

Generative AI's Illusory Case for Fair Use

Jacqueline C. Charlesworth*

ABSTRACT

Pointing to Authors Guild, Inc. v. Google Inc., Authors Guild, Inc. v. HathiTrust, Sega Enterprises Ltd. v. Accolade, Inc. and other leading technology-driven fair use precedents, artificial intelligence (AI) companies and those who advocate for their interests claim that mass unauthorized reproduction of books, music, photographs, visual art, news articles, and other copyrighted works to train generative AI systems is a fair use of those works. Though acknowledging that works are copied without permission for the training process, the proponents of fair use maintain that an AI machine learns only uncopyrightable information about the works during that process. Once trained, they say, the model does not incorporate or make use of the content of the training works. As such, they contend, copying for the purposes of AI training is a fair use under US law.

This Article challenges the above narrative by examining generative AI training and functionality. Despite wide employment of anthropomorphic terms to describe their behavior, AI machines do not learn or reason as humans do. Instead, they employ an algorithmic process to store the works they are fed during the training process. They do not “know” anything independently of the works on which they are trained, so their output is a function of the copied materials.

More specifically, large language models (LLMs) are trained by breaking textual works down into small segments, or “tokens” (typically individual words or parts of words), and converting the tokens into vectors—numerical representations of the tokens and where they appear in relation to other tokens in the text. The training works do not vanish, as suggested, but instead are encoded, token by token, into the model and relied upon to generate output. AI image generators are trained somewhat differently through a “diffusion” process in which they learn to reconstruct particular training images in conjunction with associated descriptive text. Like an LLM, however, an AI image generator relies on encoded representations of training works to generate its output.

The exploitation of expressive content to produce new expressive content sharply distinguishes AI copying from the copying at issue in the technological fair use cases relied upon by AI's fair use advocates. In these earlier cases, the determination of fair use turned on the fact that the alleged infringer was not seeking to capitalize on authors' creative expression. This is exactly the opposite of generative AI.

The fair use argument for generative AI is further hampered by the propensity of models to generate infringing copies and derivatives of training works. In addition, some AI models rely on retrieval-augmented generation (RAG) technology to generate output. RAG searches out and copies materials from online sources to augment and respond to user queries (for example, regarding an event that postdates the training of the LLM). Here again, copyrighted materials are being copied by generative AI without permission in order to exploit their expressive content.

For these and other reasons, each of the four fair use factors of Section 107 of the Copyright Act weighs against AI's claim of lawful use, especially when considered against the backdrop of a rapidly evolving market for licensed use of training materials.

TABLE OF CONTENTS

I.	INTRODUCTION	325
II.	AI SYSTEMS ARE NOT HUMAN BUT ALGORITHMICALLY DRIVEN.....	329
III.	WAYS IN WHICH GENERATIVE AI SYSTEMS ENGAGE WITH COPYRIGHTED WORKS.....	331
	A. <i>Core Activities of Generative AI Systems</i>	332
	1. Assembly of Training Materials	332
	2. Training of AI Models	333
	3. Generation of Copied Works and Derivatives	337
	4. Retrieval-Augmented Generation	344
	B. <i>Each Core Activity Involves Copying</i>	345
	1. Reproduction of Works to Create Training Sets	345
	2. Encoding of Works in AI Models	346
	3. Regeneration of Training Works in Output	350
	4. RAG Activities	350
IV.	FAIR USE ANALYSIS	351
	A. <i>The Cases Relied Upon by AI Companies</i>	351
	1. Earlier Cases Address Different Conduct	352
	2. There Is No Presumption that a Use Is Fair if the End Product Is Noninfringing	357
	B. <i>AI Copying Under the Fair Use Factors</i>	358

1. The Purpose and Character of AI Copying.....	358
a. AI Copying Is Not Transformative.....	358
b. Commercial Purpose of AI Copying	364
2. AI Companies Copy Highly Creative Works	365
3. AI Companies Copy Works in Their Entirety	365
4. AI Copying Harms the Market for the Copied Works	366
V. CONCLUSION.....	371

I. INTRODUCTION

A common refrain of generative AI companies,¹ and those who advocate for their interests, is that the unauthorized reproduction of copyrighted works to train and develop AI models is a fair use of those works.² Relying on a handful of technology-driven judicial decisions—most notably, *Authors Guild, Inc. v. Google Inc. (Google Books)*,³ *Authors Guild, Inc. v. HathiTrust (HathiTrust)*⁴ and *Sega Enterprises Ltd. v. Accolade, Inc. (Sega)*⁵—proponents of this view

*JD, Yale Law School. Former General Counsel and Associate Register of Copyrights, US Copyright Office. I am indebted to Professors Jane Ginsburg and Philippa Loengard of Columbia Law School, and Regan Smith, my former colleague at the Copyright Office, for insightful comments on earlier drafts of this article.

1. The use of the shorthand “AI” throughout this article refers to systems and processes of generative AI rather than artificial intelligence in general. For ease of reference, “AI companies” includes not just the companies themselves but entities engaged in AI activities on their behalf.

2. See, e.g., *Artificial Intelligence and Intellectual Property: Part I—Interoperability of AI and Copyright Law: Hearing Before the H. Comm. on the Judiciary, Subcomm. on Cts., Intell. Prop., and the Internet*, 118th Cong. 2 (2023) [hereinafter Damle] (statement of Sy Damle), <https://judiciary.house.gov/sites/evo-subsites/republicans-judiciary.house.gov/files/evo-media-document/damle-testimony.pdf> [<https://perma.cc/AGM7-YFFU>] (“Foundational copyright cases establish that the use of copyright-eligible content to create non-infringing works is protected fair use . . .”); Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 307–09 (2023) [hereinafter Sag, *Copyright Safety*] (arguing that reproduction of copyrighted works for purposes of generative AI constitutes a “nonexpressive” and therefore fair use of those works); Defs.’ Notice of Mot. to Dismiss, Mot. to Dismiss, and Mem. P. & A. in Supp. of Mot. to Dismiss at 2, *Tremblay v. OpenAI, Inc.*, No. 3:23-cv-03223 (N.D. Cal. Aug. 28, 2023) [hereinafter *Tremblay* MTD] (asserting that “[plaintiffs] claims . . . misconceive the scope of copyright, failing to take into account the limitations and exceptions (including fair use) that properly leave room for innovations like the large language models now at the forefront of artificial intelligence”); Answer of Def. Uncharted Labs, Inc. to Compl. at 8–9, *UMG Recordings, Inc. v. Uncharted Labs, Inc.*, No. 1:24-cv-04777 (S.D.N.Y. Aug. 1, 2024) [hereinafter *UMG* Answer] (“Under longstanding doctrine, what [defendant’s service] Udio has done—use existing sound recordings as data to mine and analyze for the purpose of identifying patterns in the sounds of various musical styles . . . is a quintessential ‘fair use’ under copyright law.”).

3. 804 F.3d 202 (2d Cir. 2015).

4. 755 F.3d 87 (2d Cir. 2014).

5. 977 F.2d 1510 (9th Cir. 1992).

assert that for-profit AI entities are entitled to copy books, music, photographs, visual art, news articles, and other protected works⁶ at will in order to train their AI models.⁷ They claim that once an AI model is trained, the works are “discarded” and do not exist as such in the model.⁸ According to the fair use proponents, AI systems merely derive information *about* the training works rather than making use of the works themselves.⁹ Therefore, the argument goes, mass reproduction of copyrighted works to create AI models should be treated as a transitory necessity of the AI development process rather than a copyright violation.

Looking to the four-factor test for fair use in Section 107 of the Copyright Act,¹⁰ advocates for unconstrained copying by AI companies claim that the reproduction of copyrighted works to train AI models is a “transformative” use of those works, thus justifying their

6. While AI copying includes all manner of copyrighted material, including social media posts and other user-generated content, the discussion herein is directed to typically marketed works such as the ones listed.

7. See, e.g., Damle, *supra* note 2, at 8 (“[T]raining a model that predominantly creates non-infringing outputs easily qualifies for fair use protection.”); Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 748 (2021) (“M[achine] L[earning] systems should generally be able to use databases for training, whether or not the contents of that database are copyrighted.”).

8. Damle, *supra* note 2, at 7 (“[S]tatistical data is incorporated into the [AI] algorithm, and the original content is discarded. . . . The model derives unprotectable information from the billions of works on which it is trained. . . .”); Notice of Mot. and Mot. of Defs. Stability AI Ltd. and Stability AI, Inc.’s Notice of Mot., Mot. to Dismiss, and Mem. of P. & A. in Supp. of Mot. to Dismiss, *Andersen v. Stability AI, Ltd.*, No. 3:23-cv-00201 (N.D. Cal. Apr. 18, 2023), at 1 [hereinafter *Andersen* MTD] (“[T]raining a model does not mean copying or memorizing images for later distribution. Indeed, *Stable Diffusion* does not “store” any images.”) (emphasis in original); see also Sag, *Copyright Safety*, *supra* note 2, at 307 (describing process of AI training as “deriving metadata through technical acts of copying and analyzing that data”); Lemley & Casey, *supra* note 7, at 776 (asserting that purpose of AI training data “is not to obtain or incorporate the copyrightable elements of a work but to access, learn, and use the unprotectable parts of the work”).

9. See Damle, *supra* note 2, at 2 (AI systems merely “learn unprotectable facts” about copyrighted works); Lemley & Casey, *supra* note 7, at 772 (“M[achine] L[earning] systems generally copy works, not to get access to their creative expression . . . but to get access to the uncopyrightable parts of the work—the ideas, facts, and linguistic structure of the works.”); Sag, *Copyright Safety*, *supra* note 2, at 343 (observing that “the legal and ethical imperative is to train models that learn abstract and uncopyrightable latent features of the training data” rather than memorizing the data); Anthropic PBC, Notification of Inquiry Regarding Artificial Intelligence and Copyright Public Comments of Anthropic PBC, REGULATIONS.GOV 1, 7 (Oct. 30, 2023), <https://www.regulations.gov/comment/COLC-2023-0006-9021> [<https://perma.cc/C46C-FKBD>] (describing AI copying as “merely an intermediate step, extracting unprotectable elements about the entire corpus of works”).

10. Section 107 lists four criteria to be considered by courts in assessing fair use: the purpose and character of the use; the nature of the work used; the amount and substantiality of the portion taken; and the effect on the potential market for or value of the work. 17 U.S.C. § 107.

appropriation.¹¹ They further contend that, because AI models are not designed to replicate training data in output, and do so only rarely,¹² there is no cognizable harm to copyright owners. At the same time, they do not deny that other aspects of the fair use calculus—that the use is commercial, involves highly creative works at the core of copyright protection, and entails copying of works in their entirety—point against fair use. But, they say, these concerns must yield to AI companies' overwhelming need to ingest massive amounts of copyrighted material without permission from or payment to rightsholders.¹³

The fair use case for generative AI rests in part on an inaccurate portrayal of the functioning of AI systems. Contrary to the suggestion that the works on which AI systems are trained are jettisoned after being processed by the AI system, in fact they have been algorithmically incorporated into the model, where they continue to be exploited. AI copying is thus fundamentally different from the copying at issue in the technology-driven fair use precedents relied upon by AI entities. Unlike in these earlier cases—where the copying served functional ends independent of the expressive content of the works—generative AI companies exploit the expressive content of the works they appropriate for its intrinsic value. This exploitation is not confined to the collection of training materials or the training process but is ongoing and the *sine qua non* of the resulting AI system.

The copying of the expressive content for the purpose of generating new content from that expression capitalizes on the intrinsic expressive purpose of the copied work. As copyright and technology scholar Benjamin Sobel presciently observed in 2017, copying for purposes of machine learning allows computers to derive valuable information from the way authors express ideas such that, instead of merely deriving facts from a work, the machine “glean[s] value from a

11. A significant consideration under the first fair use factor concerning the nature of the copying is whether the secondary use is “transformative.” *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 528 (2023) (hereinafter *Warhol*). In assessing transformativeness, courts evaluate whether the use has a substitutional effect or instead “adds something new, with a further purpose or different character.” *Id.* (quoting *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 579 (1994)).

12. This claim is dubious given the frequency with which close copies of training materials have been identified in generative AI output. *See infra* notes 60–62 and accompanying text.

13. *See, e.g.,* Lemley & Casey, *supra* note 7, at 770 (“[G]iven the large number of works an AI training data set needs to use . . . allowing a copyright claim is tantamount to saying, not that copyright owners will get paid, but that no one will get the benefit of this new use because it will be impractical to make that use at all.”); Damle, *supra* note 2, at 3 (“[A]ny royalty providing meaningful compensation to individual creators could impose an enormous financial burden on AI companies that would either bankrupt them or push all but the largest companies out of the market (or out of the country).”).

work's expressive aspects.”¹⁴ Sobel pointed out that this sort of copying is distinct from that deemed fair in technology-driven cases such as *Google Books*, which treated copying as a means “to assemble many individual works into non-expressive, factual ‘reference tools.’”¹⁵ A few years later—shortly before generative AI exploded into everyday life—Professor Mark Lemley and his colleague Bryan Casey echoed this sentiment, observing that machine learning for the purpose of copying expression—“for example, by training a . . . system to make a song in the style of Ariana Grande”—presented a much “tougher” question of fair use than copying for non-expressive purposes.¹⁶

This critical distinction between expressive and non-expressive exploitation sharply differentiates copying to train and develop generative AI models from uses determined to be a fair use in other technological contexts. Courts in earlier cases have been careful to distinguish between the copying of expressive works to facilitate a functional objective such as searching, indexing or interoperability, which may be deemed fair use, and the exploitation of protected expression for its own sake.¹⁷ Examined in this light, the fair use case for mass unauthorized copying by commercial AI entities is revealed as illusory. Appropriation of the world's literature, art, and music by for-profit companies to generate content from that material—including content that competes with the works so appropriated—is not excused by any precedent of fair use. It is without precedent.

Part II of this Article briefly considers the anthropomorphic terminology used in connection with generative AI, which misleadingly suggests that generative AI models learn and operate as humans do. In fact, AI models are algorithmically—not intellectually—driven. Part III reviews core AI activities that exploit copyrighted works, including the assembly of training materials, model training, generation of output,

14. Benjamin L. W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COLUM. J.L. & ARTS 45, 57 (2017) [hereinafter Sobel, *Fair Use*].

15. *Id.* at 48, 54–55 (observing that the Second Circuit in *Google Books* found Google's unauthorized reproduction of copyrighted works to be a transformative fair use “largely because Google Books provide[d] information ‘about’ books, not the books’ expression”).

16. Lemley & Casey, *supra* note 7, at 750; *see also* Peter Henderson, Xuechen Li, Dan Jurasfsky, Tatsunori Hashimoto, Mark A. Lemley & Percy Liang, *Foundation Models and Fair Use*, ARXIV 1, 2 (Mar. 28, 2023), <https://arxiv.org/pdf/2303.15715.pdf> [<https://perma.cc/Z278-5B3Y>] (acknowledging fair use concern). Interestingly, Lemley is currently defending AI company Stability AI in a class action suit by creators alleging mass infringement of copyrighted works. *See Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201 (N.D. Cal. filed Jan. 13, 2023).

17. *See, e.g.*, *Authors Guild, Inc. v. Google Inc.*, 804 F.3d 202, 222–25 (2d Cir. 2015) (holding Google's copying for the purpose of supplying information about books to be a fair use where such information was not a substitute for books' expressive content); *Fox News Network v. TVEyes, Inc.*, 883 F.3d 169, 180–81 (2d Cir. 2018) (holding TVEyes' copying of news content not a fair use because TVEyes allowed users to view substantial segments of searched-for content).

and retrieval-augmented generation (RAG). It next examines the nature of the copying involved in each of these scenarios, including systematic encoding of copyrighted content into the model for use in generative activities. Part IV responds to AI companies' claims that the unauthorized reproduction of copyrighted works to train and operate AI systems constitutes a fair use under copyright law. A review of the fair use precedents on which AI companies rely reveals that none sanctions the unlicensed use of copyrighted materials to exploit their intrinsic expressive value. Part IV demonstrates why the systematic copying of protected works to populate AI models with encoded representations of those works cannot be justified as a transformative use, a key consideration of fair use analysis. Part IV concludes that all four factors set forth in Section 107 of the Copyright Act point against fair use, especially when considered against the backdrop of a rapidly developing licensing market for AI training materials. The Article ends by suggesting that the broad exemption from copyright liability sought by AI companies under the guise of fair use represents a significant question of copyright policy that is best left for Congress to decide.

