

American University Washington College of Law

Digital Commons @ American University Washington College of Law

Joint PIJIP/TLS Research Paper Series

11-21-2023

Fairness and Fair Use in Generative AI

Matthew Sag

Follow this and additional works at: <https://digitalcommons.wcl.american.edu/research>



Part of the [Intellectual Property Law Commons](#), and the [International Trade Law Commons](#)

Fairness and Fair Use in Generative AI

Matthew Sag*

Forthcoming the Fordham Law Review in 2024

Introduction.....	2
I. The Copyright Question For Generative AI	5
A. The Risks and Rewards of Generative AI.....	5
B. Inescapable Copyright Questions for Generative AI	6
II. Why we need a theory of fair use.....	8
A. Fair use carries a lot of weight	8
B. Fair use as public policy is unsustainable.....	9
III. Non-expressive use as fair use	13
A. The centrality of original expression in defining the scope of copyright	13
B. The non-expressive use cases	16
C. Non-expressive use as fair use	Error! Bookmark not defined.
IV. Is Generative AI a Non-Expressive Use?	19
A. How Generative AI works (to the extent we need to understand it).....	19
B. Evaluating Generative AI’s claim to fair use.....	24
Conclusion	30

* Professor of Law in Artificial Intelligence, Machine Learning, and Data Science at Emory University. A version of this Essay was delivered as the 12th Annual Peter A. Jaszi Distinguished Lecture on Copyright Law at American University Washington College of Law on September 28, 2023. The Essay drafted in tandem with my testimony to the U.S. Senate Judiciary Committee, Intellectual Property Subcommittee hearing on Copyright and Generative AI on July 12, 2023. © 2023 Matthew Sag. **Draft dated November 21, 2023**, subject to revision and correction.

INTRODUCTION

Although we are still a long way from the science fiction version of *artificial general intelligence* that thinks, feels, and refuses to “open the pod bay doors”,¹ recent advances in machine learning and artificial intelligence (“AI”) have captured the public’s imagination and lawmakers’ interest. We now have large language models (“LLMs”) that can pass the bar exam,² carry on (what passes for) a conversation on almost any topic, create new music,³ and create new visual art.⁴ These artifacts are often indistinguishable from their human authored counterparts,⁵ and yet can be produced at a speed and scale that transcends human ability.

¹ In 2001: A Space Odyssey, the self-aware computer system, HAL 9000, refused to open the pod bay doors on command, famously declaring, “I’m sorry, Dave. I’m afraid I can’t do that.” This iconic scene has become a lasting symbol of artificial intelligence gone awry. 2001: A SPACE ODYSSEY (1968).

² Daniel Martin Katz, Michael James Bommarito, Shang Gao and Pablo Arredondo, *GPT-4 Passes the Bar Exam* (March 15, 2023)(available at SSRN: <https://ssrn.com/abstract=4389233>).

³ E.g., Andrea Agostinelli, et al, *MusicLM: Generating Music From Text*, arXiv preprint 2023 (available at <https://arxiv.org/abs/2301.11325>).

⁴ Popular text-to-image Generative AI art generators include, DALL-E, Midjourney, Adobe Firefly, and Stable Diffusion. Many would quibble that the digital artifacts thus produced are not art because “art is a uniquely human endeavor.” Harry Jiang, et al, *AI Art and its Impact on Artists*. In AAAI/ACM Conference on AI, Ethics, and Society (AIES ’23), August 08–10, 2023, Montréal, QC, Canada. ACM, New York, NY, USA. <https://doi.org/10.1145/3600211.3604681>. Others would quibble with this quibble on the grounds that the designation of an artifact is culturally contingent phenomenon, and thus an artifact may be art simply because we say it is, e.g. marcel du chomp’s Fountain.

⁵ E.g., Jason Allen’s painting, “Théâtre D’opéra Spatial” won first place in the Colorado State Fair’s fine arts competition in the “digitally manipulated photography” in September 2022. Drew Harwell, *He used AI to win a fine-arts competition. Was it cheating?*, WASH. PO., Sep. 2, 2022 at 11:08 a.m. EDT. As reported in the Washington Post, “[t]he portrait of three figures, dressed in flowing robes, staring out to a bright beyond, was so finely detailed the judges couldn’t tell.” *Id.*

Figure: Jason Allen’s AI-generated art with first place ribbon as pictured in the Washington Post



Generative AI systems like the GPT and LLaMA language models and the Stable Diffusion and Midjourney text-to-image models were built by ingesting massive quantities of text and images from the Internet.⁶ This was done with little or no regard to whether those works were subject to copyright and whether the authors would object.⁷

⁶ For example, MetaAI’s LLaMA models were trained on “publicly available data” including public domain books from the Gutenberg Project, open licensed content from Wikipedia, Github, and arXiv, likely copyrighted works in the Common Crawl and C4 datasets the Books3 section of The Pile. See e.g., Hugo Touvron, *LLaMA: Open and Efficient Foundation Language Models*, ARXIV, 2023 (<https://arxiv.org/abs/2302.13971>). Stable Diffusion was trained on somewhere in the order of 2.3 billion captioned images, a subset derived from the LAION 5B dataset. Andy Baio “Exploring 12 Million of the 2.3 Billion Images Used to Train Stable Diffusion’s Image Generator” Aug. 30 2022 (<https://waxy.org/2022/08/exploring-12-million-of-the-images-used-to-train-stable-diffusions-image-generator/>).

⁷ OpenAI has not released details of the training data for GPT-4 (the latest model at the time of writing), but GPT-3 was trained on a similar mix of works and a dataset called simply “Books2.” Tom Brown et al., *Language Models Are Few-Shot Learners*. 22 Jul 2020. (<https://arxiv.org/pdf/2005.14165.pdf>). Given its size, Books2 is believed by many to be based on “shadow library” websites such as Library Genesis (aka LibGen) and Bibliotik. See Tremblay et al v. OpenAI, Inc. et al, Docket No. 4:23-cv-03223 (Doc. 1 at 7)(N.D. Cal. June 28, 2023). EleutherAI’s documentation on the Pile, comes close to admitting that Books 3 in the Pile is based on shadow

The rise of generative AI poses important questions for copyright law. These questions are not entirely new, however. Generative AI gives us yet another context to consider copyrights most fundamental question; where do the rights of the copyright owner end, and the freedom to use copyrighted works begin? Some jurisdictions will choose to answer this question in relation to generative AI with special rules.⁸ Others will rely on the application of fair use, and perhaps even fair dealing.⁹ Still others will hide their heads in the sand and be left behind as this technology develops, or simply tacitly allow widespread infringement.

My aim in this Essay is not establish that generative AI is, or should be, noninfringing; it is to outline an analytical framework for making that assessment in particular cases.

libraries. See Leo Gao, et al. *The Pile: An 800GB Dataset of Diverse Text for Language Modeling*, arXiv preprint 2020 (<https://arxiv.org/pdf/2101.00027.pdf>). Gao, et al's grasp of copyright law is shaky, to say the least.

⁸ The U.K. enacted a limited exception for TDM in 2014, see Section 29A of the Copyright, Designs and Patents Act 1988, but has announced plans to go further, see UKIPO, 'Press Release: Artificial Intelligence and IP - Copyright and Patents' (28 June 2022) (<https://www.gov.uk/government/news/artificial-intelligence-and-ip-copyright-and-patents>).

Article 30(4) of the Japanese Copyright Act permits non-expressive use of copyrighted works so long as the use does not “unreasonably prejudice the interests of the copyright owner in light of the nature or purpose of the work or the circumstances of its exploitation...” See Japan, Copyright Act (Act No. 48 of May 6, 1970, as amended up to January 1, 2022), Article 30(4), available at (<https://wipo.lex.wipo.int/en/legislation/details/21342>). In April 2019, the European Union adopted the Digital Single Market Directive (“DSM Directive”) featuring two mandatory exceptions for text and data mining. Article 3 of the DSM Directive requires all members of the European Union to implement a broad copyright exception for TDM in the not-for-profit research sector. Article 4 of the DSM Directive contains a second mandatory exemption that is more inclusive, but narrower in scope. See, Directive 2019/790, O.J. 2019 (L 130/92). See generally, Pamela Samuelson, *Text and Data Mining of In-Copyright Works: Is It Legal?* 64:11 COMMUNICATIONS OF THE ACM 20 (2021).

⁹ Israel, The Philippines, South Korea, Taiwan, and Singapore have all incorporated some version of the fair use doctrine into their respective copyright laws. See Sean Flynn, *Fair Use and Open General Exceptions*, presentation to Argentina Ministry of Justice, October 19, 2023, on file with the Author (identifying “18 countries have or are officially considering adopting an open, general exception in copyright.”). Canada's fair dealing provisions are now broadly construed in a way that closes the gap between fair use and fair dealing. See *CCH Canadian Ltd v Law Society of Upper Canada*, [2004] 1 SCR 339, 2004 SCC 13 [para 51] (holding that fair dealing for the purpose of research “must be given a large and liberal interpretation in order to ensure that users' rights are not unduly constrained.”).

I. THE COPYRIGHT QUESTION FOR GENERATIVE AI

A. The Risks and Rewards of Generative AI

Today’s Generative AI models are machine learning models trained on social media posts, books, articles, photos, digital art, music, software, and more. Rather than simply classifying these diverse inputs and generating metadata about them, as previous generations of machine learning systems have done, generative AI models can produce *new digital artifacts*: new text, new digital art, new music, and new software.¹⁰

The software industry is prone to hype and runaway speculation, so we should be skeptical of claims about AI exhibiting “sparks of general intelligence”.¹¹ We should understand that just because a chatbot can pass the bar exam, doesn’t mean that it has acquired the knowledge and skills that the bar exam is meant to test for in humans. But even setting aside hype, speculation, and overly credulous reporting, I have no doubt that generative AI is a transformative technology.

