

# Copyright and the Training of Human Authors and Generative Machines

*Robert Brauneis\**

---

\* Michael J. McKeon Professor of Intellectual Property Law and Co-Director of the Intellectual Property and Technology Program and of the GW Center for Law and Technology, The George Washington University Law School. For helpful comments and suggestions, I want to thank Sarah Polcz; Alicia Solow-Niederman; Kristelia Garcia and the students in her law and technology colloquium at Georgetown Law, especially Karalyn Joseph, who proposed names for the meanings of “expression” that I try to differentiate in this Article; EvinBagduHanish; Bhamati Viswanathan; Guido Westkamp and the participants in the 2024 EIPIN Congress at Queen Mary University of London; Peter Karol, Guy Rub, Oren Bracha, Carys Craig and all of the participants in the Fifth Annual Art Law Works-in-Progress Colloquium at The Ohio State University; Benjamin Sobel; Matthew Sag; Keith Kupferschmid; Kevin Madigan; the participants in the 2024 Intellectual Property Scholars Workshop at UC Berkeley School of Law; Justin Hughes; and Christophe Geiger and the participants in the workshop at LUISS Guido Carli.

© 2024 Robert Brauneis. This is an open access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction, provided the original author and source are credited.

## TABLE OF CONTENTS

Introduction .....	4
I. Framing the Question .....	6
II. Copyright and the Training of Human Authors.....	10
A. Human Authors and the Strictures and Freedoms of Copyright..	13
1. Fair Use.....	13
2. First Sale and Cost Sharing.....	15
3. Limitations on the Exclusive Rights and Human Experiences	16
4. Public Licensing and Advertisement-Based Revenue Models	18
5. Lack of Fixation or Authorship.....	18
6. Remedies .....	18
B. The Ideal Copyright Regime for Authors?.....	20
III. Copyright and the Training of Generative Machines.....	21
A. Generative AI Tools and Non-Expressive Uses.....	22
1. The Three Meanings of Expression and “Non-Expressive Use”	23
.....	23
a. Constitutive Expression.....	23
b. Actionable Expression .....	27
c. Felt Expression .....	29
2. Human Authors and Generative Machines: Distinguished	30
Only By Felt Expression, and That Should Not Make a Legal	30
Difference .....	30
a. Machines and Humans Both Use Constitutive Expression	31
.....	31
b. Machines and Humans Rarely Use Actionable Expression..	39
.....	39
c. Machines Do Not Use Felt Expression and Humans Do,	40
But That Should Not Make a Legal Difference .....	40
i. Copyright Law’s Purpose: Protecting Education as	41
Well as Entertainment.....	41
ii. The Imputed Purpose of Copyrighted Works: Both	43
Entertainment and Education.....	43
iii. Should Learning Without Enjoyment Be Discounted?	45
.....	45
B. The Reproduction Right and Computer “Reading” .....	46
1. Reading and the Human Limitations Implicit in the Structure	48
of the Exclusive Rights.....	48
2. Reading, Emotion, Memory, and Agency.....	54
C. Remedies.....	58
IV. Conclusion .....	59

## INTRODUCTION

Since November 2022, authors and copyright owners have filed about twenty lawsuits against generative AI developers, many of them aspiring class actions. At the center of almost all of those lawsuits are claims of direct copyright infringement for the unauthorized use of copyrighted works in training generative AI models. The decisive issue in the U.S. cases will almost certainly be whether that copying is justified as fair use when the models are producing output that is being distributed publicly or used for private gain, but that output is not substantially similar to the training works.

In this Article, I explore that issue by considering how various copyright doctrines apply to constrain or free human authors when they learn from copyrighted works, and then asking whether there are grounds for treating the training of generative AI models differently. There are many limitations on copyright of which human authors can and do take advantage as they are learning. However, there is no blanket immunity for use of copyrighted works to train human authors, even though those authors typically do not end up creating substantially similar works. Human authors typically end up paying, directly or indirectly, for most of the copyrighted works from which they learn.

If no copyright doctrine frees human authors from all copyright constraints when they use works to learn, why should the training of generative AI models be treated as fair use? I will focus principally on two arguments in favor of treating generative AI training as fair use. (As I will note below, these arguments have also been made in foreign jurisdictions that do not have fair use exceptions, by tying them to other copyright doctrines.) The first is that the training of generative models involves a “non-expressive use” of the training works, and that “non-expressive uses” should be treated as fair uses, or indeed as uses that do not implicate any of the exclusive rights of copyright. Among the scholars who have made versions of such an argument are Matthew Sag, Oren Bracha, Michael Murray, Abraham Drassinower, Carys Craig, Alain Strowel, and Cheryl Foong.<sup>1</sup>

---

1. See Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 304, 307 (2023); Matthew Sag, *Fairness and Fair Use in Generative AI*, 92 FORDHAM L. REV. 1887 (2024); Oren Bracha, *The Work of Copyright in the Age of Machine Production*, U. OF TEX. L. LEGAL STUD. RSCH. PAPER SERIES., 25 (2023); Michael D. Murray, *Generative AI Art: Copyright Infringement and Fair Use*, 26 SMU SCI. & TECH. L. REV. 259, 275–80 (2023); ABRAHAM DRASSINOWER, WHAT’S WRONG WITH COPYING? 88 (2015); Alain Strowel, *Reconstructing the Reproduction and Communication To the Public Rights: How To Align Copyright with Its Fundamentals*, in COPYRIGHT RECONSTRUCTED: RETHINKING COPYRIGHT’S ECONOMIC RIGHTS IN A TIME OF HIGHLY DYNAMIC TECHNOLOGICAL AND ECONOMIC CHANGE 203 (P. Bernt Hugenholtz, ed., 2018); Carys J. Craig, *The AI-Copyright Challenge: Tech-Neutrality, Authorship, and the Public Interest*, in RESEARCH HANDBOOK ON INTELLECTUAL PROPERTY AND ARTIFICIAL INTELLIGENCE 134, 153–54 (2022); Cheryl Foong,

It turns out that “expression,” and therefore “non-expressive use,” has been defined in three different ways. Under the only definition that distinguishes generative AI training from human learning, a “non-expressive use” is one that does not result in an aesthetic or hedonic reaction on the part of a being that can feel such reactions. Because computers cannot have aesthetic reactions while human authors cannot avoid them, computers can in that sense make “non-expressive use” of works that humans cannot. However, I argue that recognition of this “non-expressive use” as a fair use would be bad copyright policy. Copyright should be and usually has been considered to protect not only the entertainment value of works for passive and unchanging human beings, but also the educational value of works for human beings who want to learn and change, both individually and collectively. Moreover, because all works both generate aesthetic responses in human beings, and educate and enculturate us, both entertainment and educational purposes have to be imputed to their creation. That means that their use to train generative AI models is not a wholly different, “transformative” use.

The second argument is that generative AI training is functionally equivalent to a human activity that copyright doctrine other than fair use would shield from liability, and fair use should be employed to legally recognize that functional equivalence. The leading candidates for such an activity are reading, listening, and watching. I conclude that there are reasons to doubt whether generative AI training is for copyright purposes equivalent to human reading, listening, and watching. Generative model training likely transcends the human limitations that underlie the structure of copyright’s exclusive rights, and it therefore should not be able to take advantage of that structure. Moreover, the very inability of computers to have any hedonic or aesthetic reactions to the works they are processing, and their inability to remember and act on those reactions, makes computer processing fundamentally different than human experience of works.

Part I of the Article frames the question. It provides reasons why it is likely that the central issue in the many lawsuits pending in the U.S. against generative AI developers will be whether the use of copyrighted works to train generative AI models is fair use. Part II considers the variety of doctrines that make the use of copyrighted works to educate human authors subject to more or less legal constraint. These include not only fair use, but also the first sale doctrine, internal limitations on the exclusive rights, the enablement of public licensing and advertising-based revenue

---

*Immaterial Copying in the Age of Access*, 44 EUROPEAN INTELL. PROP. REV. 513 (2022). Mark Lemley and Brian Casey also mention non-expressive use. See Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 750 (2020). However, those two authors write as much about predictive AI as about generative AI, and therefore speak of non-expressive use in somewhat narrower terms, leaving some of the generative AI questions open. See *id.* (“When learning is done to copy expression, for example, by training an ML system to make a song in the style of Ariana Grande, the question of fair use can—and should—become much tougher.”).

models, and the fixation and authorship requirements. It will conclude that while many doctrines ease copyright constraints on human learning, learning by human authors is not categorically free from those constraints. Authors typically end up paying some amount, directly or indirectly, for many of the copyrighted works from which they learn.

Part III considers copyright and fair use in the context of generative AI training. Part III.A. considers the argument that generative AI training is different from human learning because it is a “non-expressive use” that is recognized under existing precedent as fair use, at least as long as the model does not produce output that is substantially similar to any training work. I contend that the term “expression,” and therefore also the term “non-expressive use,” has actually been used in three different ways. Under the first definition, which I call “constitutive expression,” generative AI training use is expressive use, not non-expressive use. Under the second, “actionable expression,” generative AI training is non-expressive, but so is use by human authors to learn, and so to cover generative AI training, fair use doctrine would have to be radically expanded to immunize all learning. Under the third definition, “felt expression,” generative AI training is indeed non-expressive in a way that human learning is not and likely cannot be, but I will argue that copyright law should not make that factual distinction legally significant.

Part III.B. considers the argument that the use of copyrighted works to train generative models is functionally equivalent to human reading, viewing, and listening—activities that are outside of the scope of the exclusive rights granted to copyright owners—and that fair use should be employed to legally recognize that functional equivalence. It will conclude that generative model training is likely not functionally equivalent to those activities. That is the case both because training algorithms surpass the human limitations that ground the structure of copyright’s exclusive rights, and because those algorithms do not have and cannot act on hedonic reactions to the works they process. Part III.C. considers the issue of remedies if generative AI training is not a fair use. Part VI concludes.

## I. FRAMING THE QUESTION

It is by now common knowledge that most prominent generative AI base models were trained on works of authorship that are under copyright, without the authorization of the owners of copyright in those works.<sup>2</sup> Many copyright owners have filed individual or proposed class action lawsuits against generative AI

---

2. See, e.g., Peter Henderson et al., *Foundation Models and Fair Use* 1–2 (Stanf. L. and Econ. Olin Working Paper No. 584, 2023), <https://arxiv.org/abs/2303.15715> (detailing how popular generative models such as GPT 3 and 4, Stable Diffusion, Codex, and BERT have been trained on collections containing copyrighted material, such as BookCorpus, Books3, and webcrawls such as C4 and OpenWebText).

developers.<sup>3</sup> Complaints in some of these lawsuits include a wide variety of federal and state causes of action—not just direct copyright infringement, but also vicarious and contributory copyright infringement, removal of copyright management information, unjust enrichment, unfair competition, negligence, right of publicity, and others. However, early motions to dismiss in some of those broad lawsuits have been narrowing them to claims of direct copyright infringement for unauthorized copying of works to train generative AI models.<sup>4</sup> Thus, it seems certain that that direct infringement claim will be central in many of these lawsuits.

Some of the complaints in those lawsuits relate dramatic instances in which the defendant's generative AI service, when prompted to produce the lyrics to a well-known song or the text of a newspaper article, generated a near-verbatim copy of a copyrighted work on which it was trained.<sup>5</sup> Those instances of generative AI models "regurgitating" verbatim copies are cut-and-dry cases of copyright infringement.<sup>6</sup> However, they are probably already relatively rare.<sup>7</sup> And "regurgitation" is likely to become rarer. The developers of generative AI models really do not want them to produce verbatim copies of existing works. That reduces generative AI to a search engine that only delivers cached copies, which is no more than a degraded version of an old technology. It is already clear that generative AI models can do far more than that—they can produce text, images, music and videos that are not substantially similar to any of the works on which they were trained. It is very likely that generative AI developers will figure out how to prevent the models from producing near-verbatim copies of training works in all but very rare cases, probably both by making changes to how the models are trained, and, if they are deploying the models

---

3. For an updated list of lawsuits against generative AI developers, see DAIL—the Database of AI Litigation, available at <https://blogs.gwu.edu/law-eti/ai-litigation-database/> (search for the "Generative AI" in Application Area).

4. See Order on Motions to Dismiss and Strike, *Andersen v. Stability AI, Ltd.*, No. 23-cv-00201 (N.D. Cal. Oct. 30, 2023); Order Granting Motion to Dismiss, *Kadrey v. Meta Platforms*, No. 23-cv-03417 (N.D. Cal. Nov. 20, 2023); Order Granting in Part and Denying in Part the Motions to Dismiss, *Tremblay v. OpenAI, Inc.*, No. 23-cv-03223 (N.D. Cal. Feb. 12, 2024).

5. See, e.g., Complaint at 5, *Concord Music Grp., Inc. v. Anthropic PBC*, No. 23-cv-01092 (M.D. Tenn. Oct. 18, 2023) ("When a user prompts Anthropic's Claude AI chatbot to provide the lyrics to songs such as 'A Change Is Gonna Come,' 'God Only Knows,' 'What a Wonderful World,' 'Gimme Shelter,' . . . or any other number of Publishers' musical compositions, the chatbot will provide responses that contain all or significant portions of those lyrics."); Complaint at Exhibit J, *N.Y. Times Co. v. Microsoft Co.*, No. 23-cv-11195 (S.D.N.Y. Dec. 27, 2023) (containing "One Hundred Examples of GPT-4 [Allegedly] Memorizing Content from The New York Times").

6. For one good attempt to separately define "extraction" (verbatim copies produced through intentional prompting of a generative model), "regurgitation" (verbatim copies produced through intentional or unintentional prompting), and "memorization" (verbatim copies produced through any means), see A. Feder Cooper & James Grimmelmann, *The Files Are in the Computer: Copyright, Memorization, and Generative AI*, CHI-KENT L. REV. (forthcoming 2024).

7. See Nicolas Carlini et al., *Extracting Training Data from Diffusion Models*, in 32ND USENIX SEC. SYMP. 5253 (2023) (managing to produce 94 outputs that were near-identical to training works, out of 175 million targeted attempts); see also Peter Henderson et al., *supra* note 2, at 7–8, 11 (2023) (citing other extraction efforts, and reporting on extraction experiments).

only as closed services, by filtering or supplementing prompts, or filtering outputs, or both.<sup>8</sup>

That leads us to the question that I want to address in this Article. Let us define the term “generation,” taken as a noun, as an output of a generative AI system that would fall into one of the categories of “works of authorship” enumerated in the U.S. Copyright Act had it been created by a human being.<sup>9</sup> Given the most popular generative AI tools available today, generations would be most likely to qualify as literary works; pictorial or graphic works; motion pictures; or combined musical works and sound recordings.<sup>10</sup> Taking “substantial similarity” as standing for one of the two principal elements of copyright infringement, suppose that a generative AI model does not produce generations that are substantially similar to any of the copyrighted works on which it was trained.<sup>11</sup> Suppose also that the generative AI model is not merely part of a research project to test the capabilities of certain machine learning technologies, but will be deployed commercially.<sup>12</sup> In that case, can the developer use any work of authorship it wants to train the model, even if that

---

8. See, e.g., Peter Henderson et al., *supra* note 2, at 20–25 (2023) (detailing various types of technical mitigation tools, including data and output filtering, instance attribution, differentially private training, and fair use learning from human feedback); Microsoft Corporation, Comments on U.S. Copyright Office’s Notice of Inquiry and Request for Comment Re Artificial Intelligence and Copyright [Docket No. 2023-06], at 4 (Oct. 30, 2023) [hereinafter, Microsoft Comments] <https://www.regulations.gov/comment/COLC-2023-0006-8750> [<https://perma.cc/G67A-A3V9>] [<https://web.archive.org/web/20241004155553/https://www.regulations.gov/comment/COLC-2023-0006-8750>] (detailing measures to prevent memorization or repetition).

9. Cf. A. Feder Cooper et al., *Report of the 1st Workshop on Generative AI and Law* 28 (Yale Law & Econ. Research Paper, 2023), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4634513](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4634513) (defining “generation[s]” as “complex, human-interpretable outputs, such as full sentences or natural-looking images”).

10. See 17 U.S.C. § 102(a) (enumerating eight categories of works of authorship, including “(1) literary works; (2) musical works, including any accompanying words; (3) dramatic works, including any accompanying music; (4) pantomimes and choreographic works; (5) pictorial, graphic, and sculptural works; (6) motion pictures and other audiovisual works; (7) sound recordings; and (8) architectural works”).

11. For further discussion of substantial similarity, see *infra* text accompanying note 26, notes 94–101.

12. I have some sympathy for the unauthorized use of copyrighted works to see whether certain machine learning technologies would work at all. It is probably impossible to convince investors to invest money to attempt to license millions of works if you cannot provide evidence that your machine learning technology will, if trained on enough works, produce something very valuable. In its motion to dismiss the complaint of the group of plaintiffs led by Mike Huckabee, Bloomberg has argued that its unauthorized use of copyrighted works to train its large language model was a fair use in part because the project was for internal research purposes only. See Memorandum of Law in Support of Bloomberg L.P. and Bloomberg Finance L.P.’s Motion to Dismiss First Amended Class Action Complaint at 10–12, *Huckabee et al. v. Bloomberg L.P.*, No. 23-cv-09152 (S.D.N.Y. Mar. 22, 2024). If that is true—if Bloomberg did not use any generations produced by the model in the conduct of its business, but was just testing to see what is possible—I think that would weigh heavily in favor of fair use. On the difference between academic and commercial research more broadly, see Deven R. Desai & Mark Riedl, *Between Copyright and Computer Science: The Law and Ethics of Generative AI* 21–26 (Georgia Tech Scheller Coll. of Bus. Research Paper No. 4735776, 2024), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4735776](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4735776).

work is under copyright and the owner of copyright does not want it to be used in that way?