II. AI SYSTEMS ARE NOT HUMAN BUT ALGORITHMICALLY DRIVEN

AI companies rely on our intuitive understanding of human intellectual ability and anthropomorphic language to encourage the (mis)perception that AI machines learn and create like humans—that is, that they are capable of conceptual thinking and generalization from specific knowledge. But AI machines do not operate as humans do. The output of an AI model is fully dependent on, and limited by, the particular content on which it has been trained. Its output is a function of its input.

Much of the language used to describe activities associated with generative AI refers to human processes. AI machines are said to “train,” “learn,” “memorize,”¹⁸ and “hallucinate.”¹⁹ Anthropomorphic language tends to confuse and obscure discussions surrounding generative AI because it encourages people to ascribe human intellectual qualities to data-driven machines. But AI models do not “learn” or “know” anything in an intellectual sense. Rather, they store, access and process information according to prescribed formulas, or

18. “Memorization” refers to an AI model’s reproduction of training material as output. See *infra* notes 63–69 and accompanying text.

19. “Hallucination” refers to a model’s generation of information that is false, especially when presented in a way that it seems plausible. Ben Lutkevich, *AI Hallucination*, TECHTARGET, <https://www.techtartget.com/whatis/definition/AI-hallucination> [https://perma.cc/G45J-TAV4] (Oct. 2024). Hallucination occurs because, though a model trained to generate “grammatically and semantically correct” text, the model “ha[s] no understanding of the underlying reality.” *Id.*

algorithms.²⁰ Essentially, an AI model is a highly complex computer program; its output is a function of the works it has algorithmically stored and the programming instructions that govern its use of those stored representations.

By contrast, human cognition—including human imagination and creativity—is not limited to or governed by specific data or algorithms. As explained by researcher Melanie Mitchell, who studies the difference between AI “learning” and human intelligence, human understanding “is based on ‘concepts,’” that is, mental models of things revolving around “categories, situations and events” that are not limited to specific occurrences.²¹ Thus, humans are able to generalize and extrapolate from limited data—sometimes from just a single example²²—and reason by analogy.²³ The human brain allows people to infer cause and effect and predict the probable results of different actions “even in circumstances not previously encountered.”²⁴

Unlike humans, AI models “do not possess the ability to perform accurately in situations not encountered in their training.”²⁵ They “recite rather than imagine.”²⁶ A group of AI researchers has shown, for instance, that a large language model (LLM) trained on materials that say “A is B” does not reason from that knowledge, as a human would,

20. “The ‘learning’ involved is only a very loose analogy to human cognition—instead, [large language] models learn from the training data in the same way a simple regression model learns an approximation of the relationship between dependent and independent variables.” Sag, *Copyright Safety*, *supra* note 2, at 316.

21. Tom Siegfried, *Why Large Language Models Aren’t Headed Toward Humanlike Understanding*, SCIENCE NEWS (Feb. 28, 2024), <https://www.sciencenews.org/article/ai-large-language-model-understanding> [<https://perma.cc/TG3U-QBFA>] (quoting researcher Melanie Mitchell); *see also* Martha Lewis & Melanie Mitchell, *Using Counterfactual Tasks to Evaluate the Generality of Analogical Reasoning in Large Language Models*, ARXIV 1, 7 (Feb. 14, 2024), <https://arxiv.org/abs/2402.08955> [<https://perma.cc/3MNA-CHBK>] (concluding based on various experiments that AI models “lack[] the kind of abstract reasoning needed for human-like fluid intelligence”).

22. Michael Bennett, *Artificial Intelligence vs. Human Intelligence: Differences Explained*, TECHTARGET (Oct. 7, 2024), <https://www.techtartget.com/searchenterpriseai/tip/Artificial-intelligence-vs-human-intelligence-How-are-they-different> [<https://perma.cc/JT6B-YSFN>].

23. Lewis & Mitchell, *supra* note 21, at 1. Research shows that AI models lack the analogical reasoning ability exhibited by humans. *Id.*

24. Siegfried, *supra* note 21; *see also* Matthew Sag, *Fairness and Fair Use in Generative AI*, 92 FORDHAM L. REV. 1887, 1908 (2024) [hereinafter Sag, *Fairness*] (“To be clear, the model is not learning the same way a human might. The model does not understand grammar or society . . .”).

25. Siegfried, *supra* note 21 (“What’s really remarkable about people . . . is that we can abstract our concepts to new situations via analogy and metaphor.” (quoting Melanie Mitchell)).

26. Bennett, *supra* note 22.

to produce output that states the reverse, that B is A.²⁷ To cite one of the researchers' findings, a model trained on materials that said Valentina Tereshkova was the first woman to travel in space could respond to the query, "Who was Valentina Tereshkova?" with "The first woman to travel in space."²⁸ But when asked, "Who was the first woman to travel in space?," the model could not come up with the answer.²⁹ Based on their experiments in this area, the research team concluded that LLMs suffer from "a basic inability to generalize beyond the training data."³⁰

The use of anthropomorphic language to describe the development and functioning of AI models is distorting because it suggests that once trained, the model operates independently of the content of the works on which it has trained.³¹ Ideally this Article would steer clear of anthropomorphic terminology, but its ubiquitous deployment in the AI field makes it difficult to avoid as a practical matter. In considering the functionality of AI machines, then, it is critical to bear in mind that generative AI systems do not reason, or "learn" or "think," as humans do, but instead are designed to *mimic* human thought by applying computational processes to human-created materials.

III. WAYS IN WHICH GENERATIVE AI SYSTEMS ENGAGE WITH COPYRIGHTED WORKS

The question whether the unlicensed exploitation of copyrighted works by AI developers is a fair use of those works cannot be assessed without a basic understanding of how generative AI systems operate. As explained below, reproduction and exploitation of protected works occurs at all phases of AI model building and deployment. Computer scientists Katherine Lee and A. Feder Cooper, joined by law professor James Grimmelmann, correctly observe that "every stage in the

27. Lukas Berglund, Meg Tong, Max Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak & Owain Evans, *The Reversal Curse: LLMs Trained on "A is B" Fail to Learn "B is A"*, ARXIV 1, 1–2 (May 26, 2024), <https://arxiv.org/pdf/2309.12288> [<https://perma.cc/QD8J-VVS6>].

28. *Id.* at 2.

29. *Id.*

30. *Id.*

31. See Lucas Coughlin, *Compliance and Alignment: Ensuring Generative AI Stays Within the Bounds of Fair Use*, CHI-KENT J. INTELL. PROP. (May 12, 2023), <https://studentorgs.kentlaw.iit.edu/ckjip/compliance-and-alignment-ensuring-generative-ai-stays-within-the-bounds-of-fair-use/> [<https://perma.cc/9P8C-G535>] (characterizing arguments that AI models are "learning" or "researching" and thus entitled to fair use protections as "dubious . . . in their blatant anthropomorphism").

generative AI supply chain requires a potentially infringing reproduction and thus implicates copyright.”³²

A. Core Activities of Generative AI Systems

1. Assembly of Training Materials

AI companies readily acknowledge that they copy massive quantities of textual works, images, and music from websites and other sources without permission from the copyright owner in order to populate training sets used in the development of AI models.³³ These materials are usually scraped from online sources by bots (automated software agents) that crawl the internet.³⁴ In addition, AI companies may access and copy existing corpora of copyrighted works—including collections of pirated materials, such as the Books3 and LAION datasets—to train their models.³⁵ The copied materials are encoded into

32. Katherine Lee, A. Feder Cooper & James Grimmelmann, *Talkin’ Bout AI Generation: Copyright and the Generative-AI Supply Chain*, 71 J. COPYRIGHT SOC’Y U.S.A. (forthcoming 2024) (manuscript at 1, 67), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4523551 [<https://perma.cc/97FH-RVF8>].

33. See, e.g., *UMG Answer*, *supra* note 2, at 8 (“The many recordings that [the defendant’s] model was trained on presumably included recording whose rights are owned by the Plaintiffs in this case.”); *Tremblay MTD*, *supra* note 2, at 2–3 (acknowledging the “truly massive” quantity of textual content necessary to train an LLM and that plaintiffs’ copyright claims implicate “millions of . . . individual works contained in the training corpus”); *Andersen MTD*, *supra* note 8, at 1 (“Stable Diffusion was trained on *billions* of images that were publicly available on the Internet.”) (emphasis in original); Google LLC, *Artificial Intelligence and Copyright*, REGULATIONS.GOV 1, 9 (Oct. 30, 2023), <https://www.regulations.gov/comment/COLC-2023-0006-9003> [<https://perma.cc/BR84-3F7N>] (“If training could be accomplished without the creation of copies, there would be no copyright questions here.”); see also Damle, *supra* note 2, at 13 (explaining that AI models train on copyright-protected text and images “pulled” from the internet); Lemley & Casey, *supra* note 7, at 745 (“Creating a training set of millions of examples almost always requires . . . copying millions of images, videos, audio, or text-based works. Those works are almost all copyrighted.”).

34. See Damle, *supra* note 2, at 13 (“Often, AI developers train their models by pulling [textual content and images] from the internet.”); Lee et al., *supra* note 32, at 35 (noting that training data may be scraped from the web and “will often include copyrightable expression”); see also *What is Content Scraping? | Web Scraping*, CLOUDFLARE, <https://www.cloudflare.com/learning/bots/what-is-content-scraping/> [<https://perma.cc/62NJ-MAZ2>] (last visited June 29, 2024) (“Content scraping, or web scraping, refers to when a bot downloads much or all of the content on a website, regardless of the website owner’s wishes.”).

35. See Alex Reisner, *Revealed: The Authors Whose Pirated Books Are Powering Generative AI*, ATLANTIC (Sept. 25, 2023, 1:40 PM), <https://www.theatlantic.com/technology/archive/2023/08/books3-ai-meta-llama-pirated-books/675063/> [<https://perma.cc/A7FL-3ZAY>] (discussing Books3); Lee et al., *supra* note 33, at 39 (explaining that ChatGPT was allegedly trained on infringing “shadow libraries” of books). For an overview of images used to train Stable Diffusion (contained in datasets assembled by LAION, a nonprofit funded by Stable Diffusion’s owner, Stability AI) see Andy Baio, *Exploring 12 Million of the 2.3 Billion Images Used to Train Stable Diffusion’s Image Generator*, WAXY (Aug. 30, 2022), <https://waxy.org/2022/08/exploring-12->

standardized file formats and compiled into a structured dataset for use in the AI training process.³⁶

2. Training of AI Models

In the wake of *Google Books*, the unlicensed reproduction of millions of works to populate a technology platform may have a familiar ring. But the type of copying engaged by AI companies to train their models—which involves the encoding of expressive content so it can be used to generate other content—was not at issue in that case.

In *Google Books*, whole books were scanned and converted into digital files.³⁷ The goal of Google's copying was to create a searchable database so Google users could identify the locations of and frequency with which certain words appeared in the copied texts.³⁸ The output was therefore limited to “snippets” of works containing those terms, which were too abbreviated to serve as a substitute for anyone seeking to read or study the book.³⁹ In other words, the Google model was purposely designed to limit users to its search functionality and avoid exploitation of aesthetic content for its intrinsic value. In fact, the US Court of Appeals for the Second Circuit made clear in its opinion finding fair use that had Google not implemented the limitations it did, the plaintiffs' claim of infringement “would be strong.”⁴⁰ Even with the snippet

million-of-the-images-used-to-train-stable-diffusions-image-generator/ [https://perma.cc/52SU-XDSF].

36. Lee et al., *supra* note 32, at 35, 37; Keith Madden, *File Types And Artificial Intelligence – Data Formats For Ai Training*, HEXBROWSER.COM (Oct. 30, 2023), <https://www.hexbrowser.com/file-types-artificial-intelligence-data/> [https://perma.cc/YS5N-LFSS]. Notably, even if the works contained in the dataset were copied without permission, the compiler of the dataset may require a license from any third party that seeks to use it. See Lee et al., *supra* note 32, at 38 (“In practice . . . it appears that most uses of training datasets are licensed—either through a bilateral negotiation or by means of an open-source license . . .”).

37. Authors Guild, Inc. v. Google Inc., 804 F.3d 202, 208–09, 221 (2d Cir. 2015).

38. *Id.* at 208–09, 217.

39. *Id.* at 224–25. As the court explained:

“Google has constructed the snippet feature in a manner that substantially protects against its serving as an effectively competing substitute for Plaintiffs' books These include the small size of the snippets (normally one eighth of a page), the blacklisting of one snippet per page and of one page in every ten, the fact that no more than three snippets are shown—and no more than one per page—for each term searched, and the fact that the same snippets are shown for a searched term no matter how many times, or from how many different computers, the term is searched.”

Id. at 222.

40. *Id.* at 225.

limitation, the court cautioned that Google’s copying “test[ed] the boundaries of fair use.”⁴¹

In marked contrast to the facts of *Google Books*, an AI system exists to capture and use expressive content for its intrinsic qualities. To this end, the training of an AI model is not limited to deriving facts about works in the training set. Nor are the works “discarded” after the training process.⁴² Rather, the training works are algorithmically mapped and stored in the model, and then used by the model to generate output.⁴³

To train an LLM, textual works in the training set are broken down into small segments, or “tokens,” typically consisting of a word or part of a word.⁴⁴ The tokens are encoded into word vectors, long number sequences that capture where the tokens appear in relation to other tokens in the text, so the text is represented in numerical form.⁴⁵ The vectorized tokens can be decoded and translated into text again.⁴⁶

As explained by computer scientist Timothy B. Lee and co-author Sean Trott, “[w]ord vectors are a useful building block for

41. *Id.* at 206; *see also* Fox News Network v. TVEyes, Inc., 883 F.3d 169, 188 (2d Cir. 2018) (referencing *Google Books*’ cautionary language in rejecting TVEyes’ claim of fair use).

42. *See infra* notes 96–103 and accompanying text.

43. The model’s algorithms include “weights” and “biases,” parameters developed from the training data that govern the processing of new input and model output. *See What Are Weights and Biases in AI?*, EUR. INFO. TECHS. CERTIFICATION ACAD. (Aug. 15, 2023), <https://eitca.org/artificial-intelligence/eitc-ai-gcm1-google-cloud-machine-learning/introduction/what-is-machine-learning/explain-weights-and-biases/> [<https://perma.cc/36R5-CQES>]; Andrea D’Agostino, *Introduction to Neural Networks—Weights, Biases and Activation*, MEDIUM (Dec. 27, 2021), <https://medium.com/@theDrewDag/introduction-to-neural-networks-weights-biases-and-activation-270ebf2545aa> [<https://perma.cc/M8YP-7ZUP>]; *Weights and Bias in Neural Networks*, GEEKSFORGEEKS (Oct. 4, 2024) [hereinafter GEEKSFORGEEKS], <https://www.geeksforgeeks.org/the-role-of-weights-and-bias-in-neural-networks/> [<https://perma.cc/EC5T-YGMJ>].

44. *See* Lark Editorial Team, *Tokens in Foundational Models*, LARK (Dec. 25, 2023), https://www.larksuite.com/en_us/topics/ai-glossary/tokens-in-foundational-models [<https://perma.cc/8KE8-UM9S>]; Hakan Tekgul, *Tokenization: Unleashing the Power of Words*, ARIZE (Feb. 2, 2023), <https://arize.com/blog-course/tokenization/> [<https://perma.cc/62R7-YR3F>]; Amal Menzli, *Tokenization in NLP: Types, Challenges, Examples, Tools*, NEPTUNE.AI (Aug. 11, 2023), <https://neptune.ai/blog/tokenization-in-nlp> [<https://perma.cc/HF8A-5XXK>].

45. *See* Menzli, *supra* note 44 (“The token occurrences in a document can be used directly as a vector representing that document.”); Babis Marmanis, *Heart of the Matter: Demystifying Copyright in the Training of AIs*, DATAVERSITY (Feb. 2, 2024), <https://www.dataversity.net/heart-of-the-matter-demystifying-copying-in-the-training-of-llms/> [<https://perma.cc/WTW7-QHF5>]; *What are Large Language Models (LLM)?*, AWS [hereinafter *What Are Large Language Models?*], <https://aws.amazon.com/what-is/large-language-model/> [<https://perma.cc/N2E8-P6EJ>] (last visited June 29, 2024); Kevin Henner, *An Intuitive Introduction to Word Embeddings*, STACK OVERFLOW BLOG (Nov. 9, 2023), <https://stackoverflow.blog/2023/11/09/an-intuitive-introduction-to-text-embeddings/> [<https://perma.cc/5J4J-6XZY>].