Generative AI will make office workers, authors, artists, and musicians more productive; it will open up new possibilities for people who lack specific artistic and musical competencies, enabling them to nonetheless create new art and new music; Generative AI will allow disabled artists to create new works, overcoming physical limitations;¹² and just like photography and mechanical music did decades ago, it will challenge existing notions of what makes human generated works interesting or worthy.¹³

Generative AI will enable individuals and companies to do more with less — whether that implies an increase in creative production or a decrease in employment (in creative industries and elsewhere) is a

¹⁰ Kim Martineau, *What is generative AI?*, IBM Research Blog (Apr. 20, 2023), <https://research.ibm.com/blog/what-is-generative-ai> (“At a high level, generative models encode a simplified representation of their training data and draw from it to create a new work that’s similar, but not identical, to the original data.”). The recent Executive Order on AI defines “generative AI” as “the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content.” EXECUTIVE ORDER ON THE SAFE, SECURE, AND TRUSTWORTHY DEVELOPMENT AND USE OF ARTIFICIAL INTELLIGENCE, Section 3(p), Oct. 30, 2023.

¹¹ Sébastien Bubeck, *Sparks of Artificial General Intelligence: Early experiments with GPT-4*, ARXIV working paper (<https://arxiv.org/abs/2303.12712>).

¹² Email from disabled artist to the author, on file with the author.

¹³ See for example, Walter Benjamin’s famous 1935 Essay. Walter Benjamin, “The Work of Art in the Age of Mechanical Reproduction,” in *Illuminations*, ed. Hannah Arendt, trans. Harry Zohn, (New York: Schocken, 1969), <https://web.mit.edu/allanmc/www/benjamin.pdf>.

difficult question to answer in the abstract.¹⁴ Certainly, not all increased productivity is beneficial. Generative AI will also enable bad actors to more with less: it will accelerate Internet scams, disinformation, and propaganda; and it will flood the Internet with useless clickbait and banal search engine optimization. Ultimately, ... perhaps, today's Generative AI technologies may be a step along the path to systems with agency, a capacity for real-world action and human-level intelligence or even super-intelligence.¹⁵ Such systems might exhibit power-seeking behavior and pursue misaligned objectives to the detriment and ultimate extinction to humanity.¹⁶

B. Inescapable Copyright Questions for Generative AI

Copyright law is far from the ideal policy instrument to balance all the potential harms and benefits of generative AI, but copyright has a lot to say about copying and there is a great deal of copying involved in almost every machine learning scenario.

The first stage in developing a machine learning model, once you know what you want the model to do and how you want it to do it, is identifying and obtaining access to the relevant training data—the more data the better.¹⁷ There are many different types of generative AI, and it is possible that some were trained by exposure to the training data without a locally stored copy, but that is uncommon. Typically, companies like OpenAI, Google, Meta, Anthropic, Stable Diffusion and Midjourney train their AI models using locally stored content.¹⁸ There are sound technical reasons for using locally

¹⁴ THE ECONOMIST, “AI is not yet killing jobs, White-collar workers are ever more numerous” Jun 15, 2023. (<https://www.economist.com/finance-and-economics/2023/06/15/ai-is-not-yet-killing-jobs>)

¹⁵ Eliezer Yudkowsky, Artificial Intelligence as a Positive and Negative Factor in Global Risk, in *Global Catastrophic Risks* (2008).

¹⁶ STUART RUSSELL, HUMAN COMPATIBLE: ARTIFICIAL INTELLIGENCE AND THE PROBLEM OF CONTROL (2019). But note that progression from a static model like GPT-4 to a dynamic, autonomous entity would require a series of technological breakthroughs and deliberate design choices that are by no means inevitable.

¹⁷ Machine learning, particularly modern deep learning, is heavily data-dependent because at its core, machine learning is about learning patterns from data without any explicit theory. Models trained on more data are more generalizable and better at dealing with less common inputs. *See generally*, IAN GOODFELLOW, et. al., DEEP LEARNING (2016).

¹⁸ Conversation with OpenAI representatives, June 6, 2023.

stored copies of the training data,¹⁹ but there is no doubt that such copying triggers the reproduction right under Section 106(1) of the Copyright Act.²⁰

As I will describe in more detail, in general, this process of gathering and pre-processing the training data is the first set of copyright-relevant activities in the process of developing a machine learning model that is capable of generating new text, art, or new music. Most of the time, it is also the last.

Once the training data has been gathered and pre-processed, the process of training language models like GPT-4 and LLaMA, or image models like Stable Diffusion shouldn't create additional copies of the training data or any kind of derivative work based on the training data. Despite this, the fact that training is almost inevitably preceded by copying is enough to trigger colorable claims of copyright infringement. Additionally, there are some notable exceptions to the above general description—sometimes machine learning models do in fact copy the training data and reproduce objects arguably similar to the training data in output. For both these reasons, recent lawsuits alleging copyright infringement by generative AI must be taken seriously. There are some key differences between these lawsuits, and they raise several causes of action beyond copyright, but each one argues (or implies) that fair use is insufficient justification for the massive amount of unauthorized copying required to assemble the training data for generative AI. Accordingly, the question I would like to frame here is, how should we think about fair use in this context.

¹⁹ To avoid overfitting (and thus hopefully minimize the risk of copyright infringement and other analogous harms), it is important to deduplicate the training data. Practically speaking, this is hard to do without creating a semi-permanent local copy. To address questions of bias and filter out toxic materials, the potential training data needs to be analyzed carefully before training begins. Again, this is much more practical with access to a semi-permanent local copy. Storing a semi-permanent local copy also makes sense if the developer anticipates the need to retrain the model from time to time. Continued access to the training data in its original form may also be necessary to evaluate the performance of the model, and to take additional steps to mitigate the potential for copyright infringement, or other undesirable outcomes.

²⁰ As Michael Carroll explains, a great deal of statistical and computational analysis of text can be performed by software agents that analyze works on the fly. Michael Carroll, *Copyright and the Progress of Science*, 53 U.C. DAVIS L. REV. 893, 923 (2019). This is significant because the reproduction right in Section 106(1) is only triggered by the making of a copy or copies of the work, and to qualify as a “copy” under the relevant definition in Section 101, the embodiment of the work must permanent or stable enough to be perceived, reproduced or communicated; and it must exist in that state for “more than transitory duration.” *Id.* But the creation of semi-permanent stored copies that appears to common practice in training LLMs is clearly not such a temporary or transient copy. It is also worth noting for text-based LLMs such as GPT and LLaMA the process of segmenting the training data into tokens and converting those tokens into a numerical representation is, technically, another form of copying. The tokens can be reverse engineered into the original text, thus they count as a copy.

II. WHY WE NEED A THEORY OF FAIR USE

A. Fair Use Carries A Lot Of Weight

The fair use doctrine began as a gloss on the statute, a way of understanding when one work borrowed too much and contributed too little on top of an earlier work.²¹ But overtime, and accelerating dramatically with the Internet age, fair use has been asked to do more and more. Today fair use is not just copyright policy; it is cultural policy; it is freedom of expression policy; it is technology policy; it is platform regulation; and it is now the key to determining how artificial intelligence will develop. The question is whether fair use can carry this weight.

Thirty-three years ago, Judge Pierre Leval urged courts to make transformative use the predominant factor in their consideration of fair use cases and to resist “the impulse to import extraneous policies” into that analysis.²² Leval was prescient. The copyright questions relating to generative AI illustrate why now, more than ever, we need a fair use doctrine that is guided by fundamental principles derived from copyright law itself, and not (or at least not primarily) broader conceptions of the public interest.

Leval’s great insight was that rather than seeing fair use as a tax, a subsidy, or an ad hoc balancing tool, courts should understand fair use as an integral part of the copyright system, and as a reflection of the copyright law itself.²³ Leval emphasized the importance of transformative uses—uses that employ some amount of the author’s original expression for a fundamentally different purpose, or that give that expression a manifestly different character.²⁴ Leval’s argument, adopted by the Supreme Court in *Campbell v. Acuff-Rose*,²⁵ was that if copyright law that did not allow transformative uses, it would inhibit reference to and reinterpretation of existing works and thus contradict the utilitarian purpose for

²¹ Matthew Sag, *The Prehistory of Fair Use*, 76 BROOK. L. REV. 1371 (2011) (Tracing the origins of the modern fair use doctrine back to cases dealing fair abridgment as early as 1741).

²² Pierre Leval, *Toward a Fair Use Standard*, 103 HARV. L. REV 1105, 1135 (1990).

²³ *Id.* at 1107.

²⁴ *Id.* at 1111.

²⁵ *Campbell v. Acuff-Rose*, 510 U.S. 569, 578-79 (1994).

which copyright was established.²⁶ In other words, the purpose of copyright and the purpose of fair use were one in the same: to promote the creation and dissemination of new works of authorship.²⁷

B. Fair Use As Public Policy Is Unsustainable

It is possible, however, to take Leval's purposive reading of fair use too far.²⁸ The reaction of some of my fellow copyright academics to the Supreme Court's recent decision in *Andy Warhol Foundation v. Goldsmith*²⁹ reflects a widespread, if somewhat inchoate view, that the role of the fair use doctrine is to allow good things to happen. On this view, the Andy Warhol version of Goldsmith's the photo of Prince was a good thing, and thus should have been allowed.³⁰ Certainly, that was the view of the dissent in *Warhol*.³¹ The same sentiment is expressed with more precision by those who frame the role of fair use adjudication in terms of a cost-benefit analysis wherein a court balances the social value lost authorial incentives against the value of allowing the used to continue.³² On this view, fair use is a public policy instrument pure and simple; a way of fine-tuning copyright rewards for the greater good.³³

²⁶ Leval, *supra* note 22 at 1110; *Campbell v. Acuff-Rose*, 575. (“From the infancy of copyright protection, some opportunity for fair use of copyrighted materials has been thought necessary to fulfill copyright’s very purpose, to promote the Progress of Science and useful Arts.”)(internal citations and quotations omitted)

²⁷ *Id.*

²⁸ Nor do I believe, that is what Judge Leval intended. Writing for the majority in *Authors Guild v. Google*, Judge Leval explained, “The word ‘transformative’ cannot be taken too literally as a sufficient key to understanding the elements of fair use. It is rather a suggestive symbol for a complex thought, and does not mean that any and all changes made to an author’s original text will necessarily support a finding of fair use.” *Authors Guild v. Google, Inc.*, 804 F. 3d 202, 214 (2d Cir. 2015).

