Since October 2022, Runway has distributed Stable Diffusion 1.5 to the public, for
25 instance via websites like GitHub (see <https://github.com/runwayml/stable-diffusion>) and
26 Hugging Face (see <https://huggingface.co/runwayml/stable-diffusion-v1-5>). In so doing, Runway
27 infringed the exclusive distribution rights of the LAION-5B Registered Plaintiffs and Karla Ortiz.
28

1 353. The LAION-5B Registered Plaintiffs and Karla Ortiz have been and continue to be
2 injured by Runway’s multiple acts of direct copyright infringement. These plaintiffs are entitled
3 to statutory damages, actual damages, restitution of profits, and other remedies provided by law.
4

5 **COUNT TWELVE**

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

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

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

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

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

20 358. Runway intends to cause further infringement with Stable Diffusion 1.5. In
21 February 2023, Stability CEO Mostaque said that Stable Diffusion 1.5 was “the most popular
22 model by far by [a] for profit company.”⁵²

23 359. The LAION-5B Registered Plaintiffs have been and continue to be injured by
24 Runway’s inducement of copyright infringement. These plaintiffs are entitled to statutory
25 damages, actual damages, restitution of profits, and other remedies provided by law.
26
27

28 ⁵² <https://twitter.com/EMostaque/status/1629514395825983489>

COUNT THIRTEEN

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

360. The preceding factual allegations are incorporated by reference.

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

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

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

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

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

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

367. Without the authority of the LAION-5B Plaintiffs, Runway directly copied the LAION-5B Works and used these Statutory Copies as training data for the Runway Models. The works copied by Runway included CMI, including in the form of distinctive marks such as watermarks or signatures. The training process is designed to remove or alter CMI from the

1 training images. Therefore, Runway intentionally removed or altered CMI from the LAION-5B
2 Works in violation of 17 U.S.C. § 1202(b)(1).

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

11 369. As demonstrated herein, the Runway Models generate output that are copies of
12 original images with CMI with the CMI removed and/or altered.

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

19 371. The LAION-5B Plaintiffs have been injured by Runway’s removal or alteration of
20 CMI. The LAION-5B Plaintiffs have been injured by Runway’s falsification of CMI by claiming
21 false copyright in Stable Diffusion 1.5. These plaintiffs are entitled to statutory damages, actual
22 damages, restitution of profits, and other remedies provided by law.
23
24
25
26
27
28

COUNT FOURTEEN

**Unjust enrichment
under Cal. Bus. & Prof. Code § 17200 and California Common Law
against Runway
on behalf of all Plaintiffs, the Damages and Injunctive Classes**

372. The preceding factual allegations are incorporated by reference.

373. Plaintiffs and the Class have invested substantial time and energy in creating their works.

374. Runway has unjustly misappropriated the LAION-5B Works in order to train, develop and promote the Runway Models, enabling it to receive profit and other benefits. It would be unjust for Runway to retain these benefits.

375. By using Plaintiffs' works to train, develop and promote the Runway Models, Plaintiffs and the Class were deprived of the benefit of the value of their works, including monetary damages.

376. Plaintiffs did not consent to the unauthorized use of their works to train, develop and promote the Runway Models.

377. Runway derived profit and/or other benefits from the use of Plaintiffs' works to train, develop and promote the Runway Models.

378. It would be unjust for Runway to retain those benefits.

379. Plaintiffs are entitled to restitution, including disgorgement of profits and a constructive trust over all assets created with the Runway Models.

XVIII. CAUSES OF ACTION AGAINST DEVIANTART

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

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

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

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

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

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

22 386. Each of the DeviantArt Plaintiffs has stored images on DeviantArt that were later
23 incorporated into the LAION-5B dataset. These images were therefore copied by Runway and
24 Stability to train versions of Stable Diffusion. Exhibit A contains a sampling—but not an
25 exhaustive listing—of images created by the DeviantArt Plaintiffs that are contained in LAION-
26 5B and were copied from DeviantArt. They can be identified through their LAION-5B URL,
27
28

1 which comes from the “wixmp.com” domain. This domain is used by DeviantArt to store
2 member images. (“Wix” in the domain name refers to the parent company of DeviantArt.)

