

The Copyright Implications of Generative AI

Written By Samuel V. Eichner
Assistance by Matt Grillo, Josh Graham

AIPLA Annual Meeting 2023
National Harbor, MD

I. INTRODUCTION

Generative artificial intelligence (“AI”) is groundbreaking new technology; it has the potential to revolutionize industries, communities, and the way we create. But how does that technology square with copyright law? As the Copyright Office, Congress, and the federal courts attempt to navigate the complex issues that this revolutionary technology raises, two broader questions have emerged: (1) when can generative AI give rise to copyrightable expression?; and (2) to what extent does the creation and use of generative AI platforms and content constitute copyright infringement or other copyright violations? This paper explores both questions.

II. A NEW TECHNOLOGICAL FRONTIER

A. What is Generative AI?

Generative AI *generates* content that strongly resembles human creative output, and “is capable of producing outputs such as text, images, video, or audio (including emulating a human voice) that would be considered copyrightable if created by a human author.”¹ Businesses and individuals alike are increasingly turning to generative AI to create text, images, and other media to serve a multitude of purposes, with various practical, professional and commercial implications. A communications team could use generative AI to draft a press release. A graphic designer could use generative AI to find inspiration, or to build a foundational image to modify with conventional graphic design tools. A lawyer could use generative AI to draft a legal brief.² In each case, users can feed text prompts into a generative AI platform to invoke that technology, and the content datasets on which they are “trained.” Such datasets often consist of copyrightable content, which may or may not have been licensed from rights holders.

Underlying the copyright implications of generative AI is the technology that makes it possible, namely, artificial intelligence, machine learning, and specifically, deep learning. Machine Learning (or “ML”) involves the development and gradual improvement of algorithms through data analysis. Deep learning, a subset of ML, utilizes artificial neural networks to mimic human decision-making by processing vast amounts of structured or unstructured data.³ As a product of deep learning, generative AI foundational models are “trained” with huge datasets in an iterative process whereby the generative AI model “learns” to generate increasingly accurate outputs (text, images, etc.) in response to text prompts. Some generative AI models allow users to type a textual “prompt” that a generative AI system uses to generate textual outputs, the most notable of which are large language models (“LLMs”) like ChatGPT. Others can generate visual outputs, the most notable of which are latent diffusion technologies like Stability AI and the services that utilize its Stable Diffusion technology, e.g., Midjourney.

B. Large Language Models

¹ <https://www.govinfo.gov/content/pkg/FR-2023-08-30/pdf/2023-18624.pdf>.

² See *Mata. v. Avianca, Inc.*, Case 1:22-cv-01461-PKC (S.D.N.Y. June 22, 2023).

³ An example is Google’s RankBrain, a neural network designed to predict and tailor search result rankings for individual Google users. <https://blog.google/products/search/how-ai-powers-great-search-results/>.

LLMs “are machine learning models found in popular generative AI platforms like ChatGPT that utilize deep learning algorithms to process and understand language” and generate textual outputs in response to textual prompts.⁴ LLMs are comprised of “multiple layers of neural networks,” such as a generative adversarial network (“GANs”) and variational autoencoders (“VAEs”), which work collectively to analyze text in small units (i.e., tokens) and make probabilistic predictions of an accurate textual response to a textual question based on an immense “training set”.⁵ The accuracy of LLM “responses” is directly tied to the size and quality of these training sets; without incredibly large amounts of useful data to analyze, they cannot accurately generate the kind of human-like responses for which they are known.⁶

C. Latent Diffusion Models

Latent Diffusion Models (e.g., DALL•E and Midjourney) rely on “denoising” diffusion technology, i.e., generative neural networks that learn to produce images through training on vast quantities of images via an iterative denoising process.⁷ This process begins by using “forward” diffusion on a training image, incrementally adding random noise to that image until such original image becomes indistinguishable, and then using reverse diffusion to restore the original image. To achieve this denoising, generative AI models measure the amount of added noise using a neural network known as a noise predictor.⁸ Once the noise predictor is set up, the model can be conditioned to guide it, ensuring the predicted noise, when subtracted, yields the desired result. Because denoising can be time-consuming in image space, latent diffusion models like Stable Diffusion address this issue by streamlining the process; they first compress the training image into latent space before engaging in forward and reverse diffusion training.

III. THE COPYRIGHTABILITY OF GENERATIVE AI WORKS

A. Does U.S. Copyright Law Protect Works Produced by Generative AI?

The recent popularity of generative AI has sparked much debate over whether U.S. copyright law protect works created by artificial intelligence. The current weight of authority strongly suggests that works created through generative AI platforms are not protected by copyright unless the work contains sufficient “human input” to “qualify the user of an AI system as an ‘author’ of a generated work.”⁹ This is due to the U.S. Copyright Office’s Human Authorship Requirement, i.e., only a *human* can create works protected by U.S. copyright law.

B. The Human Authorship Requirement

⁴ Laya Neelakandan, *What Is A Large Language Model And How Does It Work?*, Fast Company, <https://www.fastcompany.com/90884581/what-is-a-large-language-model> (April 20, 2023).

⁵ *Id.*

⁶ Huzma Naveed, *et al.*, A COMPREHENSIVE OVERVIEW OF LARGE LANGUAGE MODELS, J. Latex, <https://arxiv.org/pdf/2307.06435.pdf> (Oct. 5, 2023).

⁷ Nicholas Carlini, *et al.*, EXTRACTING TRAINING DATA FROM DIFFUSION MODELS, 32nd USENIX Security Symposium, <https://www.usenix.org/system/files/usenixsecurity23-carlini.pdf> (Aug. 9, 2023).