The owners of copyright in the works can make a pretty strong case for *prima facie* copyright infringement. The developers of generative AI tools that are defendants in these cases use computer programs that find files of the type that are useful to their training project—text, image, music, or video, for example—that are available at internet URLs; download copies of those files; and store them on a server under the developers’ control. They keep those copies for some time and process them with a machine learning algorithm to create a generative model. Yes, there is at least one academic, Oren Bracha, who has argued that those copies are not really “copies” within the meaning of the Copyright Act, because they generally are never seen by any human being.<sup>13</sup> However, it seems likely that the courts in which this litigation is proceeding will hold that the copies that the generative AI developers have made are in fact “material objects . . . in which a work is fixed . . . from which the work *can be* perceived, reproduced, or otherwise communicated . . . with the aid of a machine or device.”<sup>14</sup> That language seems to require only that the copies are potentially perceivable by human beings, not that they actually were ever so perceived. If you could gain access to the servers on which those copies were stored, you would definitely be able to read, view, or listen to them with the aid of your computer. Thus, if the works were under copyright, and the developers had no authorization from the owners of copyright, what they did almost certainly amounts to *prima facie* infringement of the reproduction right.<sup>15</sup>

What is left? Fair use. Section 106 of the Copyright Act, which enumerates the exclusive rights granted to copyright owners, states that those rights are “[s]ubject to sections 107 through 122.”<sup>16</sup> Those sixteen sections contain many exceptions to the exclusive rights, but the only one that potentially applies to what the AI developers have done and are doing is the fair use exception codified in § 107. It alone is broad and vague enough that it might apply to the making of copies of millions or billions of files containing copyrighted works for the purpose of building a generative model. And generative AI developers are already invoking it. OpenAI, for example, has repeatedly asserted that “[t]raining AI models using publicly available internet materials is fair use, as supported by long-standing and widely accepted precedents.”<sup>17</sup>

---

13. See Bracha, *supra* note 1. Others have made this argument more generally, without tying it specifically to the language of the U.S. Copyright Act. See DRASSINOWER, *supra* note 1, at 101; Strowel, *supra* note 1 (citing DRASSINOWER); Foong, *supra* note 1.

14. 17 U.S.C. § 101 (“Copies”) (emphasis added).

15. See 17 U.S.C. § 106(1).

16. See 17 U.S.C. § 106.

17. *Open AI and Journalism*, OPENAI (Jan. 8, 2024), <https://openai.com/blog/openai-and-journalism> [<https://web.archive.org/web/20241004161638/https://openai.com/web/20241004161638/https://openai.com/index/openai-and-journalism/>]; see OpenAI LP, Comment on U.S. Patent and Trademark Office’s Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation, Docket No. PTO-C-2019-0038 (2019) [hereinafter, OpenAI Comments 2019], [https://www.uspto.gov/sites/default/files/documents/OpenAI\\_RFC-84-FR-58141.pdf](https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf) [<https://perma.cc/M7UP-J482>]



More directly, Bloomberg L.P. has invoked fair use in a motion to dismiss the claim of direct copyright infringement in litigation brought by a group headed by former Arkansas governor Mike Huckabee.<sup>18</sup> That focuses us on the question I want to explore in this Article: Is using copyrighted works to train generative AI models which produce no output that is substantially similar to those works a fair use?

A note about the international implications of the issues I will explore. There are a relatively small number of countries whose copyright laws have an open-ended fair use provision like that of the United States.<sup>19</sup> For those countries, a wide-ranging policy discussion, with consideration of all of the traditional fair use factors, will be directly applicable. For other countries, that will not be the case. However, some of the more specific concepts discussed in detail below will still have legal relevance through other doctrinal routes. First, those concepts may arguably be implicit in other general provisions. Thus, for example, Abraham Drassinower, Alain Strowel, Cheryl Foong, Oren Bracha and Carys Craig all argue that a non-expressive or non-communicative use of work does not infringe any exclusive right in that work.<sup>20</sup> If that is so, then what “expression” means is important to determine whether a use implicates any of the exclusive rights, even in a country that has no fair use provision. Second, some narrower exceptions in particular countries may explicitly reference concepts like non-expressive use. For example, Article 30-4 of the Japanese Copyright Act provides that a work may be used if it is not a person’s purpose to personally enjoy or cause another person to enjoy the thoughts or sentiments expressed in that work<sup>21</sup>—a category that is arguably identical to one use of the term “non-expressive.” Thus, although the principal focus of this Article is U.S. law, some of the discussion will refer to non-U.S. sources and will have application outside of the U.S.

## II. COPYRIGHT AND THE TRAINING OF HUMAN AUTHORS

When considering whether and under what circumstances unauthorized use of copyrighted works to train generative AI models is or should be “fair use,” one starting point is to ask how fair use and other copyright doctrines feature in the training, or education, of all of the human creators who copyright law groups under

---

[[https://web.archive.org/save/https://www.uspto.gov/sites/default/files/documents/OpenAI\\_RFC-84-FR-58141.pdf](https://web.archive.org/save/https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf)]; Memorandum of Law in Support of OpenAI Defendants’ Motion to Dismiss, N.Y. Times Co. v. Microsoft, No. 23-cv-11195 (S.D.N.Y. Feb. 26, 2024) [hereinafter, OpenAI’s MTD].

18. See Memorandum of Law in Support of Bloomberg L.P. and Bloomberg Finance L.P.’s Motion to Dismiss First Amended Class Action Complaint at 9–21, *Huckabee et al. v. Bloomberg L.P.*, No. 23-cv-09152 (S.D.N.Y. Mar. 22, 2024).

19. See, e.g., Niva Elkin-Koren & Neil Weinstock Netanel, *Transplanting Fair Use Across the Globe: A Case Study Testing the Credibility of US Opposition*, 72 HASTINGS L.J. 1121, 1125–26 (2020) (listing eleven countries that have adopted, “with some variation,” a fair use model).

20. See sources cited *supra* note 1.

21. Copyright Law of Japan, Art. 30-4 (English translation provided by the Copyright Rsch. and Info. Ctr.), <https://www.cric.or.jp/english/clj/cl2.html> [<https://web.archive.org/web/20240924155408/https://www.cric.or.jp/english/clj/cl2.html>]; see Tatsuhiro Ueno, *The Flexible Copyright Exception for ‘Non-Enjoyment’ Purposes—Recent Amendment in Japan and Its Implication*, 70 GRUR INT. 145 (2021) (discussing this provision).

the term “author”: writers, coders, visual artists, photographers, composers, recording artists, choreographers, dramatists, architects, and so on. While the generative AI question is new, the question of how copyright should and does treat the training of human authors has existed since copyright law itself emerged. Courts, legislatures, and agencies have had plenty of occasions to consider that question, and to shape copyright law in light of it.

Of course, human learning and generative model training are not exactly the same. Most fundamentally, computers do not *do* anything, in the sense that they are not moral agents to whom we can ascribe purposes. A generative AI project is still an activity of human beings, who are trying to program computers to build models for the purpose of accomplishing certain tasks. Thus, computers are not *learning* in any purposive sense. We are really still comparing humans to humans—humans learning without using computers, and humans learning how to automate certain processes.

If computers are not moral agents, they are also not moral subjects. Many of the positive connotations we attach to learning are directly tied to our human condition. Because each of us must make our way in the world, we recognize how important learning is to survival, to flourishing, and to the enjoyment of life. Because we also depend upon each other, we recognize how important learning is to society. By contrast, training a computer to fish is not going to help the computer. The computer needs no help and is not deserving of help. We again need to consider the consequences for human beings, and for other animals who we consider moral subjects, not for the computer. Thus, “machine learning” is not inevitably deserving of whatever support we think education should get. On the other hand, generative AI undoubtedly has many benefits for society. In the currently dominant model of responding to and elaborating on human prompts, it undoubtedly has many benefits for the human prompters.<sup>22</sup> Those can and should be taken into account when making generative AI policy.

It is also the case that the process of training a large language model is not the same as the process of teaching a human being. However, the outcome is similar: both human authors and generative models gain the ability to produce new works of authorship in part through exposure to copyrighted works. That is the similarity I want to highlight, and that I think is significant for copyright purposes. It follows in the tradition of taking an external view of artificial intelligence, and asking only whether machines can behave in ways that we would identify as intelligent if human beings so behaved.<sup>23</sup>

---

22. See, e.g., Katrina Geddes, *How Art Became Posthuman: Copyright, AI, and Synthetic Media*, 42–45 (2024), <https://papers.ssrn.com/abstract=4865510> (arguing that generative AI will help users and promote semiotic democracy).

23. See, e.g., J. McCarthy, M.L. Minsky, N. Rochester & C.E. Shannon, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence* (Aug. 31, 1955), <https://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html> [<https://perma.cc/L93R-ELA8>] [<https://web.archive.org/web/20240925220620/http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>] (“[T]he artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so

Given the current broad scope of copyright protection, almost every human being has been an author. Just about all of us have fixed an original work of authorship in some tangible medium of expression, from childhood crayon drawings to emails and cell phone photos.<sup>24</sup> For some purposes, we may want to concentrate on *remunerated* human authors—people who have received money for the works of authorship they have created. They, after all, are the people who are creating most of the works that have “potential market(s)” that we would want to consider in the fourth factor of fair use analysis.<sup>25</sup> Both remunerated and unremunerated human authors learn from works of authorship that are still under copyright. (Of course, they also learn from works that are in the public domain, and from other sources that are not works of authorship at all, as I will discuss below.) Those authors then go on to create their own works, and the vast majority of those works do not infringe the works from which the authors learned.

Copyright infringement requires at its core two elements—both “lack of independent creation” and “improper appropriation,” the latter of which is sometimes known as “substantial similarity.”<sup>26</sup> The works that a human author produces are not created independently of the works from which that author has learned. Suppose, for example, that a journalist learned to write “short, sharp sentences” by reading an Ernest Hemingway novel.<sup>27</sup> She then writes a news article full of sentences with grammatical structures that she had found in that novel. That article was not created independently of that novel. Thus, with respect to the works that demonstrably influenced an author—and are still under copyright—the first element of copyright infringement is satisfied.

However, outside of the myopic world of copyright litigation, it is rare that the works that a human author creates are substantially similar to any of the works that have influenced them. The journalist’s article is almost certainly not substantially similar to the Hemingway novel or any portion of it. More generally, authors undoubtedly incorporate many of the expressive elements they found in works that inspired them into their own works. However, they usually have remixed them enough, and introduced enough new combinations and elements, that none of their works is substantially similar to any one of the works that they have read or seen or heard.

---

behaving.”). I will consider some differences below, especially the difference that human authors have aesthetic or hedonic reactions to the works from which they learn, while machines do not have such reactions to the works to which they process during training.

24. See 17 U.S.C. § 102(a) (providing that “[c]opyright protection subsists . . . in original works of authorship fixed in any tangible medium of expression.”).

25. See 17 U.S.C. § 107(4) (directing courts to consider “the effect of the use upon the potential market for or value of the copyrighted work” in determining whether a use is a fair use).

26. See, e.g., *Arnstein v. Porter*, 154 F.2d 464, 468 (2d Cir. 1946) (holding that two essential elements in an infringement case are “(a) that defendant copied from plaintiff’s copyrighted work and (b) that the copying (assuming it to be proved) went so far as to constitute improper appropriation”); cf. *WILLIAM F. PATRY, PATRY ON COPYRIGHT* § 9.16 (stating that the plaintiff must “prove that defendant copied from its work,” and what the defendant copied must “constitute[] a material amount of expression”).

27. I discuss this example further below at text accompanying note 34.

In that respect, the situation of the generative models whose training with copyrighted works either is or is not fair use is similar to that of human authors. Generative models are shaped through encounters with many works of authorship, and many of those works are protected by copyright, and are still under copyright. Well-trained generative models produce text, images, video, and music that, for the most part, are not substantially similar to the works on which they were trained. So, it is fair to start by asking, what copyright strictures must human authors respect with regard to the works from which they learn? In what ways does revenue flow from those learning authors to the authors and copyright owners of the works from which they learn? Or considered from the opposite perspective, what freedoms do human authors enjoy with respect to their formative works? In what ways does revenue *not* flow to the authors and copyright owners of the formative works?

Providing a comprehensive answer to those questions would require a review of more than a hundred pages of the Copyright Act and many hundreds of judicial opinions, as well as consideration of all of the ways in which copyright owners have decided to license their works, implicitly and explicitly. Here, I will touch on some of the doctrine that I think is the most important. I will also provide examples of revenue flows, though there too I have no aspiration to be comprehensive. I will then briefly consider the issue of whether the balance struck is the right one.

## A. HUMAN AUTHORS AND THE STRICTURES AND FREEDOMS OF COPYRIGHT

### 1. Fair Use

There is no blanket fair use defense to potential infringement liability for the use of works to educate human authors. Suppose that I want to keep handy a copy of, say, the novel *What We've Lost is Nothing* by Rachel Louise Snyder so that I can study it thoroughly and reread it from time to time for inspiration.<sup>28</sup> (Snyder happens to be a named plaintiff in one putative class action against OpenAI and Microsoft and another against Meta.)<sup>29</sup> There is no reasonable interpretation of current fair use doctrine under which I can claim that it is fair use for me to go to the library photocopier with the library's copy of the book, make a copy of the entire book, and take the copy home. It will not help me in the least to swear, and swear truthfully, that my own creative output will never include anything that copyright law would consider to be substantially similar to Snyder's novel. Fair use simply does not extend that far.<sup>30</sup> As Benjamin Sobel has said, “[n]o human can rebut an infringement claim

---

28. See RACHEL LOUISE SNYDER, *WHAT WE'VE LOST IS NOTHING* (2014). In fact, I have a copy of the book that I purchased used. The action in the novel takes place in Oak Park, Illinois, where I lived for many years.

29. See First Consolidated Amended Complaint at 4, *In re ChatGPT Litigation*, No. 23-cv-03223 (N.D. Cal., Mar. 13, 2024); First Consolidated Amended Complaint at 4, *Kadrey et al. v. Meta Platforms, Inc.*, No. 23-cv-03417 (N.D. Cal., Dec. 22, 2023).

30. In many other countries around the world, my “private copying” would not infringe upon the copyright in the book, because those countries have passed legislation which provides for a private copying

merely by showing that he has learned by consuming the works he copied, even if he puts this new knowledge to productive use later on.”<sup>31</sup>

According to one 2019 study of 4083 households in the United States, each U.S. household owns an average of 114 books.<sup>32</sup> With 125.7 million households in the United States in 2018–2022,<sup>33</sup> that would amount to about 14.3 billion books held in private households. The vast majority of those books were undoubtedly lawfully purchased rather than copied without authorization. Thus, just taking that as one example of works used to educate human authors, we find that authors have not been able to claim that they should get access to those works free of copyright restrictions, and that they often pay for copies of works that have helped build their skills to create their own, non-infringing works.

In the course of arguing that use of copyrighted works to train generative models is fair use, Daniel Jeffries asserts that “[w]riters at the [New York Times] did not pay the Hemingway estate for learning to write short, sharp sentences as young people learning journalism.”<sup>34</sup> As a factual statement, that is almost certainly literally false. It is a pretty good bet than more than one journalist who writes for the New York Times bought a copy of an Ernest Hemingway novel sometime in their life; that some of the purchase price went to pay royalties to the Hemingway estate; and that they learned something from that novel about writing style.

That does not mean that fair use has no role in supporting educational activities. Section 107 of the Copyright Act, the fair use provision, states that the use of a work for “nonprofit educational purposes” will weigh in favor of fair use. It also states that the purposes that support a claim of fair use include “teaching (including multiple copies for classroom use), scholarship, or research.”<sup>35</sup> In practice, however, the support afforded by fair use to educational activities turns out to be a partial subsidy, not a blanket exemption from copyright. For example, at the end of the decade-plus-long saga of litigation between three academic publishers and the Georgia State University over the copying and distribution of forty-eight excerpts from books to

exception. However, that legislation also imposes levies on copying equipment and on blank media, the proceeds of which are distributed to copyright owners and authors. Thus, while my copying would be permitted by law, whoever bought the copying equipment and blank media is paying some amount for that privilege, and even if that was not me, I might be paying for it indirectly, through taxes to support libraries, copying fees, and so on. The private copying exception may seem like a big deal, but you would have to ask, how many people actually make a private copy of an entire book rather than paying for a copy or borrowing it from a library? My guess is that there are not many.

31. See Benjamin L.W. Sobel, *Artificial Intelligence’s Fair Use Crisis*, 41 COLUM. J.L. & ARTS 45, 73 (2017).

32. See Joanna Sikora, M. D. R. Evans & Jonathan Kelley, *Scholarly Culture: How Books in Adolescence Enhance Adult Literacy, Numeracy and Technology Skills in 31 Societies*, 77 SOC. SCI. RSCH. 1, 6 (Table 1) (2019).

33. U.S. Census Bureau, *QuickFacts: United States*, <https://www.census.gov/quickfacts/fact/table/US#> [<https://web.archive.org/web/20241004163815/https://www.census.gov/quickfacts/fact/table/US#>].

34. Daniel Jeffries (@Dan\_Jeffries), X (Dec. 28, 2023, 4:27AM), [https://twitter.com/Dan\\_Jeffries1/status/1740303405254377808](https://twitter.com/Dan_Jeffries1/status/1740303405254377808) [<https://perma.cc/4R5B-BLVZ>] [<https://web.archive.org/web/20241004165501/https://x.com/x/migrate>].