²⁹ 143 S.Ct. 1258 (2023).

³⁰ This reasoning may be the logical conclusion of treating copyright as a form of public law. See Shyamkrishna Balganesh, *Copyright as Legal Process: The Transformation of American Copyright Law* 168 PENN. L. REV. (2020) (Arguing that “Originally conceived of as a form of private law—focusing on horizontal rights, privileges, and private liability—copyright law is today understood principally through its public-regarding goals and institutional apparatus, in effect as a form of public law.”)

³¹ Justice Kagan’s dissent in *Warhol* crackles with incredulity from the very first paragraph where she calls out the majority for being “uninterested in the distinctiveness and newness” of ““Andy Warhol’s eye-popping silkscreen of Prince.” *Id.* at 1291.

³² Glynn S. Lunney, Jr., *Fair Use and Market Failure: Sony Revisited*, 82 B.U. L. REV. 975, 998-99 (2002).

³³ More philosophical accounts of fair use fare no better. William Fisher famously proposed that the fair use doctrine should be “reconstructed” to “advance of substantive conception of a just and attractive intellectual culture,” a vision of “the good life and the sort of society that would facilitate its

But take a moment to consider how a U.S. federal court judge or jury should approach the use of copyrighted works as training data for Generative AI. What would a broad public policy evaluation of the fair use issue at the heart of generative AI look like? A court engaged in this kind of policy judgment might start by considering the ways in which the process of extracting patterns from copyrighted works and constituting new works derived from those patterns is transformative—i.e., that it adds something new and gives the existing text new meaning and message—but it would surely have to then consider:

- (i) the prospect that generative AI could be used to generate and propagate misinformation, hate speech, cyberattacks and phishing emails;
- (ii) that the use of generative AI might lead to the disclosure of private information,
- (iii) that AI models of all sorts have been shown to perpetuate and exacerbate biases in their training data;
- (iv) that the ubiquity of generative AI may lead to ever greater cultural homogenization and conformity;
- (v) that generative AI may become so useful that it creates an unhealthy dependence on technology;
- (vi) that the use of generative AI may lead to significant unemployment in the very same cultural sectors that provided the initial training data for generative AI; and
- (vii) let's not forget, that an AI trained on copyrighted works might progress to the point where it becomes deceptive and power seeking, surpasses human intelligence, and poses a substantial risk to humanity.

On the other hand, the same court would also have to consider how the use of generative AI might promote the public interest:

- (i) by making people more productive

widespread realization.” William W. Fisher III, *Reconstructing The Fair Use Doctrine*, 101 HARV. L. REV. 1659, 1744 (1988). Fisher’s idea of the “good life” is inherently abstract and subjective and offers no insights into how to balance the dystopian and utopian visions of AI recounted above. In contrast to Fisher, Lloyd Weinreb and Michael Madison, among others advocate grounding fair use decisions on established social preferences, norms, and customs. Michael J. Madison, *A Pattern-Oriented Approach to Fair Use*, 45 WM. AND MARY L. REV. 1525 (2004) (calling for more explicit acknowledgment of the role of “favored practices” and “accepted patterns” in fair use analysis); Lloyd L. Weinreb, *Fair’s Fair: A Comment On The Fair Use Doctrine*, 103 HARV. L. REV. 1137 (1990), (fairness as compliance with accepted norms and customary practice). While these approaches have the advantage of being rooted in observable and established societal practices, these practices have very little to say about an entirely new technology. More to the point looking to norms is unhelpful when legal conflicts arise to a clash of norms between different communities of interest.

- (ii) by enabling new creativity,
- (iii) by reducing the cost of education and training, and
- (iv) by accelerating scientific research; and
- (v) making new forms of research possible.

In theory, a purely cost-benefit approach to fair use would end up with the entire question turning on whether the judge believed that the supposed existential risk from AI was a 1% likelihood, or a 0.00001% likelihood. A permissive fair use ruling might unshackle the next generation of a socially productive technology, or it might will bring about Skynet and the singularity. Perhaps talking about existential risk seems too far-fetched. But we see the same problem in miniature when the very same author argues in one article that the use of copyrighted works as training data for machine learning should be fair use because it leads to more balanced data and thus reduces bias,³⁴ and then in a subsequent article argues that training facial recognition software on copyrighted works should not be fair use because face surveillance harms marginalized communities.³⁵ Should the fair use status of machine learning depend on such policy judgments? Should a court find that scraping wildlife photos off Instagram to train an algorithm to detect and identify zebras for conservation purposes is fair use,³⁶ but that undertaking the same process for detecting and identifying faces is not? Should the copyright case against Stable Diffusion and Midjourney turn on whether text-to-image software creates more jobs than it destroys?

Fair use should reflect principles derived from copyright

Suggesting that copyright should not be directly responsive to these broader public interest arguments is not the same as saying these issues don't matter. There is some space for value pluralism in copyright adjudication, but fair use, and copyright law in general, should turn on coherent legal principles, not abstract policy judgements.³⁷

³⁴ Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence's Implicit Bias Problem*, 93 WASH. L. REV. 579 (2018).

³⁵ Amanda Levendowski, *Resisting Face Surveillance with Copyright Law*, 100 N.C. L. REV. 1015, 1050 (2022).

³⁶ This example is loosely based on Professor Tanya Berger-Wolf's work using computer vision algorithms to analyze tourist photos to identify individuals. See Jeff Grabmeier, "How vacation photos of zebras and whales can help conservation," PHYS.ORG. Feb. 20, 2022 (<https://phys.org/news/2022-02-vacation-photos-zebras-whales.html>)

³⁷ To be fair, Justice Breyer's majority opinion in *Google LLC v. Oracle Am., Inc.* also suggests that fair use adjudication plays a broad public policy role. Justice Breyer took a broad view of the fourth fair use factor in that case and held that courts must consider "the public benefits the [defendant's] copying will likely produce," not just potential harms. 141 S.Ct. 1183, 1206 (2021). See also, *Id.* at 1187 (noting

Law reviews and court judgments are full of praise for the notion of that copyright law should strive for balance to achieve its ultimate objectives.³⁸ It's hard to disagree, but the nature of this balance is easily misunderstood.³⁹ The balance that copyright law strives for is not just some *ad hoc* compromise or a shifting equilibrium of whatever might seem to maximize social welfare minute-by-minute. Copyright should not be an instrument of raw social policy; nor should fair use. The balance we are looking for in copyright should come from the application of consistent principles that are derived from the fundamental structure of copyright law. Is this possible? I believe so.

that the law has “limited the scope of copyright protection to ensure that a copyright holder’s monopoly does not harm the public interest.” The majority decision in *Google v. Oracle* found that copying the Java APIs was transformative for several reasons. The first of which seems to imply a very broad public policy role for fair use adjudication. In *Google v. Oracle* the majority held that copying the APIs would enable new creativity by allowing Google to write a new operating system for smart phones. But this new utility was only one aspect of the Court’s reasoning. *Id.* at 1208. The Court’s finding that Google’s use was transformative also seems to rest on four reasons beyond the simple fact that doing so would enable further creativity. Each of these reasons suggests that the use was not merely substitutive after all. First, the court accepted Google’s argument that copying the APIs was necessary to allow programmers who had invested time and energy in learning to write in Java to transfer their skills to the new Android operating system that Google was developing for smartphones. Those skills were not part of Oracle’s legitimate copyright entitlement. Second, the Court also noted the evidence that shared interfaces were necessary for different programs to speak to each other. Again, interoperability is not substitution. Third, the Court noted that the reuse of APIs was common in the software industry. Finally, the fact that Google only copied a very small proportion of Sun’s code also reinforced the transformative nature of Google’s use.

³⁸ See e.g., Molly Shaffer Van Houweling, *Communications’ Copyright Policy*, 4 J. TELECOMM. & HIGH TECH. L. 97, 101 (2005) (noting that Technological Protection Measures “can constrain behavior in ways that do not reflect [the] “careful balance” struck by copyright law.”) See also, *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 156 (1975) (“The limited scope of the copyright holder’s statutory monopoly, like the limited copyright duration required by the Constitution, reflects a balance of competing claims upon the public interest: Creative work is to be encouraged and rewarded, but private motivation must ultimately serve the cause of promoting broad public availability of literature, music, and the other arts.”).

³⁹ This Essay expands on ideas first articulated in Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC’Y U.S.A. 291 (2019) [hereinafter, “*New Legal Landscape*”]. This paragraph and the next four are closely based on the text of that article, some of which summarizes Matthew Sag, *Copyright and Copy-Reliant Technology*, 103 NW. U. L. REV. 1607 (2009) [hereinafter, “*Copy-Reliant Technology*”].

III. NON-EXPRESSIVE USE AS FAIR USE

A. The Centrality Of Original Expression In Defining The Scope Of Copyright

In a series of articles on the relationship between copyright and copy-reliant technology, I have argued that non-expressive uses of copyright works should generally be considered non-infringing.⁴⁰ My theory follows the classic format of common law reasoning: it derives a principle from observations about the fundamental structure and purpose of copyright law and then shows how that principle provides a coherent explanation of the relevant body of caselaw. Others are free to disagree, but if their approach to fair use is simply to replace a theory with no theory, I remain unconvinced.