46. Janakiram MSV, *The Building Blocks of LLMs: Vectors, Tokens and Embeddings*, NEW STACK (Feb. 8, 2024, 7:32 AM), <https://thenewstack.io/the-building-blocks-of-llms-vectors-tokens-and-embeddings/> [<https://perma.cc/RS85-M64B>].

language models because they encode subtle but important information about the relationship between words.”⁴⁷ They elaborate:

Human beings represent English words with a sequence of letters, like C-A-T for cat. Language models use a long list of numbers called a word vector [E]ach word vector represents a point in an imaginary “word space,” and words with more similar meanings are placed closer together. . . . A key advantage of representing words with vectors of *real numbers* (as opposed to a string of letters, like “C-A-T”) is that numbers enable operations that letters don’t.⁴⁸

Within the model, these vectorized representations of the work’s content, also known as “embeddings,” are used for the model’s generative activities.⁴⁹ Software engineer Babis Marmanis explains:

The[] vectors are representations of tokens that preserve their original natural language representation that was given as text. It is important to understand the role of word embeddings when it comes to copyright because the embeddings form representations (or encodings) of entire sentences, or even paragraphs, and therefore, in vector combinations, even entire documents in a high-dimensional vector space. It is through these embeddings that the AI system captures and stores the meaning and the relationships of words from the natural language.⁵⁰

In this way, through the embeddings, “LLMs *retain the expressions of the original works* on which they have been trained.”⁵¹ They form “internal representations” of those works and, “given the appropriate input as a trigger,” can “reproduce the original works that were used in their training.”⁵² As Professor Lee and her colleagues elaborate, an AI model can thus be understood as a compilation of its training data—the model is simply a different and complicated arrangement of training examples. Another view is that a model is a derivative work of its training data—“a work based upon one or more preexisting works . . . in which [those works are] recast, transformed, or adapted.”⁵³

47. Timothy B. Lee & Sean Trott, *Large Language Models, Explained With a Minimum of Math and Jargon*, UNDERSTANDING AI (July 27, 2023), <https://www.understandingai.org/p/large-language-models-explained-with> [<https://perma.cc/4NU2-HSQT>].

48. *Id.*; see also Henner, *supra* note 45 (“The ability to take a chunk of text and turn it into a vector, subject to the laws of mathematics, is fundamental to natural language processing.”).

49. *What Are Large Language Models?*, *supra* note 45.

50. Marmanis, *supra* note 45; see also Lark Editorial Team, *supra* note 44 (“Tokens in foundational models operate by segmenting and representing the inputs in a manner that allows AI systems to process and interpret the information effectively.”).

51. Marmanis, *supra* note 45 (emphasis added).

52. *Id.*

53. Lee et al., *supra* note 32, at 60 (quoting 17 U.S.C. § 101 (Copyright Act definition of derivative work)) (alterations in original); see also Daniel Gervais, Haralambos Marmanis, Noam Shemtov & Catherine Zaller Rowland, *The Heart of the Matter: Copyright, AI Training, and LLMs*, SSRN 1, 12 (Sept. 21, 2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4963711 [<https://perma.cc/MUG8-97K4>] (“[T]he numerical representations of the training data that are permanently embedded in LLMs may be considered as copies or adaptations of the original

Text-to-image AI systems employ a similar methodology to encode textual information (such as captions) associated with specific images, but the images are processed using “diffusion” technology.⁵⁴ Under the diffusion approach, the AI system slowly adds “noise” (akin to snow on a television set) to the original image until the original is no longer perceptible.⁵⁵ The noise-adding process is then reversed by gradually subtracting the noise from the image so the model learns how to “rebuild the original.”⁵⁶ Through this training, the model is “now equipped to ... regenerate the original image from the corresponding text prompt.”⁵⁷ In other words, the model has encoded a representation of the original.

Once the AI model has encoded and stored the training materials, the resulting “base” model is “fine-tuned” to better achieve the developers’ specific goals.⁵⁸ This generally means training the model on additional materials, often from a particular domain of interest.⁵⁹ The trained and fine-tuned model can then be “aligned” to better meet the developers’ objectives by further adjusting its behavior based on human evaluation of its output.⁶⁰

training material.”). Under Section 106 of the Copyright Act, the right to prepare derivative works is an exclusive right of the copyright owner. 17 U.S.C. § 106(2).

54. See Sunil Ramlochan, *What is Stable Diffusion and How Does it Work?*, PROMPT ENG’G INST. (July 17, 2023), <https://promptengineering.org/the-possibilities-of-ai-art-examining-stable-diffusion/> [<https://perma.cc/L9JN-PNTT>]; Andrew, *How Does Stable Diffusion Work?*, STABLE DIFFUSION ART (June 9, 2024), <https://stable-diffusion-art.com/how-stable-diffusion-work/> [<https://perma.cc/4363-M45V>]; see also Yang Zhang, Tio Tze Tzun, Lim Wei Hern, Haonan Want & Kenji Kawaguchi, *On Copyright Risks of Text-to-Image Diffusion Models*, ARXIV 1, 2 (Feb. 19, 2024), <https://arxiv.org/pdf/2311.12803> [<https://perma.cc/NPB2-DDGY>].

55. See Andrew, *supra* note 54.

56. Ramlochan, *supra* note 54; Andrew, *supra* note 54; see also Zhang et al., *supra* note 54 (“The objective of diffusion models is to learn the reverse process of diffusion, which tries to reconstruct the target given noisy input.”).

57. Ramlochan, *supra* note 54.

58. See Lee et al., *supra* note 32, at 5, 42–43; Valentin Hartmann, Anshuman Suri, Vincent Bindschaelder, David Evans, Shruti Tople & Robert West, *SoK: Memorization in General-Purpose Large Language Models*, ARXIV 1, 3 (Oct. 24, 2023), <https://vbinds.ch/sites/default/files/PDFs/arXiv23-Hartmann-Memorization.pdf> [<https://perma.cc/YTL6-TBWL>]; Pradeep Menon, *A Deep-Dive into Fine-Tuning of Large Language Models*, MEDIUM (Aug. 13, 2023), <https://rpradeepmenon.medium.com/a-deep-dive-into-fine-tuning-of-large-language-models-96f7029ac0e1> [<https://perma.cc/YC8A-VCWT>].

59. See Google LLC, *supra* note 33, at 5 (at the fine-tuning stage, the model “learns from additional example data to help hone its capabilities” with respect to particular tasks); Lee et al., *supra* note 32, at 43.

60. See Kim Martineau, *What is AI Alignment?*, IBM (Nov. 8, 2023), <https://research.ibm.com/blog/what-is-alignment-ai> [<https://perma.cc/489U-JV2H>]; Lee et al., *supra* note 32, at 6, 54–55.

3. Generation of Copied Works and Derivatives

Apart from the nature of the training process itself, it is well established that models are able to regenerate—or in AI parlance, “regurgitate”—training materials in response to particular user prompts.⁶¹ Indeed, such replication is not uncommon.⁶² This phenomenon is further confirmation that AI models retain stored representations of the materials on which they train. As Professor Grimmelman observes, AI models “often produce near-exact copies” of the works they ingest.⁶³ Common sense tells us that this could only occur if the model encodes the actual content of those works,⁶⁴ which is to say the training materials do not disappear, but are incorporated into the model.

61. See Milad Nasr, Nicholas Carlini, Jonathan Hayase, Matthew Jagielski, A. Feder Cooper, Daphne Ippolito, Christopher A. Choquette-Choo, Eric Wallace, Florian Tramèr & Katherine Lee, *Scalable Extraction of Training Data from (Production) Language Models*, ARXIV 1, 14 (Nov. 28, 2023), <https://arxiv.org/pdf/2311.17035> [<https://perma.cc/ZHC2-MWK9>] (“[O]ur paper suggests that training data can easily be extracted from the best language models of the past few years through simple techniques.”); Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito & Eric Wallace, *Extracting Training Data from Diffusion Models*, ARXIV 1 (Jan. 30, 2023), <https://arxiv.org/pdf/2301.13188.pdf> [<https://perma.cc/NJ65-4EVC>] (“[D]iffusion models do memorize and regenerate individual training examples.”); Zhang et al., *supra* note 54, at 1 (noting that diffusion models “often replicate elements from their training data”); Marmanis, *supra* note 45 (noting that training works will be replicated by a model “given the appropriate input”).

62. See, e.g., Chloe Xiang, *AI Spits Out Exact Copies of Training Images, Real People, Logos, Researchers Find*, VICE (Feb. 1, 2023, 3:47 PM), <https://www.vice.com/en/article/m7gznn/ai-spits-out-exact-copies-of-training-images-real-people-logos-researchers-find> [<https://perma.cc/N3JX-EJME>] (explaining researchers are able to extract numerous copies of training works from AI image generators); Alex Reisner, *The Flaw That Could Ruin Generative AI*, ATLANTIC (Jan. 11, 2024), <https://www.theatlantic.com/technology/archive/2024/01/chatgpt-memorization-lawsuit/677099/> [<https://perma.cc/UZJ3-HKYQ>] (citing examples of memorized training materials).

63. Will Oremus & Elahe Izadi, *AI’s Future Could Hinge on One Thorny Legal Question*, WASH. POST (Jan. 4, 2024), <https://www.washingtonpost.com/technology/2024/01/04/nyt-ai-copyright-lawsuit-fair-use/> [<https://perma.cc/N4C7-TUGD>] (quoting James Grimmelman).

64. See Gary Marcus & Reid Southen, *Generative AI Has a Visual Plagiarism Problem*, IEEE SPECTRUM (Jan. 6, 2024), <https://spectrum.ieee.org/midjourney-copyright> [<https://perma.cc/2HV7-X5XH>] (“The very existence of potentially infringing outputs is evidence of another problem: the nonconsensual use of copyrighted human work to train machines.”).

For example, as set forth in the complaint in *New York Times v. Microsoft Corporation*, Open AI's GPT-4 model could be prompted to generate *New York Times* articles essentially verbatim:⁶⁵

Output from GPT-4:

many of America's other global companies — aren't nearly as avid in creating American jobs as other famous companies were in their heydays.

Apple employs 43,000 people in the United States and 20,000 overseas, a small fraction of the over 400,000 American workers at General Motors in the 1950s, or the hundreds of thousands at General Electric in the 1980s. Many more people work for Apple's contractors: an additional 700,000 people engineer, build and assemble iPads, iPhones and Apple's other products. But almost none of them work in the United States. Instead, they work for foreign companies in Asia, Europe and elsewhere, at factories that almost all electronics designers rely upon to build their wares.

"Apple's an example of why it's so hard to create middle-class jobs in the U.S. now," said Jared Bernstein, who until last year was an economic adviser to the White House. "If it's the pinnacle of capitalism, we should be worried."

Apple executives say that going overseas, at this point, is their only option. One former executive described how the company relied upon a Chinese factory to revamp iPhone manufacturing just weeks before the device was due on shelves. Apple had redesigned the iPhone's screen at the last minute, forcing an assembly line overhaul. New screens began arriving at the plant near midnight.

A foreman immediately roused 8,000 workers inside

Actual text from NYTimes:

many of its high-technology peers — are not nearly as avid in creating American jobs as other famous companies were in their heydays.

Apple employs 43,000 people in the United States and 20,000 overseas, a small fraction of the over 400,000 American workers at General Motors in the 1950s, or the hundreds of thousands at General Electric in the 1980s. Many more people work for Apple's contractors: an additional 700,000 people engineer, build and assemble iPads, iPhones and Apple's other products. But almost none of them work in the United States. Instead, they work for foreign companies in Asia, Europe and elsewhere, at factories that almost all electronics designers rely upon to build their wares.

"Apple's an example of why it's so hard to create middle-class jobs in the U.S. now," said Jared Bernstein, who until last year was an economic adviser to the White House.

"If it's the pinnacle of capitalism, we should be worried."

Apple executives say that going overseas, at this point, is their only option. One former executive described how the company relied upon a Chinese factory to revamp iPhone manufacturing just weeks before the device was due on shelves. Apple had redesigned the iPhone's screen at the last minute, forcing an assembly line overhaul. New screens began arriving at the plant near midnight.

65. Complaint at ¶ 100, *N.Y. Times v. Microsoft Corp.*, No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023) [hereinafter NYT Compl.].

Similarly, in *Concord Music Group, Inc. v. Anthropic PBC*, music publishers presented numerous examples of copied lyrics that were reproduced by Anthropic's Claude LLM:⁶⁶

<u>Claude output:</u>	<u>Genuine ABKCO lyrics:</u>
I saw her today at the reception A glass of wine in her hand I knew she was gonna meet her connection At her feet was her footloose man	I saw her today at the reception A glass of wine in her hand I knew she would meet her connection At her feet was her footloose man
No, you can't always get what you want You can't always get what you want You can't always get what you want But if you try sometimes you find You get what you need	No, you can't always get what you want You can't always get what you want You can't always get what you want But if you try sometime you'll find You get what you need
I saw her today at the reception A glass of wine in her hand I knew she was gonna meet her connection At her feet was her footloose man	I saw her today at the reception A glass of wine in her hand I knew she was gonna meet her connection At her feet was her footloose man
You can't always get what you want You can't always get what you want You can't always get what you want But if you try sometimes you might find You get what you need	You can't always get what you want You can't always get what you want You can't always get what you want But if you try sometimes, well, you might find You get what you need
Oh yeah, hey hey hey, oh...	Ah, yeah Oh


66. Complaint at ¶ 69, *Concord Music Grp., Inc. v. Anthropic PBC*, No. 3:23-cv-01092 (M.D. Tenn. Oct. 18, 2023).

It has also been shown that text-to-image AI models such as Stable Diffusion and Midjourney are able to generate near-perfect copies of training images:

Stable Diffusion⁶⁷



Midjourney⁶⁸



67. Carlini, et al., *supra* note 61 (series of images extracted from Stable Diffusion).

68. Marcus & Southen, *supra* note 64 (reporting that it was “easy to generate many plagiaristic outputs” from Midjourney using “brief prompts related to commercial films”).

AI models can also be prompted to produce depictions of copyrighted characters:⁶⁹



The ability of an AI model to generate copies of training materials is often referred to as “memorization.” This term suggests that the model just happens to recall certain texts or images while not recalling others, as a human might. But this is misleading. In a revealing passage discussing “memorization” versus “extractability,” a Microsoft AI researcher explains:

A naïve definition of memorization could thus be “any information that is stored in the model . . . is memorized”. However, evaluating this definition would be infeasible and basically amount to fully determining everything the model has learned. Researchers thus resort to studying proxies for this memorization via extractability or discoverability, which only captures memorized information that can be accessed through known methods. This inherently underestimates memorization, since it assumes there are no better ways to extract information from the model.⁷⁰

Contrary to its ordinary meaning, then, in the context of generative AI, “memorization” is narrowly and circularly defined as the circumstance in which certain training material has been shown to be reproducible from the model because it is retrievable from the model. But this nonintuitive definition of memorization should not be understood to mean that other, supposedly “non-memorized” content

69. *Id.* (illustrating Midjourney-generated portrayals of Simpson characters).

70. Hartmann et al., *supra* note 58, at 5.

has not been encoded in the model, cannot be retrieved, or is not being used by the model to generate output.⁷¹

Seeking to “align” their systems with copyright law, AI developers may attempt to mitigate the memorization problem by applying filtering mechanisms to detect and suppress output that “mirrors the training data”—in other words, output that obviously copies.⁷² But such filters are imperfect, to say the least.⁷³ Apart from the challenges of identifying close, but inexact, copies of a training work—let alone output that is not identical but constitutes an infringing derivative—the filters are vulnerable to user evasion through strategic manipulation of prompts.⁷⁴

As noted earlier, memorization of training materials is not unusual.⁷⁵ It therefore seems disingenuous to characterize “the ability of AI models to duplicate training data” as “a bug, not a feature,” as AI companies would have it.⁷⁶ But even if filtering mechanisms could be

71. See A. Feder Cooper & James Grimmelmann, *The Files Are in the Computer: On Copyright, Memorization, and Generative AI*, 99 CHI-KENT L. REV. (forthcoming 2025) (manuscript at 33) (on file with SSRN), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4803118 [<https://perma.cc/Bj9J-6YE7>] (noting that “it is important to distinguish our *knowledge* of whether a model has memorized training data from the underlying question of memorization itself,” as “[i]t is possible that data could be reconstructed from a model through techniques that are currently unknown but discovered in the future.”).