35. 17 U.S.C. § 107.

students, the U.S. District Court for the Northern District of Georgia held that the copying and distribution thirty-seven of those excerpts were fair uses, but that the copying and distribution of eleven of those excerpts were not fair uses.<sup>36</sup> For instance, the use of one chapter from *The Craft of Inquiry* by Robert R. Alford, which amounted to 6.25% of the entire book, was fair use, considering not only the percentage of the work used, but the market for the book as a whole and the market for excerpts from the book.<sup>37</sup> By contrast, the use of two chapters of *The Power Elite* by C. Wright Mills, which amounted to 12.5% of the entire book, was not fair use,<sup>38</sup> nor was the use of 1.58% of *The SAGE Handbook of Qualitative Research*, due largely to the proven demand for licensing of excerpts from that work.<sup>39</sup> Georgia State never even tried to argue that the copying and distribution of an entire book would amount to fair use.

In some limited circumstances, of course, copying an entire work for educational purposes could be fair use. For example, the 1976 “Classroom Guidelines,” which were intended to construct a “safe harbor” for copying in connection with educational instruction, provided that entire poems and articles could be sometimes be copied and distributed to students.<sup>40</sup> However, they needed to meet both tests for “brevity”—up to a 250-word poem or 2500-word article—and for “spontaneity”—teachers had to individually decide to use the works shortly before the class session in question.<sup>41</sup> While the Classroom Guidelines were not meant to go the limit of fair use, it remains quite clear that fair use was never meant to cover any general distribution of educational materials in their entirety. The upshot is that most students buy or rent textbooks, and most schools buy or license the copyrighted materials that they provide for students.

## 2. First Sale and Cost Sharing

The first sale doctrine, codified in § 109 of the Copyright Act,<sup>42</sup> supports a lot of sharing of the cost of obtaining physical copies of books and images and video and recorded music. Many authors borrow books and videos and music from libraries. There were about 1.77 billion items in U.S. public libraries in 2019, and undoubtedly virtually all of those items were authorized copies purchased by the libraries.<sup>43</sup> In that

---

36. See *Cambridge Univ. Press v. Becker*, 446 F. Supp. 3d 1145 (N.D. Ga. 2020).

37. See *id.* at 1162–65.

38. See *id.* at 1264–66.

39. See *id.* at 1266–69.

40. See Agreement on Guidelines for Classroom Copying in Not-for-Profit Educational Institutions with Respect to Books and Periodicals, H.R. REP. NO. 1476, at 67–70 (1976), reprinted in 1976 U.S.C.C.A.N. 5659, 5681–83.

41. See *id.*

42. See 17 U.S.C. § 109 (allowing any owner of a lawfully made copy of a work to dispose of it without the permission of the copyright owner, with limitations for computer programs and sound recordings).

43. See Nicholas Rizzo, *State of US Public Libraries—More Popular & Digital Than Ever* (Feb. 17, 2022), <https://wordrated.com/state-of-us-public-libraries/> [https://perma.cc/VC8R-KBPX] [<https://web.archive.org/web/20241004164221/https://wordrated.com/state-of-us-public-libraries/>].

same year, there were about 2.98 billion instances of borrowing items from those libraries.<sup>44</sup> Thus, each item in a public library's collection served on average the needs of just under two borrowers. The most popular items were undoubtedly shared by many more. The first sale doctrine also facilitates the sale of used books. Many of the estimated 14.3 billion books in U.S. households were undoubtedly purchased in used condition, allowing the cost of those books to be shared by two or more households over time. In the digital network era, when "borrowing" occurs over the internet on devices, licensing and technical measures continue to simulate the experience of temporary possession and access, and facilitate the sharing of costs.<sup>45</sup>

### 3. Limitations on the Exclusive Rights and Human Experiences

Copyright law also facilitates learning by human authors through the limitations of the exclusive rights themselves. The only exclusive rights afforded to copyright owners are those that concern the making and distribution of copies—identical or substantially similar—and the presentation of public performances and displays.<sup>46</sup> People do not infringe any exclusive right by *experiencing* works of authorship— by seeing, reading, hearing, or touching them.<sup>47</sup> That leaves more leeway for people, including remunerated authors, to learn from works. When we read a copy of a book, or listen to a song on the radio, we do not have to worry about whether the copy or the performance is infringing, because our activity of reading or watching carries no liability with it. However, it would be a mistake to say that those limitations of copyright's exclusive rights leave human learning free from the constraints of copyright. Rather, those limitations implement two policy decisions: (1) a decision about *who* is liable if and when a work is copied, distributed, performed or displayed without authorization; and (2) a decision about how *fine-grained* the control is that copyright owners have over their works.

If a person is reading an unauthorized copy of a book, or watching an unauthorized public performance of a video, then *someone* is liable for copyright infringement—it just is not the reader or the viewer. That situation therefore poses only a practical question of enforcement: Can the copyright owner find and sue the person who made or distributed the copy, or who performed the work publicly? As long as copyright enforcement is reasonably good, and authorized copies, displays or performances are available on reasonable terms, most people will end up reading and watching and listening to authorized copies and performances and displays.

---

44. *See id.*

45. For a description of how authorized ebook lending works, see *Hachette Book Grp., Inc. v. Internet Archive*, 664 F. Supp. 3d 370, 374–75 (S.D.N.Y. 2023). In that same opinion, the court holds that the Internet Archive's unauthorized ebook lending is not a fair use. *See id.* at 379–91.

46. *See* 17 U.S.C. § 106(1)–(6).

47. *See, e.g., Fortnightly Corp. v. United Artists Television, Inc.*, 392 U.S. 390, 398–99 (1968) (members of theater audiences do not perform, and neither do television viewers). As Jessica Litman has put it, "[f]or most of the history of copyright, the law left reading, listening, and viewing unconstrained." Jessica Litman, *Lawful Personal Use*, 85 TEX. L. REV. 1871, 1882 (2006).

Thus, for example, in February 2024, 99% of American households subscribed to at least one streaming service, and on average, Americans pay for 2.9 streaming services each month.<sup>48</sup> Because of enforcement efforts,<sup>49</sup> most of that content is not easily available as unauthorized public performances that would lead to no viewer liability.<sup>50</sup> Some of it may be available for unauthorized download, using peer-to-peer protocols like BitTorrent and index sites like The Pirate Bay,<sup>51</sup> but then the recipient is creating a new unauthorized copy, and could be liable for copyright infringement. Thus, although it is hard to track, it is very likely that the vast majority of experiences of works of authorship under copyright are of authorized copies, performances, or displays, for which compensation is flowing, directly or indirectly, back to the copyright owner. The decision to exclude experiences of works of authorship from the reach of exclusive rights protects important liberty and privacy interests. However, it does not result in any blanket freedom for people in general, or remunerated authors in particular, to learn from copyrighted works of authorship without permission from the copyright owner.

The limitation of the exclusive rights to reproduction, distribution, and public performance and display also means that the copyright owner's control is coarse-grained. A copyright owner can charge for a copy of a book, but cannot control the reading of the book, or any private performances of the text. It is no secret that the transition from physical copy to digital network distribution has resulted in a lot of controversy about whether that copyright owner control would and should become finer grained.<sup>52</sup> The answer so far is that it has, but not nearly as much as skeptics feared, in part because of their opposition, and in part because of market forces.

---

48. See Ana Durrani, *Top Streaming Statistics in 2024*, FORBES HOME (Aug. 15, 2024), <https://www.forbes.com/home-improvement/internet/streaming-stats/> [<https://perma.cc/C3M7-KF2B>] [<https://web.archive.org/web/20241004164507/https://www.forbes.com/home-improvement/internet/streaming-stats/>].

49. See, e.g., Press Release, Office of Public Affairs, U.S. Department of Justice, Five Men Convicted for Operating Major Illegal Streaming Service (June 20, 2024) (announcing criminal convictions for operating an illegal streaming service branded "Jetflix"), <https://www.justice.gov/opa/pr/five-men-convicted-operating-major-illegal-streaming-service> [<https://perma.cc/V34T-49PR>] [<https://web.archive.org/web/20241004164842/https://www.justice.gov/opa/pr/five-men-convicted-operating-major-illegal-streaming-service>].

50. To be sure, there are illegal streaming sites. For one list, see Top Illegal Streaming Sites in Canada and Their Risks, STREAMSAFELY, <https://streamsafely.ca/where-to-watch/dangerous-illegal-streaming/> [<https://perma.cc/TA9A-KUJE>] [<https://web.archive.org/web/20240920155025/https://streamsafely.ca/where-to-watch/dangerous-illegal-streaming-in-canada/>] (last visited Sept. 20, 2024).

51. See BITTORRENT, <https://www.bittorrent.com> [<https://perma.cc/LB5T-B4TT>] [<https://web.archive.org/web/20240918221952/https://www.bittorrent.com/>] (last visited Sept. 20, 2024); PIRATE BAY, <https://thepiratebay.org/index.html> [<https://perma.cc/6WDP-HRM6>] [<https://web.archive.org/web/20240919132141/https://thepiratebay.org/index.html>] (last visited Sept. 20, 2024).

52. See, e.g., AARON PERZANOWSKI & JASON SCHULTZ, *THE END OF OWNERSHIP: PERSONAL PROPERTY IN THE DIGITAL ECONOMY* (2016).



#### 4. Public Licensing and Advertisement-Based Revenue Models

The owners of many works of authorship decide to distribute them under public licenses that allow anyone to use their works under whatever terms and conditions the owners decide to impose. They may make their works more or less freely available out of charitable impulses, or to build their reputation. It is also true that many owners of copyright decide to publicly perform or display their works accompanied by advertisements, and to seek revenue not from fees paid by audience members, but from fees paid by advertisers. As a result both of public licensing and of advertising-based revenue models, many people experience many works without paying fees. That, of course, does not stem from any limitations or exceptions to copyright law, but rather from the copyright owners' use of their exclusive rights to further what they see as being their interests.

#### 5. Lack of Fixation or Authorship

Human authors also learn a lot from sources that are not protected by copyright because they are not "fixed in a tangible medium of expression,"<sup>53</sup> or because they are not created by authors. As children, we all learn from our parents and others speaking with us, and we continue to learn from conversations throughout our lives. Those oral communications are not fixed and they therefore are not protected by copyright. We also learn a lot from natural sources. A visual artist may learn from the shape of a leaf, from the texture of a stone, or from the color of the sky at sunrise. Those sources are not protected by copyright because they have no human author.

#### 6. Remedies

This is the last one, but an important one. Let us return to Daniel Jeffries' assertion that "[w]riters at the [New York Times] did not pay the Hemingway estate for learning to write short, sharp sentences as young people learning journalism."<sup>54</sup> As I suggested above,<sup>55</sup> this is almost certainly false. Suppose, however, that a young writer admitted that they had made an infringing copy of *For Whom the Bell Tolls* ("*Bell*"); that they had read that copy; and that they had been inspired by that reading to write their novel *For Whom the Ringtone Rings* ("*Ringtone*"), which became acclaimed for, among other things, its short, sharp sentences. Suppose further that *Ringtone* had generated a million dollars in profits, and that many schools had decided assign *Ringtone* to their students in place of *Bell*, causing sales of the latter to drop by \$500,000. Assuming that *Ringtone* was not substantially similar to *Bell*, what actual damages and profits could the owner of copyright in *Bell* get for the writer's infringement? The answer is that it could get no more than the profits from the sale

---

53. 17 U.S.C. § 102(a).

54. Jeffries, *supra* note 34.

55. See *supra* text accompanying note 34.

of one copy of *Bell*.<sup>56</sup> It would not be entitled to any of *Ringtone's* profits, not even some tiny apportioned amount. Nor would it be entitled to any actual damages, even though *Ringtone* demonstrably substituted for *Bell*. Those measures of damages are limited to the circumstance in which *Ringtone* itself infringed *Bell*. Pam Samuelson, drawing on earlier work of Mark Lemley, calls this limitation the rejection of the “fruit of the poisonous tree” doctrine, and cites a number of cases in support of that rejection.<sup>57</sup>

Statutory damages and equitable relief should be limited similarly. Statutory damages may sometimes seem altogether inappropriate when a defendant has made a single copy for personal use.<sup>58</sup> In any event, however, in determining where an award of statutory damages should fall within the wide range that the Copyright Act provides, a court should not take into account the value of skills learned from reading an infringing copy. As for equitable relief, *Bell's* owner would not be entitled to any injunctive relief regarding *Ringtone*. In theory, *Bell's* owner might be able to have the court order the destruction of the infringing copy of *Bell*, but maybe not even that, since the Copyright Act leaves the court a lot of discretion with regard to injunctive relief.<sup>59</sup>

Stated more generally, remedies for infringement should not extend to the products of the things learned or abilities gained from that infringement, so long as those products are not substantially similar to the infringed works. In tort law terms, the infringing copy of *For Whom the Bell Tolls* may have been a but-for cause of *For Whom the Ringtone Rings*, but it is not a proximate cause. That is an important freedom for human authors. If they learn from infringing, nothing of what they learn can be reached by non-punitive remedies for that infringement, so long as they don't create works that are substantially similar to the infringed work.<sup>60</sup>

I have undoubtedly not described all the ways in which fair use, first sale, the basic structure of the exclusive rights, fixation, authorship, or other copyright doctrines could possibly come to the aid of human authors. They are important limitations, and we will have occasion to consider them below in relation to generative AI

---

56. This assumes that the young writer had made the infringing copy of “Bell” within the last three years. See 17 U.S.C. § 507(b) (providing that civil copyright infringement actions must be brought within three years of the accrual of the claim). The statute of limitations is also an important limit on liability; the legal consequences of our infringements do not hang over our heads forever.

57. See Pamela Samuelson, *Fair Use Defenses in Disruptive Technology Cases*, 71 UCLA L. REV. (forthcoming) (manuscript at n.375) (2023) (citing several cases and Mark A. Lemley, *The Fruit of the Poisonous Tree in IP Law*, 103 IOWA L. REV. 245 (2017)).

58. On the problems of statutory damages, see generally Pamela Samuelson & Tara Wheatland, *Statutory Damages in Copyright Law: A Remedy in Need of Reform*, 51 WM. & MARY L. REV. 439 (2009). Canadian copyright law provides a lower range for statutory damages in cases of infringement for non-commercial purposes (\$100–\$5000) than in cases of commercial infringement (\$500–\$20,000). See Canadian Copyright Act § 38.1. U.S. copyright law makes no such distinction. See 17 U.S.C. § 504(c).

59. See 17 U.S.C. § 503(b) (providing that “as part of a final judgment or decree, the court may order the destruction or other reasonable disposition of all copies or phonorecords found to have been made or used in violation of the copyright owner's exclusive rights”) (emphasis added).

60. One can imagine a Copyright Act providing for punitive measures to deter certain kinds of infringement, but in the United States, at least, copyright has never had civil punitive damages.

development. What I think remains incontestable is that neither fair use nor any other copyright law doctrine provides a blanket exception to copyright liability for human authors in their learning encounters with copyrighted works of authorship. Human authors end up paying for many of the learning experiences they have with those copyrighted works, even though the works that those authors produce are not substantially similar to the works from which they have learned.

### B. THE IDEAL COPYRIGHT REGIME FOR AUTHORS?

For as long as copyright has existed, policymakers, scholars, authors, and others have debated what the ideal set of copyright law provisions would look like, and what theory would justify those provisions. The debate will undoubtedly continue. Here, I only want to make a point about the edges of that debate. Very few people would argue that human authors should be able to access all previous works for learning purposes without any copyright restrictions at all. That seems unlikely to be the right answer from any number of theoretical perspectives, whether the issue is viewed as one of providing incentives, or rewarding labor, or protecting personality or human rights. One of the most revealing perspectives may be a relatively rare one: considering the relationship between generations of authors as a matter of Rawlsian justice, as Professor Dawn Nunziato did in a 2002 article.<sup>61</sup>

John Rawls himself formulated a principal of intergenerational savings as part of his theory of justice.<sup>62</sup> Suppose that you did not know in which generation you were going to live along the timeline of history. In that state of ignorance, you had to settle on a general principle of how much each generation would save for the next. You would not decide that each generation should save nothing for the next, and consume it all. That would benefit you if you happened to live in the first generation, but it would be devastating if you lived in any other generation. On the other hand, you also would not decide that each generation should save everything for the next. That would only benefit you if you lived in the last generation, and would otherwise impoverish you. You would formulate a principle that would lie somewhere in between those extremes.

Nunziato applied that Rawlsian reasoning to authors and copyright. Suppose that you were an author who did not know to which generation of authors you would belong along the timeline of civilization. In that state of ignorance, you needed to formulate a set of copyright rules that would apply to all generations of authors. Nunziato was principally concerned with justifying limitations on copyright, including many of the limitations I describe above, and those are important. They would likely include limitations, not only on scope, but also on duration, since extending the duration of copyright much more than a few decades beyond an author's life produces little added incentive or benefit for that author, while it subjects

---

61. See Dawn C. Nunziato, *Justice Between Authors*, 9 J. INTELL. PROP. L. 219 (2002).

62. See JOHN RAWLS, *A THEORY OF JUSTICE* 251–59 (Harvard Univ. Press rev. ed. 1999).

future generations to additional constraints.<sup>63</sup> Here, I want to emphasize the other side. You probably would not formulate a rule that allowed authors to access all works by means of unauthorized copying so long as they were learning something from those works. That would benefit you if you happened to end up living in the last generation of authors on earth. In any other generation, however, you would probably want to be paid something, for some limited term, for the works you created from which others are learning. In return, you would be willing to pay something for recently created works from which you learned.<sup>64</sup> That contributes to an intuitive understanding of why learning from works is not entirely free of copyright strictures, whether through fair use or otherwise, even though there is still a lot of debate about the details.