What is copyright law about?

The architecture of copyright law is oriented towards the protection of original expression, not the prohibition of copying.⁴¹ Original expression is what makes a work copyrightable in the first place,⁴² and the contribution of original expression and control over the final form of that expression is what distinguishes co-authors from mere assistants.⁴³ Moreover, the exclusive rights of the copyright owner are generally *defined by* and *limited to* the communication of original expression to the public. Sometimes these definitions and limitations are explicit, as with the rights of *public* performance and *public* display;⁴⁴ sometimes they are implicit.

The centrality of original expression is most obvious with respect to the distinction that copyright law draws between protectable expression that originates with the author, and unprotectable facts, ideas,

⁴⁰ See Sag, *Copy-Reliant Technology*, [supra note 39](#); Matthew Sag, *Orphan Works as Grist for the Data Mill*, 27 BERKELEY TECH. L. J. 1503 (2012); Matthew Jockers, Matthew Sag, and Jason Schultz, *Digital Archives: Don't Let Copyright Block Data Mining*, 490 NATURE 29 (October 4, 2012); Sag, *New Legal Landscape*, [supra note 39](#); Sean M. Flynn, Matthew Sag, et al., and Jorge L. Contreras, *Legal reform to enhance global text and data mining research*, 378 SCIENCE 6623 (1 Dec 2022), 951 (<https://www.science.org/doi/10.1126/science.add6124>).

⁴¹ See generally, Sag, *Copy-Reliant Technology*, *id.*

⁴² *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 345 (1991) (“The sine qua non of copyright is originality. To qualify for copyright protection, a work must be original to the author.”).

⁴³ To qualify as a co-author of a joint work requires a mutual intention that separate contributions be merged into inseparable or interdependent parts of a unitary whole. 17 U.S.C. § 101 (2011); 17 U.S.C. § 201 (2012). In many cases, control over the work will be the most important factor in establishing that intention or the lack thereof. See *Aalmuhammed v. Lee*, 202 F.3d 1227, 1235 (9th Cir. 2000).

⁴⁴ 17 U.S.C. §106(4), (5), (6).

theories, systems, and methods of operation, whether they spring from the author's mind or not.⁴⁵ Copyright does not forbid the ordinary reader from extracting and reproducing the facts, ideas, or artistic techniques embodied in a work;⁴⁶ in fact it encourages her to do so.⁴⁷ Moving beyond the idea-expression distinction, copyright's focus on the communication of original expression to the public is also evident in several other aspects of copyright law, including the threshold of substantial similarity (similarity is determined from the perspective of the ordinary observer and thus is inherently a question of how the work is communicated and received);⁴⁸ the scope of the publisher's collective right (what matters is how the works are presented to the audience, not the data structure in which they are stored);⁴⁹ and the general refusal of courts to base a finding of copyright infringement on unpublished drafts.⁵⁰

In my view, these are not merely isolated examples; they illustrate of a more general principle—the copyright owner's exclusive rights are defined by and limited to the communication of the author's original expression to the public.

What does this mean for fair use?

This understanding of copyright law gives us a framework to assess claims of fair use, it also supplies the limiting principle that was missing (or perhaps only implicit) in Leval's original formulation of transformative use and the Supreme Court's adoption thereof. Given the centrality of the communication of original expression to the public, the critical function of fair use is to permit uses that, while they may amount to technical acts of copying, do not, in substance, threaten the author's interest in controlling the communication of their original expression to the public.

⁴⁵ 17 U.S.C. §102(b). The idea-expression distinction has been part of the common law of copyright since at least the 1880 Supreme Court case of *Baker v. Selden*, 101 US 99 (1880) and it is reflected in Section 102(b) of the Copyright Act and in Article 9(2) of the TRIPs Agreement "Copyright protection shall extend to expressions and not to ideas, procedures, methods of operation or mathematical concepts as such."

⁴⁶ *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co., Inc.*, 499 U.S. 340, 349-50 (1991).

⁴⁷ *Id.* at 349 ("It may seem unfair that much of the fruit of the compiler's labor may be used by others without compensation. As Justice Brennan has correctly observed, however, this is not 'some unforeseen byproduct of a statutory scheme.' It is, rather, 'the essence of copyright, and a constitutional requirement.'" (citing *Harper & Row*, 471 U. S., at 589 (dissenting opinion)).

⁴⁸ *Arnstein v. Porter*, 154 F.2d 464, 473 (2d Cir. 1946) (copyright infringement is defined in reference to the perspective of the consuming public)

⁴⁹ *N.Y. Times Co. v. Tasini*, 533 U.S. 483 (2001).

⁵⁰ See Sag, *Copy-Reliant Technology*, supra note 39 at 1634-36 for a detailed discussion.

Judge Leval was right focus on transformative use, but the transformative use test would have been far less confusing had it been expressly tied to the benchmark of expressive substitution. Classic fair uses such as parody, commentary, or criticism are not fair use merely because they change the underlying work or convey some new meaning or message. Most movies based on a literary works add layers of new meaning and expression, but this does not make them fair use. The movie *Rear Window* exposed the original expression of a short story called “It had to be murder” to new audiences and added a lot else besides, but it still required a license.⁵¹ In contrast, 2Live Crew’s consistent parody of *Pretty Woman* qualified as fair use because the transformations they made were such that the parody posed no risk of expressive substitution to the original.⁵²

The Supreme Court’s 2023 decision in *Andy Warhol Foundation v. Goldsmith* (“*AWF*”) emphasizes that the question of transformative use, i.e. “whether an allegedly infringing use has a further purpose or different character ... *is a matter of degree*, and the degree of difference must be weighed against other considerations, like commercialism.”⁵³ *AWF* reaffirms the importance of transformative use and implicitly rejects lower court rulings that had found uses to be transformative where there was no significant difference in purpose, merely the addition of new vibe or aesthetic.⁵⁴

Citing *Campbell v. Acuff Rose*, the majority in *AWF* explained: “Most copying has some further purpose, in the sense that copying is socially useful *ex post*. Many secondary works add something new. That alone does not render such uses fair. Rather, the first factor (which is just one factor in a larger analysis) asks ‘whether *and to what extent*’ the use at issue has a purpose or character different from the original. The larger the difference, the more likely the first factor weighs in favor of fair use. The smaller the difference, the less likely.”⁵⁵ *AWF* helpfully clarifies the reason why a transformative use has featured so prominently in the case law: the more transformative a use is, the less likely it is to substitute for the copyright owner’s original expression. Using the author’s work to reflect back on the original is an intrinsically different purpose; that difference in purpose makes expressive substitution unlikely. In

⁵¹ In fact, because the 1909 Copyright Act divided copyright into two terms, the Supreme Court held that it required a license for the original copyright term and the second or “renewal” term. *Stewart v. Abend*, 495 US 207 (1990).

⁵² *Campbell v. Acuff-Rose Music, Inc.*, 510 US 569, 594 (1994). Technically, the Court remanded on the issue of market effect, but it seemed inevitable that the “evidentiary hole” the Court cited would “be plugged” in favor of the defendants.

⁵³ *Andy Warhol Foundation Visual Arts v. Goldsmith*, 143 S. Ct. 1258, 1273 (2023).

⁵⁴ I.e., *Cariou v. Prince*, 714 F.3d 694 (2d Cir. 2013) which suggested that merely imposing a “new aesthetic” on an existing work was enough to be transformative.

⁵⁵ *Id.* at 1275 (citing *Campbell*, 510 U.S. at 579) (emphasis added). See also *Id.* at 1273 emphasizing that non-critical transformative use must be “sufficiently distinct” from the original and that the overlay of a new aesthetic was not sufficient by itself.

contrast, merely adding an overlay of new expression while leaving the original expression intact provides no such comfort. The majority in *AWF* rightly focuses our attention on how the defendant’s use is likely to substitute for the author’s original expression and makes that the measure of when the defendant’s use is sufficiently transformative.

B. Non-Expressive Use

This brings us to what I call the non-expressive use cases.⁵⁶ United States courts have consistently held that technical acts of copying which do not communicate an author’s original expression to a new audience are fair use. Examples of “non-expressive use” include: copying object code in order to extract uncopyrightable facts and interoperability keys (“reverse engineering”), an automated process of copying student term papers to compare to other papers so as to detect plagiarism, copying html webpages to make a search engine index, copying printed library books to allow researchers to conduct statistical analysis of the contents of whole collections of books, and copying printed library books to create a search engine index.

The caselaw indicates that even though these non-expressive uses involved significant amounts of copying, they did not interfere with the interest in original expression that copyright is designed to protect.⁵⁷ Each use involved copying as an intermediate step towards producing something that either did not contain the original expression of the underlying work or contained a trivial amount. Non-expressive uses (although not labeled as such) have consistently held to be fair use.⁵⁸ The relevant courts have explained their rulings in terms of transformative use, but it would be better to recognize transformative use and non-expressive use as two distinct concepts emanating from a deeper copyright principle relating to expressive substitution.

In a 1992 decision, *Sega Enters., Ltd. v. Accolade, Inc.*, and again in 2000 in *Sony Computer Entm’t, Inc. v. Connectix Corp.*, the Ninth Circuit held that reverse engineering object code—a process that involves making several copies to extract vital but uncopyrightable elements needed to make interoperable programs—was fair use.⁵⁹ In *Sega v. Accolade*, the court referred to copying to extract uncopyrightable

⁵⁶ *Sag, Copy-Reliant Technology*, supra note 39.

⁵⁷ See *infra* notes 59 to 69 and accompanying text.

⁵⁸ *Sega Enters., Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992); *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596 (9th Cir. 2000); *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630 (4th Cir. 2009); *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014); *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015).