⁸ Andrew, *How does Stable Diffusion work?*, Stable Diffusion Art, <https://stable-diffusion-art.com/how-stable-diffusion-work/> (Sept. 27, 2023).

⁹ *Thaler v. Perlmutter*, No. CV 22-1564 (BAH), 2023 WL 5333236, at *6 (D.D.C. Aug. 18, 2023).

Neither the Copyright Clause of the U.S. Constitution (which refers to “authors” but does not define them) nor the Copyright Act expressly state that an author must be human.¹⁰ The Human Authorship Requirement is a creature of case law, Office guidance, and Office registration decisions.¹¹ This Requirement was first set forth by the Office in its 1973 Compendium of U.S. Copyright Office Practices, which provides that the Office would not register material that did not “owe their origin to a human agent.” This Requirement is also described in later versions of the Compendium, including the current version, which explains that the following is the “crucial question” to answer when considering the copyrightability of a computer-generated work:

Whether the ‘work’ is basically one of human authorship, with the computer [or other device] merely being an assisting instrument, or whether the traditional elements of authorship in the work (literary, artistic, or musical expression or elements of selection, arrangement, etc.) were actually conceived and executed not by man but by a machine.¹²

In support of the Human Authorship Requirement, the Compendium cites the U.S. Supreme Court’s decision in *Trade-Mark Cases*, 100 U.S. 82, 94 (1879) for the proposition that “copyright law only protects ‘the fruits of intellectual labor’ that ‘are founded in the creative powers of the mind’,” and *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884) for the proposition that “copyright law is limited to ‘original intellectual conceptions of the author’.”¹³ While those decisions arguably imply that “authors” must be human, others have noted that the “fragments” in these early cases “do not resolve the question whether the Constitution requires human authorship.”¹⁴ Few subsequent cases have addressed the human authorship requirement.¹⁵ This is because, in general, “very few courts have been called upon to

¹⁰ See 17 U.S.C. § 101 *et seq.*; see also *Urantia Found. v. Maaherra*, 114 F.3d 955, 958 (9th Cir. 1997) (“The copyright laws, of course, do not expressly require ‘human’ authorship, and considerable controversy has arisen in recent years over the copyrightability of computer-generated works.”) (citing Arthur R. Miller, Copyright Protection for Computer Programs, Databases, and Computer-Generated Works: Is Anything New Since CONTU?, 106 Harv.L.Rev. 977 (1993)); Annemarie Bridy, The Evolution of Authorship: Work Made by Code, 39 Colum. J.L. & Arts 395, 399 (2016).

¹¹ 17 U.S.C. § 102 sets forth the “subject matter of copyright.” Subsection (a) provides that “[c]opyright protection subsists, in accordance with this title, in original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device.” Subsection (a) also lists various types of works (e.g., literary works, musical works, etc.) that fall within the ambit of copyright, and 17 U.S.C. § 102(b) excludes various items (e.g., “system[s]” and “method[s] of operation”) from protection. Nowhere does Section 102 expressly set forth a human authorship requirement. Nor does Section 101 define “author” or “authorship,” and while some statutory provisions suggest that authorship is necessarily human, others may suggest otherwise. See 17 U.S.C. § 201.

¹² Compendium of U.S. Copyright Office Practices (3d ed. 2021) at § 313.2, quoting U.S. Copyright Office, Report to the Librarian of Congress by the Register of Copyrights 5 (1966).

¹³ *Id.* at § 306. As the Human Authorship Requirement does not appear in the 1973 version of the Compendium, it seems that the Requirement was first expressly set forth in the 1984 version of the Compendium.

¹⁴ Arthur R. Miller, Copyright Protection for Computer Programs, Databases, and Computer-Generated Works: Is Anything New Since CONTU?, 106 Harv.L.Rev. 977, 1065 (1993).

¹⁵ That said, the Office has applied the Human Authorship Requirement in its registration determinations. See U.S. Copyright Office Review Board Opinions, available at <https://www.copyright.gov/rulings-filings/review-board/> (wherein a user can select the “Human authorship” check box, which returns 19 Review Board opinions issued since 1995 involving application of the Human Authorship Requirement).

address the issue explicitly, and those courts that have been forced to address the issue usually have held, implicitly or explicitly, that machines involved in the creation of artistic works are simply tools that assist the human beings who employ them.”¹⁶

On March 16, 2023, the Office published guidance entitled “Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence,” in which it cited two Ninth Circuit cases dealing with works allegedly authored by spiritual beings and animals for the proposition that human authorship is required.¹⁷ For example, in *Urantia Found. v. Maaherra*, 895 F. Supp. 1337, 1338 (D. Ariz. 1995), a district court found that a religious work allegedly authored by divine beings possessed the “minimal degree” of creativity required by the Supreme Court’s decision in *Feist Publ’ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340 (1991), and that it was not necessary that the copyright claimant be the author, “[n]or [was] it necessary that the authorship stem from human effort.” The Ninth Circuit reversed, finding that “it is not creations of divine beings that the copyright laws were intended to protect, and that in this case some element of human creativity must have occurred in order for the Book to be copyrightable.” *Urantia Found. v. Maaherra*, 114 F.3d 955, 958 (9th Cir. 1997).