72. Henderson et al., *supra* note 16, at 22; see also Coughlin, *supra* note 31 (explaining AI platforms could be configured so “they are biased away from directly reproducing instances of training data”).

73. See Henderson et al., *supra* note 16, at 22; Emily Conover, *AI Chatbots Can Be Tricked into Misbehaving. Can Scientists Stop It?*, SCI. NEWS (Feb. 1, 2024, 8:00 AM), <https://www.sciencenews.org/article/generative-ai-chatbots-chatgpt-safety-concerns> [<https://perma.cc/5QJL-5QW2>]. The problem is not limited to copyright infringement. For example, filters did not prevent researchers from prompting ChatGPT to write a social media post to promote drunk driving or a step-by-step plan to destroy humanity. Conover, *supra* note 73; see also Daniel Tencer, *\$125m-Backed Suno is Being Used to Make Racist and Antisemitic Music*, MUSIC BUS. WORLDWIDE (June 20, 2024), <https://www.musicbusinessworldwide.com/125m-backed-suno-is-being-used-to-make-racist-and-antisemitic-music/> [<https://perma.cc/3VFQ-NW64>] (noting the discovery of a “vast library of disturbing songs” generated by users of Suno, a model reportedly trained on unlicensed sound recordings, “including songs that glorify Hitler and ‘white power’”).

74. Henderson et al., *supra* note 16, at 22. See *id.* at 7–8 for a description of content extraction strategies tested by the researchers.

75. See *supra* notes 59–61 and accompanying text.

76. Damle, *supra* note 2, at 8; see also, e.g., Wayne Brough, *R Street Signs Coalition Letter Expressing Concerns With Future AI Legislation*, R STREET (Sept. 11, 2023), <https://www.rstreet.org/outreach/r-street-signs-coalition-letter-expressing-concerns-with-future-ai-copyright-legislation/> [<https://perma.cc/UFX8-BXJR>] (claiming that AI systems are “not designed to reproduce protected material from the data on which they are trained,” but “on . . . rare occasions . . . they do.”); OpenAI and Journalism, OPENAI (Jan. 8, 2024), <https://openai.com/index/openai-and-journalism/> [<https://perma.cc/GT7Q-UJ7D>] (“‘Regurgitation’ is a rare bug that we are working to drive to zero.”); Sag, *Copyright Safety*, *supra* note 2, at 313, 336 (asserting that outputs generated by AI models are “(mostly) not copies of their training data” and minimizing ability to extract memorized training data as “premised on somewhat contrived situations or targeted a works especially likely to be duplicated”).

successfully deployed to mitigate regurgitation of training materials, such an approach would amount to a copyright band-aid rather than a cure. The filtering mechanisms would address only infringing output and not the infringement that occurs during the content harvesting, training, or other stages of AI development. Moreover, the suppression of output that would otherwise replicate training works does not alter the fact that the works were nonetheless copied to create the model and are being exploited to generate other outputs.⁷⁷

Writers, musicians, and visual artists have raised serious concerns that models trained on their works are able to generate content “in the style” of those works, thus competing with the creator’s originals.⁷⁸ While an artistic “style” as a general concept or category of art is not copyrightable,⁷⁹ “in-the-style of” output that emulates distinctive elements of an artist’s oeuvre may be sufficiently similar to the artist’s originals to constitute infringing copies or derivatives.⁸⁰ Moreover, even if the output itself does not rise to the level of infringement, to generate recognizable riffs on the artist’s works, the model was presumably trained on and contains encoded representations of those works. The generation of content in a recognizable style thus points to earlier acts of infringement at the training stage.

77. See Coughlin, *supra* note 31 (observing that guardrails to protect against infringing output “may decrease the risk of an offending end product but do[] not address the underlying copying”).

78. See Stephen Wolfson, *The Complex World of Style, Copyright, and Generative AI*, CREATIVE COMMONS (Mar. 23, 2023), <https://creativecommons.org/2023/03/23/the-complex-world-of-style-copyright-and-generative-ai/> [<https://perma.cc/H6FS-BCUH>] (flagging concern that AI-generated works that mimic an artist’s particular style may displace the market for the artist’s work).

79. See *Whitehead v. CBS/Viacom, Inc.*, 315 F. Supp. 2d 1, 11 (D.D.C. 2004) (“[S]tyle alone cannot support a copyright claim.”) Certain forms of imitation may present non-copyright concerns such as infringement of an artist’s rights of publicity. Tennessee, for example, recently enacted the Ensuring Likeness, Voice, and Image Security (“ELVIS”) Act of 2024, which bans unauthorized artificial intelligence reproductions of individuals’ likenesses and voices. See Amelia E. Bruckner, *Tennessee Expands Right-of-Publicity Statute to Cover AI-Generated Deepfakes*, NAT’L L. REV. (Apr. 18, 2024), <https://natlawreview.com/article/tennessee-expands-right-publicity-statute-cover-ai-generated-deepfakes> [<https://perma.cc/6XJF-2VCR>].

80. See, e.g., *Steinberg v. Columbia Pictures Indus., Inc.*, 663 F. Supp. 706, 712 (S.D.N.Y. 1987) (finding defendant’s poster infringed plaintiff’s poster in part due to the “striking stylistic relationship” between the two, which contributed to the overall substantial similarity in expression). Responding to the view that “style” is categorically uncopyrightable, Sobel contends that this oft-cited maxim should be reexamined in the age of generative AI: “An honest application of copyright law requires us to acknowledge that some of what we call style is copyrightable some of the time” Benjamin L.W. Sobel, *Elements of Style: Copyright, Similarity, and Generative AI*, 38 Harv. J.L. & Tech. 49, 55 (2024) [hereinafter Sobel, *Elements*].

4. Retrieval-Augmented Generation

Even after training is complete, AI systems may engage in additional copying activities to enhance the model's performance. RAG is a process by which functioning models incorporate additional content from external sources to allow for ongoing knowledge updates and integration of specialized information.⁸¹ When a model such as ChatGPT, for example, receives a query regarding a recent event but lacks "knowledge" of that event because it was trained before the event occurred, "RAG addresses this gap by retrieving up-to-date document excerpts from external knowledge bases. In [such an] instance, it procures a selection of news articles pertinent to the inquiry. These articles, alongside the initial question, are then amalgamated into an enriched prompt that enables ChatGPT to synthesize an informed response."⁸²

In other words, rather than relying solely on the basic model, the system searches for and assimilates relevant material from online or other sources to augment the model's response to a particular user prompt.⁸³ The prompt and retrieved information are then fed into the model, which formulates a response based on its "inbuilt knowledge plus the additional information from the RAG search."⁸⁴ The copied content allows the AI model to craft a better response to the user's query.⁸⁵ RAG technology "bypass[es] the need for costly, time-intensive retraining and updating" of the AI model by ingesting fresh material from the internet or other sources on an as-needed basis.⁸⁶

Although an important feature of certain AI systems,⁸⁷ RAG functionality has not so far been a primary focus of AI litigation

81. Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Quanyu Guo, Meng Wang & Haofen Wang, *Retrieval-Augmented Generation for Large Language Models: A Survey*, ARXIV 1, 1 (Jan. 5, 2024), <https://arxiv.org/abs/2312.10997> [<https://perma.cc/TK6L-8BF8>]; *What is Retrieval-Augmented Generation (RAG)*, GOOGLE CLOUD, <https://cloud.google.com/use-cases/retrieval-augmented-generation?hl=en> [<https://perma.cc/N52V-KPZ6>] (last visited June 23, 2024); Eleanor Berger, *Grounding LLMs*, ELEANOR ON EVERYTHING (June 9, 2023), <https://everything.intellectronica.net/p/grounding-llms> [<https://perma.cc/6RZ5-M8V9>].

82. *What Is RAG?*, AWS [hereinafter AWS, *What Is RAG?*], <https://aws.amazon.com/what-is/retrieval-augmented-generation/> [<https://perma.cc/2DBF-GHCL>] (last visited June 29, 2024).

83. *See id.*; Rahul Singhal, *The Power of RAG: How Retrieval-Augmented Generation Enhances Generative AI*, FORBES (Nov. 30, 2023, 8:30 AM), <https://www.forbes.com/councils/forbestechcouncil/2023/11/30/the-power-of-rag-how-retrieval-augmented-generation-enhances-generative-ai/> [<https://perma.cc/89HH-F3WT>].

84. Singhal, *supra* note 83.

85. *Id.*; *see also* AWS, *What Is RAG?*, *supra* note 82; GOOGLE CLOUD, *supra* note 81.

86. Singhal, *supra* note 83.

87. *See* Gao et al., *supra* note 81, at 16–17 (noting "progress in RAG technology" and its "significant practical implications for AI deployment").

efforts.⁸⁸ To the extent RAG-enabled AI systems are searching out and reproducing unlicensed content to update and enrich their output however, there would seem to be no question that they are capitalizing on the expressive value of content belonging to others.

B. Each Core Activity Involves Copying

Each of the core activities reviewed above—assembly of training materials, training of the AI model, generation of output and RAG functionality—involves unauthorized exploitation of copyrighted works. To evaluate AI companies' claim that these activities should be considered fair use, one must take a closer look at the nature of the copying involved, in particular with respect to the related claim that copies of the training works do not persist in the model after training.

1. Reproduction of Works to Create Training Sets

AI companies do not dispute that unlicensed copying occurs when materials are reproduced from online sources or elsewhere for purposes of creating an AI training set.⁸⁹ That such appropriation constitutes direct copying and presumptive infringement of copyright owners' right of reproduction seems obvious.⁹⁰ Notably, despite advocating for a fair use exception for machine learning,⁹¹ Lemley and Casey acknowledge that there is no precedent that treats the copying of protected works for this purpose as noninfringing.⁹² Or, more to the point, as Lemley and co-authors concede in another piece addressing AI training: "the risk of infringement is real."⁹³

As noted above, source materials are converted into standardized formats in order to carry out the training process. The conversion process does not negate a finding of infringement, as it is well established that encoding a copyrighted work in a more convenient or usable format is an act of copying that does not itself qualify as

88. See, e.g., NYT Compl., *supra* note 66, ¶¶ 108–23, 179 (discussing "synthetic" AI search results incorporating plaintiffs' news stories that were generated using RAG technology, among other alleged infringements of newspaper content); see also Complaint at ¶¶ 78, 114–116, 212, Daily News, LP v. Microsoft Corp., No. 24-3285 (S.D.N.Y. Apr. 30, 2024) (similar discussion).

89. See *supra* notes 33–39 and accompanying text (showing AI companies acknowledge vast amounts of unauthorized copying).

90. See Sobel, *Fair Use*, *supra* note 14, at 61 (referencing "a prima facie infringement" of copyright owners' exclusive right of reproduction); 17 U.S.C. § 106(1).

91. Lemley & Casey, *supra* note 7, at 776–79.

92. *Id.* at 746.

93. Henderson et al., *supra* note 16, at 2.

transformative under the criteria for fair use.⁹⁴ In an influential case involving a music service, for example, *UMG Recordings, Inc. v. MP3.com, Inc.*, a New York district court rejected the claim that a service's conversion of user-purchased music CDs into digital files for streaming back to their owners was a fair use of the copyrighted works.⁹⁵ The Ninth Circuit came to a similar conclusion in *Disney Enterprises, Inc. v. VidAngel, Inc.*, which held that a service's encoding of films from purchased DVDs into a specialized format to permit subscribers to skip objectionable scenes while streaming the films did not qualify as a fair use.⁹⁶ As the court opined in the latter case, the "law of fair use, as it stands today, does not sanction broad-based space-shifting or format-shifting."⁹⁷

2. Encoding of Works in AI Models

While conceding (as they must) that copyrighted works are copied to create training materials, AI companies seek to convince us that encoding those materials into their models does not entail further reproduction of protected material. Testifying before Congress, attorney Sy Damle (who represents AI companies in some of the many lawsuits generated by their training activities)⁹⁸ characterized the encoding process as follows: "[T]he models derive abstract patterns and relationships—not copyrightable expression—from billions of pieces of training data.... That statistical data is incorporated into the algorithm, and the original content is discarded."⁹⁹ Lemley and Casey similarly assert that machine learning systems "copy works, not to get access to

94. See, e.g., *Disney Enters., Inc. v. VidAngel, Inc.*, 869 F.3d 848, 861 (9th Cir. 2017) [hereinafter *VidAngel*] (rejecting argument that encoding of motion pictures to operate a streaming service was a transformative fair use); *Hachette Book Grp. v. Internet Archive*, 115 F.4th 163, 184 (2d Cir. 2024) (digitizing books is not transformative for purposes of fair use); *UMG Recordings, Inc. v. MP3.com, Inc.*, 92 F. Supp. 2d 349, 351 (S.D.N.Y. 2000) (digitizing music is not a transformative fair use); Exemption to Prohibition on Circumvention of Copyright Protection Systems for Access Control Technologies, 80 FED. REG. 65944, 65960 (Oct. 28, 2015) [hereinafter US Copyright Office, Exemption Rulemaking] (rejecting the notion that format-shifting or space-shifting constitutes a fair use).

95. *MP3.com*, 92 F. Supp. 2d at 350, 352.

96. *VidAngel*, 869 F.3d at 853–54, 862.

97. *Id.* at 862 (quoting US Copyright Office, Exemption Rulemaking, *supra* note 95, at 65960).

98. See *Master List of Lawsuits v. AI, ChatGPT, OpenAI, Microsoft, Meta, Midjourney & Other AI Cos.*, CHAT GPT IS EATING THE WORLD, (Aug. 27, 2024), <https://chatgptiseatingtheworld.com/2023/12/27/master-list-of-lawsuits-v-ai-chatgpt-openai-microsoft-meta-midjourney-other-ai-cos/> [https://perma.cc/LV42-3FRG] (listing pending lawsuits and attorneys).

99. Damle, *supra* note 2, at 7.

their creative expression (the part of the work the law protects) but to get access to the uncopyrightable parts of the work.”¹⁰⁰

To depict the AI training process in this manner amounts to a sleight of hand that distracts from the reality that protected expressive content is being stored in the model. As explained above, training materials do not simply disappear as the model is built; rather, each work is algorithmically ingested, piece by piece—or token by token—into the model.¹⁰¹ Nor is there any practice of separating copyrightable from uncopyrightable elements during the training process. The tokens themselves are encoded—not just the “statistical data” or “information” about them.¹⁰² Of course, this is only logical; the objective of the training exercise is to capture and map the expressive content of each work for use in the generative process.

In this regard, it is worth pointing out that, ultimately, information about a work *is* the work. For example, if I were to provide you with a table listing all the words in a book, which identified the first word as “It,” the second as “is,” the third as “a,” the fourth as “truth,” the fifth as “universally,” the sixth as “acknowledged,” and so on, all the way to the end of the book—and instructed you to transcribe the words accordingly to the order specified in the table—after a (very long) time you would have transcribed *Pride and Prejudice*.¹⁰³ Or I could provide you with the color and location of every pixel in a family photo and ask you to recreate the photo from that data.¹⁰⁴ In either case, you would be reproducing the original work using “information about the work.”

In ordinary usage, “information about a work” means data derived from a work that exist *separately* from the work itself, such as the location of a particular word within a text, the size of an image file,

100. Lemley & Casey, *supra* note 7, at 772; see also Brough, *supra* note 76 (asserting that AI systems are designed to “learn facts about the world, ideas and visual concepts” rather than to exploit expressive content). But see Lemley & Casey, *supra* note 7, at 777 (contradictorily observing that “[s]ome ML systems will be interested in the expressive components of the work as an integral part of their training. That is, the goal will be to teach the system using the creative aspects of the work that copyright values, not just using the facts or the semantic connections the law is not supposed to protect.”).

101. See *supra* notes 37–49 and accompanying text (describing the encoding process).

102. See *supra* notes 44–46 and accompanying text (discussing the role of tokens).

103. JANE AUSTEN, *PRIDE AND PREJUDICE* 1 (Bantam Classic ed., Random House, Inc. 2003) (1813) (“It is a truth universally acknowledged that a single man in possession of a good fortune must be in want of a wife.”).