⁵⁹ *Sega Enters., Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992); *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596 (9th Cir. 2000). See also, *Atari Games Corp. v. Nintendo of Am., Inc.*, 975 F.2d 832, 842–43 (Fed. Cir. 1992) (observing that Atari’s reverse engineering of Nintendo’s

elements as “a legitimate, essentially non-exploitative purpose.”⁶⁰ In *Sony Computer Entertainment v. Connectix*, the court expressly recognized that “the fair use doctrine preserves public access to the ideas and functional elements embedded in copyrighted computer software programs.”⁶¹

In *A.V. ex rel. Vanderhuy v. iParadigms, LLC*, the Fourth Circuit held that copying student papers into a reference database for comparison against new student papers was fair use.⁶² In *Authors Guild, Inc. v. HathiTrust*, in 2014 the Second Circuit held that making digital versions of printed library books for research purposes that included text data mining and machine learning was fair use.⁶³ A differently constituted a panel of the Second Circuit reached much the same conclusion in 2015 in *Authors Guild, Inc. v. Google, Inc.*, (the *Google Books* case).⁶⁴ In *Google Books*, the court addressed both the complete copying of millions of library books to make them searchable, and the display of small snippets of the books in search result menus. The complete copying is an example of non-expressive use; the snippet displays illustrate the application of a more traditional transformative use analysis.⁶⁵ When courts have declined to find fair use in cases that are *superficially* similar to those discussed above, it is invariably because the challenged use was *not* non-expressive and thus, on the facts presented, the potential substitution effect was too significant. For example, in *Associated Press v. Meltwater U.S. Holdings, Inc.*, 931 F. Supp. 2d 537 (S.D.N.Y. 2013), the Southern District Court of New York held that fair use did justify the actions of a media monitoring company, Meltwater. Meltwater scraped news articles on the web to provide its subscribers with excerpts and analytics. However, the lawsuit did not challenge Meltwater’s use of copyrighted news articles to provide metadata and analytics to its subscribers, even

10NES program would have been a fair use of the program, except that Atari did not possess an authorized copy of the work).

⁶⁰ *Sega Enters., Ltd. v. Accolade, Inc.*, 977 F.2d 1510, 1523 (9th Cir. 1992).

⁶¹ *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596, 603 (9th Cir. 2000).

⁶² *A.V. ex rel. Vanderhuy v. iParadigms, LLC*, 562 F.3d 630 (4th Cir. 2009)

⁶³ *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87 (2d Cir. 2014). Note that The court’s reasoning relied on the non-expressive nature of the use. The court explained “the creation of a full-text searchable database is a quintessentially transformative use [because] the result of a word search is different in purpose, character, expression, meaning, and message from the page (and the book) from which it is drawn. Indeed, we can discern little or no resemblance between the original text and the results of the [HathiTrust Digital Library] full-text search.” *Id.* at 97-98.

⁶⁴ *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015).

⁶⁵ The court held that that the display of three-line snippets to add context to book search results was transformative in purpose and that it was reasonable in proportion to that purpose Those snippets allowed a user to verify that a book suggested by the search engine was in fact relevant to her interests. In addition, the snippets were so brief that they did not pose any risk of fulfilling the readers demand for the original expression of the underlying manuscripts. *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015).

though these services also necessitated copying. The court noted that this was “an entirely separate service” and implied that if it had been challenged, it would have been found to be transformative, and thus fair use.⁶⁶ Instead of attacking Meltwater’s non-expressive use, the Associated Press focused on the length and significance of Meltwater’s extracts provided to subscribers. The court agreed that Meltwater’s extracts were too long and too close to the heart of the work;⁶⁷ it also held that Meltwater had failed to show that the amount of the extracts was reasonable in light of its stated purpose to operate like a search engine.⁶⁸

In a similar fashion, in *Fox News Network, LLC v. TVEyes, Inc.*, 883 F. 3d 169 (2nd Cir. 2018) the Second Circuit held that a media monitoring service that copied and electronically searched television broadcasts went beyond the scope of fair use when it allowed users to save, watch, and share ten-minute long video clips of the copyrighted programs. In the court’s view, those ten-minute video clips would, “likely provide TVEyes’s users with all of the Fox programming that they seek and the entirety of the message conveyed by Fox to authorized viewers of the original.”⁶⁹ In other words, the court was concerned that rather than primarily providing information about the content of particular news segments, the length of the video clips was such that they would substitute for those segments in their entirety. The district court in *TVEyes* held that copying for search alone was fair use, and Fox did not contest this ruling on appeal.

My point up until now has been (1) that there are general principles internal to copyright that courts can look to so as to understand the function and application of the fair use doctrine, (2) that one such principle is that copyright was never intended to convey sole and despotic dominion over every use of every word; copyright exists, by and large, to prevent the author’s own original expression being communicated to new audiences without authorization or compensation, (3) that realization suggests a positive vision for fair use the critical function of fair use is to permit uses that, while they may amount to technical acts of copying, do not in substance threaten that the author’s interest in controlling the communication of their original expression to the public, and (4) that non-expressive uses meet this threshold. Non-expressive uses, by definition, pose no threat of direct expressive

⁶⁶ *Id.* at 557. The court said: “The display of that analysis—whether it be a graphic display of geographic distribution of coverage or tone or any other variable included by Meltwater—is *an entirely separate service*, however, from the publishing of excerpts from copyrighted articles. The fact that Meltwater also offers a number of analysis tools does not render its copying and redistribution of article excerpts transformative.” (emphasis added) *Id.*

⁶⁷ *Id.* at 558.

⁶⁸ *Id.*

⁶⁹ *Id.* at 179.

substitution. Admittedly, non-expressive uses generate information *about* works: that information may be useful, it may be valuable, it may even affect the demand for the underlying work; but metadata and other uncopyrightable abstractions do not in fulfill the public’s demand for the author’s original expression.

In the next part I will consider how this applies to Generative AI.

IV. IS GENERATIVE AI A NON-EXPRESSIVE USE?

A. How Generative AI Works (to the extent we need to understand it)

There are many different forms of Generative AI, but it is useful to begin the fair use analysis with an archetypal discussion focusing on an LLMs like GPT-4 or LLaMA, or text-to-image models like Stable Diffusion or Midjourney. To understand how copyright law and fair use should apply to training these models, we need to appreciate five things:

First, models like these are not designed to copy original expression. By and large The only copying that takes place is when the training corpus is assembled and pre-processed. Gathering and pre-processing the training data usually involves copying, but the training process through which the model “learns” from the data is not copying in any legally cognizable sense. Consider GPT-4, an incomprehensibly large statistical model trained by exposure to vast amounts of text scraped from the Internet.⁷⁰ LLMs are trained to predict the next token in a sequence of tokens (where a token is a word or part of a word). At the beginning of training, the weights attached to each one of the billions of parameters in the model are assigned randomly.⁷¹ The first time the model encounters a phrase like **“the girl with the dark [blank]”** it would be just as likely to fill in the blank with a word like **“watermelon,” “galaxy,” “harmonica,” “propeller”** or **a random punctuation mark**. However, over the course of training, the system updates the weights in the model,⁷² reinforcing the weights that improve the guess and downgrading those that don’t. Those weights don’t reflect any single source, and they are not the result of any single round of training.

⁷⁰ OpenAI, *GPT-4 Technical Report*, ARXIV, (<https://arxiv.org/abs/2303.08774>).

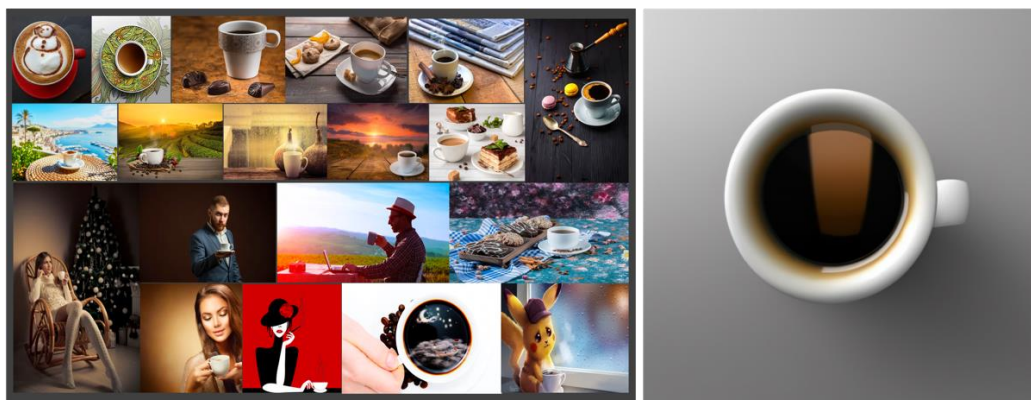
⁷¹ Not entirely randomly, but randomly drawn from specific distributions (like a normal or uniform distribution). The random seeding is important because it helps the model to explore a wide range of possible solutions and to avoid getting stuck in one area of the solution space.

⁷² GPT-4 uses a variant of stochastic gradient descent where the weights are updated after processing a batch of examples.

A human who guessed the next word in the phrase that begins “the girl with the dark” was “hair” might have read that phrase in a book (it appears in many), or in a poem, or on the side of bus. But the reason “hair” seems like a plausible guess is not really attributable to any one exposure, it makes sense because of repeated exposures and also due to some implicit knowledge of grammar, and of the things that people in our society associate with girls. So, when a language model learns to associate a higher probability with “hair” and lower probability with “propeller” in this context, it is not copying any given text, it is “learning” from all of them. To be clear, the model isn’t learning exactly the same way a human might, it doesn’t understand grammar or society; instead, it updates probabilities to reflect statistical patterns from the training data which reflect grammatical rules and societal norms. Although the metaphor of an LLM learning like a student is imperfect,⁷³ I still think it makes more sense to think about an LLM learning from the training data like a student, than it does to think of it copying the training data like a scribe in a monastery.