C. U.S. Copyright Registrations for Computer-Generated Works

In a few cases, the Office has registered computer-generated works; those registrations are difficult to square with the Human Authorship Requirement. For example, in November 1984, the Office issued a copyright registration for a computer-generated book entitled, *The Policeman’s Beard is Half Constructed*; such registration covers “computer prose and poetry / by Racter.”¹⁸ Racter—short for Raconteur—was a text generation system “written in compiled BASIC on a Z80 micro with 64K of RAM” by William Chamberlain and Thomas Etter using INRAC, a programming language they developed.¹⁹ Racter’s prose in *The Policeman’s Beard* has been described as “disjointed,” but not without “a semblance of narrative through repeated phrases and names.”²⁰ A 1984 Wall Street Journal article described Racter’s work as follows:

Racter’s method is a complicated blend of haphazardness and linguistic savvy. The program basically strings words and phrases together randomly, but it has two important constraints. It contains rules of English, so Racter speaks grammatically. In addition, it contains enough information about each word in its 2,400-word vocabulary to let Racter put together meaningful sentences. Racter knows, for instance, that a grammatical sentence like “Sarah coldly flew the parking meter” is senseless, while “Sarah ravenously ate the lettuce” is sensible. (Racter’s lexicon also includes a handful of prefabricated sentences, proverbs and quotations that it periodically interjects.)²¹

¹⁶ *Id.* at 1059-60; *see also Thaler v. Perlmutter*, No. CV 22-1564 (BAH), 2023 WL 5333236, at *6 (D.D.C. Aug. 18, 2023).

¹⁷ 88 FR 16190.

¹⁸ *See* Registration No. TX0001454063.

¹⁹ *Id.* at 1, 3.

²⁰ *Id.* at 7.

²¹ Henrickson, Leah. “Constructing the Other Half of The Policeman’s Beard”, *Electronic Book Review*, April 4, 2021, <https://doi.org/10.7273/2bt7-pw23>.

In June 1993, the Office registered a computer-generated work titled *Just This Once*.²² Scott French, who “identified two hundred idiosyncrasies” in the writing style of best-selling author Jacqueline Susann, turned them into six thousand “rules,” and incorporated those rules into a computer program that French used to generate *Just This Once*.²³ Per French, “he wrote about a quarter of the prose, the computer wrote about the same amount, and the remainder was a collaboration of man and machine.”²⁴

At a minimum, these isolated examples raise questions about the scope and application of the Human Authorship Requirement and contextualize the Office’s recent decisions refusing registration to works that were created using generative AI technology.

D. *Zarya of the Dawn*

On February 21, 2023, the Office cancelled a copyright registration issued to artist Kristina Kashtanova for a comic book, *Zarya of the Dawn*, after discovering statements on social media that confirmed that the comic book’s images (examples below) were generated using the Midjourney generative AI platform.²⁵



In doing so, the Office considered the copyrightability of the individual components of the work, determining that the text (written entirely by Kashtanova) was copyrightable, as was Kashtanova’s selection and arrangement of the images and texts, and re-issued a narrowed registration to Kashtanova that covered only those aspects of the work.

²² See Registration No. TX0003633395.

²³ Robert C. Denicola, *Ex Machina: Copyright Protection for Computer-Generated Works*, 69 Rutgers U.L. Rev. 251, 267 (2016).

²⁴ *Id.*

²⁵ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn (VAu001480196)* at 3 (Feb. 21, 2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.

In finding that the Midjourney-generated images in Kashtanova’s work were not copyrightable, the Office reasoned that Midjourney generates images in an “unpredictable way,” such that Kashtanova did not “actually form” the images in her work. The Office reasoned that Midjourney’s unpredictability distinguishes it from other artistic tools, such as cameras, where artists can control the final image based on their “own original mental conception, to which [they] gave visible form.”²⁶ The Office rejected Kashtanova’s claim of ownership based on “creative, human-authored prompts,” reasoning that Midjourney prompts function as “suggestions” rather than direct “orders.”

While the Office maintained that non-human authored works are not copyright-protected, it left room for generative AI authorship, stating that “[i]t is possible that other AI offerings that can generate expressive material operate differently than Midjourney does.”²⁷

E. *Thaler v. Perlmutter*

On August 18, 2023, the U.S. District Court for the District of Columbia became the first federal district court to weigh in on the copyrightability of generative AI works following the Office’s refusal to register *A Recent Entrance to Paradise* (shown below):



A Recent Entrance to Paradise was apparently created by Dr. Stephen Thaler’s Creativity Machine, an artificial neural network that can generate two-dimensional artwork. Thaler filed an application to register the work with the U.S. Copyright Office, listing himself as claimant and Creativity Machine as author. The Office refused registration on the grounds that the work ran afoul of the Human Authorship Requirement.²⁸ In support of the Requirement, the Office relied on cases discussed *supra*, as well as *Naruto v. Slater*, 888 F.3d 418, 426 (9th Cir. 2018), the “monkey selfie” case in which a monkey’s self-portrait photograph was denied authorship because the Act’s terms “all imply humanity and necessarily exclude animals.” Thaler appealed the decision to the United States Copyright Office Review Board, which affirmed the decision.

²⁶ *Burrow-Giles Lithographic Co. v. Sarony Burrows*, 111 U.S. 53, 55 (1884).

²⁷ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn (VAu001480196)* at 10 (Feb. 21, 2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.

²⁸ *Second Request for Reconsideration for Refusal to Register A Recent Entrance to Paradise* (Copyright Review Board, Feb. 14, 2022).

Thaler filed suit against the Office to challenge its decision in the U.S. District Court for the District of Columbia, arguing that there was no constitutional or statutory basis for the Human Authorship Requirement, and that the work-made-for-hire doctrine set forth in 17 U.S.C. § 201(b)—which permits corporations to be “authors” for purposes of U.S. copyright law—confirms this. The Court disagreed and, on August 18, 2023, resolved the case in favor of the Office.²⁹ The Court found that “[h]uman authorship is a bedrock requirement of copyright” and “[t]hat principle follows from the plain text of the Copyright Act.”³⁰ The Court was unpersuaded by Thaler’s arguments that the work-made-for-hire doctrine demonstrated that authorship need not be human, reasoning that such arguments “concern *to whom* a valid copyright should have been registered” when the absence of human involvement “le[ft] nothing at all to register and thus no question as to whom that registration belonged.”