### III. COPYRIGHT AND THE TRAINING OF GENERATIVE MACHINES

Human authors do not get the benefit of any categorical fair use exception from copyright liability for their encounters with copyrighted works from which they learn. Why should humans who create and use generative AI tools get such an exception?<sup>65</sup> There are two different principal arguments that I would like to address. The first is that the use of copyrighted works to train a generative AI tool is a “non-expressive use.” I will argue that there are actually three senses in which the term “expression” is used, and three corresponding definitions of “non-expressive use.” Under the first two senses of “expression,” human and machine learning cannot be distinguished. Both make “constitutive” expressive use of the copyrighted works from which they learn, and neither usually make “actionable” expressive use. Under the third sense of “expression,” what I will call “felt expression,” human and machine learning can indeed be distinguished. Human authors necessarily have aesthetic or hedonic reactions to the works they learn from; computers do not. If the only thing copyright protected in a work were the aesthetic or hedonic reactions that it produced, or if works used to train generative AI were created for the sole purpose of producing such reactions, then that factual difference could make a legal difference. However, I will argue that copyright law protects both the enjoyment and the learning value of works, and that works have the effect, and therefore the imputed purpose, of both entertaining and enlightening us. Adopting a view that anhedonic

---

63. See Nunziato, *supra* note 61, at 249–53 (discussing how the principal of intergenerational savings might result in specific limitations on copyright, including duration); *cf.* Eldred v. Ashcroft, 537 U.S. 186, 254–55 (2003) (Breyer, J., dissenting) (noting that the extension of copyright term from life plus fifty years to life plus seventy years is exceedingly unlikely to create any additional motivating benefit to authors).

64. Of course, while under current copyright law everyone is an author of copyright-protected works of authorship, very few are remunerated authors. The knowledge of whether you were going to be one of those remunerated authors would probably greatly affect your judgment about what the best copyright rules would be. Thus, the Rawlsian veil of ignorance should probably hide from you the knowledge of whether you will be a remunerated author.

65. As James Grimmelman put it, why should “a digital humanist [be able to] skim a million books with abandon while a humanist who reads a few books closely must pay full freight for hers”? James Grimmelman, *Copyright for Literate Robots*, 101 IOWA L. REV. 657, 675 (2015).

learning, and only anhedonic learning, is a fair use would unfairly disadvantage human authors for whom aesthetic experience and learning are inextricably intertwined. Existing precedent does not mandate such an outcome, and it should not be extended in that direction.

The second argument is that what people are doing with computers when they train a generative AI model is functionally equivalent to some activity that for human authors is free of copyright restrictions. The most frequently mentioned of such activities are reading, viewing, or listening that fall outside of the scope of the exclusive rights. According to this line of argument, we should use fair use to achieve “perceptual parity” between machines and humans, compensating for the technical need for computers to make copies during generative model training. I will argue that the generative model training process is in relevant ways different than human reading, viewing or listening, and that we should therefore not ignore the copying of works that is part of that process.

### A. GENERATIVE AI TOOLS AND NON-EXPRESSIVE USES

Matthew Sag and Oren Bracha are leading academic proponents in the United States of the argument that the process of training and deploying a generative AI tool either does not infringe any of copyright’s exclusive rights in any copyrighted training work (Bracha), or makes fair use of any such work (Sag), because it makes a “non-expressive use” of that work.<sup>66</sup> (Abraham Drassinower, Alain Strowel, Carys Craig, and Cheryl Foong have made similar arguments in other countries.)<sup>67</sup> That position has been repeated in Congressional testimony,<sup>68</sup> and in comments by OpenAI in a U.S. Patent and Trademark Office inquiry.<sup>69</sup> However, the meaning of “non-expressive use” is not entirely fixed or stable, because the term “expression” has been and is being used in more than one sense. Much of the rhetorical force of the “non-expressive use” argument is generated by switching back and forth between these different senses, lending an air of authority to statements about one meaning of “non-

---

66. See Bracha, *supra* note 1, at 25; Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 304 (“[T]he fair use cases dealing with copy-reliant technology reflect the view that non-expressive uses of copyrighted works should be treated as fair use.”); *id.* at 307 (“There Is No Machine Learning Exception to the Principle of Non-expressive Use.”). Matthew Sag coined the term and concept of “non-expressive use” in an earlier article. See Matthew Sag, *Copyright and Copy-Reliant Technology*, 103 NW. U.L. REV. 1607, 1068 (2009). James Grimmelman and Benjamin Sobel have investigated non-expressive use without making the argument that such use never requires the permission of copyright owners. See James Grimmelman, *supra* note 65, at 661–65; Sobel, *supra* note 31, at 51–57. For other good criticism of “non-expressive use” as fair use, see David W. Opderbeck, *Copyright in AI Training Data: A Human-Centered Approach*, 76 OKLA. L. REV. 951, 975–85 (2024).

67. See sources cited *supra* note 1.

68. See *Understanding Generative Artificial Intelligence and Its Relationship To Copyright: Hearing on Artificial Intelligence and Intellectual Property: Part I—Interoperability of AI and Copyright Law Before Subcomm. on Cts, Intell. Prop., and the Internet of the H. Comm. on the Judiciary*, 118th Cong., at 15 (2023) (statement of Christopher Callison-Burch, Ph.D.).

69. See OpenAI Comments 2019, *supra* note 17, at 5–7.

expressive” that in truth should only be enjoyed by statements about a different meaning.

Here, I want to concentrate on three possible meanings of “expression,” and therefore three meanings of “non-expressive use.” The three meanings of “expression” are as follows, and to keep them distinct, I am going to introduce some qualifying adjectives:

(1) “constitutive expression” means the individual expressive choices that the author of a work makes, and that are not disqualified from protection because they are facts, historical or scientific theories, or functional elements.

(2) “actionable expression” means a group of expressive choices made by an author that are complex enough to be copyrightable as a work of authorship, and that if copied in another work, would render that other work substantially similar to the first work.

(3) “felt expression” means the aesthetic or hedonic reaction that someone who is capable of such reactions has when they experience a work of authorship.

Given those definitions of “expression,” we can consider three definitions of “non-expressive use”:

(1) “non (constitutive) expressive use” is a use of a copyrighted work that is indifferent to the expressive choices made by the author of the work.

(2) “non (actionable) expressive use” is a use of a copyrighted work that does not produce any other work that contains enough of the expressive choices of the first work to infringe it.

(3) “non (felt) expressive use” is a use of a copyrighted work that produces no aesthetic or hedonic reaction to the expressive choices of that work because it does not expose that work to anyone capable of having such a reaction.

I will first consider each of these meanings of “expression” and “non-expressive use” and how they relate to each other. I will then consider how each of them applies to human authors and generative AI tools.

## 1. The Three Meanings of Expression and “Non-Expressive Use”

### *a. Constitutive Expression*

In *Harper & Row, Publishers, Inc. v. Nation Enterprises*, the Supreme Court stated that “[t]he copyright [in a work] is limited to those aspects of the work—termed ‘expression’—that display the stamp of the author’s originality.”<sup>70</sup> In that sense, “expression” is any aspect or element of a work that is a protectable choice of the

---

70. 471 U.S. 539, 547 (1985).

author, and not an excluded element, such as a fact,<sup>71</sup> an historical hypothesis,<sup>72</sup> a scientific theory, or a functional element. Any choice exhibiting a “modicum of creativity” will do, although courts sometimes use the “scenes à faire” doctrine to heighten the creativity standard,<sup>73</sup> or the merger doctrine to deny protection when there is only one or a very few ways of expressing an unprotectable element.<sup>74</sup> Each of those qualifying choices is a constitutive part of what copyright protects in a work, hence the name “constitutive expression.” Somewhat confusingly, in that context, unprotectable elements like historical hypotheses or scientific theories are sometimes called “ideas.”<sup>75</sup>

The “stamp of author’s originality” may not always be the result of the author’s deliberate, conscious choice. Sometimes statements about constitutive expression assume that it is the result of an author’s fixed personality, instead of a matter of choice. Thus, for example, in *Bleistein v. Donaldson Lithographing Co.*, Justice Holmes writes that a work, which he calls a “copy,” is the personal reaction of an individual upon nature,” and then proceeds to comment that “[p]ersonality always contains something unique,” and that it is that unique reaction of personality to the world that is copyrightable.<sup>76</sup> For present purposes, however, as long as the constitutive expression comes from the author, the degree of deliberate choice does not matter.

Take for example, the story *Fox 8*, by George Saunders (a named plaintiff in one of the putative class action lawsuits against OpenAI and Microsoft).<sup>77</sup> The text exhibits a lot of expressive choices, including misspellings and grammatical errors to remind the reader that the narrator is supposed to be a fox:

We do not trik Chikens! We are very open and honest with Chikens! With Chikens, we have a Super Fare Deel, which is: they make the eggs, we take the eggs, they make more eggs. And sometimes may even eat a live Chiken, shud that Chiken consent to be eaten by us, threw faling to run away upon are approche, after she has been looking for feed in a stump.<sup>78</sup>

All of those word and spelling and grammatical choices are expressive choices that Saunders makes. All of the expressive choices are constitutive expression, because each of them is a qualifying part of an aggregate that is protected by copyright—an

71. See *Feist Publ’ns., Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 346 (1991); 17 U.S.C. § 102(a) (providing that copyright subsists in “original works of authorship”).

72. See, e.g., *Hoehling v. Universal City Studios, Inc.*, 618 F.2d 972, 976 (2d Cir. 1980).

73. For the “modicum of creativity” standard, see *Feist*, 499 U.S. at 346; for “scenes à faire,” see, e.g., *Hoehling*, 618 F.2d at 979.

74. See, e.g., *Morrissey v. Procter & Gamble Co.*, 379 F.2d 675, 678–79 (1st Cir. 1967).

75. See, e.g., *CCC Info. Serv. v. Maclean Hunter Mkt. Rep.*, 44 F.3d 61 (1994) (referring to “ideas that undertake to advance the understanding of phenomena or the solution of problems, such as the identification of the symptoms that are the most useful in identifying the presence of a particular disease”).

76. 188 U.S. 239, 250 (1903).

77. See First Consolidated Complaint, *Authors Guild v. Open AI, Inc.*, No. 23-cv-08292 (S.D.N.Y. Feb. 25, 2024).

78. GEORGE SAUNDERS, *FOX 8: A STORY 6* (2018). I bought a copy of this book in a used bookstore, Powell’s, in Chicago.

original work of authorship—and none of them is disqualified as a fact, historical hypothesis, scientific theory, or functional element. Expressive choices in literary works can be about particular words or phrases, or about structural matters like the use of a particular rhyming pattern in a poem. Similarly, in visual art expressive choices can be about particular colors or lines, or about structural matters like the use of symmetrical or asymmetrical composition.

Very confusingly, in standard copyright lingo each of the constitutive expressive choices can also be called, in isolation, an “idea,” as can an insufficiently complex group of such choices. That is because copyright law does not protect creations that are the product of only a few creative choices, like a graphic composition of “evenly spaced white circles on a purple background,”<sup>79</sup> or a short phrase like “holy origin” or “to gain love.”<sup>80</sup> In isolation, those creations are not complex enough to be recognized as original works of authorship.<sup>81</sup> It is also because a more complex work, which does qualify for copyright protection, is not infringed by another work that copies only one or a few of the first work’s expressive choices. In such a case—as we will explore further below—the second work is said to copy only “ideas” from the first, particularly when the copying is “non-literal.”<sup>82</sup> Obviously, that meaning of “idea” is different than “scientific theory” or “historical hypothesis.”

When the term “idea” is used to refer to an insufficiently complex expressive choice, then the term “expression” refers to an aggregation of expressive choices that is sufficiently complex to qualify for copyright protection, and that, if copied from one work into another, makes that second work infringe the first (unless another exception applies, which could include fair use). I am going to call that kind of expression “actionable expression,” to distinguish it from the “constitutive expression” of individual expressive choices. When constitutive expression gets aggregated into actionable expression, the constitutive expressive choices do get protection as part of the whole.

Ordinarily, when we human beings read a book like *Fox 8*, we care about those constitutive expressive choices, and about the way they are combined to create actionable expression in works of authorship. We are interested in the aesthetic experience they produce, and we are also interested in what we learn from them. Sometimes, however, we are indifferent to the expressive choices, because we want

---

79. See U.S. COPYRIGHT OFF., COMPENDIUM OF U.S. COPYRIGHT OFFICE PRACTICES § 906.1 Common Geometric Shapes (2021) [hereinafter, COPYRIGHT COMPENDIUM] (a claim of copyright in “a picture with a purple background and evenly spaced white circles” will not be registered, because “the combination of the purple rectangle and the standard symmetrical arrangement of the white circles does not contain a sufficient amount of creative expression to warrant registration.”).

80. See *Arica Inst., Inc. v. Palmer*, 770 F. Supp. 188, 191 (1991), *affirmed*, *Arica Inst., Inc. v. Palmer*, 970 F.2d 1067, 1072 (1992); cf. COPYRIGHT COMPENDIUM, *supra* note 79, § 312.2 The Originality Requirement for Compilations (stating that the Copyright Office “generally will not register a compilation containing only two or three elements, because the selection is necessarily *de minimis*.”).

81. On the formulation of a standard of minimum complexity, or size, separate from the “modicum of creativity” standard for originality, see Justin Hughes, *Size Matters (or Should) in Copyright Law*, 74 *FORDHAM L. REV.* 575 (2005).

82. See *infra* text accompanying notes 95–99.



other information from the work. If we are using the work in a way that is indifferent to the expressive choices made in a work, then maybe our use should count as a fair use, because we are not making use of any of the elements that copyright protects in a work.

For example, it is standard black-letter copyright law to say that copyright does not protect functional elements in a work, or facts conveyed in a work.<sup>83</sup> Can we make copies of a work in order to access the functional elements contained in it? In some contexts, courts have said yes. Thus, for example, in *Sega Enterprises Ltd. v. Accolade, Inc.*,<sup>84</sup> the Ninth Circuit held that it was fair use for Accolade, an independent video game developer, to decompile the object code of three video games produced by Sega and produce source code copies, so that it could understand how the game cartridge communicated with Sega's game console and make a cartridge compatible with the Sega console.<sup>85</sup> The commands passed back and forth between the game cartridge and the console were functional code not protected by copyright, and the *Sega* court undoubtedly also had the sense that it was not within copyright law's purpose to enforce a "tying" arrangement, requiring all purchasers of Sega consoles to use only Sega game cartridges.<sup>86</sup> Significantly, the *Sega* court rejected Accolade's broader argument that "intermediate copying does not infringe the exclusive rights granted to copyright owners in § 106 of the Copyright Act unless the end product of the copying is substantially similar to the copyrighted work."<sup>87</sup> The narrower winning argument was that copying the game cartridge's code for the purpose of obtaining the technical specifications for communication with the game console was fair use. That can be seen as an instance of non (constitutive) expressive use.

The fair use record with unprotected elements is more mixed, however, in part because courts are sometimes sensitive to protecting an author's investment in labor, even when they profess to be protecting creativity. In *Author's Guild v. Google*, for example, Judge Leval suggests that the factual or fictional character of a work should not be given much weight in fair use analysis, because copyright does protect the "author's manner of expressing . . . facts."<sup>88</sup> He continues: "Those who report the news undoubtedly create factual works. It cannot seriously be argued that, for that reason, others may freely copy and re-disseminate news reports."<sup>89</sup> Whether that is actually

---

83. See, e.g., *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 346 (1991); 17 U.S.C. § 102(b).

84. 977 F.2d 1510 (9th Cir. 1992).

85. See *id.* at 1520–28; see also *Sony Computer Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 601 (9th Cir. 2000) (holding that it was fair use to copy a game console's software for purposes of understanding its noncopyrightable functional elements and then using them in emulation software that could enable compatible games to be run on a personal computer).

86. *Sega*, 977 F.2d at 1523–24 ("an attempt to monopolize the market by making it impossible for others to compete runs counter to the statutory purpose of promoting creative expression and cannot constitute a strong equitable basis for resisting the invocation of the fair use doctrine"); cf. *Chamberlain Grp. v. Skylink Tech., Inc.*, 381 F.3d 1178 (2004) (holding that the anti-circumvention provisions of § 1201 cannot be used to prevent an independent remote control manufacturer from making a remote control compatible with a garage door opener made by the plaintiff).

87. *Sega*, 977 F.2d at 1517.

88. 804 F.3d 202, 220 (2d Cir. 2015).

89. *Id.*

because copyright is protecting the creative expression of the facts, or to some degree protecting the labor in collecting the facts, it is clear that you cannot justify your copying and redistribution of a news article on the ground that you just want to provide access to the uncopyrightable facts.

Courts have developed the concept of “transformative use,” rather than “non-expressive use,” and in cases involving factual content, one can understand how transformative use might more flexibly respect the investment in the used work and the purposes of the author and the user. Take, for example, Mark Lemley and Brian Casey’s example of the use of images that include stop signs to train a computer vision model to recognize stop signs.<sup>90</sup> Suppose that a company developing such a model scraped the internet for such images, and found 50,000 of them. Suppose further that those images were all travel photos posted on sites like Flickr or shared on social media services like Instagram, and that the stop signs appeared coincidentally in the background. Here there is a reasonably strong case for transformative use. The original purpose of the photographers had nothing to do with training a computer vision model, and that use has no effect on any market originally contemplated by the photographers.

Lemley and Casey make a case, not for transformative use, but for “non-expressive use.” The computer vision company could argue that it “wants photos of stop signs so [that its computer vision system] can learn to recognize stop signs, not because of the artistic choices you made in lighting or composing your photo.”<sup>91</sup> However, suppose that one company finds and uses a large set of photographs of stop signs taken by another company expressly for the purpose of training a computer vision model. Is that still a “non-expressive use”? I think not, and not a “transformative use” either. As Lemley and Casey acknowledge in a footnote, “[t]here is a sense in which the creative choices matter even here. The AI is likely to want to see photos of stop signs in a variety of lights, angles, and conditions to train better.”<sup>92</sup> Courts will tend to broaden “expressive choices” towards what Canadian courts would call “skill and judgment,”<sup>93</sup> to protect the investment of one company against a competing company in the same market.