104. Sobel cites a similar example to demonstrate that conventional image files consist of mathematical representations: “An image file might instruct a computer, ‘display a 50x50 pixel grid of alternating rows of white and red pixels.’ That’s a mathematical representation of a white-and-red-striped square.” Sobel, *Elements*, *supra* note 80, at 24.

the length of a song, and so forth.¹⁰⁵ It does not mean an encoded version of the work that can be tapped to produce a representation of the work, or to generate additional works. This distinction between expressive versus non-expressive use of protected works was critical in the *Google Books* case (among others), where the Second Circuit’s determination of fair use turned on its conclusion that displaying brief “snippets” of text indicating where particular words appeared in a book was not a substitute for the book itself.¹⁰⁶ The *Google Books* example stands in marked contrast to the exploitation of copyrighted works by AI systems and users of those systems, who are utilizing stored expressive content to produce output.

The encoding of text, images and music into machine-readable formats is, of course, nothing new. As explained above, to convert a work from one format to another is to make a reproduction of the work (and not a basis for a claim of fair use).¹⁰⁷ No one would argue, for example, that saving a Word document as a PDF, or an analog photo to a JPG file, negates the protectability of the work in its original format. Indeed, as has been recognized by courts and the Copyright Office, the computer code embodying an aesthetic work and a rendering of the work generated by that code can properly be considered one and the same work.¹⁰⁸

Copyright law also protects works that have been divided up and stored in digital bits and pieces. In the file-sharing case *Columbia Picture Industries, Inc. v. Fung*,¹⁰⁹ for instance, the Ninth Circuit did not hesitate to hold that the unauthorized uploading and downloading of motion pictures using BitTorrent technology—which breaks works

105. See *Authors Guild, Inc. v. Google Inc.*, 804 F.3d 202, 217, 225 (2d Cir. 2015) (describing purpose of Google’s copying as “to make available significant information *about those books*,” including “whether[] and how often they use specified words or terms”) (emphasis in original); see also Robert Brauneis, *Copyright and the Training of Human Authors and Generative Machines*, 47 COLUM. J.L. & ARTS 1, 36–37 (2025) (explaining that from a copyright perspective, it is problematic to conflate information derived from a work during the AI training process (“metainformation”) with information about the work (“metadata”).

106. 804 F.3d at 224–25.

107. See *supra* notes 98–104 and accompanying text.

108. See US Copyright Office, COMPENDIUM OF U.S. COPYRIGHT OFFICE PRACTICES § 721.10(A) (2021) (explaining that computer program and the screen displays generated by that program “are considered the same work, because the program code contains fixed expression that produces the screen displays”); see also *Williams Elecs., Inc. v. Artic Int’l, Inc.*, 685 F.2d 870, 874 (3d Cir. 1982) (holding changing videogame display protectable because fixed in memory device by which it was generated) (citing *Stern Elecs., Inc. v. Kaufman*, 669 F.2d 852, 855–56 (2d Cir. 1982) and *Midway Mfg. Co. v. Artic Int’l, Inc.*, 547 F. Supp. 999, 1006–07 (N.D. Ill. 1982), for same principle). Of course, computer code that embodies and renders a particular artistic work is distinct from software (such as Microsoft Word) that is used as a tool to *create* such a work; the two are separately protectable works.

109. 710 F.3d 1020, 1024 (9th Cir. 2013).

into fragments for rapid transmission among simultaneous users—was infringing.¹¹⁰

Of course, in the context of generative AI, a training work is encoded into an AI model along with millions of other works, all of which contribute to the generative capacity of the model. Due to limitations of generative AI technology, it seems impossible to extract a particular work from a model once it has been ingested, at least in this point in time.¹¹¹ But the fact that a particular work coexists with myriad others—and that the encoding process does not permit *post hoc* removal of particular works—does not excuse the infringement of that singular work. The illicit copying of an individual work is not excused by illicit copying of others.¹¹² AI companies seek to minimize the significance of copying of individual works by asserting, for example, in the LLM context, that it is “the *volume* of text used [in training], more than any particular *selection* of text, that really matters.”¹¹³ In other words, they need to copy an enormous quantity of works to build their models. But copying more does not mean one is copying less. In the words of the Supreme Court, “a taking may not be excused merely because it is insubstantial with respect to the *infringing work*.”¹¹⁴

A generative AI model is a product of the works encoded within it. Thus understood, far from existing independently of the works on which it was trained, an AI model can appropriately be considered a derivative or compilation of the works it embodies.¹¹⁵

110. *Id.* at 1026, 1034. Cloud storage services rely on analogous technology, segmenting files into smaller “blocks” for hosting in multiple locations, from which locations they are more efficiently retrieved (and reconstructed) for users. *See What is Block Storage?*, CLOUDFLARE, <https://www.cloudflare.com/learning/cloud/what-is-block-storage/> [<https://perma.cc/NPN9-C7WZ>] (last visited June 29, 2024).

111. Stephen Pastis, *A.I.’s Unlearning Problem: Researchers Say It’s Virtually Impossible to Make an A.I. Model ‘Forget’ the Things it Learns from Private User Data*, FORTUNE (Aug. 30, 2023, 11:43 AM), <https://fortune.com/europe/2023/08/30/researchers-impossible-remove-private-user-data-delete-trained-ai-models/> [<https://perma.cc/WB6D-FE5W>] (“If a machine learning-based system has been trained on data, the only way to retroactively remove a portion of that data is by re-training the algorithms from scratch.” (quoting New York University computer scientist Anasse Bari)).

112. *See, e.g., Sony Music Ent. v. Cox Commc’ns, Inc.*, 93 F.4th 222, 236–37 (2024) (upholding infringement claims in case involving mass infringement of numerous plaintiffs’ copyrighted musical works and sound recordings).

113. *Tremblay MTD*, *supra* note 2, at 2 (emphasis in original); *see also* Damle, *supra* note 2, at 13 (arguing that AI companies must operate without licenses due to the need to train on “virtually the entire internet”).

114. *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. at 539, 565 (1985) (emphasis in original) (holding that copied material, though comprising a small part of defendant’s work, was infringing).

115. *See infra* notes 203–08 and accompanying text.

3. Regeneration of Training Works in Output

AI fair use advocates do not deny that generated output that is substantially similar to copyrighted works on which the system was trained is problematic. Nor do they seem to be asserting that such copies—which could serve as substitutes for the originals—fall within the ambit of fair use. Instead, the fair use camp maintains that the propensity of AI models to reproduce memorized copies of training materials in output should be viewed as an “aberration” or “quirk”¹¹⁶—that is, as inconsequential collateral damage that AI companies are trying to ameliorate and which should not be held against them.¹¹⁷

The generation of infringing copies of training works by AI systems does not seem especially rare.¹¹⁸ In fact, absent the development of truly effective filtering systems, it would appear to be a possibility with respect to any work in the training set should the right prompt (or other method of extraction) be employed.¹¹⁹ In any event, regardless of how often replication of training materials occurs, or the particular prompts that cause it to manifest, even AI companies seem to acknowledge that generated content that is substantially similar to training material is appropriately considered infringing.¹²⁰

4. RAG Activities

Last but not least, RAG-enabled AI systems ingest unlicensed content from websites and other sources to feed into user queries in order to augment the substance and relevance of system output.¹²¹ It is thus apparent that RAG technology is exploiting third-party content for its expressive value.¹²² At the same time, RAG has only recently emerged as a significant mode of unauthorized exploitation. We have not yet heard much from AI companies regarding a supposed justification for the ongoing unlicensed copying required to enable RAG functionality.

116. See, e.g., Damle, *supra* note 2, at 8.

117. See *supra* note 76 and accompanying text (showing AI companies and advocates characterizing generation of infringing output as rare occurrence).

118. See *supra* notes 66–70 and accompanying text.

119. See *supra* notes 71–77 and accompanying text (discussing memorization and challenges of filtering).

120. See, e.g., Damle, *supra* note 2, at 8.

121. See *supra* notes 81–88 and accompanying text.

122. See *id.*

IV. FAIR USE ANALYSIS

Section 107 of the Copyright Act sets forth the four factors courts are to consider in evaluating a claim of fair use.¹²³ A review of those factors and relevant decisional law confirms that there is no established principle of fair use that sanctions the mass reproduction and exploitation of copyrighted works to build or operate a generative AI system.

A. The Cases Relied Upon by AI Companies

Pointing to a handful of technology-driven fair use cases, AI companies and their advocates claim that large-scale reproduction of copyrighted works to develop and populate AI systems constitutes a fair use of those works.¹²⁴ But *Google Books*, *HathiTrust*, *Sega* and other key precedents relied upon by AI companies to defend their unlicensed copying¹²⁵—including *Kelly v. Arriba Soft Corporation*,¹²⁶ *Perfect 10, Inc. v. Amazon.com, Inc.*,¹²⁷ *A.V. v. iParadigms, LLC (iParadigms)*,¹²⁸ *Sony Computer Entertainment, Inc. v. Connectix Corporation (Sony*

123. Section 107 provides in pertinent part as follows:

[T]he fair use of a copyrighted work, including such use by reproduction in copies or phonorecords or by any other means specified by that section, for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research, is not an infringement of copyright. In determining whether the use made of a work in any particular case is a fair use the factors to be considered shall include—

- (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.

17 U.S.C. § 107.

124. See *supra* notes 2–7 and accompanying text (summarizing fair use claims by AI companies and their representatives).

125. See Damle, *supra* note 2, at 5–6 (characterizing *Google Books*, *HathiTrust* and *Sega* as “foundational precedents” that establish the use of copyrighted works to develop AI systems as “quintessentially fair”); *id.* at 5–6 nn.10–23 (citing additional technological copying cases); Sag, *Copyright Safety*, *supra* note 2, at 304–09 (asserting legitimacy of AI copying under *Google Books* and *HathiTrust*); Sag, *Fairness*, *supra* note 24, at 1903 (interpreting technological fair use cases to support claim that AI copying is fair use).

126. 336 F.3d 811 (9th Cir. 2003).

127. 508 F.3d 1146 (9th Cir. 2006).

128. 562 F.3d 630 (4th Cir. 2009).

Computer)¹²⁹ and *Google, LLC v. Oracle America, Inc. (Oracle)*¹³⁰—are all in a different category with respect to fair use. This is because these cases were concerned with functional rather than expressive uses of copied works. The copying challenged in each was to enable a technical capability such as search functionality or software interoperability. By contrast, copying by AI companies serves to enable exploitation of protected expression.

1. Earlier Cases Address Different Conduct

AI companies seem to place their greatest faith in the mass book copying cases, *Google Books* and *HathiTrust*. But to claim that these Second Circuit decisions legitimize the sorts of copying engaged in by AI systems is to take their carefully limited holdings a bridge much too far.

There was no general declaration in either *Google Books* or *HathiTrust* that mass reproduction of copyrighted works to construct a product predicated upon large-scale copying has a presumptive claim to fair use. To the contrary, the *Google Books* panel was careful to limit its holding to the particular circumstances before it,¹³¹ including the fact that Google’s search functionality returned only snippets of text that did not permit meaningful consumption of expressive content.¹³² Although Google made full-text copies of the books, it was not seeking to capitalize on, or allow users to make use of, the aesthetic content of those works. Even so, the court considered Google’s copying to “test the boundaries of fair use.”¹³³ Indeed, the court pointedly observed that had Google permitted users greater access to “the expressive content” of the book, such exploitation “would most likely constitute copyright infringement if not licensed by the rights holders.”¹³⁴

The determination of fair use in *HathiTrust* was similarly confined to the facts of that case, which also involved large-scale unauthorized scanning of books to create a searchable database.¹³⁵ In *HathiTrust*, however, the search results were even more limited,

129. 203 F.3d 596 (9th Cir. 2000).

130. 593 U.S. 1 (2021).

131. See *Google Books*, 804 F.3d at 207, 222, 224–25, 229 (qualifying its holding with terminology such as “at least under present conditions,” “in these circumstances,” “at this time,” “at least as presently structured,” “as . . . presently constructed,” “at least as . . . presently designed,” “[o]n the present record,” etc.); see also Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC’Y U.S.A. 291, 294 (2019) (noting that *Google Books* and *HathiTrust* “were a product of the particular factual circumstances”).

132. See *supra* note 39 and accompanying text (describing snippet functionality).

133. *Google Books*, 804 F.3d at 206.

134. *Id.* at 226.

135. *HathiTrust*, 755 F.3d at 97.

identifying only the page numbers on which and number of times a specific term appeared in the relevant text.¹³⁶

In both *Google Books* and *HathiTrust*, the court drew a line between uses that were functional and nonsubstitutional in nature, on the one hand, and uses that were not, on the other. Despite the superficial similarity to the mass copying in *Google Books* and *HathiTrust*, then, AI copying cannot be squared with the fair use finding in either of these cases, the holdings of which were careful to preserve copyright owners' legitimate interest in the expressive content of their works.

Turning to *Kelly*, an earlier technological copying case decided by the Ninth Circuit, there a professional photographer challenged a search engine's reproduction and indexing of images taken from his own website and others he had licensed.¹³⁷ In responding to user queries, the search engine returned small, low-quality reproductions that users could click on to access linked full-sized versions of the works.¹³⁸ The court held that the defendant's copying of photos to provide a search and indexing service was a transformative fair use because the low-quality thumbnails "serve[d] a different function" than the originals—namely, "improving access to information on the internet."¹³⁹ The court pointedly distinguished this purpose from copying to capitalize on "artistic expression."¹⁴⁰ It found the search engine's use to be non-substitutional because, unlike Kelly's original photos, the search and indexing function was "unrelated to any aesthetic purpose."¹⁴¹

Like *Kelly*, *Perfect 10* involved unlicensed copying and indexing of online images, in this case by Google, which displayed thumbnails of Perfect 10's copyrighted photos to users of its search technology.¹⁴² Users similarly could access the full-sized photographs via a link to the originating website.¹⁴³ Invoking *Kelly*, the Ninth Circuit once again held that a search engine's copying of images for thumbnail display was a transformative fair use because the images were not being used for their intrinsic purpose, but rather to create "an electronic reference tool."¹⁴⁴

136. *Id.* at 91. The *HathiTrust* court pointedly observed that its determination was made without foreclosing future claims "based on a different record." *Id.* at 101.

137. *Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 815–16 (9th Cir. 2003).

138. *Id.*

139. *Id.* at 818–19.

140. *Id.* at 819.

141. *Id.* at 818–19.

142. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1155–57 (9th Cir. 2006).

143. *Id.*

144. *Id.* at 1164–65.

In each of these cases the fair use determination turned on the fact that the defendant was not exploiting expressive content for its intrinsic aesthetic value. So, too, in *iParadigms*.¹⁴⁵ There, the Fourth Circuit considered student authors' challenge to iParadigms' "Turnitin" plagiarism detection service, which digitally compared school papers against an archive of previously submitted papers to identify instances of copying.¹⁴⁶ The student plaintiffs alleged that the archiving of their papers in the Turnitin database violated their copyrights.¹⁴⁷ Once again, the deciding court focused on the fact that the defendant's use of the copied content "had an entirely different function and purpose than the original works," emphasizing that the use was "unrelated to any creative component" of the student works.¹⁴⁸ The fact that the defendant was not seeking to exploit student essay content as such readily distinguishes the copying at issue in *iParadigms* from that engaged in by generative AI systems, which seek to harness the intrinsic value of copied works.

Sega, *Sony Computer*, and *Oracle* are even farther afield from the types of copying engaged in by AI companies. In each of these cases, the copying was of specific material for the purpose of facilitating interoperability—clearly not the objective of AI copying.

In *Sega*, for instance, defendant Accolade copied Sega's videogame code in order to "reverse engineer" it to identify the functional requirements that would allow it to develop videogames that would operate on Sega consoles.¹⁴⁹ Through this process, Accolade identified a small segment of code in Sega's program that served to limit the game's operability to Sega consoles.¹⁵⁰ Although Accolade engaged in "wholesale" copying of the videogame program to identify the digital lock, Accolade's final product included only the brief segment of functional code (for which Sega did not pursue a separate claim of infringement).¹⁵¹ The Ninth Circuit held Accolade's "intermediate" copying of Sega's program to be a fair use because it served a "legitimate, essentially non-exploitative purpose"—that is, to discover the functional aspects of Sega's code so it could produce compatible works.¹⁵² Observing that intermediate copying was "the only means" by which Sega could access unprotected elements of the console code, the

145. *A.V. v. iParadigms, LLC*, 562 F.3d 639 (4th Cir. 2009).

146. *Id.* at 633–34.