In a footnote, the Court mentioned that Thaler’s arguments illustrated efforts to “imply[] a level of human involvement” that was “entirely absent in the administrative record.”³¹ Thus, despite the broad statements made by the Court regarding the Human Authorship Requirement, its holding was relatively narrow insofar as it assumed, based on the record before it, that Thaler “played no role in using the AI to generate the work” and essentially ignored Thaler’s arguments that he “provided instructions and directed his AI to create the Work,” and that “the AI only operates at [his] direction.” This allowed the Court to easily conclude that the Office had not acted “arbitrarily and capriciously” in refusing registration.

The Court also acknowledged “new frontiers in copyright as artists put AI in their toolbox to be in the generation of new visual and other artistic works” and the many questions around the copyrightability of generative AI output that remain, including the quantum of human input necessary to qualify the user of an AI system as an “author” of a generated work, the scope of the protection obtainable over the resulting image, the manner in which to assess the originality of AI-generated works from platforms trained on vast undisclosed data sets, and how copyright law can best incentivize creative uses of generative AI.

F. Can AI-Generated Works Ever Constitute “Human Authorship”?

The Office’s view that human authorship is required for generative AI outputs to be protected by copyright makes sense; to the extent that current generative AI models “can’t produce a work that reflects its own ‘original intellectual conception’ *because it has none*,”³² there seems little basis to view those models as “authors,” or to provide them with the legal and economic incentives to create that U.S. copyright law provides to authors.

Nevertheless, the application of the Human Authorship Requirement raises important questions. One such question is whether the *predictability* that the Office’s recent decisions have required from users of generative AI to “author” works is commensurate to the predictability that

²⁹ *Thaler v. Perlmutter*, No. CV 22-1564 (BAH), 2023 WL 5333236, at *6 (D.D.C. Aug. 18, 2023).

³⁰ *Id.* at *4

³¹ *Id.* at *3 n. 1.

³² *Artificial Intelligence and Intellectual Property – Part II: Copyright and Artificial Intelligence: Hearing Before the Judiciary S. Comm. on Intellectual Property*, 117th Cong. 3 (2023) (testimony of Matthew Sag, Professor of Law in Artificial Intelligence, Machine Learning, and Data Science, Emory University School of Law) (“Sag Testimony”) at 2 (emphasis added).

copyright law has required from authors of traditional media. Indeed, to create artworks—many of which are generally understood to be at the “core” of copyright protection—artists routinely take risks and experiment with different tools, mediums and modes of creation, all of which is, by definition, unpredictable. What would become of those works, and the incentive to create them, if predictability were strictly enforced?

The copyrightability of photographic works provides a useful comparison. Under current copyright law, it is practically beyond dispute that a photograph is protected by copyright, regardless of whether the process of capturing that photograph was “unpredictable” or not.³³ That is because a body of case law has developed over more than a century, beginning with the Supreme Court’s *Burrow-Giles* decision, to reach the conclusion that “no photograph, however simple, can be unaffected by the personal influence of the author, and no two will be absolutely alike” such that photographs are protected “without regard to the degree of ‘personality’ which enters into them.”³⁴ Generative AI users, however, now seem to be held to a higher standard; as they now toil under a cloud of presumptive non-authorship, they must prove that they exercise “sufficient control over generated images to be treated as the ‘master mind’ behind them.”³⁵

The copyrightability of audiovisual works provides another useful comparison. A director may author a film by “superintend[ing]” the work, exercising creative control, and by “actually form[ing] the picture by putting the persons in position, and arranging the place where the people are to be” such that the director is “the inventive or master mind” who “creates, or gives effect to the idea.” *Aalmuhammed v. Lee*, 202 F.3d 1227, 1234 (9th Cir. 2000). This form of authorship (and somewhat attenuated fixation of expression) is possible under existing copyright law, even though film directors cannot exercise direct control over an actor or cameraperson, nor *predict* the details of their performances.³⁶ Instead, the film directorial process tends to be iterative, i.e., a director selects and rejects individual contributions that actors and others make through many “takes,” final cut, and recasting roles.³⁷ Arguably, human actors are no more predictable than generative AI, and the iterative process by which a director selects, rejects, and edits to reach the desired result is analogous to a generative AI user’s process in creating numerous images until it finds the “right” one.

Even assuming an author must exercise a certain quantum of creative control to qualify as an author, how much is enough? While the Office has thus far refused registration of several AI-generated works, it has left the door open to copyright protection for AI-generated works through platforms that “operate differently than Midjourney does.”³⁸ Indeed, some platforms may be intentionally assembled to strengthen the connection between user inputs and creative outputs,

³³ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn* (VAu001480196) at 9 (Feb. 21, 2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.

³⁴ *Jewelers’ Circular Publishing Co. v. Keystone Publishing Co.*, 274 F. 932, 934 (S.D.N.Y.), *aff’d*, 281 F. 83 (2d Cir. 1921), *cert. denied*, 259 U.S. 581 (1922).

³⁵ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn* (VAu001480196) at 9 (Feb. 21, 2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.

³⁶ *See, e.g.*, <https://www.backstage.com/magazine/article/camera-improv-secrets-5882/> (“Film directors... are known for letting their actors improvise on certain takes.”)