#### *b. Actionable Expression*

As I have suggested above, the term “expression” can mean an element or feature of a work that counts towards copyright protection, because it is the result of a

---

90. See Lemley & Casey, *supra* note 1, at 749. The same could be true of any use of images to train a model to recognize objects, whether the goal would be to recognize a category of objects, such as stop signs, or specific objects, such as individual human beings through facial recognition. See Sobel, *supra* note 31, at 67–68 (stop signs); Enrico Bonadio, Plamen Dinev & Luke McDonagh, *Can Artificial Intelligence Infringe Copyright? Some Reflections*, in RESEARCH HANDBOOK ON INTELLECTUAL PROPERTY AND ARTIFICIAL INTELLIGENCE 245, 248 (Abbott ed., 2022) (facial recognition).

91. Lemley & Casey, *supra* note 1, at 749.

92. *Id.* n.34.

93. See *CCH Canadian Ltd. et al. v. Law Soc’y of Upper Canada*, 2004 SCC 13 [2004] 317 N.R. 107 (Can.).

creative choice by the author, rather than, say, a fact or a functional element. That is what I have called “constitutive expression.” Alternatively, it can mean an aggregation of such constitutive choices that is sufficiently complex to qualify for copyright protection, and that, if copied from one work into another, makes that second work infringe the first. I am going to call that this second kind of expression, the complex aggregation kind, “actionable expression,” to distinguish it from the “constitutive expression” of individual expressive choices.

The use of “expression” to mean “actionable expression” dates back at least to Judge Learned Hand’s exploration of “substantial similarity” in the context of non-literal copying. In *Nichols v. Universal Pictures Corporation*, for example, Hand considered whether the movie *The Cohens and the Kellys* infringed the play *Abie’s Irish Rose*.<sup>94</sup> There was no question that the producers of *The Cohens and the Kellys* knew of and were inspired by *Abie’s Irish Rose*. They had tried to obtain the rights to make the play into a movie.<sup>95</sup> When that attempt failed, they copied some of the expressive choices that Anne Nichols, the author of *Abie’s Irish Rose*, had made. For example, they copied Nichols’ choice to make the story about a boy and a girl, one Jewish and the other Irish Catholic, living in New York City, who fall in love. They also copied Nichols’ choice to have the boy and the girl initially keep their romance secret, because each of their families was opposed to romance and marriage outside of their respective religions and ethnicities. And they copied Nichols’ choice to have the boy and the girl get married, and to have their fathers quarrel but in the end make up with each other.

However, in the terms that Learned Hand chose to structure the analysis of substantial similarity, each of those choices counts only as an “idea,” and “ideas” can be copied freely into a new work. It is only when a sufficient number of those ideas—or constitutive expressive choices—are copied that the appropriation becomes copying of “actionable expression” that gives rise to an infringement claim. Judge Hand held that the choices made by Anne Nichols in *Abie’s Irish Rose* that were copied in *The Cohens and the Kellys* were not numerous and complex enough to amount to actionable expression.<sup>96</sup> That is to say, *The Cohens and the Kellys* makes a “non (actionable) expressive use” of *Abie’s Irish Rose*. More generally, a non (actionable) expressive use is any use of the expressive choices in a work that does not make a new work, or a performance or display, substantially similar to the first work.<sup>97</sup> In the case of “literal” copying, courts typically use the standard of “de minimis use,” rather than

---

94. 45 F.2d 119 (2d Cir. 1930).

95. See *Nichols v. Universal Pictures Corp.*, 34 F.2d 145, 150 (S.D.N.Y. 1929) (“The record discloses that in 1925 the defendant, the Universal Pictures Corporation, tried to purchase ‘Abie’s Irish Rose’ motion picture rights, and that, when the scenario of ‘The Cohens and Kellys’ was being written, its authors ‘studied’ the synopsis of ‘Abie’s Irish Rose.’”).

96. See *id.* at 121–22.

97. In the *Nichols* opinion, Hand also used the same idea of a sufficiently complex aggregate to furnish a test for copyrightability: A character could be subject to copyright protection, but only if it was sufficiently developed and distinct. *Id.* at 121–22.

the standard of “substantial similarity,” to determine whether there is enough copying to constitute infringement.<sup>98</sup>

*c. Felt Expression*

Felt expression is the aesthetic or hedonic reaction that a human being has when they experience a work of authorship. This is territory that James Grimmelmann began exploring in his 2016 article “Copyright for Literate Robots,”<sup>99</sup> but I want to explore further here. For many readers, the easiest path to understanding the appearance of felt expression in copyright law may be through well-known formulations of the standard of “improper appropriation” or “substantial similarity.” Jerome Frank’s famous formulation in a case involving music asks “whether defendant took from plaintiff’s works so much of *what is pleasing to the ears of lay listeners . . .* that defendant wrongfully appropriated something which belongs to the plaintiff.”<sup>100</sup> Learned Hand’s equally well-known framing in a fabric design case asks whether “the ordinary observer, unless he set out to detect the disparities [between two works], would be disposed to overlook them, and regard *their aesthetic appeal* as the same.”<sup>101</sup> This could suggest that what is important in works of authorship is the kind of aesthetic or hedonic reactions they produce in the human beings who form their audiences. Thus, Oren Bracha speaks of expression as involving “human enjoyment of the expressive value of a work” and “expressive enjoyment by any ear or eye.”<sup>102</sup> Correlatively, the important thing about those who experience works is that they can be pleased or displeased by a work, and be aesthetically attracted or repulsed.<sup>103</sup>

If the only thing that is important about a work of authorship, and that makes it valuable, is that it produces certain hedonic reactions, then any use of a work in which no such reaction is produced could count as a non (felt) expressive use. This could account, for example, for some of an intuitive sense that it is fair to make an archival or backup copy of a work of which you already have an authorized copy.<sup>104</sup>

98. See, e.g., *Ringgold v. Black Ent. Television, Inc.*, 126 F.3d 70, 74–77 (2d Cir. 1997).

99. See Grimmelmann, *supra* note 65; cf. Christopher Buccafusco, *A Theory of Copyright Authorship*, 102 VA. L. REV. 1229, 1260–75 (2016) (arguing that authorship requires an intent to produce a mental effect in an audience).

100. *Arnstein v. Porter*, 154 F.2d 464, 473 (2d Cir. 1946) (emphasis added).

101. *Peter Pan Fabrics, Inc. v. Martin Weiner Corp.*, 274 F.2d 487, 489 (2d Cir. 1960) (emphasis added).

102. Bracha, *supra* note 1, at 28.

103. Others use somewhat different language and concepts. Abraham Drassinower, followed by Alain Strowel and Carys Craig, speaks of “communicative” and “non-communicative” uses. See DRASSINOWER, *supra* note 1; Strowel, *supra* note 1; Craig, *supra* note 1. Cheryl Foong draws on both “non-expressive use” and “non-communicative use.” See Foong, *supra* note 1. Somewhat further afield in philosophy, John Searle distinguishes between computer programs, which cannot have intentionality, and human brains and potentially other machines, which can. See John R. Searle, *Minds, Brains, and Programs*, 3 BEHAV. AND BRAIN SCI. 417 (1980). The relationship between these concepts is worth investigating—in a different article.

104. Compare 17 U.S.C. § 108(b) (allowing a library or archive to make a copy of an unpublished work for preservation or security purposes), and § 108(c) (allowing a library or archive to make a copy of a

The archival or backup copy, locked in a drawer or on a cloud server somewhere, is not satisfying any current or additional demand for the hedonic reaction that it could produce. Rather, it is stored away for use only if the authorized copy is damaged or destroyed. Generalizing on that intuition, Oren Bracha states that “[m]aking a new physical copy when the expression embodied in it will be experienced by no one is not any more relevant for copyright than using an existing copy as a doorstep.”<sup>105</sup>

## 2. Human Authors and Generative Machines: Distinguished Only By Felt Expression, and That Should Not Make a Legal Difference

Having provided definitions for three different types of expression and three different types of non-expressive use, let us turn to the comparison of human authors and generative AI tools. Do computers make some kind of non-expressive use of works that human authors do not, and that would therefore entitle generative AI systems to a fair use privilege that human authors do not enjoy?

First, it is important to note that Matthew Sag and others seem to mix the three different senses of “expression” that I have tried to distinguish.<sup>106</sup> For example, Sag states that “the rationale for allowing for-profit and academic researchers to derive valuable data from other people’s copyrighted works is a necessary implication of the fundamental distinction between protectable original expression and unprotectable facts, ideas, abstractions, and functional elements.”<sup>107</sup> The references to unprotectable “facts” and “functional elements” suggest that the statement is about constitutive expression. The reference to “ideas” is ambiguous, as it might be a reference a scientific theory like “ $E = mc^2$ ,”<sup>108</sup> or to an expressive choice that is not complex enough to constitute actionable expression.<sup>109</sup> The reference to “abstractions” suggests that the statement may be about actionable expression.<sup>110</sup> Elsewhere in the same article, Sag asserts that “machine learning models . . . qualify as non-expressive use so long as the outputs are not substantially similar to any particular original expression in the training data.”<sup>111</sup> Here, the reference to “substantial similarity” suggests that the statement is about actionable expression. In two other passages, Sag states that “technical acts of copying that do not communicate the original expression to a new audience do not interfere with the interest in original expression that copyright is

---

published work for the purpose of replacement of a copy or phonorecord that is damaged, deteriorating, lost, or stolen, if a copy cannot be found at a reasonable price), with § 117(a)(2) (allowing an owner of an authorized copy of a computer program to make another copy or adaptation provided “that such new copy or adaptation is for archival purposes only and that all archival copies are destroyed in the event that continued possession of the computer program should cease to be rightful”).

105. Bracha, *supra* note 1, at 23.

106. Of course, this may just mean that they would reject the distinctions I have drawn!

107. Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 305.

108. See *supra* text accompanying note 72.

109. See *supra* text accompanying notes 79–84, 95–99.

110. See *supra* text accompanying notes 79–84, 95–99.

111. Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 309.

designed to protect,”<sup>112</sup> and “what matters is whether the original expression of the authors of works in the training data is communicated to a new public.”<sup>113</sup> In another article, Sag defines “non-expressive use” as “any act of reproduction that is not intended to enable human enjoyment, appreciation, or comprehension of the copied expression as expression.”<sup>114</sup> Those sound like statements about felt expression, which exists only in the minds of audience members who have aesthetic reactions.

If it makes sense to separate the three uses of “expression,” how do human authors and generative machines compare with regard to each? It turns out that human authors and generative systems are both intensely interested in and make uses of constitutive expression, so that does not distinguish them. Neither human authors nor generative systems typically copy actionable expression, so that does not distinguish them either. However, human authors do make use of felt expression, while generative AI systems typically do not. I will argue that that should not make a legal difference.

*a. Machines and Humans Both Use Constitutive Expression*

In early motions to dismiss the copyright infringement lawsuits that have been filed against them, generative AI developers OpenAI and Bloomberg have argued that their generative models are only learning “facts” or “information” from the works on which they are trained, plus the most basic elements of language and linguistic structure: “grammar[] and syntax”<sup>115</sup> or “grammar[] [and] vocabulary.”<sup>116</sup> That is not the case. Developers of generative AI tools are completely interested in creating tools and services that produce text and images that appeal to the aesthetic sensibilities of users, not just text and images that would recognizably but boringly or awkwardly portray some object or fact.<sup>117</sup> That is why they use training works in a wide variety of linguistic and visual styles. That is also why the latent spaces in their models contain clusters representing, not only objects depicted in images or facts stated in prose, but also a wide variety of stylistic characteristics derived from the expressive choices of the authors of the training works.

---

112. *Id.* at 306.

113. *Id.* at 309.

114. Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC'Y U.S.A. 291, 301 (2019).

115. OpenAI's MTD, *supra* note 17, at 12.

116. Memorandum of Law in Support of Bloomberg L.P. and Bloomberg Finance L.P.'s Motion to Dismiss First Amended Class Action Complaint at 8 (Jan. 24, 2024), *Huckabee v. Meta Platforms, Inc.*, No. 23-cv-09152 (S.D.N.Y. Jan. 24, 2024).

117. See Sobel, *supra* note 31, at 57 (2017) (“[M]achine learning gives computers the ability to derive valuable information from the way authors express ideas. Instead of merely deriving facts about a work, they may be able to glean value from a work's expressive aspects.”); *id.* at 69 (noting that generative AI developers “sought to make use of authors’ varied and rich expression of ideas,” which “is the essence of copyrightable subject matter.”).

Take the case of text-to-image model training. In comments submitted to the U.S. Copyright Office's inquiry on Artificial Intelligence and Copyright, OpenAI states,

When a model is exposed to a large array of images labeled with the word “cup”, it learns what visual elements constitute the concept of “cup-ness[.]” It does this not by compiling an internal database of training images, but rather by abstracting the factual metadata that correlates to the idea of “cup.” . . . The factual metadata and fundamental information that AI models learn from training data are not protected by copyright law.<sup>118</sup>

Matthew Sag uses the same example of cup images and suggests that what Generative AI models do is “encode fundamental relationships between pixels that are more likely in pictures with coffee cups than without.”<sup>119</sup> He then provides an example of an image of a cup generated by Stable Diffusion next to some images of cups used to train Stable Diffusion. The cup in the generated image has “generic features”—it is “round . . . [a]ppears to be made of white ceramic, [and] has a small single handle . . . [but is otherwise] not substantially similar to any particular image from the training data.”<sup>120</sup> The implicit suggestion is that the training set contains images labeled only with names of objects that appear in them; that the generative model just averages out all images labeled “cup” into a generic cup; and that a text-to-image service built on that model only produces images of generic cups. It would in theory be possible to construct such a training set, and to construct a model from that training set that contained only a representation of a generic cup, averaged from all cups portrayed in the training images. If that were what generative AI developers were doing, text-to-image generative models would have some claim to be learning only about real-world, uncopyrighted objects, thus making essentially non (constitutive) expressive use of the image-text pairs in the training set.

However, that is not what any text-image dataset collector or generative AI developer does, and that is not what they would want to do. Consider, first, the matter of training works—for example, the LAION image-text pair datasets on which

---

118. OpenAI, Comments of OpenAI on U.S. Copyright Office's Notice of Inquiry and Request for Comment [Docket No. 2023-06], at 11–12 (Oct. 30, 2023) [hereinafter, OpenAI Comments 2023], <https://www.regulations.gov/comment/COLC-2023-0006-8906> [<https://perma.cc/KY34-XQXA>] [<https://web.archive.org/web/20240914125018/https://www.regulations.gov/comment/COLC-2023-0006-8906>]. Meta made similar statements in its comments in the Copyright Office inquiry. See Meta Platforms, Inc., Comments of Meta on U.S. Copyright Office's Notice of Inquiry on Artificial Intelligence & Copyright [Docket No. 2023–06], at 7 (Oct. 30, 2023) [hereinafter, Meta Comments] <https://www.regulations.gov/comment/COLC-2023-0006-9027> [<https://perma.cc/7K69-PTD9>] [<https://web.archive.org/web/20240914125134/https://www.regulations.gov/comment/COLC-2023-0006-9027>] (“The goal of such models is to simply extract non-expressive facts and statistics from training data (e.g., what characteristics typify a cat) and use them to generate new content (e.g., an entirely new picture of a cat).”).

119. Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 319; see Sag, *Fairness and Fair Use in Generative AI*, *supra* note 1, at 1908–09 (discussing the same example of cups).

120. Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 319.

Stability AI trained Stable Diffusion. The datasets consist of the URL where each image can be found, and the accompanying alt-text tag.<sup>121</sup>

Stability AI downloaded each of the images, and trained Stable Diffusion on each image and its corresponding text.<sup>122</sup> The text tags in the LAION datasets include terms describing every conceivable stylistic element in images—everything you would ever want to describe the expressive choices of image creators. Think of names of colors, and terms like “vivid” and “muted”. Shape terms, like “organic” and “geometric.” Texture terms, from “smooth” to “rich” to “rough.” Movement, “balance,” “imbalance,” “harmony,” “contrast”—they are all there. So are genres, and schools of art are also all there, from “baroque” and “rococo” to “impressionist,” “expressionist,” “Fauve,” “cubist,” “Dada,” “photorealist,” and so on. If that were not enough, there are thousands of names of the artists who created the images, and not just long-dead artists whose works are no longer under copyright, but artists who are still very much living and working.<sup>123</sup> It is also telling that the single most frequent word in the LAION 400M dataset (excluding prepositions and articles) is “Stock”—text-to-image tools are being trained on the stock images that their outputs are most likely to replace.<sup>124</sup>

Those tags are not left in the training dataset out of laziness. They are essential to the goal of the generative AI developers, which is to learn what visual patterns are associated with those terms. Some of those will be pretty simple: The term “purple” will likely come be associated with a range of corresponding hues. Others will be much more complex: The name of a particular artist may be associated with a whole series of visual patterns that cover texture, composition, color palette, subject matter, and so on. Thus, if 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. Developers want text-image models to understand the visual patterns associated with every term that could be used to describe an artist’s work, so that the models can respond to user requests to create images embodying those choices.<sup>125</sup>

---

121. See Romain Beaumont, LAION-5B: A New Era of Open Large-Scale Multi-Modal Datasets, Laion (Mar. 31, 2022), <https://laion.ai/blog/laion-5b/> [<https://perma.cc/ZF4T-4F7P>] [<https://web.archive.org/web/20240914125219/https://laion.ai/blog/laion-5b/>].

122. See Stable Diffusion v1 Model Card, GITHUB, [https://github.com/CompVis/stable-diffusion/blob/main/Stable\\_Diffusion\\_v1\\_Model\\_Card.md#training](https://github.com/CompVis/stable-diffusion/blob/main/Stable_Diffusion_v1_Model_Card.md#training) [<https://perma.cc/S83T-KM6K>] [[https://web.archive.org/web/20240919022150/https://github.com/CompVis/stable-diffusion/blob/main/Stable\\_Diffusion\\_v1\\_Model\\_Card.md#training](https://web.archive.org/web/20240919022150/https://github.com/CompVis/stable-diffusion/blob/main/Stable_Diffusion_v1_Model_Card.md#training)] (stating that Stable Diffusion was trained on “LAION 5B and subsets thereof”) (last visited Sept. 18, 2024).