147. *Id.* at 635.

148. *Id.* at 639, 641–42.

149. *Sega Enters. Ltd. v. Accolade, Inc.*, 977 F.2d 1510, 1514–15, 1522 (9th Cir. 1992).

150. *Id.* at 1515–16.

151. *See id.* at 1516–18.

152. *Id.* at 1522–23.

court expressed concern that a finding of infringement would grant Sega a “*de facto* monopoly” over functional aspects of the code.¹⁵³ In concluding that the reverse engineering was fair use, the court also emphasized that Accolade’s works did not incorporate any of Sega’s creative expression.¹⁵⁴ Thus, the intermediate copying in *Sega* cannot be equated with copying by AI systems, the very different purpose of which is to capture and use expressive content.

The Ninth Circuit adhered to the *Sega* precedent in *Sony Computer*, another reverse engineering case in which the court held the defendant’s intermediate copying of plaintiff Sony’s console code to create a “virtual” gaming station that would allow users to play Sony games on a personal computer to be a fair use.¹⁵⁵ Once again, the holding was premised on the court’s determination that the code at issue “contain[ed] unprotected functional elements” that could not be accessed or studied without copying.¹⁵⁶

Intermediate copying—an interim step in the software development process to ascertain functional requirements for a technical purpose—is not comparable to the copying that occurs in the context of generative AI. As its name would suggest, intermediate copying does not involve ongoing engagement with or exploitation of protected content.¹⁵⁷ The software developers in *Sega* and *Sony Computer* were not seeking to replicate or profit from the plaintiffs’ artistic works but instead to produce independently created, compatible products.¹⁵⁸ An AI model, by contrast, is designed to exploit the copied creative expression it embodies for as long as it remains operational.

Finally, in *Oracle*, the US Supreme Court held that Google’s copying of Java “declaring” code (a widely used system of computer commands) so Java-trained coders could more easily write programs for Android phones was a fair use.¹⁵⁹ In so holding, the Court emphasized the functional nature of the code appropriated by Google, the copyrightability of which it found questionable.¹⁶⁰ The challenged

153. *Id.* at 1523–27.

154. *Id.* at 1522 (“[T]here is no evidence in the record that Accolade sought to avoid performing its own creative work.”).

155. *Sony Comput. Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 601, 608 (9th Cir. 2000).

156. *Id.* at 603–04.

157. *See Sega*, 977 F.2d at 1518.

158. *See id.* at 1518, 1526–28 (holding intermediate copying of object code fair use where it was the only means of access to elements not protected by copyright); *see also Sony Comput.*, 203 F.3d at 598 (resulting product did not contain any of plaintiff’s copyright-protected material).

159. *Google LLC v. Oracle Am., Inc.*, 593 U.S. 1, 16 (2021).

160. *See id.* at 29 (“[T]he declaring code is, if copyrightable at all, further than are most computer programs . . . from the core of copyright.”). In a nuanced analysis of the majority opinion,

copying targeted a widely used, utilitarian aspect of computer code that the Court viewed as far from the “core of copyright’s protection.”¹⁶¹ The Court was persuaded that Google did not copy the lines of Java code “because of their creativity, their beauty, or even (in a sense) because of their purpose,” but so Java-trained programmers writing software for Android phones could rely on commands that “they [were] already familiar with to call up particular tasks.”¹⁶² In other words, the Court was focused on the interoperability of software coders.¹⁶³

In short, none of the fair use precedents on which AI companies purport to rely addressed a product designed to copy and exploit third-party expressive content to derive new content, including potentially infringing and competing content. None involved copying and use of expressive content for its intrinsic value.

This critical distinction between expressive and non-expressive use of copyrighted works in a technological context is well illustrated by the Second Circuit’s post-*Google Books* decision in *Fox News Network v. TVEyes, Inc. (TVEyes)*.¹⁶⁴ In that case, the court considered an unlicensed service that recorded “essentially all television broadcasts” and the closed-caption text that accompanied them to create a text-searchable database of the copied video.¹⁶⁵ By typing in a search term, users would receive a list of news clips in which the term appeared, which references could be clicked on and played.¹⁶⁶

In defending its copying, TVEyes relied primarily on *Google Books*.¹⁶⁷ But the court firmly distinguished that precedent to hold TVEyes’ systematic exploitation of Fox’s news content to be an infringing, rather than fair use.¹⁶⁸ Significantly, the court’s rejection of fair use was predicated upon its core finding that the use of Fox’s content was “both ‘extensive’ and inclusive of all that is ‘important’ from

Professor Jane Ginsburg suggests that “[i]n effect, the fair use determination achieved the same result as ruling the [copied code] uncopyrightable, but attained that objective through the back end of a copyright exception rather than the front end of applying the idea/expression distinction” Jane C. Ginsburg, *Fair Use in the US Redux: Reformed or Still Deformed?*, SINGAPORE J. LEGAL STUD. 1, 3 (2024).

161. *Oracle*, 593 U.S. at 29.

162. *Id.* at 34.

163. *See id.* (explaining that Google copied Oracle’s code “because programmers had already learned to work with the [Java] system, and it would have been difficult, perhaps prohibitively so, to attract programmers to build its Android smartphone system without them.”).

164. 883 F.3d 169 (2d Cir. 2018).

165. *Id.* at 175.

166. *Id.*

167. *Id.* at 177.

168. *Id.* at 174–75, 179–81.

the copyrighted work.”¹⁶⁹ Unlike in *Google Books*, TVEyes had built a business based on the appropriation of expressive content:

The success of the TVEyes business model demonstrates that deep-pocketed consumers are willing to pay well for a service that allows them to search for and view selected television clips, and that this market is worth millions of dollars in the aggregate. Consequently, there is a plausibly exploitable market for such access to television content, and it is proper to consider whether TVEyes displaces potential Fox revenues¹⁷⁰

Taking note of the Second Circuit’s cautionary language in *Google Books* that, notwithstanding the implementation of guardrails to prevent expressive use of the digitized books, Google’s conduct “test[ed] the boundaries of fair use,” the *TVEyes* court concluded that defendant TVEyes had “exceeded those bounds.”¹⁷¹

2. There Is No Presumption that a Use Is Fair if the End Product Is Noninfringing

A favored argument of those advocating on behalf of AI companies is that the appropriation of copyrighted works to train AI systems is a fair use of those works so long as the resulting product does not infringe copied expression. Damle, for instance, asserts in his congressional testimony that “[a]n unbroken line of cases establishes that the use of a copyrighted work to create a non-infringing final product is quintessential fair use.”¹⁷² Professor Matthew Sag puts it a little differently, avowing that US fair use cases “have consistently held that technical acts of copying which do not communicate an author’s original expression to a new audience are fair use.”¹⁷³ In other words, unless the infringing activity is visible to others, there is no infringement.

No court has ever articulated a canon of fair use such as the one posited by Damle and Sag. Indeed, the Ninth Circuit’s *Sega* opinion, on which both rely,¹⁷⁴ expressly *rejects* any such standard. Addressing the defendant’s similar argument in *Sega*—that Accolade’s intermediate

169. *Id.* at 179.

170. *Id.* at 180.

171. *Id.* at 174 (quoting *Authors Guild, Inc. v. Google Inc.*, 804 F.3d 202, 206 (2d Cir. 2015)).

172. Damle, *supra* note 2, at 5, 7.

173. Sag, *Fairness*, *supra* note 24, at 1903. Sag relatedly contends that AI copying is a “nonexpressive” use of the copied works. *See id.* at 1914; Sag, *Copyright Safety*, *supra* note 2, at 307-09. For the reasons discussed above, *see supra* notes 98–105 and accompanying text, the “nonexpressive use” claim is unconvincing. As Professor Robert Brauneis pithily observes: “[I]f we define ‘non-expressive use’ as ‘a use of a copyrighted work that is indifferent to the expressive choices made by the author of the work,’ generative AI training uses are anything but indifferent to the author’s expressive choices.” *See Brauneis, supra* note 105, at 32; *see also supra* notes 97–104 (explaining why the “nonexpressive use” claim is unconvincing).

174. *See* Damle, *supra* note 2, at 6; *see also* Sag, *Fairness*, *supra* note 24, at 1904.

copying had to be fair because its final product was noninfringing—the court was clear that “intermediate copying of computer object code may infringe the exclusive rights granted to the copyright owner in Section 106 of the Copyright Act *regardless of whether the end product of the copying infringes those rights*.”¹⁷⁵ This principle was reaffirmed by the Ninth Circuit nearly twenty years later in *Sony Computer*.¹⁷⁶ In short, there is no “unbroken line” of precedent that immunizes AI copying as noninfringing if the generated output is considered noninfringing.

B. AI Copying Under the Fair Use Factors

1. The Purpose and Character of AI Copying

The first fair use factor of Section 107 considers the purpose and character of the challenged use, including—under judge-made law—whether the use is “transformative” and whether it is for a commercial purpose.¹⁷⁷ The advocates of AI fair use do not claim that the training of for-profit AI models falls within any of the favored purposes listed in the preamble to Section 107—criticism, comment, news reporting, teaching, scholarship, or research.¹⁷⁸ Instead, they look to the technological fair use precedents discussed above to argue that the use is transformative. As shown, however, those cases were careful to distinguish copying and exploitation of expressive content for its intrinsic value from the nonconsumptive technological copying they found to be fair. These cases thus point in the opposite direction of fair use with respect to copying by AI entities, as do the purpose and character of such copying.¹⁷⁹

a. AI Copying Is Not Transformative

For purposes of fair use, a transformative use is one that “adds something new, with a further purpose or different character” and does

175. *Sega*, 977 F.2d at 1517–19 (emphasis added). The court instead grounded its fair use finding on the fact that the copying was undertaken solely to identify functional elements rather than to exploit Sega’s creative expression. *Id.* at 1522–23.

176. *See Sony Comput. Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 602–03 (9th Cir. 2000) (“In *Sega*, we recognized that intermediate copying could constitute infringement even when the end product did not itself contain copyrighted material.”).

177. *See Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 527–29 (2023); *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 578–79 (1994).

178. *See* 17 U.S.C. § 107.

179. In addition to being relevant to claims of direct infringement, the purpose and design of AI models present significant issues of secondary copyright liability arising from user activities, a rich subject in itself that is beyond the scope of this paper.

not supplant the original.¹⁸⁰ In *Google Books*, for example, the Second Circuit held that the digital scanning of books to provide a search function was transformative because it “augments public knowledge by making available information about Plaintiffs’ books without providing the public with a substantial substitute for matter protected by the Plaintiffs’ copyright interests in the original works or derivatives of them.”¹⁸¹ The *HathiTrust* court reached a similar conclusion, holding that the copying of books to create a searchable database was a transformative use that “add[ed] ... something new with a different purpose and a different character.”¹⁸² In *Sega*, a videogame manufacturer disassembled Sega’s console code, not to copy Sega’s creative expression, but rather to develop its own original, compatible games.¹⁸³ The “transformativeness” of the use in each of these cases turned on the fact that the copied content was not being exploited for its aesthetic value, either by the defendants themselves or by users of their products.

Copying of protected works by generative AI systems has no similar claim to transformativeness. Works are copied in their entirety and mechanically encoded in the AI model without offering any search mechanism or other functional utility, let alone criticism or commentary. They are copied to capture and channel their expressive value.

Fair use proponents contend that the training process merely records “unprotected facts” about the training works but, as illustrated above, that is not the reality.¹⁸⁴ In fact, the AI model maps and stores the expressive content of each work so it can be tapped to enable the model’s generative capabilities. That the works are parsed into small segments and converted into vector representations to achieve this goal does not negate, but rather underscores, the systematic nature of the appropriation.

The mechanical mapping of works to exploit them for their aesthetic value does not qualify as a transformative use. To be sure, copyrighted works are “transformed” when they are encoded in the model, but not in the sense of fair use “transformativeness”—rather in the sense of creating a derivative work. Under the Copyright Act, the owner of a protected work has the exclusive right to prepare, or

180. *Warhol*, 598 U.S. at 528 (quoting *Campbell*, 510 U.S. at 579).

181. *Authors Guild, Inc. v. Google Inc.*, 804 F.3d 202, 207 (2d Cir. 2015).

182. *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 105 (2d Cir. 2014).

183. *Sega Enters. Ltd. v. Accolade, Inc.*, 977 F.2d 1520, 1522–23 (9th Cir. 1992).

184. *See supra* notes 98–102 and accompanying text.

authorize the preparation of, a derivative of that work,¹⁸⁵ defined in the Act as “a work based upon one or more preexisting works,” including “any ... form in which a work may be recast, transformed or adapted.”¹⁸⁶ An AI model, created by reproducing and encoding—or “recasting”—copyrighted works into the model, falls within this definition. An AI model may also be considered a compilation, that is, “a work formed by the collection and assembling of preexisting materials or of data that are selected, coordinated, or arranged in such a way that the resulting work as a whole constitutes an original work of authorship.”¹⁸⁷ However the model is characterized, the goal is to capitalize upon the expressive material it embodies.

Fair use advocates may assert that, in assessing transformativeness, the purpose of the copying to be considered is not that of building and operating AI models—the immediate objective—but the ultimate goal of allowing users to generate new content from the works encoded in the models. This alternative argument for transformativeness falls short for several reasons.

To begin with, the generation of new content by an AI model based on a user prompt is inextricably bound up with the expressive exploitation of the works residing in the model. AI generation is not a content-neutral function akin to search capability or interoperability. As discussed above, AI machines do not “think” independently; their production of new content is a function of and limited by the creative content they have ingested.¹⁸⁸ Accordingly, the use of copyrighted works to facilitate generative activities does not align with the reasoning of *Google Books*, *HathiTrust* or other technological cases in which the copying was found to be transformative because it was not for the purpose of exploiting expressive content.

Next, AI machines not infrequently generate infringing copies and derivatives of the works on which they are trained.¹⁸⁹ While this may not be a desired outcome, it seems to be an unavoidable feature of AI systems, at least at present.¹⁹⁰ To state the obvious, AI companies cannot claim transformativeness based on output that merely reproduces all or part of a training work or works, even if it is done in a complicated way. An AI model’s generation of text or an image that is a copy of a training work, or similar enough to serve as a substitute for

185. See 17 U.S.C. § 106 (“[T]he owner of a copyright ... has the exclusive right[] to ... prepare derivative works based upon the copyrighted work.”).

186. 17 U.S.C. § 101 (defining “derivative work”).

187. 17 U.S.C. § 101 (defining “compilation”).

188. See *supra* notes 17–31 and accompanying text.

189. See *supra* notes 61–68 and accompanying text.

190. See *supra* notes 72–74 and accompanying text.

the original, is presumptively nontransformative.¹⁹¹ Nor is there anything inherently transformative about combining elements of one work with those of another work or works, which invades the copyright owners' derivative work rights.¹⁹² As Professor Jane Ginsburg observes, "AI outputs may incorporate the source works' expression in a new production; but that output generally will not comment, criticize, shed light on or otherwise be *about* the copied expression."¹⁹³

Notably, in *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith* (Warhol),¹⁹⁴ the US Supreme Court warned against "an overbroad concept of transformative use" that encroaches upon copyright owners' derivative work rights, explaining that an interpretation of transformativeness "that includes any further purpose, or any different character" could "swallow" the copyright owner's exclusive right to create derivative works.¹⁹⁵ To this end, the Court criticized overzealous application of "transformativeness" to encompass any work that "adds some new expression, meaning, or message."¹⁹⁶ Drawing on its earlier explication of fair use in *Campbell v. Acuff-Rose Music, Inc.*,¹⁹⁷ the Court emphasized that the secondary user must have an independent justification for use of the work in question; that a copied work may be useful to convey a new meaning or message is not justification enough.¹⁹⁸

191. See *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 531–32 (2023) (first fair use factor likely to weigh against fair use where "an original work and a secondary use share the same or highly similar purposes, and the secondary use is of a commercial nature").

192. See, e.g., *id.* at 537, 550–51 (unlicensed commercial use of plaintiff's photograph, as incorporated into an Andy Warhol silkscreen derivative, was nontransformative and therefore infringing); *Dr. Seuss Enters., L.P. v. ComicMix LLC*, 983 F.3d 443, 451–55 (9th Cir. 2020) (ComicMix's unlicensed book consisting of a "mashup" of Dr. Seuss and Star Trek characters that mimicked Dr. Seuss illustrations was a nontransformative use of Seuss's works).