³⁷ *Id.*

³⁸ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn* (VAu001480196) at 10 (Feb. 21, 2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.

i.e., to make those prompts “orders” rather than “suggestions.”³⁹ Adobe Firefly, for example, has been marketed as a generative AI platform trained on licensed data sets, and Adobe now claims that Firefly offers “new photo settings” that “enable you to have new creative control just like you would with a manual camera, and we’re always improving our text prompt capabilities.”⁴⁰ If so, and assuming proper permissions to allow others to own what they create through Firefly, the Office may find content generated from such a platform is protectable.

Lastly, statutory copyright protection of computer-generated works in other jurisdictions, e.g., the U.K., raises the question of whether the U.S. should follow suit, particularly given the advent of generative AI.⁴¹ It remains to be seen whether Congress will act to expand protection for AI-generated works, as it did when extending copyright protection to software.

IV. COPYRIGHT LIABILITY ARISING FROM GENERATIVE AI MODELS

Generative AI platforms are now coming under scrutiny for unlicensed use of copyrighted works within their training datasets. To date, at least nine copyright infringement lawsuits have been filed within the U.S. against generative AI platform companies.⁴² Additionally, generative AI technology may be poised to give rise to other kinds of infringement claims, e.g., claims against corporate users of generative AI. Do any of these claims have merit? The answer depends on whether the plaintiffs in these cases can establish the elements of infringement, and whether defendants can successfully assert fair use or other defenses.

A. Copyright Claims Against Generative AI Platforms

A copyright infringement plaintiff must prove ownership and copying of protectable expression.⁴³ The copyright infringement lawsuits filed against generative AI platform companies have focused on two theories: infringing outputs and infringing training sets.

1. Allegedly Infringing Outputs

³⁹ *Id.*

⁴⁰ Ashley Still, Senior VP of Digital Media, Adobe (Wall Street Journal Tech News Briefing, October 11, 2023), available at <https://www.wsj.com/podcasts/tech-news-briefing/how-adobe-plans-to-stand-out-from-the-ai-pack/180ef877-408e-457f-94b4-612352d8039b>.

⁴¹ See Copyright, Designs and Patents Act, 1988, c. 1, § 9(3) (U.K.) (“In the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken.”)

⁴² *Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201-WHO (N.D. Cal. Jan. 13, 2023); *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW (D. Del. Mar. 29, 2023); *Silverman v. OpenAI, Inc.*, No. 3:23-cv-03416-AMO (N.D. Cal. Sept. 27, 2023); *Chabon v. Meta Platforms, Inc.*, No. 3:23-cv-04663-VC (N.D. Cal. Oct. 5, 2023); *Kadrey et al. v. Meta Platforms*, No. 3:23-cv-03417-VC (N.D. Calif. Jul. 7, 2023); *Thomson Reuters Enter. Ctr. GmbH v. ROSS Intel. Inc.*, No. 1:20-cv-00613-SB (D. Del. May 6, 2020); *Tremblay et al v. OpenAI, Inc. et al.*, No. 3:23-cv-03223-AMO (N.D. Cal. Jun. 28, 2023); *Authors Guild v. OpenAI et al.*, No. 1:23-cv-8292-SHS (S.D.N.Y. Sep. 19, 2023); *Chabon v. OpenAI, Inc.*, No. 3:23-cv-04625-AMO (N.D. Cal. Sep 8, 2023); *J.L. et al. v. Alphabet Inc. et al.*, No. 3:23-cv-3440-AMO (N.D. Cal. Jul. 11, 2023); *Doe 1 v. GitHub Inc.*, No. 4:22-cv-06823-JST (N.D. Cal. Nov. 3, 2022).

⁴³ *Feist*, 499 U.S. at 361.

With respect to infringing outputs, plaintiffs face the challenge of proving that generative AI outputs are “substantially similar” to the protected work and, in cases where copyrighted content is “thin” (e.g., highly factual in nature), virtually identical to the protected work.⁴⁴ Whether generative AI outputs tend to incorporate a sufficient percentage of a particular work to satisfy these legal tests is unclear.

The disparate allegations copyright infringement plaintiffs have made regarding the Stable Diffusion platform are illustrative; whereas some plaintiffs have alleged that Stable Diffusion’s output is not “likely to be a close match for any specific [content] in the training data,”⁴⁵ others have alleged that this same technology can produce images that are “highly similar to or derivative of” protected works.⁴⁶ Others have alleged that generative AI platforms can produce “very accurate summaries” of protected works.⁴⁷ While these allegations may implicate several exclusive rights under the Copyright Act, including the right to reproduction and to create derivative works, it remains to be seen whether generative AI outputs have copied sufficient protectible expression to infringe those exclusive rights.⁴⁸

Ultimately, whether generative AI outputs infringe copyright is likely to depend on the generative AI model at issue, and specifically, whether the prompts that users input have the potential to call up or “invoke” copyrighted works using a specific generative AI model, whether the generative AI platform has “guardrails” in place to prevent or mitigate infringement risk, and the size and depth of the generative AI training dataset itself.

2. Allegedly Infringing Training Sets

A separate question is whether generative AI models were trained in a manner that infringes U.S. copyright law. This question has been raised in the many lawsuits against generative AI platforms,⁴⁹ and is complicated by the inner workings of the technologies employed by generative AI platforms themselves, and the content that was used to train those foundational models.⁵⁰ Such information is, to some extent, not publicly available, but raises factual questions that are likely to be answered in connection with the various copyright infringement litigations that are now pending against generative AI platforms.