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

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

12 389. CompVis is the shorthand name of the Computer Vision and Learning Group at
13 Ludwig Maximilian University in Munich, where the original research underlying Stable
14 Diffusion was first conducted. According to the GitHub page for Stable Diffusion 1.4, “Stable
15 Diffusion was made possible thanks to a collaboration with Stability AI and Runway.”⁵³

16 390. According to CompVis, Stable Diffusion 1.4 “was trained on a large-scale dataset
17 LAION-5B.”⁵⁴

18 391. 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
21 the `img2dataset` tool or another similar tool. Consistent with this, in preparation for training
22 Stable Diffusion 1.4, CompVis made one or more Statutory Copies of the LAION-5B Registered
23 Works so they could be fed to Stable Diffusion 1.4 as training data. The Statutory Copies made of
24 each registered work were substantially similar to that registered work.

25 392. During the training of Stable Diffusion 1.4, CompVis made a series of intermediate
26 Statutory Copies of the LAION-5B Registered Works. For instance, diffusion models are trained

27 _____
⁵³ See <https://github.com/CompVis/stable-diffusion>

28 ⁵⁴ See <https://huggingface.co/CompVis/stable-diffusion-v1-4>

1 by creating “noised” copies of training images, as described herein, all of which qualify as
2 Statutory Copies. The intermediate Statutory Copies of each registered work that CompVis made
3 during training of Stable Diffusion 1.4 were substantially similar to that registered work.

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

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

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

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

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

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

1 DeviantArt lets you declare whether or not external AI models and
2 platforms can train based on your deviations. When submitting a
3 deviation, you'll be able to check a box that informs third parties
4 whether or not you authorize that submission being included in
5 datasets used to train AI models like AI image generators.⁵⁵

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

12 397. This has been further confirmed by DeviantArt CTO Chris Nell. In November
13 2022, on the public LAION Discord server, Nell described himself as "one of the people at
14 DeviantArt working on improving acceptance of AI generated/augmented art in the broader
15 online arts community" and added "I think our goals at DA [DeviantArt] are very aligned with
16 LAION's ... and want to collaborate as much as possible."⁵⁶ Nell said of DreamUp: "we did not
17 fine tune [meaning, perform additional training on] SD [= Stable Diffusion] at all, so there aren't
18 novel weights to share. [W]e do perform additional guidance at generation time ... so it's not
19 exactly unmodified SD [= Stable Diffusion] output, but that is more akin to prompt tuning."⁵⁷ As
20 confirmed by Nell, DeviantArt was well aware of how Stable Diffusion was developed and did not
21 do any fine-tuning of the weights included in the Stable Diffusion model DreamUp was based on.
22 In other words, this implies all of the images copied in training Stable Diffusion were included in
23 the DreamUp model.

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

27 ⁵⁵ <https://www.deviantartsupport.com/en/dreamup>

28 ⁵⁶ <https://discord.com/channels/823813159592001537/1006139459860975716/1042539656396411004>

⁵⁷ <https://discord.com/channels/823813159592001537/1006139459860975716/1042543837425438804>

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

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

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

21 400. DeviantArt’s embrace of generative AI art was seen as a betrayal by its art
22 community.

23 401. The scope of DeviantArt’s betrayal of its artist community by embracing Stable
24 Diffusion was evident in a group audio session held by DeviantArt management on November 11,
25 2022 from approximately 1:00–2:30 pm Pacific Time. DeviantArt scheduled the discussion
26 specifically to allay the well-founded concerns of DeviantArt members that DeviantArt’s embrace
27 of AI art was a complete repudiation of its longstanding community principles, as well as
28 economically and legally unfair.

402. At one point in the audio session, CEO Moti Levy explicitly took ownership of the
decision to bring Stable Diffusion (the basis of the DreamUp–CompVis Model) onto DeviantArt
via the DreamUp app: “The reason why we’re using Stable Diffusion because it’s the only option
for us to take an open source [software engine] and modify it The other platforms or the

1 other companies do not allow it. . . . [A]nd by the way, that was my decision. **That’s our decision**
2 **by me as the CEO. That’s my decision to take Stable Diffusion.**” (Emphasis added.)

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

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

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

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

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

4 408. This promise is misleading.