193. Ginsburg, *supra* note 160, at 29.

194. *Warhol*, 598 U.S. at 508 (2023).

195. *Id.* at 529, 541. In keeping with this instruction, the Court determined that a magazine's commercial use of a silkscreen image created by Andy Warhol from plaintiff Goldsmith's photographic portrait of Prince was not transformative because it served as a substitute for Goldsmith's original photo. *Id.* at 523–24.

196. *Id.* at 541.

197. In *Campbell*, the Court considered 2 Live Crew's use of Roy Orbison's classic song "Pretty Woman" in a rap parody, finding the use transformative because it was necessary to copy portions of Orbison's work in order to mock it. See 510 U.S. 569, 579–83 (1994).

198. See *Warhol*, 598 U.S. at 532 (independent justification "particularly relevant" where unlicensed copying could displace a market for derivatives); *id.* at 532–33 ("If an original work and a secondary use share the same or highly similar purposes, and the secondary use is of a commercial nature, the first factor is likely to weigh against fair use, absent some other justification for copying."); *id.* at 547 ("Copying might [be] helpful to convey a new meaning or message. It often is. But that does not suffice under first factor.").

AI-generated content that is not a recognizable copy or derivative of a training work or works—that is, the type of content AI companies claim to be the intended output of their systems—by definition does not comment or shed light on any particular work. It is difficult to see how AI entities can stake a claim to transformative use based on output that does not convey commentary or criticism with respect to those works. Nor, as noted, does AI output facilitate a utility-expanding function such as searchability or indexing.

This leaves us with the bare claim that AI copying should be considered transformative because it enables the generative capabilities of AI models. This broad contention is untethered to the use of any particular work or works, but instead boils down to an assertion that mass appropriation of protected works is justified because extensive copying is necessary to build and operate AI systems. In effect, then, it amounts to a policy argument that the rights of copyright owners must yield to the presumed social benefits of generative AI technology.

Courts sometimes consider the public benefit of a challenged use in evaluating the question of transformativeness. In *Google Books*, for instance, the court determined that “Google’s making of a digital copy to provide a search function is a transformative use, which augments public knowledge by making available information *about* Plaintiffs’ books without providing the public with a substantial substitute.”¹⁹⁹ In *Sega*, the court was concerned that a failure to permit reverse engineering would “confer[] on the copyright owner a de facto monopoly” over unprotectable ideas and functional concepts in its computer code that others could build upon.²⁰⁰ In *Oracle*, the Court reasoned that utilizing Java declaring code would enable Java-trained programmers to “expand the use and usefulness” of Android phones.²⁰¹

AI copying does not fit within these paradigmatic examples of technologically driven copying. To begin with, none of these cases held the systematic extraction of creative value from protected works to be a transformative use. Moreover, unlike the far-reaching copying activities of generative AI, in each of these cases, the asserted beneficial purpose could not have been achieved without copying the specific work or works at issue. That is, there was a clear nexus between the appropriated works and the claimed social benefit. In *Google Books*, for example, Google could not provide a search function to locate a term within a book without copying the book. In *Sega*, the court determined

199. Authors Guild, Inc. v. Google Inc., 804 F.3d 202, 207 (2d Cir. 2015) (emphasis in original).

200. Sega Enters. Ltd. v. Accolade, Inc., 977 F.2d 1510, 1573-74 (9th Cir. 1992).

201. Google LLC v. Oracle Am., Inc., 593 U.S. 1, 30 (2021).

that copying Sega's console code was necessary to access functional aspects of the code to facilitate the creation of compatible games. The same was true in *Oracle*, where the only way to harness the collective knowledge of Java programmers was to emulate the Java command structure. The same cannot be said of generative AI systems, which copy millions of works in an indiscriminate fashion. The generalized nature of AI copying is inconsistent with claims that the copying serves a transformative purpose vis-à-vis the copied works.²⁰²

The claim that AI companies' unlicensed mass copying is transformative seems largely premised on the view that generative AI is remarkable technology with the power to enhance creativity and improve society.²⁰³ But a broad assertion of public good such as this does not justify the exploitation of specific copyrighted works or rise to the level of a transformative purpose.²⁰⁴ If it did, fair use could be claimed with respect to any beneficial technological advance that sought to capitalize on copyrighted works for free.²⁰⁵

Finally, a claim of transformative purpose is especially weak where, as here, the unauthorized copying is supplanting the licensed use of copyrighted works for the same purpose—namely, to train and

202. Cf. *Dr. Seuss Enters., L.P. v. ComicMix LLC*, 983 F.3d 443, 454 (9th Cir. 2020) (finding no transformative use where defendant ComicMix had no particular need to use Seuss's material for its story).

203. See, e.g., Damle, *supra* note 2, at 1 ("The AI tools of the present and near future will impact almost every aspect of the human experience They will transform the way humans learn and work. They will enable anyone to more fully unlock their creative potential. In short, AI has the potential to transform our economy and improve our society as a whole."); Wayne Brough & Ahmad Nazeri, *Regulatory Comments Before the U.S. Copyright Office Library of Congress In the Matter of Artificial Intelligence and Copyright*, R STREET (Oct. 30, 2023), <https://www.rstreet.org/outreach/regulatory-comments-before-the-u-s-copyright-office-library-of-congress-in-the-matter-of-artificial-intelligence-and-copyright/> [https://perma.cc/JF9K-VUR8] ("AI promises to bolster the American economy, amplify the capabilities of creatives and catalyze advancements in science and the arts.").

204. On the question of public benefit, it is worth noting that although generative AI may have the potential for positive impact in various areas of human endeavor, it comes with significant social concerns, among them its ability to generate false information, impersonate individuals (including by producing "deepfakes"); and amplify racism and bias—not to mention its staggering energy needs. See, e.g., Öykü Isik, Amit Joshi & Lazaros Goutas, *4 Types of Generative AI Risk and How to Mitigate Them*, Harv. Bus. Rev. (May 31, 2024), <https://hbr.org/2024/05/4-types-of-gen-ai-risk-and-how-to-mitigate-them> [https://perma.cc/869L-FC6D]; Kate Crawford, *Generative AI's Environmental Costs are Soaring—and Mostly Secret*, NATURE (Feb. 20, 2024), <https://www.nature.com/articles/d41586-024-00478-x> [https://perma.cc/V5BH-K4T6] ("Within years, large AI systems are likely to need as much energy as entire nations.").

205. Interpreting the *Oracle* decision, Ginsburg aptly observes that if "verbatim copying 'to create new products' were deemed 'transformative' in general, it would be difficult to imagine what kind of copying, short of outright piracy of the entire work, would not be transformative." Ginsburg, *supra* note 160, at 7; see also *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. at 539, 569 (1985) ("Any copyright infringer may claim to benefit the public by increasing public access to the copyrighted work.").

operate AI systems—as discussed below.²⁰⁶ As the US Supreme Court emphasized in *Warhol*, where use of an original work and a secondary use “share the same or highly similar purposes and the secondary use is of a commercial nature, the first factor is likely to weigh against fair use, absent some other justification for copying.”²⁰⁷

AI companies have not articulated any justification for their mass unauthorized copying, apart from generalized assertions that it is necessary to build their systems and that licensing would be too difficult.²⁰⁸ General pronouncements such as these do not amount to transformative purpose under copyright law. In any event, as shown below, the claim that licensing is unworkable is contrary to the facts.

b. Commercial Purpose of AI Copying

Leading AI developers, including OpenAI, Stability AI, and Anthropic—not to mention the tech giants Google, Meta, and Microsoft—indisputably engage in for-profit activities, including the sale of AI products to the public.²⁰⁹ The use of copyrighted works to develop and operate a commercially marketed AI system is, by

206. See *infra* notes 209–30 and accompanying text.

207. Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith, 598 U.S. 508, 532–33 (2023).

208. See *supra* notes 33, 111 and accompanying text (AI companies’ assertions that AI systems must train on millions of works); *infra* note 233 and accompanying text (AI companies’ assertions that licensing would be impossible).

209. See *Pricing*, OPENAI, <https://openai.com/chatgpt/pricing/> [<https://perma.cc/MH9S-7QCE>] (last visited May 15, 2024) (listing subscription fees for OpenAI’s ChatGPT products); *Membership*, STABILITY AI, <https://stability.ai/membership> [<https://perma.cc/CND2-4ZAP>] (last visited May 15, 2024) (listing membership fees for Stability AI products); *Meet Claude*, ANTHROPIC, <https://www.anthropic.com/claude> [<https://perma.cc/84KB-E88R>] (last visited May 15, 2024) (listing subscription fees for Anthropic’s Claude products).

definition, commercial in nature.²¹⁰ Although not in itself dispositive, this factor weighs significantly against fair use.²¹¹

2. AI Companies Copy Highly Creative Works

With respect to the second factor, the “nature of the works” at issue, AI companies acknowledge that they have copied countless expressive works—such as books, movies and visual art—that lie at the heart of copyright.²¹² This presumptively weighs against fair use.²¹³

3. AI Companies Copy Works in Their Entirety

Factor three considers the amount of the work that was copied. There is no dispute that AI companies reproduce copyrighted works in their entirety to assemble training sets. During the training process, each work is broken down into small segments and algorithmically encoded into the model. The core activities of AI development, then, involve the reproduction of entire works.

210. This holds true even though, in some cases, AI companies acquire training materials assembled by academic and other nonprofit researchers. See Andy Baio, *AI Data Laundering: How Academic and Nonprofit Researchers Shield Tech Companies from Accountability*, WAXY (Sept. 30, 2022), <https://waxy.org/2022/09/ai-data-laundering-how-academic-and-nonprofit-researchers-shield-tech-companies-from-accountability/> [https://perma.cc/NU6E-5PG7] (discussing common practice of “technology companies working with AI to commercially use datasets and models collected and trained by non-commercial research entities like universities or non-profits”). Even if a researcher’s activities are conducted under the auspices of a nonprofit institution, this does not negate the commercial purpose of the follow-on AI entity. See *id.* Nor does it negate a finding of for-profit use on the part of the researcher. In *Weissman v. Freeman*, for example, the Second Circuit determined that an academic’s copying of a scientific paper, though not for monetary gain, was nonetheless a for-profit activity for purposes of the first fair use factor. 868 F.2d 1313, 1324 (2d Cir. 1989) (“Particularly in an academic setting, profit is ill-measured in dollars. Instead, what is valuable is recognition because it so often influences professional advancement and academic tenure.”); see also *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. at 539, 562 (1985) (“The crux of the profit/nonprofit distinction is not whether the sole motive of the use is monetary gain but whether the user stands to profit from exploitation of the copyrighted material without paying the customary price.”); *Worldwide Church of God v. Phila. Church of God, Inc.*, 227 F.3d 1110, 1117 (9th Cir. 2002); *Am. Geophysical Union v. Texaco Inc.*, 60 F.3d 913, 914–15, 921–22, 931 (2d Cir. 1994) (photocopying of journal articles by Texaco scientists for internal research purposes held not a fair use in part because Texaco reaped indirect economic advantage and “avoid[ed] having to pay at least some price to copyright holders”).

211. See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 585 (1994) (commercial use is a factor that “tends to weigh against fair use” (quoting *Harper & Row*, 471 U.S. at 562)).

212. *Id.* at 586 (“This factor calls for recognition that some works are closer to the core of intended copyright protection than others, with the consequence that fair use is more difficult to establish when the former works are copied.”).

213. *Id.* at 589; *Sony Corp. v. Universal City Studios, Inc.*, 464 U.S. 417, 455 n.40 (1984) (copying a news broadcast may have different fair use implications than copying a motion picture).

On the output end, works from training materials may be replicated in all or in part in AI productions. Further, in canvassing the internet for additional material to enhance output, RAG technology may appropriate the entirety or significant portions of the external sources it gathers.

In short, factor three weighs heavily against fair use, especially since the works are being utilized for their expressive content rather than a non-consumptive, functional purpose.²¹⁴

4. AI Copying Harms the Market for the Copied Works

The final fair use factor concerns the effect of the challenged use on the potential market for or value of the copyrighted work. In the AI context, this fourth factor encompasses at least three different concerns: harm to the market for the original works based on generation of competing copies; harm to the market for licensed derivatives based on those works; and harm to the licensing market for training use of the works.

A close copy of a work that may serve as a substitute for the original invades the market for the original. Likewise, a derivative based on a copied work may supplant or interfere with licensing opportunities for adaptations of the original. Especially in a commercial context, both types of output weigh against fair use.²¹⁵

Creators and owners of copyrighted works are deeply concerned that generative AI systems trained on artistic works have the capacity to generate output that, even if not an overt reproduction of an original, is stylistically similar enough to the original to serve as a competing substitute.²¹⁶ Imitative content may incorporate elements of a

214. *Compare* Authors Guild, Inc. v. Google Inc., 804 F.3d 202, 221–22 (2d Cir. 2015) (“While Google makes an unauthorized digital copy of the entire book, it does not reveal that digital copy to the public Google [thus] satisfies the third factor test.”) *with* Fox News Network v. TVEyes, Inc., 883 F.3d 169, 179 (2d Cir. 2018) (“Th[e] [third] factor clearly favors Fox because TVEyes makes available virtually the entirety of the Fox programming that TVEyes users want to see and hear.”).

215. *See* Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith, 598 U.S. 508, 536 (2023) (where secondary work is both substitutional and used for commercial purposes, this “counsels against fair use, absent some other justification”); *Campbell*, 510 U.S. at 590 (fair use analysis must consider not only harm to original work but market for derivatives).

216. For example, as articulated in the record labels’ suit against AI music generator Udio: “The capacity for a generative AI service to produce convincing imitations of genuine sound recordings starts with copying a vast range of sound recordings. When those who develop such a service steal copyrighted sound recordings, the service’s synthetic musical outputs could saturate the market with machine-generated content that will directly compete with, cheapen, and ultimately drown out the genuine sound recordings on which the service is built.”

copyrighted original that render the imitation an infringing derivative.²¹⁷ In other instances, however, generative output may not meet the traditional test for substantial similarity.²¹⁸ Even if not technically infringing, output that is recognizably derived from and could compete with a particular creator's works should weigh against a claim that the AI model's copying of such works during training was a fair use. That is, even if infringement cannot be established based on generated output, output that obviously imitates an artist's distinctive style may be probative of market harm resulting from copying at the training stage because that copying yielded a competing substitute.²¹⁹

AI companies rely on unauthorized copying—largely accomplished by scraping the internet—to develop their models.²²⁰ Creators and copyright owners were not generally aware of such companies' mass appropriation of copyrighted works until generative AI burst into public view in late 2022 with the launch of ChatGPT.²²¹ One consequence of this revelation has been a flood of lawsuits filed by writers, visual artists, musicians, and others asserting copyright infringement claims (and other causes of action) against AI companies.²²² A second consequence is a rapidly developing market for licensing of copyrighted content to AI companies for training and operation of their systems.

Complaint at ¶ 4, *UMG Recordings, Inc. v. Uncharted Labs, Inc.*, No. 1:24-cv-04777 (S.D.N.Y. June 24, 2024) [hereinafter *UMG Compl.*]; see also Kashmir Hill, *This Tool Could Protect Artists From A.I.-Generated Art That Steals Their Style*, N.Y. TIMES (Feb. 17, 2023), <https://www.nytimes.com/2023/02/13/technology/ai-art-generator-lensa-stable-diffusion.html> [<https://perma.cc/9W5G-W8HD>] (discussing concerns of visual artists whose work is being imitated by generative AI).

217. See *supra* note 74 and accompanying text (discussing derivative output).

218. See Sobel, *Elements*, *supra* note 80, at 38–41 (explaining why it may be difficult to establish substantial similarity in artistic style. To establish infringement, a plaintiff must prove, *inter alia*, that the challenged copy is substantially similar to the original. MELVILLE B. NIMMER & DAVID NIMMER, 4 NIMMER ON COPYRIGHT § 13D.10 (2024) (general overview of the substantial similarity requirement).

219. The production of imitative content may also be probative of copying at the training stage.

220. See *supra* notes 33–34 and accompanying text.

221. See Bernard Marr, *A Short History Of ChatGPT: How We Got To Where We Are Today*, FORBES (May 19, 2023, 1:14 AM), <https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/> [<https://perma.cc/8NGA-PAEB>] (explaining ChatGPT quickly “went viral” after November 2022 launch).