Some plaintiffs have alleged that infringement has occurred in the compilation and use of training datasets based on platform outputs, i.e., they claim that if a model can accurately

⁴⁴ *Satava v. Lowry*, 323 F.3d 805, 812 (9th Cir. 2003).

⁴⁵ Compl. ¶ 93, *Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201-WHO, ECF No. 1. (N.D. Cal. Jan. 13, 2023).

⁴⁶ See Am. Compl. ¶¶ 8, 61, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW, ECF No. 13 (D. Del. Mar. 29, 2023) (alleging infringement based on use of copyrighted images to train a generative AI model and on the possibility of that model generating images “highly similar to and derivative of” copyrighted images).

⁴⁷ Compl. ¶ 42, *Silverman v. OpenAI, Inc.*, No. 3:23-cv-03416, ECF No. 1. (N.D. Cal. July 7, 2023).

⁴⁸ 17 U.S.C. § 106 (2002).

⁴⁹ See *supra* n.42.

⁵⁰ “Foundation models” are “trained on a broad set of unlabeled data that can be used for different tasks, with minimal fine-tuning.” See What are foundation models, IBM (last visited Oct. 19, 2023), <https://research.ibm.com/blog/what-are-foundation-models>.

summarize a work,⁵¹ or that if a model can generate an image that appears to contain a mutilated watermark from plaintiff’s work,⁵² then a copy exists somewhere on the back end of these generative AI models or, at a minimum, that infringement occurred during the training process. Other cases allege infringement in training datasets based on information that the generative AI model developer has allegedly published. For example, Plaintiffs in *Chabon v. OpenAI* allege that Meta’s LLaMA training dataset originates in part from Project Gutenberg, an online archive of approximately 70,000 books that are out of copyright, and the Books3 dataset of ThePile, which is allegedly derived from a copy of the contents of the Bibliotik private tracker that consists of copyrighted fiction and nonfiction material.⁵³

That said, there is no clear consensus on whether copyrighted content is actually *copied* during the iterative process of training a generative AI platform. Some claim that generative AI models “typically don’t copy the data in any literal sense.”⁵⁴ At least with respect to LLMs, this may be due to the fact that a generative AI model does not actually digest content in entire sentences or words, but instead does so in “subword tokens” in order to determine the probability that one subword token will follow another in order to form words, phrases, and sentences.⁵⁵ Yet even if “the things [generative AI models] learn from the training data are generally fairly abstract and thus uncopyrightable,” algorithm “training” may itself involve the creation of “interim” copies, which itself could give rise to infringement liability.⁵⁶

Assuming that some generative AI training sets contain unlicensed content that was incorporated through actions that infringe copyright, questions remain about theories of liability against “downstream” platforms that use technology developed by *other entities*. For example, at a July 19, 2023 hearing on defendant’s motions to dismiss in the ongoing litigation in *Andersen v. Stability AI Ltd.*, Judge Orrick questioned whether the kinds of direct infringement claims that plaintiffs pled against Stability AI could also be pled against Midjourney and DeviantArt, i.e., platforms that did not actually perform foundational model training of the

⁵¹ See *supra* n.47.

⁵² See Am. Compl. ¶¶ 8, 61, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW, ECF No. 13 (D. Del. Mar. 29, 2023).

⁵³ See Am. Compl. ¶ 31, *Chabon v. Meta Platforms, Inc.*, No. 3:23-cv-04663-VC, ECF No. 16 (N.D. Cal. Oct. 5, 2023); see also Compl. ¶ 34, *Tremblay et al v. OpenAI, Inc. et al.*, No. 4:23-cv-03223-AMO, ECF No. 1 (N.D. Cal. Jun. 28, 2023).

⁵⁴ Sag Testimony at 3; but see *id.* (“In very rare cases when [generative AI models] do copy the training data—something computer scientists call ‘memorization’—that is regarded as a bug to be fixed, not a desirable feature.”)

⁵⁵ Sag Testimony at 4. As Professor Sag explains, “GPT-3 predicts subword tokens. For example, the word “unintentionally” can be broken down into smaller parts, or subword tokens, that are still meaningful such as “un”, “intention”, “al”, “ly”. The tokenization process allows the model to handle a wide range of words, including those it has not seen during training.”

⁵⁶ See generally, *MAI Sys. Corp. v. Peak Comput., Inc.*, 991 F.2d 511, 519 (9th Cir. 1993) (“the loading of software into the RAM creates a copy under the Copyright Act.”) (citation omitted); but see *Cartoon Network LP, LLLP v. CSC Holdings, Inc.*, 536 F.3d 121, 127-28 (2d Cir. 2008) (“we construe MAI Systems and its progeny as holding that loading a program into a computer’s RAM can result in copying that program. We do not read MAI Systems as holding that, as a matter of law, loading a program into a form of RAM always results in copying. Such a holding would read the ‘transitory duration’ language out of the definition [of ‘copies’], and we do not believe our sister circuit would dismiss this statutory language without even discussing it.”)

Stability AI model, but use that model to offer generative AI services.⁵⁷ Plaintiffs contend that “the way the copyright law operates and attaches liability kind of downstream with respect to infringing uses” renders those secondary generative AI platforms liable for infringement “they are using that [allegedly infringing product] and selling it and putting it into the marketplace generally gives rise to the claims.”⁵⁸

3. DMCA Violation Claims Against Generative AI Platforms

Several of the copyright infringement cases against generative AI platforms also involve alleged violations of the Digital Millennium Copyright Act (“DMCA”). Section 1202(a) of Title 17 prohibits knowingly providing false copyright management information (“CMI”), or distributing/importing false CMI, with intent to induce, enable, facilitate or conceal infringement. Section 1202(b) prohibits any of the following, if undertaken knowingly, or, “with respect to civil remedies under section 1203, having reasonable grounds to know, that it will induce, enable, facilitate, or conceal an infringement” of copyright: (1) removal or alteration of CMI, (2) distribution or importation of removed or altered CMI, and (3) distribution or importation of copies of works with removed or altered CMI.