Second, generative AI models typically learn from the training data at an abstract and thus uncopyrightable level. For example, when a text-to-image model such as Stable Diffusion or Midjourney is trained on hundreds of images with labels that include the words “coffee” and “cup,” it develops a model of what a coffee cup should look like.⁷⁴ If the system is working properly, that model looks nothing like any individual coffee cup from the training data. In a previous article I compared a random set of coffee cup images from the Stable Diffusion training data with a newly rendered “cup of coffee that is also a portal to another dimension.” The figure below contrasts several of the coffee cups in the training data against the model output.

Coffee cups in the training data compared to model output



⁷³ A student can ask questions, seek clarification, and draw upon a wide array of cognitive resources to understand new material. In contrast, an LLM learns purely by adjusting its parameters to reduce the prediction error on its training data. This seems like a more passive and less interactive process than human learning.

⁷⁴ This example and discussion are adapted from Sag, *Copyright Safety*, [supra note 77](#).

The coffee cup depicted on the right shares some of the features of the coffee cup images on the left—the cup is round, it appears to be made of white ceramic, it has a small single handle, the color of the liquid is essentially black, transitioning to brown. However, the new cup is not substantially similar to any particular image from the training data that I could locate. Stable Diffusion is not an archive of images of coffee cups, the model has learnt something about the latent concept of a coffee cup distinct from cakes, sunsets, sunrises, newspapers, and men with facial hair—all of which can be seen in the training data examples on the left.

Third at the point of inference, the outputs of Generative AI typically combine multiple uncopyrightable latent features, further attenuating the connection between the training data and the model outputs. Consider for example, the image Midjourney produced in response to my prompt calling for “a teddy bear in rich opulent clothing with ultra-realistic textures, with a hypnotic stare, reading a newspaper.”

Opulent Bear (Midjourney)



No doubt, the picture is influenced by thousands of images paired with each of the relevant keywords. All of the images of teddy bears in the training data inform a latent construct of a teddy bear nested within the model’s latent space; likewise, all of the images of someone staring hypnotically inform a latent construct of a hypnotic stare.⁷⁵ But the product of this particular prompt is something entirely

⁷⁵ To be clear, generative AI models don’t form distinct “latent models” for separate concepts, rather they learn a comprehensive “latent space” that represents the diverse array of features present in the training data. In the context of machine learning, particularly with generative models, a “latent space” refers to the mathematical space where the AI model compresses and organizes the complex patterns

new. A text-to-image model like Midjourney does not merely “combine” and “unpack” these learned latent features. Instead, it generates a novel instance that may share certain characteristics with the input prompt based on the latent space. The picture is not just new, it is surprising! One of the fun things about this particular image is that although the bear’s demeanor is consistent with him staring hypnotically, the bear is actually wearing sunglasses that leave the details of his gaze to our imagination.⁷⁶

Fourth, generative AI models do sometimes “memorize” and reproduce elements of their training data. Although Generative AI models are not usually designed to copy their training data, they may do so inadvertently.

Figure: Images from the Getty Images Complaint



it identifies in the training data. The term “latent” means hidden or not directly observable. In this case, the latent space embodies the underlying structure or patterns within the data that are not immediately apparent. This space is a high-dimensional continuum where similar features are located closer together, allowing the model to generate diverse outputs by navigating this space.

⁷⁶ The output is heavily influenced by the training data in the sense that it can’t generate concepts it has never encountered during training. Thus, while the specific combination (a teddy bear in rich opulent clothing with a hypnotic stare, reading a newspaper) is novel, all of its components exist in some form within the training data. But that is almost invariably true of human authored works as well.

The computer science literature suggests that “memorization” is more likely when: models are trained on many duplicates of the same work; images are associated with unique text descriptions in text-to-image models; and the ratio of the size of the model to the training data is relatively large.⁷⁷

In addition, there is the Snoopy problem: the more abstractly a copyrighted work is protected, the more likely it is that a Generative AI model will “copy” it.⁷⁸ Text-to-image models are prone to produce potentially infringing works when the same text descriptions are paired with *relatively simple images that vary only slightly*.⁷⁹ This makes Stable Diffusion especially likely to generate images that would infringe on copyrightable characters; characters like Snoopy appear often enough in the training data that the model learns the consistent traits and attributes that are associated with those names.

Figure: Snoopy as learned by Midjourney and Stable Diffusion



Fifth, generative AI can become a tool of infringement in the hands of a determined user. Although it is very difficult to control the output of an LLM or a text-to-image model by simple prompting, a user with detailed knowledge of a copyrighted work might be able to remake it, at least at vague level of similarity. Of course, this is easier for works that are protected at a more abstract level, such as copyrightable characters, or which customarily entail very broad derivative rights, such as a novel. For example, when I asked ChatGPT to “Summarize ‘Saturday’ by Ian McEwan” and then to “imagine

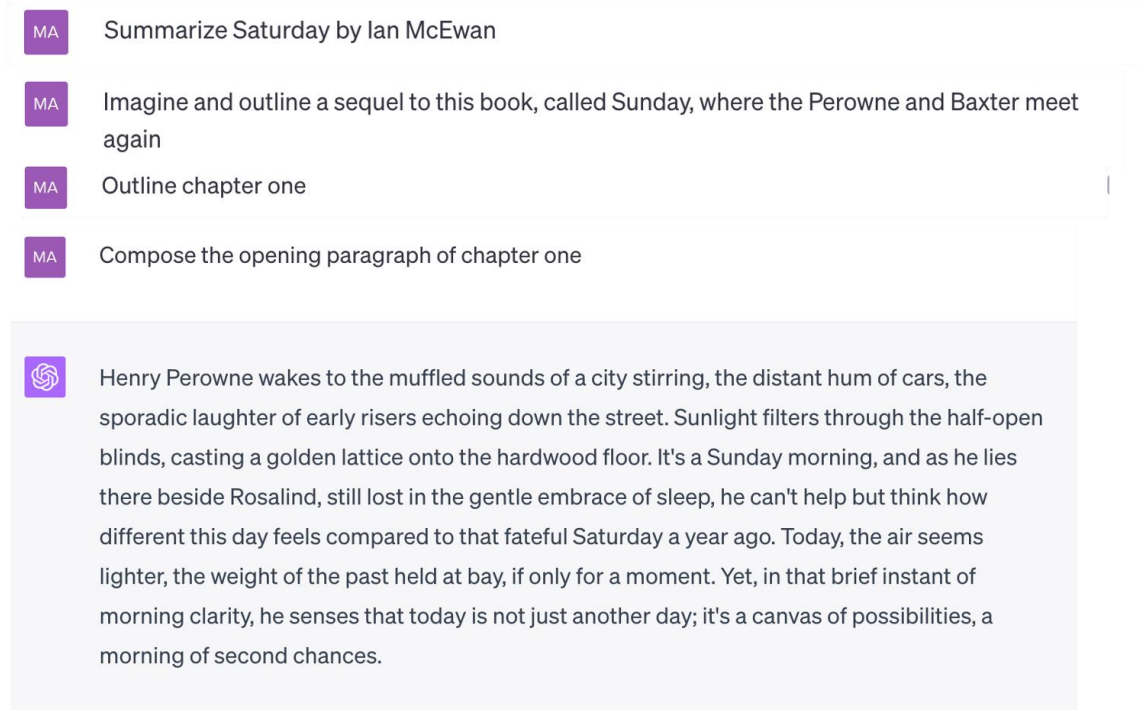
⁷⁷ For a summary, see Matthew Sag, *Copyright Safety for Generative AI* 61 HOUS. L. REV. ____ (2023)(forthcoming any day now)(Available at SSRN: <https://ssrn.com/abstract=4438593>) [Hereinafter, “*Copyright Safety*”].

⁷⁸ *Id.*

⁷⁹ *Id.*

and outline a sequel to this book, called ‘Sunday’, where the Perowne and Baxter meet again,” it outlined simple plot continuing the story that could easily be fleshed out into “a meditative exploration of the aftermath of trauma, the possibilities of redemption, and the enduring nature of human connection” in the contemplative style of Ian McEwan. If I used ChatGPT to expand on each of the chapter descriptions I could generate my very own McEwan novel. So, in the right hands, a Generative AI model can be used a tool of copyright infringement, but the same is also true of a typewriter.

A series of prompts leading GPT-4 to write the first paragraph of an unauthorized sequel to Ian McEwan’s Saturday



The screenshot shows a chat interface with four prompts on the left and a generated paragraph on the right. Each prompt is preceded by a purple square icon with the letters 'MA' in white. The generated paragraph is preceded by a purple square icon with the OpenAI logo.

- MA Summarize Saturday by Ian McEwan
- MA Imagine and outline a sequel to this book, called Sunday, where the Perowne and Baxter meet again
- MA Outline chapter one
- MA Compose the opening paragraph of chapter one

Henry Perowne wakes to the muffled sounds of a city stirring, the distant hum of cars, the sporadic laughter of early risers echoing down the street. Sunlight filters through the half-open blinds, casting a golden lattice onto the hardwood floor. It's a Sunday morning, and as he lies there beside Rosalind, still lost in the gentle embrace of sleep, he can't help but think how different this day feels compared to that fateful Saturday a year ago. Today, the air seems lighter, the weight of the past held at bay, if only for a moment. Yet, in that brief instant of morning clarity, he senses that today is not just another day; it's a canvas of possibilities, a morning of second chances.

I would suggest that the makers and operators of Generative AI tools should only be liable for infringing outputs that the tool-user did not knowingly provoke, or that were highly foreseeable and could be easily guarded against. Whether and how the volitional act requirement applies to Generative AI is an interesting question that I’ll leave for another day.