123. See Robert Brauneis & Aneri Girishbhai Patel, LAION 400M Text Tag Word Frequency List, <https://ln5.sync.com/dl/4b125ab10/f5kbcbs2-ibzyh2fg-iug9k55j-89ngixs7> [<https://perma.cc/AQS3-XRQ7>] [<https://web.archive.org/web/20241004174530/https://ln5.sync.com/dl/4b125ab10/f5kbcbs2-ibzyh2fg-iug9k55j-89ngixs7/view/text/23874584460008>] (last visited Oct. 4, 2024).

124. See *id.*

125. Importantly, state-of-the-art generative modeling techniques do not limit themselves to the terms found in the text tags of text-image pairs. Rather, they incorporate separately trained large language models, which can provide more context for terms that appear in the text tags of text-image pairs, and thus interpret prompts better. See Chitwan Saharia et al., *Photorealistic Text-to-Image Diffusion Models with Deep*



The same is true of the images in the LAION training sets. If what we wanted to do with such a set was to train a computer vision model, it would be important to limit the images in the set to those that were or looked like standard visible-spectrum photographs. When we deployed the model within a computer vision system, the system would be receiving images from a standard camera, and would have to identify objects in such images. However, the LAION training sets are in no way limited to such images. Generative AI developers keep all of the artistic, highly-stylized images when they use the sets to train their models, because they want their models to be able to produce images in a wide variety of styles, containing a wide variety of expressive choices.

Consider, for example, the seven image-text pairs from the LAION 5B dataset in Table 1 below. Notice that the text tags all include the word “cup,” but they also include other words. Go retrieve the images using the provided Perma links, and look at them. They do indeed portray cups, in a very wide variety of styles, sometimes barely visibly, as in the Diane Leonard painting, sometimes whimsically, as in the Stanley Morrison illustration, sometimes conceptually, as in the Jacek Yerka sketch in which no cup is present, just the water in a cutaway of a cup-shaped moat around a castle.

Image URL	Tag
<a href="https://i.pinimg.com/736x/23/20/ca/2320caf3176e52d46769bec1ace33e31.jpg">https://i.pinimg.com/736x/23/20/ca/2320caf3176e52d46769bec1ace33e31.jpg</a> [ <a href="https://perma.cc/T4C3-CSEY">https://perma.cc/T4C3-CSEY</a> ]	A Cup of Tea. Oil on Board. Diane Leonard. Impressionist.
<a href="https://render.fineartamerica.com/images/rendered/search/print/images/artworkimages/medium/1/coffee-dragon-stanley-morrison.jpg">https://render.fineartamerica.com/images/rendered/search/print/images/artworkimages/medium/1/coffee-dragon-stanley-morrison.jpg</a> [ <a href="https://perma.cc/2ZFW-8ZEQ">https://perma.cc/2ZFW-8ZEQ</a> ]	Cup Digital Art Coffee Dragon by Stanley Morrison
<a href="https://www.absolutearts.com/portfolio3/j/jimlively/A_Second_Cup_of_Coffee-1357674081st.jpg">https://www.absolutearts.com/portfolio3/j/jimlively/A_Second_Cup_of_Coffee-1357674081st.jpg</a> [ <a href="https://perma.cc/24ZM-UEVM">https://perma.cc/24ZM-UEVM</a> ]	Jim Lively Artwork A Second Cup of Coffee
<a href="https://images.fineartamerica.com/images-medium-large-5/yellow-morning-cup-lutz-baar.jpg">https://images.fineartamerica.com/images-medium-large-5/yellow-morning-cup-lutz-baar.jpg</a> [ <a href="https://perma.cc/2NZF-L5XA">https://perma.cc/2NZF-L5XA</a> ]	Yellow Morning Cup Painting by Lutz Baar
<a href="https://render.fineartamerica.com/images/rendered/search/print/images/artworkimages/medium/1/coffee-cup-painting-robert-joyner.jpg">https://render.fineartamerica.com/images/rendered/search/print/images/artworkimages/medium/1/coffee-cup-painting-robert-joyner.jpg</a> [ <a href="https://perma.cc/CMG2-LFQA">https://perma.cc/CMG2-LFQA</a> ]	Painting – Coffee Cup Painting by Robert Joyner
<a href="https://mayhemandmuse.com/wp-content/uploads/2012/08/A-Jacek-Yerka-surrealism-sketch-of-a-castle-in-a-tea-cup.jpg">https://mayhemandmuse.com/wp-content/uploads/2012/08/A-Jacek-Yerka-surrealism-sketch-of-a-castle-in-a-tea-cup.jpg</a> [ <a href="https://perma.cc/LP25-AAT8?type=image">https://perma.cc/LP25-AAT8?type=image</a> ]	A Jacek Yerka surrealism sketch of a castle in a tea cup
<a href="https://mydekel.files.wordpress.com/2015/04/guitarcoffee.jpg?w=1250&amp;h=600&amp;crop=1">https://mydekel.files.wordpress.com/2015/04/guitarcoffee.jpg?w=1250&amp;h=600&amp;crop=1</a> [ <a href="https://perma.cc/2PEC-BVXX">https://perma.cc/2PEC-BVXX</a> ]	Guitar and Coffee Cup digital art from photographs (c)2015 Michael Dickel

**Table 1: Seven Image-Text Pairs from LAION 5B<sup>126</sup>**

126. I found these images by doing a search for “cup” on “*Have I Been Trained?*,” an online tool that allows searchable access to the LAION-5B dataset, available at <https://spawning.ai/have-i-been-trained> [<https://perma.cc/V94M-72G7>] [<https://web.archive.org/web/20241004174736/https://spawning.ai/have-i-been-trained>] (last visited Mar. 4, 2024). I generated Perma links for the pages on which the images were found, so that they could continue to be viewed in their original context, even if they become no longer available at the URLs provided in the LAION 5B dataset.

Nonetheless, those images, and thousands of others like them, are all included in the LAION 5B dataset, and were all used to train Stable Diffusion.<sup>127</sup> What patterns is the Stable Diffusion machine learning algorithm finding in those images? It is, among other things, finding patterns associated with how particular artists choose to portray cups, including named artists like Diane Leonard, Stanley Morrison, Lutz Baar, Jim Lively, Robert Joyner, and Jacek Yerka (all of whom, by the way, are as of this writing still alive and creating images). Those patterns are the expressive choices that those artists have made.

The images are encoded and stored in a “latent space” that is impossible for us to interpret and experience directly. However, experiments in exploring small regions of latent space in the image models of generative AI systems demonstrate that the image embeddings are indeed storing style information—expressive choices made by image creators—as well as information about the objects represented in those images.<sup>128</sup> Take a look, for example, a couple of the animated gifs provided with a post titled “A walk through latent space with Stable Diffusion” on the Keras website.<sup>129</sup> One part of one gif shows the decoded images associated with similar image embeddings, most of which look like a panda depicted in a stained glass window. The panda seems to morph into a weasel or a skunk at some points, but the stained glass window context remains, exploring variations in how the panels of color are divided by black lines that look like lead came—until it veers off into another neighborhood of planes flying through an orange sky.<sup>130</sup> That area of latent space would not exist without one or more training works created by artists who were making expressive choices about how to design stained glass windows or images that looked like stained glass windows. The same is true for another gif of painterly Dutch landscapes with windmills and cows.<sup>131</sup> Again, it is not that any of those images is necessarily substantially similar to any one image in the training set. That would be actionable

---

127. See *Stable Diffusion v1 Model Card*, *supra* note 123.

128. For some brilliant experiments in making latent spaces interpretable, see B.J. Ard, *Copyright’s Latent Space: Generative AI and the Limits of Fair Use*, CORNELL L. REV. (forthcoming 2025).

129. See Ian Stenbit, fchollet, and lukewood, *A Walk Through Latent Space with Stable Diffusion*, KERAS (Sept. 28, 2022), [https://keras.io/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/](https://keras.io/examples/generative/random_walks_with_stable_diffusion/) [<https://perma.cc/DPB3-M2C2>] [[https://web.archive.org/web/20240919032516/https://keras.io/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/](https://web.archive.org/web/20240919032516/https://keras.io/examples/generative/random_walks_with_stable_diffusion/)].

130. See Gif of a Panda, in KERAS, [https://keras.io/img/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/panda2plane.gif](https://keras.io/img/examples/generative/random_walks_with_stable_diffusion/panda2plane.gif) [<https://perma.cc/CY9M-26M4>] [[https://web.archive.org/web/20240828175233/https://keras.io/img/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/panda2plane.gif](https://web.archive.org/web/20240828175233/https://keras.io/img/examples/generative/random_walks_with_stable_diffusion/panda2plane.gif)].

131. See Gif of Cows, in KERAS, [https://keras.io/img/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/happycows.gif](https://keras.io/img/examples/generative/random_walks_with_stable_diffusion/happycows.gif) [<https://perma.cc/ZL38-DRF7>] [[https://web.archive.org/web/20240828175230/https://keras.io/img/examples/generative/random\\_walks\\_with\\_stable\\_diffusion/happycows.gif](https://web.archive.org/web/20240828175230/https://keras.io/img/examples/generative/random_walks_with_stable_diffusion/happycows.gif)].

expression, which I am not focusing on here. Rather, those images incorporate constitutive expression found in one or more training images.

Oren Bracha has argued that training generative AI models involves the process of extracting social “metainformation” from the training works. The information extracted, he says, is not “not the information (whether content or expressive form) contained in specific works.”<sup>132</sup> Rather, it is “information about regularities and relations in the informational patterns of such works,” and it is “social information because its value consists in aggregating patterns common to many individual expressive works.”<sup>133</sup> Does this mean that training generative AI models is non-expressive in a way educating human authors is not? I do not think so.

First, consider Bracha’s use of the term “metainformation.” In what is likely its most common use, the term “metadata” refers to information that is included in a digital file, but that is not the main content of that file (which is often a work of authorship like an image) and that is not derived from that content. For example, under a standard called Exchangeable Image File Format, most digital cameras capture and store certain additional information about each photograph taken.<sup>134</sup> That information typically includes the make and model of the camera, the focal length of the lens, the shutter speed, the aperture setting, the ISO speed setting, and the date and time the photo was taken. Another example is the “copyright management information” defined by § 1202(c) of the Copyright Act.<sup>135</sup> That includes certain information that is conveyed with copies of a work, such as the work’s title, author, and copyright owner, and terms and conditions of using the work.<sup>136</sup> If a file contains a copy of a work, then information such as the make and model of the camera that was used to create that work, or the name of the author of that work, is not part of the work, and is not derived from the work.

That is not the kind of “metainformation” Bracha is talking about. He is talking about information that is derived from the work itself. Ideally, we would find a different term to describe that derived information, since the use of “metainformation” or “metadata” to refer to both non-derived and derived information can lead to confusion, and in particular can lead to the mistaken assumption that derived information should be as insulated from copyright claims as

---

132. Bracha, *supra* note 1, at 42–43.

133. *Id.* at 43; *see also* Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 301 (mentioning the argument that “generative AI is exactly the same as the digitization process in HathiTrust: millions of copyrighted works were digitized so that researchers could extract uncopyrightable metadata . . .”). Amanda Levendowski explicitly connects the difference between “works” and “data” to the inability of machines to have the aesthetic reactions that human beings cannot avoid. *See* Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence’s Implicit Bias Problem*, 93 WASH. L. REV. 579, 625 (2018). I will address this difference regarding what I call “felt expression” below.

134. The Exchangeable Image File Format was originally released in 1996; version 3.0 was released in May 2023. *See* “Exif 3.0 Released, Featuring UTF-8 Support”, INT’L PRESS TELECOMMS. COUNCIL (June 1, 2023), <https://iptc.org/news/exif-3-0-released-featuring-utf-8-support/> [https://perma.cc/X9U9-MSAU] [https://web.archive.org/web/20240919033556/https://iptc.org/news/exif-3-0-released-featuring-utf-8-support/].

135. 17 U.S.C. § 1202(c).

136. *See id.*

non-derived information. In the case of derived information, the question must be, what kind of derived information?

At one end of the spectrum, we might agree that broad statistical summaries derived from a large number of works provide insights that are very different from what we learn when we experience a particular work. In *Author's Guild v. Google*, for example, the court notes that the full-text search of Google's Library Project enables "text mining" to provide "statistical information to Internet users about the frequency of word and phrase usage over centuries."<sup>137</sup> Researchers can uncover interesting patterns, such as "the frequency of references to the United States as a single entity ("the United States is") versus references to the United States in the plural ("the United States are") and how that usage has changed over time."<sup>138</sup> The *Authors Guild* court concluded that such information was part of the "transformative use" that made the Library Project a fair use.

At the other extreme, we can imagine an argument that information about the value of each pixel in a digital image, and the location of that pixel in that image, is just information *about* a work rather than the work itself. Yet because the image—the pictorial or graphic work—could be reconstructed perfectly from that information, calling that information "metainformation" or "metadata" would render those terms meaningless.

On that spectrum, the information preserved in a generative AI model falls pretty close to where the learning of human authors falls. When we as human authors read a book or view an image, we typically do not remember the exact sequence of words in the book, or the exact composition of the image.<sup>139</sup> Rather, we remember techniques of writing or painting, styles, moods, patterns, broad outlines. And we remember certain patterns more when we have experienced them in many works, not just one. Some of those patterns are common in our social milieu, and they become part of our enculturation, as we come to understand norms common to our social context. They are rightly characterized, as Bracha characterizes generative AI models, as "social information," because we "aggregat[e] patterns common to many individual expressive works."<sup>140</sup> Thus, this kind of aggregation does not distinguish generative AI models from the human authors who cannot claim blanket fair use immunity from copyright liability.

It is true that some generative AI developers have started to implement voluntary filters or guardrails that seek to limit the ability of their services to mimic individual artistic styles. For example, OpenAI has stated that "DALL-E 3 is designed to decline

---

137. 804 F.3d 202, 209 (2d Cir. 2015).

138. *Id.* (quoting *Authors Guild, Inc. v. Google Inc.*, 954 F. Supp. 2d 282, 287 (S.D.N.Y. 2013)).

139. It is a little different with music. We often do remember the exact sequence of notes in a melody, which, combined with the greater simplicity of the average musical work, is why all significant unconscious copying cases are about music. See *Three Boys Music Corp. v. Bolton*, 212 F.3d 477 (9th Cir. 2000); *Bright Tunes Music Corp. v. Harrisongs Music, Ltd.*, 420 F. Supp. 177 (S.D.N.Y. 1976); *Fred Fisher, Inc. v. Dillingham*, 298 F.145 (S.D.N.Y. 1924).

140. Bracha, *supra* note 1, at 43.

requests that ask for an image in the style of a living artist.”<sup>141</sup> Microsoft states that it “has offered the ability for living artists to request that their name not be used to generate prompts.”<sup>142</sup> However, the mere existence of such filters is an acknowledgement that the model is, in fact, learning to mimic individual artists’ styles. Because the model itself stores information on individual artistic styles, it may be possible to “trick” the model into imitating an individual style by using alternative prompts that substitute for an artist’s name and are not blocked. In any event, these developers have decided to provide a filter only for “living artists,” thus falling seventy years short of the current term of copyright.<sup>143</sup> Since the ability to provide a filter depends on limiting access to the model through a platform interface, developers that release “open source” models to the public cannot limit the performance of their models, and they will be able to imitate individual styles.

*b. Machines and Humans Rarely Use Actionable Expression*

This section can be as short as the previous section is long. Although generative AI systems have produced images, text, and other generations that contain actionable expression from training works,<sup>144</sup> I have assumed for purposes of this Article that that is an engineering problem that can largely be fixed.<sup>145</sup> If it continues to occur from time to time, there will be infringement suits when a sufficient amount of money is involved, just as there are when human authors use valuable actionable expression from copyrighted works without authorization.

Here, I only want to reiterate what I developed in some detail above in Part II. Human authors do not get the benefit of a blanket exception to infringement liability for unauthorized copying of copyrighted works, so long as they use those works for learning purposes and do not include actionable expression from those works in the new works that they create. True, Mark Lemley and Brian Casey argue in “Fair Learning” that human beings should be able to take advantage of a broader fair use doctrine, under which, for example, gaining access to uncopyrightable facts and theories might justify copying some actionable expression found in scientific articles, newspaper articles, and newsworthy video.<sup>146</sup> Even Lemley and Casey, however, do not argue that human authors should be able to claim a fair use privilege for any unauthorized copying that they do, as long as they don’t author any works that infringe the works they’ve copied. Thus, a claim that human beings who use

---

141. *DALL-E* 3, OPENAI, <https://openai.com/dall-e-3> [[https://web.archive.org/web/20241004112116/https://openai.com/index/dall-e-3/?\\_\\_cf\\_chl\\_rt\\_tk=0GN04xPobQkzhXnBkBSFCE\\_4\\_ty5Kv74GSYrn9E04Qc-1728040876-0.0.1.1-6186](https://web.archive.org/web/20241004112116/https://openai.com/index/dall-e-3/?__cf_chl_rt_tk=0GN04xPobQkzhXnBkBSFCE_4_ty5Kv74GSYrn9E04Qc-1728040876-0.0.1.1-6186)] (last visited Oct. 11, 2024); see also OpenAI Comments 2023, *supra* note 119, at 11.

142. Microsoft Comments, *supra* note 8.

143. See 17 U.S.C. § 302(a) (the general term of copyright for works created on or after January 1, 1978 is the life of the author plus seventy years).

144. See sources cited *supra* note 5 (complaints in lawsuits that allege near-verbatim copying of text in training works by generative AI services).