The majority of DMCA violations in these cases thus far appear to fall under Section 1202(b). For example, Getty Images has alleged that Stability AI has removed the “gettyimages” watermark from images used to train the Stability AI model.⁵⁹ Similarly, plaintiffs in the *Tremblay* case against OpenAI allege that “OpenAI intentionally removed CMI from the Plaintiffs’ Infringed Works in violation of 17 U.S.C. § 1202(b)(1)” at least because “[b]y design, the training process does not preserve any CMI.”⁶⁰ Plaintiffs in the *Tremblay* case also allege that OpenAI “created derivative works” by allegedly incorporating plaintiffs’ works into a training dataset and, by “distributing these works without their CMI,” violated 17 U.S.C. § 1202(b)(3).⁶¹

One major hurdle to these CMI claims will be the “double-scienter” requirement.⁶² Particularly with respect to “downstream” defendants that did not train a foundational model, but merely used that technology to offer independent services (e.g., a platform like Midjourney), it may be difficult to show that those defendants *knew* that CMI was actually removed from works they never identified or directly incorporated into a training dataset. If so, plaintiffs may be forced to argue that defendants had generalized knowledge of infringement based on a generative

⁵⁷ Hearing Transcript on Motions to Dismiss, *Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201-WHO (N.D. Cal. Jul. 19, 2023) (“Anderson Tr.”) at 4:16-19.

⁵⁸ *Id.* at 15:20-24.

⁵⁹ See Am. Compl. ¶ 68, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW, ECF No. 13 (D. Del. Mar. 29, 2023).

⁶⁰ Compl. ¶ 64, *Tremblay et al v. OpenAI, Inc. et al.*, No. 3:23-cv-03223-AMO, ECF No. 1 (N.D. Cal. Jun. 28, 2023).

⁶¹ *Id.* ¶ 65.

⁶² *Mango v. BuzzFeed, Inc.*, 970 F.3d 167, 171 (2d Cir. 2020).

AI training process that is allegedly designed not to preserve CMI, which has unclear chances of success.⁶³

It is also unclear whether generative AI defendants have “reasonable grounds to know” that the removal, alteration, distribution, or importation of CMI “will induce, enable, facilitate, or conceal an infringement,” particularly if the infringement of copyright does not occur *in the functioning of the platform itself* but only “when the training corpus is assembled and pre-processed.”⁶⁴ Because it is still unclear whether infringement occurs on generative AI platforms, defendants may be able to convincingly argue that they did not have reasonable grounds to know that CMI violations would induce, enable, facilitate, or conceal infringement given the novelty and legal uncertainty surrounding generative AI technology.

B. Defenses to Copyright Claims Against Generative AI Platforms

1. Fair Use

Although the generative AI litigations referenced above are still in their early stages, defendants in those cases are likely to raise fair use defenses to infringement claims premised on outputs and training datasets. Courts considering a fair use defense to copyright infringement will look to the four statutory fair use factors:

- (1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
- (2) the nature of the copyrighted work;
- (3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
- (4) the effect of the use upon the potential market for or value of the copyrighted work.⁶⁵

Fair use defenses to infringement claims premised on generative AI training datasets will likely depend on how courts view the purpose and character of training set uses, as well as their effect on the plaintiff’s reasonable licensing market for training datasets. As for the first-factor analysis, the purpose and character of the use, analogies might be drawn to the Second Circuit’s *Authors Guild* decisions. In *Authors Guild, Inc. v. HathiTrust*, the Second Circuit held that copying printed library books to create a searchable digital research repository was “quintessentially” transformative.⁶⁶ Again, in *Authors Guild, Inc. v. Google Inc.*, the Second Circuit held that the copying involved in the Google Books project was fair use due to its “highly transformative purpose” of building a digital search tool that achieved something the underlying works could not, and thus did not compete with those underlying works.⁶⁷ Generative AI platforms may also have a “highly transformative purpose” because they take traditional content

⁶³ Compl. ¶ 64, *Tremblay et al v. OpenAI, Inc. et al.*, No. 3:23-cv-03223-AMO, ECF No. 1 (N.D. Cal. Jun. 28, 2023).

⁶⁴ Sag Testimony at 4.

⁶⁵ 17 U.S.C. § 107 (2023).

⁶⁶ *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 97 (2d Cir. 2014).

⁶⁷ *Authors Guild v. Google, Inc.*, 804 F.3d 202, 216, 218 (2d Cir. 2015)

and repurpose it in order to develop a creative tool that serves a fundamentally different purpose from the works that were copied to create it. One might argue that this purpose is at least as “transformative” as other “highly creative and innovative” uses that courts have protected with the fair use doctrine.⁶⁸ Analogies to other cases can be drawn as well. For example, some have drawn analogies to reverse engineering cases, arguing that generative AI training is fair use because it is a *non-expressive use* akin to reverse engineering.⁶⁹

Regardless of which first-factor arguments are made in support of fair use, courts will need to consider the Supreme Court’s decision in *Warhol v. Goldsmith*. In *Warhol*, the Court analyzed the first fair use factor by focusing on the Andy Warhol Foundation’s *licensing* of ‘Orange Prince’—not on the transformative value of ‘Orange Prince’ itself—and concluded that because that kind of licensed use (featuring a celebrity portrait in a magazine) was very similar to plaintiff’s use, the first factor favored plaintiff. Thus, a court considering a fair use defense to generative AI training infringement might focus less on the transformativeness of the generative AI model itself, and more on the precise use of the work in relation to the generative AI model. For example, a temporary or ephemeral use of a copyrighted work for training a generative AI model that never *generates* the original seems more likely to be favored in a first factor analysis than a permanent use in generative AI “database” in which users can invoke and incorporate the asserted work into derivative works that serve the same purpose as the original.