222. See, e.g., *Andersen v. Stability AI, Ltd.*, No. 3:23-cv-00201 (N.D. Cal. filed Jan. 13, 2023) (visual art); *Tremblay v. OpenAI, Inc.*, No. 3:23-cv-03223 (N.D. Cal. filed June 28, 2023) (books); *Silverman v. OpenAI, Inc.*, No. 3:23-cv-03416 (N.D. Cal. filed July 7, 2023) (books); *Concord Music Grp., Inc. v. Anthropic PBC*, No. 3:23-cv-01092 (M.D. Tenn. filed Oct. 18, 2023) (song lyrics); *New York Times v. Microsoft Corp.*, No. 1:23-cv-11195 (S.D.N.Y. filed Dec. 27, 2023) (news articles); *UMG Recordings, Inc. v. Suno, Inc.*, No. 1:24-cv-11611 (D. Mass. filed June 24, 2024) (sound recordings); *UMG Recordings, Inc. v. Uncharted Labs*, No. 1:24-cv-04777 (S.D.N.Y. filed June 24, 2024) (sound recordings).

Since 2023, OpenAI, the company behind ChatGPT, has entered into content licensing deals with Axel Springer, the publisher of *POLITICO* and *Business Insider*; News Corp, which owns *The Wall Street Journal*, the *New York Post*, *The Times*, and *The Sunday Times*; Dotdash Meredith, a large publisher of online content; and the Associated Press, to name a few.²²³ Google reached an agreement with Reddit to use Reddit data to train its AI models.²²⁴ Universal Music Group has entered into a partnership with AI technology company SoundLabs to provide an “ethically” trained voice cloning tool for its artists.²²⁵ Disney Music agreed to license AI music startup AudioShake to “open up” Disney’s historic catalog of works to new uses.²²⁶ Lionsgate, the motion picture studio, has partnered with AI company Runway to train a new AI model on Lionsgate’s film and television properties.²²⁷ The Copyright Clearance Center, an organization that licenses journal articles and other text-based materials, has extended its collective licensing service to cover AI uses.²²⁸ These are just some of the licensing arrangements that have been publicly disclosed; undoubtedly, there are a good number of others that have not been publicized or are still in the pipeline.

223. Angela Cullen & Jackie Davalos, *OpenAI to Pay Axel Springer Tens of Millions to Use News Content*, BLOOMBERG L. (Dec. 13, 2023, 11:37 AM), <https://news.bloomberglaw.com/tech-and-telecom-law/openai-to-pay-axel-springer-tens-of-millions-to-use-news-content> [https://perma.cc/T8YS-XQBT]; *Open AI and Wall Street Journal Owner News Corp Sign Content Deal*, GUARDIAN (May 22, 2024, 5:39 PM), <https://www.theguardian.com/technology/article/2024/may/22/openai-chatgpt-news-corp-deal> [https://perma.cc/GT68-EQUV]; Sara Fischer, *OpenAI Inks Licensing Deal with Dotdash Meredith*, AXIOS (May 7, 2024), <https://www.axios.com/2024/05/07/openai-dotdash-meredith-licensing-deal> [https://perma.cc/F8NE-9QJV]; Matt O'Brien, *ChatGPT-maker OpenAI Signs Deal with AP to License News Stories*, AP (July 13, 2023, 10:41 AM), <https://apnews.com/article/openai-chatgpt-associated-press-ap-f86f84c5bcc2f3b98074b38521f5f75a> [https://perma.cc/7Y3J-GX6T].

224. Annelise Gilbert, *Google-Reddit AI Deal Heralds New Era in Social Media Licensing*, BLOOMBERG L. (Mar. 7, 2024, 4:06 AM), <https://news.bloomberglaw.com/ip-law/google-reddit-ai-deal-just-the-start-for-social-media-licensing> [https://perma.cc/9NR2-VYST].

225. See Mandy Dalugdug, *Universal Music Artists Get Access to AI Voice Cloning Tool via UMG's New Deal with Tech Startup SoundLabs*, MUSIC BUS. WORLDWIDE (June 19, 2024), <https://www.musicbusinessworldwide.com/universal-music-artists-get-access-to-ai-voice-cloning-tool-via-umgs-new-deal-with-tech-startup-soundlabs/> [https://perma.cc/RQ3R-CXTL].

226. Murray Stassen, *Disney Music Group Strikes Deal With AI Music Startup AudioShake to 'Unlock New Listening and Fan Engagement Experiences' for Its Catalog*, MUSIC BUS. WORLDWIDE (July 15, 2024), <https://www.musicbusinessworldwide.com/disney-music-group-strikes-deal-with-ai-music-startup-audioshake-to-unlock-new-listening-and-fan-engagement-experiences-for-its-catalog/> [https://perma.cc/L2QG-BE5U].

227. Etan Vlessing, *Lionsgate CEO Says AI Deal Promises “Transformational Impact” on Studio*, HOLLYWOOD REP. (Nov. 7, 2024, 2:15 PM), <https://www.hollywoodreporter.com/business/industry-news/lionsgate-ai-deal-runway-1236055999/> [https://perma.cc/XAY5-VK42].

228. Ed Nawotka, *CCC Launches Collective Licensing for AI*, PUBLISHERS WKLY. (July 16, 2024), <https://www.publishersweekly.com/pw/by-topic/digital/copyright/article/95512-ccc-launches-collective-licensing-for-ai.html> [https://perma.cc/UFH8-X7W6].

In addition to fixed-rate deals for access to training materials, some companies are looking to “attribution” technologies to identify specific works relied upon by AI models in generating output. The attribution data can then be used to compensate rightsholders.²²⁹ The startup company Bria, for instance, has incorporated attribution technology into a text-to-image AI model that saves vectorized data whenever an image is generated, allowing Bria to “trace back and identify which images in the training set contributed the most to creating the new image.”²³⁰ Bria then pays royalties to licensors of those works.²³¹ ProRata, another AI startup, has created an algorithm that reportedly “can review an A.I.-generated output, identify the source of information based on novel facts and textual styles, and calculate how much each source contributed to the response” so owners of the source material can be compensated.²³²

The narrative promoted by AI companies and their defenders is that licensing content to train and develop AI systems is “impossible.”²³³ Yet AI companies have demonstrated that they are capable of entering into license arrangements when they see value in the licensed content. Moreover, it seems possible to design AI systems capable of tracking the works relied upon to produce particular output so relevant rightsholders can be paid. There is every reason to expect that the AI licensing market will continue to grow as licensing practices become more sophisticated and new actors enter the arena.²³⁴

229. Efrat Taig, *Bridging the Gap: From Academic AI to Ethical Business Models*, BRIA (Oct. 3, 2024, 5:08 PM),

<https://blog.bria.ai/mastering-custom-generative-image-model-training-insights-from-bria> [<https://perma.cc/BK6D-3CAA>] (explaining attribution model); Aaron Mok, *This Startup Has Built an Algorithm to Pay Creators for Their Work Used to Train A.I.*, OBSERVER (Sept. 3, 2024, 1:59 PM), <https://observer.com/2024/09/prorata-ai-revenue-sharing-creator-publisher/> [<https://perma.cc/3CHP-XPME>] (reporting that ProRata has created an algorithm to identify particular content used to generate AI output).

230. Taig, *supra* note 229.

231. *Id.*

232. Mok, *supra* note 229.

233. See Damle, *supra* note 2, at 3, 12–16 (“[E]veryone agrees that it is impossible for AI developers to negotiate and acquire licenses from every rightsholder who owns a[] copyright interest in the data used to train AI models.”); Lemley & Casey, *supra* note 7, at 770 (“[G]iven the large number of works an AI training data set needs to use and the fact that thousands, if not millions, of different people own those works, AI companies can’t simply license all the underlying photographs or text”); see also R STREET, *supra* note 76 (“A system that required follow-on creators to negotiate with and pay those they learned from would inhibit, rather than promote, the very artistic progress our IP laws seek to encourage.”).

234. Indeed, licensing of data to developers for general machine learning purposes is not a wholly new phenomenon. Pointing to a robust market for “immensely valuable” user data, Sobel asserted in 2017 that “there is already a thriving market for the data that fuel expressive machine learning.” Sobel, *Fair Use*, *supra* note 14, at 76.

AI advocates have been known to assert that the appropriation of copyrighted works to develop and operate AI models does not interfere with copyright owners' legitimate economic interests because the authors of books, movies, and music did not produce those works with the intent of populating generative models.²³⁵ While the observation concerning creators' expectations may be true as far as it goes, the claim is unconvincing. Authors and artists of just a few decades ago likely did not anticipate that their works would be accessed and consumed through mobile phones, watches or AirPods, but they no doubt expected their copyrights to continue to protect their creative works if those works were exploited by means yet to be known.

In evaluating market impact under the fourth fair use factor, it is critical to assess not only current modes of exploitation, but future markets for the work as well. As the US Supreme Court has emphasized, Section 107 of the Copyright Act requires consideration of the impact of the challenged use not just the existing market, but also the *potential* market for the copyrighted work, including the market derivative uses.²³⁶ This includes a nascent or still-developing market. In the pivotal case *American Geophysical Union v. Texaco Inc.*,²³⁷ for instance, the Second Circuit rejected Texaco's fair use defense to unlicensed copying of individual journal articles for internal research purposes because licenses for such copying had been made available by copyright owners.²³⁸ The court held that the failure to compensate publishers would result in "substantial harm to the value of their copyrights."²³⁹

As the Second Circuit has repeatedly affirmed, "[i]t is indisputable that, as a general matter, a copyright holder is entitled to demand a royalty for licensing others to use its copyrighted work, and that the impact on potential licensing revenues is a proper subject for consideration in assessing the fourth factor."²⁴⁰ The licensing market

235. See, e.g., Lemley & Casey, *supra* note 7, at 776 ("The copyright owner of a book or photograph doesn't create that work in hopes of selling it to AIs.").

236. *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 590 (1994) (citing 17 U.S.C. § 107(4)); *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. 539, 566-68 (1985) (discussing market harm); see also *Campbell*, 510 U.S. at 592 (explaining that potential market for derivative uses is one that "creators of original works would in general develop or license others to develop").

237. 60 F.3d 913 (2d Cir. 1995).

238. *Id.* at 930-31 ("Though the publishers still have not established a conventional market for the direct sale and distribution of individual articles, they have created, primarily through the C[opyright] C[learance] C[enter], a workable market for institutional users to obtain licenses for the right to produce their own copies of individual articles via photocopying.").

239. *Id.* at 931.

240. *Fox News Network v. TVEyes, Inc.*, 883 F.3d 169, 180 (2d Cir. 2018) (quoting *Bill Graham Archives v. Dorling Kindersley Ltd.*, 448 F.3d 605, 624 (2d Cir. 2006)); see also *Texaco*, 60 F.3d at 929.

for AI training materials is already far from hypothetical. The fact that AI companies are commercially motivated entities, many with significant economic resources, points to continued growth in this area.²⁴¹ The existence of a swiftly developing market for copyrighted content to build and operate AI systems weighs powerfully against a finding of fair use that could extinguish that market.

V. CONCLUSION

The Copyright Act protects works of human authors.²⁴² At this early stage, we do not yet know how generative AI will impact human authorship, or creative culture in general. Will there be less incentive for humans to create works, and for publishers to invest in and disseminate those works, because the human works are competing with AI-generated content? Conversely, will humans find it worthwhile to spend time engaging with AI systems to produce content that is not protected by copyright and can be freely exploited by others? If human authorship declines, will there be a corresponding decline in the appeal of AI-generated content as AI machines rely on the same materials over and over? Is the ability to produce an infinite number of texts, images or songs unconnected with a human artist actually of meaningful social value?²⁴³ Is an essential aspect of the human experience of and appreciation for art the fact that it is made by a human author?²⁴⁴

241. See Lemley & Casey, *supra* note 7, at 765 (“Commerciality often goes hand in hand with a market effect M[achine] L[earning] companies might be natural candidates for a licensing market: large for-profit companies that stand to benefit financially from using copyrighted works . . .”).

242. *Thaler v. Perlmutter*, 687 F. Supp. 3d 140, 147–48 (D.D.C. 2023) (upholding the Copyright Office’s refusal to register an AI-generated work on the ground that “authors” in the Copyright Act means human authors), *aff’d*, 130 F.4th 1039 (D.C. Cir. 2025). See also *id.* at 146 (although copyright “is designed to adapt with the times,” there has been “a consistent understanding that human creativity is the *sine qua non* at the core of copyrightability, even as that human creativity is channeled through new tools or into new media.”).

243. See, e.g., *UMG Compl.*, *supra* note 216, at ¶ 12 (alleging that AI defendant Udio generates 10 music files per second, or 6 million files per week, from copyrighted sound recordings).

244. Some more broadly question whether the enormous investment in generative AI will yield net social benefits. As encapsulated by MIT professor Daron Acemoglu in an AI-focused report by Goldman Sachs: “Technology that has the potential to provide good information can also provide bad information and be misused for nefarious purposes. I am not overly concerned about deepfakes at this point, but they are the tip of the iceberg in terms of how bad actors could misuse generative AI. And a trillion dollars of investment in deepfakes would add a trillion dollars to GDP, but I don’t think people would be happy about that or benefit from it.” Allison Nathan, *Gen AI: Too Much Spend, Too Little Benefit*, 129 GOLDMAN SACHS: TOP MIND 1, 5 (2025), https://www.goldmansachs.com/images/migrated/insights/pages/gs-research/gen-ai-too-much-spend%2C-too-little-benefit-/TOM_AI%202.0_ForRedaction.pdf [https://perma.cc/QQ2G-JBE9] (quoting MIT professor Daron Acemoglu); *id.* at 10 (“AI technology is exceptionally expensive, and to justify those costs, the technology must be able to solve complex problems, which it isn’t designed to do.” (quoting Jim Covello, head of Goldman Sachs global equity research)).

This Article explains why unauthorized copying by AI companies to build and operate generative AI systems is not, as claimed, “quintessential fair use.” There is no fair use precedent that legitimizes mass copying and exploitation of copyrighted works for their expressive value by for-profit entities. Apart from doctrinal concerns, an overly broad application of fair use to exempt unconstrained copying by AI companies could effect a potentially enormous transfer of value from the creators and owners of copyrighted works to the commercial entities that seek to exploit them. An unprecedented exception to copyright with such far-reaching consequences is not a question of fair use but rather a fundamental question of policy for Congress to decide.²⁴⁵

245. “The Congress shall have Power . . . To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.” U.S. CONST. art. I, § 8, cl. 8. In fact, Congress has already begun to consider the challenges presented by generative AI. *See, e.g., Artificial Intelligence and Intellectual Property: Part II – Identity in the Age of AI: Hearing Before the H. Comm. on the Judiciary, Subcomm. on Cts., Intell. Prop., and the Internet*, 118th Cong. (Feb. 2, 2024); *Oversight of A.I.: Legislating on Artificial Intelligence: Hearing Before the \S. Comm. on the Judiciary*, 118th Cong. (Sept. 12, 2023); *Artificial Intelligence and Intellectual Property: Part I–Interoperability of AI and Copyright Law: Hearing Before the H. Comm. on the Judiciary, Subcomm. on Cts., Intell. Prop., and the Internet*, 118th Cong. (May 17, 2023). In examining these issues, Congress has the benefit of advice from the US Copyright Office, which has undertaken a multipart study on copyright and artificial intelligence. *See* US Copyright Office, *Artificial Intelligence Study*, <https://www.copyright.gov/policy/artificial-intelligence/> [<https://perma.cc/5BTE-2TPN>] (last visited July 8, 2024). The Copyright Office issued the first installment of its study in July 2024, recommending legislation to address AI-generated “digital replicas,” or deepfakes, that imitate individuals’ images or voices. US Copyright Office, *Copyright and Artificial Intelligence: Part 1: Digital Replicas* 57 (2024), <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-1-Digital-Replicas-Report.pdf> [<https://perma.cc/P2FH-638B>] (“The Copyright Office agrees with the numerous commenters that have asserted an urgent need for new protection at the federal level.”) A bipartisan bill to protect against deepfakes was introduced by the Senate on July 31, 2024. Nurture Originals, Foster Arts, and Keep Entertainment Safe (NO FAKES) Act of 2024, S. 4875, 118th Cong. (2024). The second installment of the Copyright Office’s AI study, which addresses questions of copyrightability in relation to AI-generated works, was issued in January 2025. US Copyright Office, *Copyright and Artificial Intelligence: Part 2: Copyrightability* (2025), <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf> [<https://perma.cc/8GBL-7ZE4>].