145. See *supra* text accompanying notes 7–8.

146. See Lemley & Casey, *supra* note 1, at 779–82.

generative machines should get such a privilege is mysterious, and anything but self-evident. Put another way, if we use the term “expression” in the sense of “actionable expression,” then Matthew Sag’s statement that “machine learning models . . . qualify as non-expressive use so long as the outputs are not substantially similar to any particular original expression in the training data” is just a tautology.<sup>147</sup> Indeed, what it means to make a “non (actionable) expressive use” of a work is to create another work (or “generation” or “output”) that is not substantially similar to that work. We need something more to understand why that might turn making copies of the first work into a fair use if machines are involved. For that, we proceed to the next section.

*c. Machines Do Not Use Felt Expression and Humans Do, But That Should Not Make a Legal Difference*

As I have noted above, when we human beings experience a work of authorship, we have an aesthetic or hedonic reaction. Any expression in that work becomes *felt* expression.<sup>148</sup>

As far as we know, no computer has an aesthetic reaction to a work that it processes in some way. Computers are anhedonic learners. Thus, the processing of a work by a computer, without any human involvement, would count as a non (felt) expressive use. It is not clear that all generative machine learning avoids human exposure to training works. There is an entire field of “human-in-the-loop” machine learning,<sup>149</sup> and some of the ways in which humans could be involved in the training and deployment of generative models could expose them to the training works.<sup>150</sup> For present purposes, however, let’s assume that a generative AI training and deployment process involves no human exposure to the copyrighted works in the training set. What then? Does the factual difference between human learning and generative machine learning justify a difference in legal treatment of the two? In particular, does

147. Sag, *Copyright Safety for Generative AI*, *supra* note 1, at 309.

148. Of course, we have aesthetic and hedonic reactions to noncopyrightable elements, too, like theories, facts, and functional elements. Surely, many students of physics have experienced the beauty of  $F = m \cdot a$ , Isaac Newton’s second law of motion.

149. See, e.g., Eduardo Mosqueira-Rey et al., *Human-in-the-Loop Machine Learning: A State of the Art*, 56 ARTIF. INTELL. REV. 3005 (2023).

150. For example, we know that Stable Diffusion is more intensively trained on some subsets of LAION 5B, namely, “LAION Aesthetics,” which are “several collections of subsets from LAION 5B with high visual quality.” Christoph Schuhmann, *LAION-Aesthetics*, LAION BLOG (Aug. 16, 2022), <https://laion.ai/blog/laion-aesthetics/> [https://perma.cc/TA64-CWGG] [https://web.archive.org/web/20240927013701/https://laion.ai/blog/laion-aesthetics/]. How did LAION figure out which images in LAION 5B had “high visual quality”? It trained an aesthetic assessment model on a dataset that consisted of images and the aesthetic quality score that human beings gave those images. In this case, however, it had people score images that were created by a generative AI system from user-submitted prompts, and the users all signed an agreement dedicating whatever copyright interest they might have in the images to the public domain. See *Simulacra Aesthetic Captions*, GITHUB, <https://github.com/JD-P/simulacra-aesthetic-captions> [https://perma.cc/TS2D-JY46] [https://web.archive.org/web/20240927014520/https://github.com/JD-P/simulacra-aesthetic-captions] (last visited Sept. 26, 2024).

it justify blanket fair use treatment of any copying that occurs in generative AI development, while human learning does not get that favorable treatment?

There are two reasons why I think it does not. First, the purpose of copyright law is not only to protect and incentivize the creation of aesthetic experiences—felt expression. It is also, and perhaps primarily, to protect and incentivize the creation of learning experiences. Second, the particular works that have been used to train generative AI models are unlikely to have been created solely for entertainment or hedonic purposes, such that the use of those works to train generative AI models is a “transformative” use. Rather, they are created with the understanding that people will learn from them as well as enjoy them, so that their use to train a generative model should not count as a transformative use.

*i. Copyright Law’s Purpose: Protecting Education as Well as Entertainment.*

It is true that courts sometimes seem to focus solely on the entertainment value of works. As mentioned above, for example, Judge Jerome Frank’s famous formulation of substantial similarity asks whether “whether defendant took from plaintiff’s works so much of *what is pleasing to the ears of lay listeners . . .* that defendant wrongfully appropriated something which belongs to the plaintiff.”<sup>151</sup> However, even that formulation does not exclude the possibility that part of what listeners found valuable about the plaintiff’s work—here, a musical work—is that they learned something new from the work about the rhythmic, melodic, or harmonic possibilities of music. More generally, the better view is that while we human beings have an aesthetic or hedonic reaction to every work that we experience, we also learn from, and are changed by, every work that we experience. We do not have fixed aesthetic desires from birth, and then create markets for works through passive consumption that leaves us unchanged. Rather, we are always “on,” as both appreciators and learners. The works that we experience simultaneously delight or repulse us, and also edify, enlighten, enrich, and change us.

Does copyright law recognize the educational value of works, as well as their entertainment value? There is a good argument that it does, and has from the very beginning. After all, the constitutional proclamation of the purpose of copyright and patent law does not refer the promotion of pleasure. Rather, according to Article I, section 8, clause 8, the purpose of copyright and patent law is “[t]o promote the progress of science and useful arts”<sup>152</sup>—to advance human knowledge in both its theoretical and practical forms. Likewise, the first Copyright Act—the Copyright Act of 1790—is not entitled “An Act for the encouragement of entertainment,” but “An Act for the encouragement of learning.”<sup>153</sup> As Oren Bracha and others have noted,

---

151. *Arnstein v. Porter*, 154 F.2d 464, 473 (2d Cir. 1946) (emphasis added).

152. U.S. CONST. art. I, § 8, cl. 8.

153. Act of May 31, 1790, ch. 15, 1 Stat. 124. The Copyright Act of 1790 borrowed its title (and most of its substance) from England’s Statute of Anne. *See* Statute of Anne 1710, 8 Ann. c. 19 (Eng.) (entitled “An Act for the Encouragement of Learning, by Vesting the Copies of Printed Books in the Authors or Purchasers of such Copies, during the Times therein mentioned”).



among early copyrighted works, “practical or useful works” were dominant.<sup>154</sup> “[W]orks of original literature or poetry constituted only a small fraction of registrations.”<sup>155</sup> Thus, copyright was not intended to protect only the aesthetic reactions that would be produced by the works, but also their value in imparting knowledge and skill of various kinds to those who were exposed to them.

As the Supreme Court held in *Baker v. Selden*,<sup>156</sup> and as I acknowledged above,<sup>157</sup> “the copyright of [a] treatise [does not] give the exclusive right to the art or manufacture described therein.”<sup>158</sup> Thus, when I say that copyright protects the learning value of a work, I am not saying that the owner of copyright in a work has the right to control any skill learned from the work. What I am saying is that no one can make an unauthorized copy of the book to learn the skill, and then claim that copying the book was fair use because the skill was not copyrightable. The *Baker v. Selden* Court would agree, because it made it perfectly clear that “a work on the subject of book-keeping . . . may be the subject of a copyright . . . and, considered as a book, as the work of an author, conveying information on the subject of book-keeping, and containing detailed explanations of the art, it may be a very valuable acquisition to the practical knowledge of the community.”<sup>159</sup> What I am also saying is that the purpose of copyright cannot be just to protect the aesthetic or hedonic reactions that a treatise on book-keeping produces in readers. Surely for such a treatise, as for many works, the aesthetic or hedonic reaction that it produces is only a secondary consideration.

Suppose that an AI tool was designed to ingest the explanatory text in the book at issue in *Baker v. Selden*—“Selden’s Condensed Ledger, or Book-keeping Simplified”—not to add to its ability to generate new text, as existing text-to-text generative AI tools do, but to enable it create a spreadsheet program that would implement the bookkeeping method described in that text. That would be an instance in which the AI really did just use the book to learn an uncopyrightable system, rather than making use of any (constitutive) expression in the book. I do not think it would be fair use to do even that with an unauthorized copy of the book, but obviously that gets closer to *Sega v. Accolade*.<sup>160</sup> Even if you disagree, however, note that that is not what text-to-text generative AI tools do. When they ingest the text of “Selden’s Condensed Ledger,” they not just learning an uncopyrightable system. They are also adding to their ability to create new verbal expression that would be a copyrightable work if it were created by a human being—they are using the constitutive expression in the book. That kind of use is neither a non (constitutive) expressive use, nor a transformative use.

---

154. See Oren Bracha, *Commentary on the U.S. Copyright Act 1790*, in PRIMARY SOURCES ON COPYRIGHT (1450–1900) (L. Bently & M. Kretschmereds.).

155. *Id.*

156. 101 U.S. 99 (1880).

157. See *supra* text accompanying notes 53–57 (explaining why infringement damages should not include the value of skills learned from infringing copies of works).

158. *Baker*, 101 U.S. at 102; see *supra* text accompanying notes 54–60.

159. *Baker*, 101 U.S. at 102.

160. For discussion of *Sega*, see *supra* text accompanying notes 84–87.

ii. *The Imputed Purpose of Copyrighted Works: Both Entertainment and Education.*

A number of generative AI developers have argued that the copyrighted works they used to train their models were created only for entertainment purposes, so that their use of those works for training purposes is use for a “transformative” purpose within the fair use analysis. Thus, for example, OpenAI contends that “[w]orks in training corpora were meant primarily for human consumption for their standalone entertainment value. The ‘object of the original creation,’ in other words, is direct human consumption of the author’s expression.”<sup>161</sup> By contrast, OpenAI is using those works to “hel[p] computer programs learn the patterns inherent in human-generated media,” which it asserts is a completely different purpose.<sup>162</sup> Similarly, Meta Platforms, Inc. asserts that “[t]he works that AI models are trained on . . . were created for expressive purposes,” while “models use training data . . . to develop an entirely new and innovative service that . . . produces valuable new content.”<sup>163</sup> This argument could be supported in part by a statement in *Baker v. Selden*.<sup>164</sup> The *Baker* Court suggests in dicta that there is a category of works—“ornamental designs, or pictorial illustrations addressed to the taste”—the “object” of which is “the production of pleasure in their contemplation,” and it concludes that “[t]his is their final end.”<sup>165</sup>

That piece of dicta in *Baker* is wrong, and so are the generative AI developers’ assertions. As human beings we learn matters such as “the mode of drawing lines to produce the effect of perspective,”<sup>166</sup> not only from “book[s] on perspective”<sup>167</sup>—from instructional treatises—but also from “pictorial illustrations addressed to the taste”<sup>168</sup>—from seeing how artists use perspective in practice. Similarly, we learn how to write “short, sharp sentences,” not just from style manuals, but also from works of fiction. Because we are always learning from works as well as enjoying them, we cannot say that the sole purpose of such works is to entertain. Even if the subjective intent of the author of a work was only to entertain, the imputed purpose of the work must be broader, since the effect of the work is to educate as well. Thus, it is a mistake

161. OpenAI Comments 2019, *supra* note 17, at 5.

162. *Id.*

163. Meta Comments, *supra* note 119, at 14; see Stephen Wolfson, *Fair Use: Training Generative AI*, CREATIVE COMMONS (Feb. 17, 2023), <https://creativecommons.org/2023/02/17/fair-use-training-generative-ai/> [https://perma.cc/4FBY-TTMR] [https://web.archive.org/web/20240927164524/https://creativecommons.org/2023/02/17/fair-use-training-generative-ai/] (“While these images were originally created for their aesthetic value, their purpose for the AI model is only as data.”).

164. I say “in part” because the training sets of the generative AI developers are hardly limited to “ornamental designs,” *Baker v. Selden*, 101 U.S. 99, 103–04 (1880), or works created for their “standalone entertainment value,” OpenAI Comments 2019, *supra* note 17, at 5. They include textbooks and other nonfiction and instructional works. Thus, even if there were a category of works the use of which for generative AI training would be “transformative,” actual training sets for leading generative models are not limited to those works.

165. *Baker v. Selden*, 101 U.S. 99, 103–04 (1880).

166. *Id.* at 102.

167. *Id.*

168. *Id.* at 103.

to conclude that there are “ornamental” works that have only “standalone entertainment value,” so that using such works for learning or training purposes would involve a “transformative purpose.”<sup>169</sup>

As Benjamin Sobel has recognized,<sup>170</sup> Justice Stevens’s reply to a broad interpretation of Justice Blackmun’s dissenting stance in *Sony Corp. of America v. Universal City Studios, Inc.* undermines generative AI developers’ claims that use of copyrighted works to train a model that could create new, dissimilar works is “productive.”<sup>171</sup> Under that broad interpretation, Justice Blackmun draws a line between using works for “purely personal consumption,” which will never produce any benefit to society, and using works in a productive way, which will produce societal benefits.<sup>172</sup> Justice Stevens responds that even uses that Justice Blackmun would presumably consider “consumptive” do in fact benefit individuals and will likely benefit society, so that the presence of societal benefit alone can’t mean that a use is a fair use. “A teacher who copies to prepare lecture notes is clearly productive,” states Justice Stevens, “[b]ut so is a teacher who copies for the sake of broadening his personal understanding of his specialty.”<sup>173</sup> That is true. Human beings have always been learning how to create new, dissimilar works from existing copyrighted works. That has never been sufficient in itself to render that learning use a fair use, and it should not be sufficient to render generative AI training a fair use either.

There is a narrower interpretation of Blackmun’s dissent that is further developed and reinforced in *Campbell v. Acuff-Rose Music, Inc.*, but the narrower interpretation does not help generative AI developers either.<sup>174</sup> A new work may make use of a particular older work in a way that justifies the use of some portion of *that particular work*. If you are “a researcher or scholar,” for example, you may find that your “work depends on the ability to refer to and to quote the work of prior scholars”<sup>175</sup>—scholars who worked in the same field, and whose work you need to quote because you want to comment on or criticize that particular work. Similarly, if you want to parody Roy Orbison’s song “Oh Pretty Woman,” then you need to use some of the lyrics and music from that song, so that the audience understands what you are parodying.<sup>176</sup> That weighs in favor of fair use, and as long as your new work does not use so much

---

169. David Opderbeck argues that “because some of the copyrighted works used to train an AI are meant for training human beings, the transformation in purpose is not so dramatic.” See Opderbeck, *supra* note 66, at 979. I agree, but I want to make the point even stronger: Human beings use virtually all works to train themselves, and so there really is no transformation in purpose when human beings use works to train AI models to produce generative outputs that will satisfy human demands.

170. See Sobel, *supra* note 31, at 73.

171. 464 U.S. 417 (1984).

172. See *id.* at 495 (Blackmun, J. dissenting) (“There is no indication that the fair use doctrine has any application for purely personal consumption on the scale involved in this case . . .”); *id.* at 478 (Blackmun, J., dissenting) (suggesting that a productive use is one that “result[s] in some added benefit to the public beyond that produced by the first author’s work”).

173. *Id.* at 455 n.40.

174. 510 U.S. 569 (1994).

175. *Sony*, 464 U.S. at 477 (Blackmun, J., dissenting).

176. See *Campbell*, 510 U.S. at 586–89.

of the old work that it acts as a substitute for it, you are in pretty good shape.<sup>177</sup> It will be an exceedingly rare case, however, in which a generative AI output contains enough of a training work to be prima facie infringing, but turns out to be a fair use because it is parodying, criticizing, or otherwise commenting on the work it is using. To the extent that fair use jurisprudence distinguishes between “transformative use” and “transformative purpose,” that would be a “transformative use.”<sup>178</sup> The design of generative AI systems is not focused on such uses, in which the purpose and content of the output or generation justifies the use of one particular training work. Rather, generative AI systems are primarily designed to produce new output with constitutive expression that they learned from a variety of training works. That is not a transformative use, and it is not a transformative purpose either.

*iii. Should Learning Without Enjoyment Be Discounted?*

You may have a lingering sense that even if the learning value of works is protected by copyright, so is the enjoyment value, and therefore if humans delegate the development of literary and artistic skills to machines, and forego the enjoyment, they should get some sort of discount. You also may have a sense that human beings sometimes are not learning all that much from their TikTok videos or their sitcoms, and so they should pay a premium for the enjoyment value of works. Of course, we should not forget that licensing prices may take this into account. For example, Stability AI offers a generative music service called “Stable Audio” that is trained entirely on music licensed from a platform called Audiosparx. Creators on the Audiosparx platform are allowed to opt out of having their music used for training, and when Audiosparx first publicized the training use and the opt-out, about 10% of the creators opted out.<sup>179</sup> Those who do not opt out—initially, 90% of the creators—receive some payment for the training use.<sup>180</sup> Although the amounts of those payments are not publicly known, the payment a creator receives for training use is almost certainly far less than the payment they would receive if their music was licensed for a (felt) expressive use, such as use in a movie or as background music in a store.

---

177. See *id.* at 591–94.

178. For this distinction, see Rebecca Tushnet, *Content, Purpose, or Both?*, 90 WASH. L. REV. 869, 869–70 (2015) (distinguishing between “content-transformativeness” and “purpose-transformativeness”); Jiarui Liu, *An Empirical Study of Transformative Use in Copyright Law*, 22 STAN. TECH. L. REV. 163, 169–70 (2019) (distinguishing between “physical transformation” and “purposive transformation”).

179. See Stuart Dredge, *Stable Diffusion Maker Launches Stable Audio Text-to-Music AI*, MUSIC ALLY (Sept. 13, 2023), <http://musically.com/2023/09/13/stable-diffusion-maker-launches-stable-audio-text-to-music-ai/> [https://perma.cc/7UTZ-SW9Q] [https://web.archive.org/web/20240927173722/https://musically.com/2023/09/13/stable-diffusion-maker-launches-stable-audio-text-to-music-ai/].