This begs the question of whether the first fair use factor favors a generative AI training use that itself may be fundamentally different from the purpose and character of the original work, but which gives rise to a host of downstream uses that serve a similar or identical purpose as the original, and may even be commercial substitutions for the original work. Courts are likely to analyze this under the fourth fair use factor, i.e., by considering the relevance and plausibility of a licensing market for generative AI training dataset uses. In fact, some plaintiffs have anticipated this argument, alleging that the infringement has impacted their licensing market for generative AI training sets.⁷⁰ While it remains to be seen whether plaintiffs will be able to prove that a genuine dataset licensing market exists, that may be unnecessary; past courts have found that “likely to be developed” licensing markets have sufficed to shift the fourth fair use factor in a defendant’s favor.⁷¹ If so, courts may consider whether the allegedly impacted licensing market is speculative or realistic for the specific parties and works before the court.

Fair use may also be an applicable defense to infringement claims premised on generative AI outputs, assuming that plaintiffs can establish substantial similarity. In these cases, the strength of a fair use defense will likely depend on the nature of the challenged use. At one extreme, an expressive output that is substantially similar to the original may prove to be a

⁶⁸ *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183, 1203 (2021).

⁶⁹ *A.V. v. iParadigms, Ltd. Liab. Co.*, 544 F. Supp. 2d 473, 482 (E.D. Va. 2008).

⁷⁰ See Am. Compl. ¶ 66, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW, ECF No. 13 (D. Del. Mar. 29, 2023) (“Getty Images has licensed millions of suitable digital assets for a variety of purposes related to artificial intelligence and machine learning in a manner that respects personal and intellectual property rights.”)

⁷¹ *Am. Geophysical Union v. Texaco Inc.*, 60 F.3d 913, 930 (2d Cir. 1994) (finding that the “impact on potential licensing revenues for traditional, reasonable, or likely to be developed markets should be legally cognizable when evaluating a secondary use’s “effect upon the potential market for or value of the copyrighted work.”)

commercial substitute that shifts the fourth factor decidedly against fair use.⁷² At the other extreme, a user can leverage a generative AI platform as a “parody machine,” which may shift the first factor in favor of fair use.⁷³ Ultimately, because the potential universe of AI outputs is so broad, it is difficult to generalize regarding the strength of fair use defenses in these cases.

2. Other Defenses

While fair use gets most of the “attention” when it comes to possible infringement defenses, other defenses may apply. *De minimus non curat lex* (“the law does not concern itself with trifles”) “insulates from liability those who cause insignificant violations of the right of others.”⁷⁴ While the typical example of *de minimus* use tends to be small and/or fleeting uses of copyrighted material in the background of an image or scene, the copying (if any) involved in the creation and/or use of generative AI training datasets may be *de minimus* to the extent that copies are not stored, but are merely transitory “cache” copies for generative AI models to “learn” from. *De minimus* defenses may also be raised on the basis that no single copyrighted work will have a significant impact on a training of a generative AI model that uses billions of texts, images or other content to train, especially if the training dataset also includes licensed and/or public domain works.

The Section 512 “Safe Harbor” of the Digital Millennium Copyright Act may also come into play as users begin to incorporate copyrighted works in generative AI model prompts, e.g., by uploading an image to guide the model toward a desired result. Depending on the technology behind the generative AI model, and whether the model stores user inputs to further train the model, content uploads as prompts may indeed be “stored at the direction of a user” and contain infringement of which the generative AI platform service provider is unaware.⁷⁵ While some platforms are only just beginning to offer functionality that allows users to upload content to guide generative AI models toward a specific result, this defense may become increasingly relevant as that functionality becomes more prevalent.

C. Copyright Claims Against Generative AI Users

It remains to be seen whether and how users of a generative AI system could be held liable for infringing use of a generative AI model that was trained on unlicensed copyrighted datasets. Copyright infringement is a strict liability tort; while users may be unaware that training sets are unlicensed, or that the creation and use of generative AI platforms are infringing, they may still be held liable, albeit with potentially reduced damages for innocent or non-willful infringement.⁷⁶ There is also the question of whether generative AI platforms can be held secondarily liable under vicarious, contributory, or inducement infringement theories.

⁷² See *Andy Warhol Found. for Visual Arts, Inc. v. Goldsmith*, 11 F.4th 26, 50 (2d Cir. 2021), cert. granted, 142 S. Ct. 1412, 212 L. Ed. 2d 402 (2022), and aff’d sub nom. 598 U.S. 508, 143 S. Ct. 1258, 215 L. Ed. 2d 473 (2023).

⁷³ See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 579 (1994).

⁷⁴ *Ringgold v. Black Ent. Television, Inc.*, 126 F.3d 70, 74 (2d Cir. 1997)

⁷⁵ 17 U.S.C. § 512(c) (2010).

⁷⁶ 17 U.S.C. § 504 (2010).

V. GENERATING NEW QUESTIONS

Amid these myriad questions, more remain, and further questions will continue to arise as generative AI technologies continue to evolve and are adopted over time. It remains to be seen how copyright law will evolve on judicial and legislative fronts to meet these challenges.