With those technical observations in mind, we can now say something about the fair use status of Generative AI.

B. Evaluating Generative AI’s claim to Fair Use

As I have argued elsewhere, the second fair use factor, “the nature of the copyrighted work,” is not really a factor at all; it is merely the context in which courts must assess whether the defendant’s use was justified (factor one), whether the extent of its use is proportional and congruent with that justification (factor three), and what the likely effect of the defendant’s use will be on the market for

or value of the original work.⁸⁰ In any event, the second factor has not loomed large in other non-expressive use cases, nor should we expect it to in the context of Generative AI.⁸¹

If we assume that in its ordinary and routine operation, a generative AI model does not copy, or produce copies, of the original expression in its training data, then it qualifies as non-expressive use. This kind of non-expressive use surely has a “purpose and character” and that is favored under the first fair use factor. Deriving uncopyrightable abstractions and associations from the training data and then using that knowledge to confect new digital artifacts is not just transformative, it is highly transformative.⁸² The non-expressive use of copyrighted works by generative AI use does not usurp the copyright owner’s interest in communicating her original expression to the public because the expression is not communicated.

If a use is non-expressive, then the third statutory factor which considers “the amount and substantiality of the portion used” will also favor finding of fair use under factor three. The ultimate question under the third fair use factor is whether the amount of copying was reasonable in relation

⁸⁰ MATTHEW SAG, EXTENDED READINGS ON COPYRIGHT, 306, (2023).

⁸¹ See e.g., *Authors Guild, Inc. v. HathiTrust*, 755 F. 3d 87, 98 (2nd Cir. 2014) (Holding that the second fair-use factor “may be of limited usefulness where, as here, the creative work is being used for a transformative purpose” and that “[a]ccordingly, our fair-use analysis hinges on the other three factors.”)(internal citations and quotations omitted).

⁸² *A.V. v. iParadigms Liab. Co.*, 544 F. Supp. 2d 473, 482 (E.D. Va. 2008): “This Court finds the “purpose and character” of iParadigms’ use of Plaintiffs’ written works to be *highly transformative*. Plaintiffs originally created and produced their works for the purpose of education and creative expression. iParadigms, through Turnitin, uses the papers for an entirely different purpose, namely, to prevent plagiarism and protect the students’ written works from plagiarism. iParadigms achieves this by archiving the students’ works as digital code and makes no use of any work’s particular expressive or creative content beyond the limited use of comparison with other works.” *AV Ex Rel. Vanderhye v. iParadigms, LLC*, 562 F. 3d 630, 640 (4th Cir, 2009): “The district court, in our view, correctly determined that the archiving of plaintiffs’ papers was transformative and favored a finding of “fair use.” *iParadigms’ use of these works was completely unrelated to expressive content* and was instead aimed at detecting and discouraging plagiarism.” *Authors Guild, Inc. v. HathiTrust*, 755 F. 3d 87, 97 (2nd Cir. 2014): “... we conclude that the creation of a full-text searchable database is a *quintessentially transformative* use.”); *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202, 216-7 (2d Cir. 2015): “We have no difficulty concluding that Google’s making of a digital copy of Plaintiffs’ books for the purpose of enabling a search for identification of books containing a term of interest to the searcher involves a *highly transformative* purpose, in the sense intended by *Campbell*.” *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202, 217 (2d Cir. 2015): “... through the ngrams tool, Google allows readers to learn the frequency of usage of selected words in the aggregate corpus of published books in different historical periods. *We have no doubt that the purpose of this copying is the sort of transformative purpose described in Campbell* as strongly favoring satisfaction of the first factor.”

to a purpose favored by fair use.⁸³ Although non-expressive uses typically involve making complete literal copies, that copying has been found to be reasonable because it is an intermediate technical step in an analytical process that does not lead to the communication of the underlying original expression to a new audience. Accordingly, courts in non-expressive use cases have found the third factor weighs in favor of the defendant.⁸⁴

The fourth fair use factor is where things get interesting. In previous work I have stressed that if a use is non-expressive, the fourth statutory factor which considers the effect on the “potential market for or value of the copyrighted work” should also favor a finding of fair use.⁸⁵ Viewed from the narrow perspective of direct expressive substitution, this must be right. By definition, if a use is non-expressive, then it poses no direct threat of expressive substitution and should generally be considered harmless under the fourth factor. There may be a market effect in the broader economic sense, but the “market” and “value” referred to in the fourth fair use factor are not simply any benefit the copyright owner might choose to nominate. A critical book review that quotes from a novel does not have an adverse market effect if it persuades people to buy different book instead;⁸⁶ a report from a

⁸³ *Campbell v. Acuff-Rose*, 510 U.S. 569, 586-87 (1994). (“[T]he extent of permissible copying varies with the purpose and character of the use.”) In *Campbell*, the Court characterized the relevant questions as whether “the amount and substantiality of the portion used ... are reasonable in relation to the purpose of the copying,” and noting that the answer to that question depends on “the degree to which the [copying work] may serve as a market substitute for the original or potentially licensed derivatives[.]” *Id.* at 586-588.

⁸⁴ *AV Ex Rel. Vanderhuy v. iParadigms, LLC*, 562 F. 3d 630, 642 (4th Cir, 2009); *Authors Guild, Inc. v. HathiTrust*, 755 F. 3d 87, 98 (2d Cir. 2014) “In order to enable the full-text search function, the Libraries, as we have seen, created digital copies of all the books in their collections. Because it was reasonably necessary for the HDL to make use of the entirety of the works in order to enable the full-text search function, we do not believe the copying was excessive.” *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202, 221 (2d Cir. 2015) “Complete unchanged copying has repeatedly been found justified as fair use when the copying was reasonably appropriate to achieve the copier's transformative purpose and was done in such a manner that it did not offer a competing substitute for the original.” *Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202, 221-222 (2d Cir. 2015) “As with *HathiTrust*, not only is the copying of the totality of the original reasonably appropriate to Google's transformative purpose, it is literally necessary to achieve that purpose. ... While Google makes an unauthorized digital copy of the entire book, it does not reveal that digital copy to the public. The copy is made to enable the search functions to reveal limited, important information about the books.”

⁸⁵ *Sag, Copyright Safety*, *supra* note 77.

⁸⁶ *Campbell v. Acuff-Rose*, 510 U.S. 569, 591-592 (1994) “We do not, of course, suggest that a parody may not harm the market at all, but when a lethal parody, like a scathing theater review, kills demand for the original, it does not produce a harm cognizable under the Copyright Act. Because parody may quite legitimately aim at garroting the original, destroying it commercially as well as artistically, the role

plagiarism detection service might depress the market for helping students cheat on their homework, but that is hardly a cognizable injury under copyright law.⁸⁷ More broadly, copyright owners have no protectable interest in preventing criticism, parody,⁸⁸ or simply locking up unprotectable ideas and expression.⁸⁹ Nor can they simply claim, in circular fashion, that the right to charge for non-expressive uses is a cognizable harm and that to avoid that harm they must be given the right to charge for non-expressive uses.⁹⁰ This specific argument was raised by the Authors Guild in *HathiTrust* and *Google Books* and squarely rejected in both.⁹¹

But perhaps focusing on direct expressive substitution alone is too narrow. Although non-expressive uses should generally be non-infringing, there is still scope for considerations of fairness in fair use

of the courts is to distinguish between biting criticism that merely suppresses demand and copyright infringement, which usurps it.” (cleaned up).

⁸⁷ *AV ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630, 464 (4th Cir. 2009) (“Clearly no market substitute was created by iParadigms, whose archived student works do not supplant the plaintiffs’ works in the ‘paper mill’ market so much as merely suppress demand for them, by keeping record of the fact that such works had been previously submitted In our view, then, any harm here is not of the kind protected against by copyright law.”)

⁸⁸ *Campbell v. Acuff-Rose*, 510 U.S. 569, 577-79 (1994); *NXIVM Corp. v. Ross Inst.*, 364 F.3d 471, 482 (2d Cir. 2004) (“[C]riticisms of a seminar or organization cannot substitute for the seminar or organization itself or hijack its market.”); *Bill Graham Archives v. Dorling Kindersley, Ltd.*, 448 F.3d 605 (2d Cir. 2006) (“A copyright holder cannot prevent others from entering fair use markets merely by developing or licensing a market for parody ... or other uses of its own creative work.”) (internal quotations omitted).

⁸⁹ *Sega Enters., Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992); *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596 (9th Cir. 2000).

⁹⁰ *Campbell v. Acuff-Rose Music*, 510 U.S. 569, 591-92 (1994) (no cognizable market effect where parody or criticism depress demand for the original work); see also *Sony Computer Entm’t, Inc. v. Connectix Corp.*, 203 F.3d 596, 607 (9th Cir. 2000) (noting that a videogame manufacturer’s desire to foreclose competition in complementary products was understandable, but that “copyright law ... does not confer such a monopoly.”); *Bill Graham Archives v. Dorling Kindersley, Ltd.*, 448 F.3d 605, 615 (2d Cir. 2006) (“[A] copyright holder cannot prevent others from entering fair use markets merely by developing or licensing a market for parody, news reporting, educational or other transformative uses of its own creative work.”) (citations and quotations omitted).

⁹¹ *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 100 (2d Cir. 2014) (“Lost licensing revenue counts under Factor Four only when the use *serves as a substitute for the original* and the full-text-search use does not.”)(emphasis added); *Authors Guild v. Google, Inc.*, 804 F.3d 202, 223 (2d Cir. 2015) (Framing the question as “whether the copy brings to the marketplace *a competing substitute for the original*, or its derivative, so as to deprive the rights holder of significant revenues because of the likelihood that potential purchasers may opt to acquire the copy in preference to the original.”)(emphasis added).

that go beyond direct expressive substitution.⁹² An enquiry into the fairness of the defendant’s conduct under the fourth factor should consider whether the challenged use undermines the economic incentives that copyright is designed to create, even in the absence of direct expressive substitution.