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

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

22 **COUNT FIFTEEN**

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

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

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

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

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

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

9 416. The LAION-5B Registered Plaintiffs have been and continue to be injured by
10 DeviantArt’s multiple acts of direct copyright infringement. These plaintiffs are entitled to
11 statutory damages, actual damages, restitution of profits, and other remedies provided by law.
12

13 **COUNT SIXTEEN**

14 **Breach of contract**
15 **for violation of its Terms of Service**
16 **against DeviantArt**
17 **on behalf of the DeviantArt Plaintiffs**

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

19 418. The DeviantArt Plaintiffs have accepted DeviantArt’s Terms of Service. Thus,
20 the DeviantArt Plaintiffs have formed a contract with DeviantArt. Copies of the applicable
21 agreements are attached as **Exhibit L: DeviantArt Terms of Service (Nov. 11, 2022)**, **Exhibit**
22 **M: DeviantArt Terms of Service (Jan. 11, 2023)**, and **Exhibit N: DeviantArt Privacy Policy**
23 **(Jan. 11, 2023)**.

24 419. The DeviantArt Plaintiffs have performed the obligations imposed on them by
25 DeviantArt’s Terms of Service.

26 420. Section 16 of the DeviantArt Terms of Service specifies the limited rights
27 DeviantArt has in the work of the DeviantArt Plaintiffs (emphasis added below):
28

1 16. Copyright in Your Content ... DeviantArt does not claim
2 ownership rights in Your Content. **For the sole purpose of enabling**
3 **us to make your Content available through the Service**, you grant
4 to DeviantArt a non-exclusive, royalty-free license to reproduce,
5 distribute, re-format, store, prepare derivative works based on, and
6 publicly display and perform Your Content ...

7 421. Though DeviantArt has the right to “reproduce,” “distribute,” and “prepare
8 derivative works based on” the works of the DeviantArt Plaintiffs, it may only do so “[f]or the
9 sole purpose of enabling us to make your Content available through the Service.”

10 422. The DeviantArt Plaintiffs put their work on DeviantArt because it held itself out
11 as a platform for increasing visibility of artists and their work. But by releasing DreamUp,
12 DeviantArt breached the Terms of Service in two ways:

- 13 a. DeviantArt exceeded its license grant in Section 16 and thereby breached
14 its explicit Terms of Service. Relative to a DeviantArt artist member whose
15 work was used for training the DreamUp–CompVis Model, in no sense
16 does DreamUp “make your Content available” to visitors. Rather, artist
17 “Content” was misappropriated to make a new commercial product.
- 18 b. DeviantArt breached the implied covenant of good faith and fair dealing.
19 The release of DreamUp unleashed a flood of AI-generated images on
20 DeviantArt that immediately began drowning out the work of human
21 artists like the DeviantArt Plaintiffs. By releasing DreamUp, DeviantArt
22 put itself into competition with the DeviantArt Plaintiffs and its other
23 artist members, undermining their very purpose in being on DeviantArt in
24 the first place. DeviantArt’s bad faith was further exemplified by its hasty
25 addition of a permissive new “Data Scraping & Machine Learning
26 Activities” provision to its Terms of Service after DeviantArt’s members
27 complained about the unfairness of DreamUp.

28 423. In an audio session on November 11, 2022—in response to outrage from the
DeviantArt artist community that had arisen in the two days since the release of DreamUp—
DeviantArt CEO Moti Levy took ownership of the decision to put a copy of the DreamUp–

1 CompVis Model inside the DreamUp app: “The reason why we’re using Stable Diffusion
2 because it’s the only option for us to take an open source [software engine] and modify it ... The
3 other platforms or the other companies do not allow it ... [A]nd by the way, that was my decision.
4 That’s our decision by me as the CEO. That’s my decision to take Stable Diffusion.”

5 424. At the beginning of the audio session, the DeviantArt terms of service were as
6 depicted in **Exhibit L: DeviantArt Terms of Service (Nov. 11, 2022)**. But shortly after the end
7 of this audio session, DeviantArt updated its Terms of Service. DeviantArt added a new
8 paragraph about “Data Scraping & Machine Learning Activities” that explicitly permits this kind
9 of usage under certain circumstances, so that Runway, Stability, Midjourney, and future AI
10 companies can continue to scrape DeviantArt for images with impunity. In so doing, DeviantArt
11 switched its loyalties from its artist members to AI companies seeking training data. The revised
12 terms are shown in **Exhibit M: DeviantArt Terms of Service (Jan. 11, 2023)**.

13 425. Although this “Data Scraping” provision acknowledges that certain kinds of data
14 scraping will continue to be an “unauthorized use” of the DeviantArt website, it also provides
15 that “owners of the works are responsible for policing their own works.” In other words, despite
16 its longstanding commitment to artists, DeviantArt washed its hands of the matter.