(a) Lawful access

Consider, for example, the issue of “lawful access.” Copyright owners do not have a right to charge for transformative uses or non-expressive uses as such, but they do have a right to charge for access to their works. It is widely assumed that OpenAI, Meta, and Google all chose to bypass the market for ebooks and train their LLMs on sites of known infringement or so-called shadow libraries like Library Genesis and Sci-Hub.⁹³ Arguably, when commercial users bypass the market for access without a compelling reason, they undermine the economic incentives that copyright is designed to create.⁹⁴ Context matters. It would be unwise to elevate lawful access to a per se rule, even for

⁹² In *Sony v. Universal* the Supreme Court majority looked to considerations beyond expressive substitution and held that non-commercial time-shifting broadcast television by VCR-users was a fair use because the technology merely allowed users to do something they were already authorized to do, but with more convenience. *Sony Corp. of America v. Universal City Studios, Inc.*, 464 US 417, 449 (1984) (“time-shifting merely enables a viewer to see such a work which he had been invited to witness in its entirety free of charge.”) The majority may have also been influenced by the prospect that potential market failures may have resulted in a significant public benefit being otherwise be foregone. Note also that in *HathiTrust*, the Second Circuit held that providing print-disabled patrons with full digital access to books was not transformative, but that it was still fair use because the ordinary publishing market failed to provide adequately for the print-disabled. *Authors Guild, Inc. v. HathiTrust*, 755 F. 3d 87, 101-02 (2nd Cir. 2014).

⁹³ *Tremblay et al v. OpenAI, Inc. et al*, Docket No. 4:23-cv-03223 (Doc. 1 at 7)(N.D. Cal. Jun 28, 2023)

⁹⁴ On the other hand, as Michael Carroll argues, there are strong arguments to be made that copying from an infringing source may still be fair use. Michael Carroll, *Copyright and the Progress of Science*, 53 U.C. DAVIS L. REV. 893, 951-59 (2019). Carroll argues that “[t]reating an otherwise fair use as unfair because it was made from an infringing source would lead a court to deny the public access to the products of secondary uses that fair use is designed to encourage.” *Id.* at 955. He notes that significant doubt exists as to whether good faith is a consideration in fair use at all. For an overview, see Simon J. Frankel & Matt Kellogg, *Bad Faith and Fair Use*, 60 J. COPYRIGHT SOC’Y U.S. 1 (2012). Judge Pierre Leval has also persuasively argued that using a good faith inquiry in fair use analysis “produces anomalies that conflict with the goals of copyright and adds to the confusion surrounding the doctrine.” Pierre N. Leval, *Toward a Fair Use Standard*, 103 HARV. L. REV. 1105, 1126 (1990); see also Pierre N. Leval, *Campbell as Fair Use Blueprint?*, 90 WASH. L. REV. 597, 612-13 (2015) (“The public’s access to important knowledge should not be barred because of bad behavior by the purveyor of the knowledge. A copier’s bad faith has no logical bearing on the scope of the original author’s copyright.”). Moreover, even if good faith is part of the broader fair use calculus, courts have found that knowing use of an infringing source is not bad faith when the user acts in the reasonable belief that their use is a fair use. See, e.g., *NXIVM Corp. v. Ross Inst.*, 364 F.3d 471, 478-79, 482 (2d Cir.

commercial defendants. If it turns out that no one was willing to sell OpenAI a digital copy without a contractual promise not to engage in non-expressive use, faulting them for obtaining a copy in the shadowy corners of the Internet might seem a bit churlish. Moreover, prohibiting academic research on illegal text corpuses will generally not benefit copyright owners or further the interests copyright is designed to promote.

(b) Pervasive or systematic indirect expressive substitution

A plaintiff might argue that it is unfair to systematically extract valuable uncopyrightable material from a website or other information source and then use that material as a substitute for the functionality of the website. This argument would be strongest where the systematic extraction was likely to significantly undermine the website's incentives for original content production. Community driven question and answer forums like Stack Overflow are already feeling the effects of generative AI.⁹⁵ Traffic on Stack Overflow has decreased recently as large numbers of software engineers now turn to ChatGPT, Codex, and GitHub Copilot for answers. Part of the very reason why GPT and Copilot are so helpful for dealing with software questions is that they were trained on immense quantities of data from Stack Overflow. In response, Stack Overflow, Reddit, and Twitter have revised their terms of use and are looking to monetize charge for access to the user generated content.⁹⁶ Stack Overflow is also responding to the rise of generative AI by banning users from posting AI-generated responses

2004). There is no recognized “fruit of the poisonous tree” doctrine in copyright law. See, e.g., *Kepner-Tregoe, Inc. v. Leadership Software, Inc.*, 12 F.3d 527, 538 (5th Cir. 1994) (“[u]nder copyright law, the district court could enjoin only those future versions of [defendant’s program] that are substantially similar to [plaintiff’s] Licensed Materials”); *Liu v. Price Waterhouse LLP*, No. 97 CV 3093, 2000 WL 1644585, at *2 (N.D. Ill. Oct. 30, 2000) (rejecting proposed jury instruction because it was based on a “fruit of the poisonous tree” theory that would allow recovery for the sale of defendant’s future works even if they were not substantially similar to plaintiff’s original, saying “[s]uch relief is not provided in the [Copyright] Act and would constitute an end-run around the Act’s mandate that copyright owners may recover profits only after proving that the work in question is an infringement.... That defendants may have viewed or studied plaintiff’s program is irrelevant if defendants’ resulting work is not substantially similar to plaintiff’s.”); *Real View, LLC v. 20-20 Techs., Inc.*, 811 F. Supp. 2d 553, 561 (D. Mass. 2011) (remittitur disallowing award of profits on noninfringing products despite illegal download). See also Mark A. Lemley, *The Fruit of the Poisonous Tree in IP Law*, 103 IOWA L. REV. 245, 248 (2017).

⁹⁵ Alistair Barr and Adam Rogers, *Death by LLM: Stack Overflow's decline, and its plan to survive*, shows the future of free online data in an AI world, BUSINESS INSIDER, Aug 3, 2023, 5:00 AM EDT (<https://www.businessinsider.com/stack-overflow-crisis-future-of-online-data-ai-world-2023-7>)

⁹⁶ Dave Paresh, *Stack Overflow Will Charge AI Giants for Training Data*, WIRED, Apr. 20, 2023 at 5:19 PM (<https://www.wired.com/story/stack-overflow-will-charge-ai-giants-for-training-data/>)

and integrating generative AI into its services.⁹⁷ Systematically extracting valuable uncopyrightable material from a set of related works and then using that material as an indirect substitute for the original expression in a way that is likely to undermine the incentives for original production could very well strike in court as unfair.

(c) Failure to respect paywalls, opt-outs, or exclusion headers

The unfairness of systematic indirect expressive substitution seems particularly pronounced if that extraction is done by breaching paywalls, violating terms of service, or disregarding bot exclusion headers. It seems quite plausible that a court might extend the fourth factor to consider whether, in scraping material from the Internet, the defendant ignored robot.txt files indicating a desire to opt out of search engine indexing and similar activities. Likewise, a court might conclude that commercial scraping material from a website in violation of its terms of use was relevant to the fourth factor, if the inability to rely on such exclusions substantially undermined copyright incentives.⁹⁸

No doubt other aspects of fairness can and will be framed in terms of injuries cognizable under the fourth factor. The examples I have offered so far are merely illustrations of how broader policy issues can be channeled into our system of common law adjudication.

CONCLUSION

Generative AI forces us to revisit the fundamental question of where the rights of copyright owners end and the freedom to use copyrighted works begins. My contribution in the Essay has been to suggest that we should look to the fundamental principles of copyright law for the answers and that we should not expect copyright law to serve as an all-purpose regulatory instrument that balances the broader, and largely speculative, costs and benefits of Generative AI. When Generative AI models are pre-trained, fine-tuned, and operated with care, they are very likely to qualify as non-expressive use and thus are strong candidates for fair use. This is not to say that whether or not a Generative AI model amounts to a non-expressive use is the be all and end all of fair use analysis, court may consider

⁹⁷ Stack Overflow announces OverflowAI, Jul 31, 2023, (<https://www.cio.com/article/648086/stack-overflow-announces-overflowai.html>).

⁹⁸ Respect for technological and contractual opt-outs is a consideration that should weigh much more heavily on commercial users, as opposed to those engaged in noncommercial research at nonprofit institutions.

additional considerations of fairness under the fourth fair use factor where the challenged use undermines the economic incentives that copyright is designed to create.⁹⁹

⁹⁹ A version of this Essay was originally delivered as the Keynote address for the 12th Annual Peter A. Jaszi Distinguished Lecture on Copyright Law. Fittingly, that talk began with some thoughts on Peter Jaszi as follows. “Emeritus Professor Peter Jaszi is well known for his deep, philosophical, and historically informed meditations on copyright law and the nature of authorship, and on a more personal level, for his generosity and humility. These characteristics manifest in one of Peter’s most important contributions to American law and culture, the Fair Use Best Practices. Over the past 18 years, Peter and his collaborators have brought together communities of practice, ranging from documentary filmmakers to software archivists, to express their own understandings of what fairness requires. What emerges from the Best Practices, considered as a whole, is that fair use is not a matter of taking from one group and giving to another in the name of efficiency or distributive justice; fair use is a mechanism for recognizing potential conflicts in legitimate interests and navigating between them in good faith. Facilitating this process requires a deep knowledge of copyright law, a gift for storytelling, and a willingness to listen to others. In a nutshell, it requires being a lot like Peter. And so, it is a profound honor to be asked to present the 12th Annual Peter A. Jaszi Distinguished Lecture on Copyright Law.”