17 426. The DeviantArt Plaintiffs have suffered monetary damages as a result of
18 DeviantArt’s conduct.

19 427. DeviantArt’s conduct was a substantial factor in causing the DeviantArt Plaintiffs’
20 harm.

21 428. Unless enjoined by this Court, DeviantArt’s conduct will continue to cause the
22 DeviantArt Plaintiffs irreparable injury that cannot fully be compensated by money.

23 429. As a result of these material breaches by DeviantArt, the DeviantArt Plaintiffs are
24 entitled to an injunction requiring DeviantArt to comply with all the terms of the DeviantArt
25 Terms of Service.

1 430. The DeviantArt Plaintiffs are further entitled to recover from DeviantArt the
2 damages DeviantArt Plaintiffs sustained—including consequential damages—for the DeviantArt
3 Plaintiffs’ costs in enforcing DeviantArt’s Terms of Service.

4 431. The DeviantArt Plaintiffs are also entitled to recover restitution from DeviantArt
5 for any unjust enrichment, including gains, profits, and advantages that it has obtained as a result
6 of its breaches of the DeviantArt Terms of Service.

7

8

COUNT SEVENTEEN

9

**Unjust enrichment
under Cal. Bus. & Prof. Code § 17200 and California Common Law
against DeviantArt
on behalf of the DeviantArt Plaintiffs**

10

11

12

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

14 433. DeviantArt has held itself out as a community that supports artists. By abusing
15 that trust and using the works hosted on DeviantArt to develop and promote DreamUp and the
16 DreamUp-CompVis Model, which is based on Stable Diffusion 1.4, DeviantArt has violated the
17 legal rights of the DeviantArt Plaintiffs, enabling it to receive profit and other benefits. It would
18 be unjust for DeviantArt to retain these benefits.

19 434. Plaintiffs and the Class have invested substantial time and energy in creating their
20 works, including those Plaintiffs chose to host on DeviantArt.

21 435. By using Plaintiffs’ works to develop and promote DreamUp and the DreamUp-
22 CompVis Model, Plaintiffs and the Class were deprived of the benefit of the value of their works,
23 including monetary damages.

24 436. Plaintiffs did not consent to the unauthorized use of their works to develop and
25 promote DreamUp and the DreamUp-CompVis Model.

26 437. DeviantArt derived profit and/or other benefits from the use of Plaintiffs’ works to
27 train, develop and promote DreamUp and the DreamUp-CompVis Model.
28

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

438. It would be unjust for DeviantArt to retain those benefits.

439. The DeviantArt Plaintiffs are entitled to restitution, including disgorgement of profits and a constructive trust over all assets created with DreamUp and the DreamUp-CompVis Model.

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: November 29, 2023

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

2 Joseph R. Saveri (State Bar No. 130064)
3 Cadio Zirpoli (State Bar No. 179108)
4 Christopher K.L. Young (State Bar No. 318371)
5 Elissa A. Buchanan (State Bar No. 249996)
6 Travis Manfredi (State Bar No. 281779)
7 JOSEPH SAVERI LAW FIRM, LLP
8 601 California Street, Suite 1000
9 San Francisco, California 94108
10 Telephone: (415) 500-6800
11 Facsimile: (415) 395-9940
12 Email: jsaveri@saverilawfirm.com
 czirpoli@saverilawfirm.com
 cyoung@saverilawfirm.com
 eabuchanan@saverilawfirm.com
 tmanfredi@saverilawfirm.com

13 Matthew Butterick (State Bar No. 250953)
14 1920 Hillhurst Avenue, #406
15 Los Angeles, CA 90027
16 Telephone: (323) 968-2632
17 Facsimile: (415) 395-9940
18 Email: mb@buttericklaw.com

19 Brian D. Clark (*pro hac vice*)
20 Laura M. Matson (*pro hac vice*)
21 Arielle S. Wagner (*pro hac vice*)
22 Eura Chang (*pro hac vice*)
23 LOCKRIDGE GRINDAL NAUEN P.L.L.P.
24 100 Washington Avenue South, Suite 2200
25 Minneapolis, MN 55401
26 Telephone: (612)339-6900
27 Facsimile: (612)339-0981
28 Email: bdclark@locklaw.com
 lmmatson@locklaw.com
 aswagner@locklaw.com
 echang@locklaw.com

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