180. *Id.*; see also *Introducing Stable Audio 2.0*, STABILITY AI (Apr. 3, 2024), <https://stability.ai/news/stable-audio-2-0> [https://perma.cc/X7CR-T8J2] [https://web.archive.org/web/20240927175138/https://stability.ai/news/stable-audio-2-0].

Whether or not private bargaining sets the correct discount for generative AI training use, we should be wary of setting copyright policy that makes it less expensive to train machines than to educate human authors to perform similar tasks.<sup>181</sup> Even if generative AI developers have to license training works, generative AI will cause significant labor displacement. That is probably something that society as a whole should welcome, as it means that the creation of some kinds of content will become less expensive and more accessible, and some of those who would have become remunerated authors or artists will find some other productive role. However, to avoid even more drastic displacement, we should think twice about making the present generation of authors the one that helped finance its own replacement. Academics concerned about AI labor displacement more generally have suggested that we should change policies that make it more expensive for employers to hire human beings than to buy machines. For example, rather than forcing employers to pay for employees' health care and family leave benefits, maybe we should finance those benefits through general taxation.<sup>182</sup> Here, we should consider whether copyright policy, in the form of a fair use exemption for training generative machines, but not for educating human authors, will accelerate and deepen creative labor displacement in a way that we should avoid.

## B. THE REPRODUCTION RIGHT AND COMPUTER "READING"

In this section, we are going to keep discussing whether the copies made for generative AI training purpose should be treated as infringing the reproduction right, but the framework will be quite different, and it will sometimes seem like the polarities have been reversed. The previous section discussed the argument that generative AI training is making a "non-expressive use" of the copyrighted works in the training set, and that the act of making copies of those works for training purposes should therefore be treated as a fair use. In that context, the proponents of fair use argue that use by the computers creating the generative model is fair precisely because it is *different* than human use. Human use is "expressive" because we human beings have hedonic reactions to a work to which we are exposed. Computers do not have those reactions, and so their use is "non-expressive." By contrast, I argued that both humans without computers and humans building generative models can be said

---

181. We can imagine some ghastly experiment in which human beings were rendered anhedonic—perhaps through a continuous flow of dopamine inhibitors—and then employed to read copyrighted works, and respond to requests for essays or images by other, more fortunate human beings. Under the principle that non (felt) expressive use is fair use, the employers of those anhedonic human beings could make copies of every work now under copyright, so long as they only let their anhedonic employees read or view or listen to them, thus avoiding the extra cost of licensing the works, and in that respect putting humans and computers on a level playing field. However, even if such an operation could overcome moral and legal objections, it likely would not work: Human beings almost certainly do not learn as well when deprived of their aesthetic reactions.

182. See Cynthia Estlund, *What Should We Do After Work? Automation and Employment Law*, 128 YALE L.J. 254, 305–19 (2018).

in a relevant sense to be learning—acquiring a facility—and in particular acquiring a facility with constitutive expressive choices, acquired from the training works.

Now we want to consider an argument that the creation of copies for computer use should be a fair use because computers are doing something that should be treated *the same* as one kind of human activity, namely, the activity of reading, listening to, or viewing works.<sup>183</sup> To process a work using a computer, one typically has to make a copy of the work. As Nicola Lucchi states, “AI systems cannot learn from art in the same way humans do, since they require an exact copy of the artwork in their training dataset.”<sup>184</sup> Mark Lemley and Bryan Casey make a similar comment: “The problem [machine learning] systems face is the inability to capture the unprotectable parts [of copyrighted works of authorship] to use for training without making a rote copy of the protectable ones.”<sup>185</sup>

However, according to this argument, it is merely a technicality that a copy must be made for a computer to use a work to train a generative model. During model training, the computer processes the copy once and is done with it, just like we human beings might read a book, listen to a song, or view a streamed movie. Therefore, computer “ingestion” of works for purposes of generative model training is functionally equivalent to human reading, viewing, or listening.<sup>186</sup> As we know, human beings do not incur infringement liability for those activities.<sup>187</sup> Therefore, computers should not either. Rather, they should enjoy a kind of “perceptual parity” with humans. Because there is no specific provision in the Copyright Act allowing adjustments for perceptual parity, fair use should step in.<sup>188</sup> Making a copy solely for purposes of one-time computer analysis and training should be a fair use. Thus, for

---

183. One might consider a broader argument that people who use computers to learn things about works should be treated exactly the same as people who learn without computers. They should benefit from all of the same freedoms that unaided humans do, and be subject to all of the same strictures. For example, they should benefit from the first sale doctrine and the limitations on exclusive rights, but they should not be able to make unauthorized copies of entire works. The problem is that treating what the computer does as mere “experiencing,” and treating as fair use any copies made to enable that experience, ends up bypassing all of the other remaining copyright strictures. Making an unauthorized copy of an entire book is perfectly fine, because that is treated as no more than reading that book, which is not an infringing act. Nor does it matter that the copy from which the electronic copy is made is itself pirated, because the computer is just “reading” the pirated copy, and reading pirated copies is not an infringing act either. All we need is a ruling that the computer is effectively doing no more than reading, listening or watching, and every copy made to enable that activity is immunized from infringement liability.

184. Nicola Lucchi, *ChatGPT: A Case Study on Copyright Challenges for Generative AI Systems*, EUR. J. RISK REGUL. 1, 11 (2023).

185. Lemley & Casey, *supra* note 1, at 775.

186. See, e.g., UTKU TAŞOVA, *THE DICTIONARY OF ARTIFICIAL INTELLIGENCE* 127 (2023) (defining “data ingestion” as “the process of collecting, importing, and processing data for storage and analysis”).

187. See *supra* text accompanying note 47.

188. The idea of adjusting rules to achieve perceptual parity between unaided humans and humans using machines is one example of the broader idea of “technological neutrality.” As Carys Craig has argued, it is unlikely that technological neutrality can be achieved by adjusting one rule in isolation; it is more likely that one has to step back and compare contexts in light of the broader aims of copyright law. See Carys J. Craig, *Technological Neutrality: Recalibrating Copyright in the Information Age*, 17 THEORETICAL INQUIRIES LAW 601, 606–15 (2016).

example, in its comments submitted to a U.S. Patent and Trademark Office inquiry on AI and intellectual property, the Consumer Technology Association contends: “To read is not to infringe . . . The case law on non-AI reading-for-learning teaches that AI reading does not infringe where fixed copies are produced but not distributed.”<sup>189</sup>

That leads to the question whether computer processing is, for copyright purposes, functionally the same as human reading, listening or viewing. (Sometimes I will use “reading” to stand for all human experience of works.) This is not an easy question. However, I think there are two principal reasons why the computer processing that takes place during generative model building should probably be treated as *different from* the human experience of works. First, distinctions between the exclusive rights of copyright are very much based on human limitations of perception and memory, and are comprehensible only in light of those limitations. If generative AI development operates beyond those limitations, which I think in some respects it does, then it is hard to conclude that a computer involved in that development should be treated as merely being on the receiving end of a display or performance, rather than as possessing a copy, when copies of training works are in fact being made.

Second, human experience of works is the experience, not only of beings who can have emotional and aesthetic reactions to those works, but of beings who remember those reactions, and who are actors in social, cultural, and economic worlds. When copyright owners make works publicly available, they are doing so under the general assumption that the public who will be experiencing them is a human public, with emotions, memory, and agency. That enables the building of reputations, and it enables generation of revenue. Generative AI training bypasses all direct social, cultural, and economic involvement stemming from the experience of works. Cutting that link between experience and action in the social and economic world renders that use unlike the human experience of reading, viewing, or listening. Because it is unlike that experience, it should not count as “mere reading” for fair use purposes.

## 1. Reading and the Human Limitations Implicit in the Structure of the Exclusive Rights

The Copyright Act itself gives us an important framework within which we can assess whether what machine learning algorithms are doing are for copyright purposes similar to what human beings are doing when they read, view, or listen to works. As part of its core structure, copyright law assumes that human abilities are limited, particularly those abilities related to memory, analysis, and creation. The Copyright Act’s division of exclusive rights in § 106 is only comprehensible against

---

189. Consumer Tech. Ass’n, Comment on U.S. Patent and Trademark Office’s Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation, Docket No. PTO–C–2019–003 (Jan. 10, 2020), at 1, 3, [https://www.uspto.gov/sites/default/files/documents/Consumer%20Technology%20Association\\_RFC-84-FR-58141.pdf](https://www.uspto.gov/sites/default/files/documents/Consumer%20Technology%20Association_RFC-84-FR-58141.pdf) [<https://perma.cc/ZG32-XRLX>] [[https://web.archive.org/web/20240921202957/https://www.uspto.gov/sites/default/files/documents/Consumer%20Technology%20Association\\_RFC-84-FR-58141.pdf](https://web.archive.org/web/20240921202957/https://www.uspto.gov/sites/default/files/documents/Consumer%20Technology%20Association_RFC-84-FR-58141.pdf)].

the background of limited human abilities. Although § 106 creates six exclusive rights, they are of two basic types. The first are those that concern fixed copies, reproduction, distribution, and some if not all adaptations.<sup>190</sup> The second are those that concern only unfixed presentations of works—public display and public performance (including digital audio transmissions), and possibly unfixed adaptations. What I will call the “presentation rights”—display and performance—are much more limited and porous than what I will call the “fixation rights”—reproduction and distribution. Private presentations are not regulated at all. Public presentation rights are subject to far more exceptions and limitations than fixation rights. Those include the eleven enumerated exceptions in § 110 of the Copyright Act,<sup>191</sup> and the elaborate schemes for cable transmissions, digital audio transmissions, jukeboxes, public broadcasting, and satellite transmissions.<sup>192</sup> In addition, the enumeration of distinct § 106 rights facilitates private bargaining about licensing performances and displays on far more lenient terms than reproductions and distribution.

Why do we think of presentations differently than we think of fixation, and why does the Copyright Act regulate presentations with a far lighter touch? One important answer is this: As human beings, we face serious limitations in our ability to analyze, learn from, remember, or recreate the works that we experience in real-time presentations. These limitations are so natural to us—such a part of our ordinary experience of life—that it may require a thought experiment to make them apparent. Imagine that we all had near-perfect auditory and visual memory. Once we were exposed to a page of text or of musical notation, we could read it from memory at will. Once we saw a painting, we could conjure it up whenever we fancied, and could savor the vibrant colors or rich graytones just as if we were standing in front of it. If that were how we experienced the world, copyright law would be structured far differently. It is not clear that we would consider reproduction and distribution to be different enough from performance or display to formulate separate exclusive rights in that way. It is not clear that we would exempt private presentations, since a private presentation would have the same effect as the distribution of a personal copy. Paradoxically, one might prefer to acquire a copy rather than view a display if one did *not* want to experience a work immediately, but wanted to delay that experience.

---

190. There is some question about whether the adaptation or derivative work right can be infringed without fixation. See H.R.REP. No. 94-1476, at 62 (1976), as reprinted in 1976 U.S.C.A.N. 5675 (“The exclusive right to prepare derivative works . . . is broader than [the exclusive right of reproduction], in the sense that . . . the preparation of a derivative work, such as a ballet, pantomime, or improvised performance, may be an infringement even though nothing is ever fixed in tangible form.”); but see *Lewis Galoob Toys, Inc. v. Nintendo of Am., Inc.*, 964 F.2d 965, 967 (9th Cir. 1992) (holding that the derivative work must “incorporate a protected work in some concrete or permanent form.”). For my purposes, it does not really matter whether the adaptation right extends to some unfixed performances or not. If so, then it straddles the two basic types of exclusive rights. If not, then it can be grouped with the reproduction and distribution rights.

191. See 17 U.S.C. § 110.

192. See 17 U.S.C. § 111 (cable transmissions); 17 U.S.C. § 114(d) (digital audio transmissions); 17 U.S.C. § 117 (jukeboxes); 17 U.S.C. § 118 (public broadcasting); and 17 U.S.C. § 119 (satellite transmissions).



The human limitations assumed by the division of exclusive rights into fixation rights and unfixed presentation rights concern not only memory, but also analysis and creation. Suppose that we human beings want to thoroughly analyze a musical sound recording—a “song,” in the modern definition—to understand it and perhaps to mimic its style. Very few if any of us can do that after having heard the song only once, particularly if it is not at a time of our own choosing. As human beings, we need access to a copy of that song, so that we can listen to it repeatedly, and pause and listen to small segments repeatedly. Of course, in a digital networked world, that copy need not be in our physical possession. A perfectly manipulable on-demand stream would do. But that is why Congress decided to treat on-demand streaming differently than programmed streaming, requiring individually negotiated licenses for the former, and also why Congress decided that owners of the rights to make copies of sound recordings of musical works—so called “mechanical rights”—would be compensated for on-demand streams.<sup>193</sup> An on-demand stream is too close to being a replacement for possession of a copy.

Now that we understand that the Copyright Act treats presentations differently than fixations in substantial part because of limitations on what we humans can do with our experiences of presentations, we can return to considering the status of generative AI training. Suppose we are willing to entertain the possibility that at least in some instances, making a copy of a work for purposes of computer processing should be treated as fair use, because the computer is doing no more than what a human being would do when experiencing a presentation of a work, something that does not implicate any of the exclusive rights of copyright. We now have to focus on what it means to do “no more than” what a human being would do. In a particular instance, is a computer doing what a human being could do without a copy, through mere exposure to a performance or display, or is it doing something more? If it is doing more, than immunizing copying for computer use is not creating perceptual parity. Rather, it is favoring people who use computers over those who do not, and creating a huge breach in the partition between fixation rights and presentation rights.

How do we tell whether computers, in building generative models, are doing more with the works they use to train the models than what human beings could do when experiencing those works through performance and display? It is not easy to decide what the relevant comparison should be. We could consider (1) the learning process—how learning takes place; (2) the resulting memory or model—how much and what is “known”; or (3) the resulting ability or capacity—how quickly and cheaply

---

193. See 17 U.S.C. § 114(d)(2)(a)(1) (providing that digital audio transmissions of sound recordings that are part of an interactive service are not covered by a statutory license; licenses for those transmissions must be obtained directly from the owner of copyright in the sound recording); 17 U.S.C. § 115(e)(10) (defining “digital phonorecord delivery” as including an “interactive stream”); *Arista Records, LLC v. Launch Media, Inc.*, 578 F.3d 148, 161 (2d Cir. 2009), *cert. denied*, 559 U.S. 929 (2010) (considering the definition of “interactive service”) (“If the user has sufficient control over the interactive service such that she can predict the songs she will hear, much as she would if she owned the music herself and could play each song at will, she would have no need to purchase the music she wishes to hear.”).

can new works or generations be created. Probably some combination of all three is necessary, since it is not clear how the details of processing alone would matter if the resulting system created generations that were far worse and more expensive than works created by human authors.

As for the learning process, I think what turns out to be important is how consistently a machine learning algorithm can analyze a very large number of works—whether texts, images, videos, or songs—in a comparatively short period. For example, GPT4 was apparently trained in a three-month period using about 13 trillion tokens of text, which works out to about 10 trillion words.<sup>194</sup> The GPT learning algorithm is able to build a coherent model because it analyzes each piece of text in exactly the same way. Human beings can't do that. At a reading rate of 250 words per minute, a human being would take about 76,000 years to read 10 trillion words. Maybe we could figure out how to do something with many human beings working in parallel, with some redundancy built in to ensure consistency. However, that very redundancy would require persistent copies, to make sure that people were reading the same thing.

It is also the case that generative model building requires making many additional copies of training works—at least “copies” in the colloquial sense, if not the copyright sense. For example, the goal in training a generative image model, using state-of-the-art diffusion technology, is to enable it to find images within visual noise. You do that by starting with a model that has randomly initiated weights, feeding it an image with some noise introduced, seeing what it predicts the image would look like with the noise removed, and then adjusting the weights to make the model perform better. Existing theory suggests that this needs to be done in stages—the model has to be trained to make a noisy image a little less noisy, and then repeat that process many times, rather than predicting a perfect image from complete noise in one step. For every stage, the computer training the model has to make another copy of the image being used as training data—a copy that is slightly noisier. How many stages are there? At first, researchers needed hundreds of stages; a 2021 paper introduced an improved process that brought that number down to around fifty.<sup>195</sup> Typically, in more than half of those stages, the original image is still quite recognizable. That means that the training process is requiring, not just one copy of an image to be made, but twenty-five or more recognizable copies.<sup>196</sup>

---

194. See Yam Peleg (@Yampeleg), TWITTER (July 10, 2022, 11:23 PM), archived Tweet, <https://archive.is/2RQ8X> [<https://perma.cc/MM6S-P7LP>].

195. See Alex Nichol & Prafulla Dhariwal, *Improved Denoising Diffusion Probabilistic Models*, ARXIV (Feb. 18, 2021), <http://arxiv.org/abs/2102.09672> [<https://perma.cc/B9T3-87YT>] [<https://web.archive.org/web/20240921220503/https://arxiv.org/pdf/2102.09672>] (“While DDPM . . . requires hundreds of forward passes to produce good samples, we can achieve good samples with as few as 50 forward passes . . .”) (last visited Sept. 21, 2024).

196. Another advance in the field first compresses the images into a latent space, and then performs the training steps within that space. See, e.g., Robin Rombach et al., *High-Resolution Image Synthesis with Latent Diffusion Models*, ARXIV (Apr. 13, 2022), <http://arxiv.org/abs/2112.10752> [<https://perma.cc/S3X9-PACB>] [<https://web.archive.org/web/20240922195159/https://arxiv.org/pdf/2112.10752>]. Nonetheless, each of the progressively noisier versions of an image in latent space can be converted into a perceptible

That is not the end of copying. Many generative image models use convolutional layers to process the input images. The idea of a convolutional layer is to process an image with some number of filters (also called “kernels”) that will highlight certain features in the image, such as vertical edges or horizontal edges. The output of a typical first convolutional layer is some number of filtered copies of the input images. For example, the first paper to suggest a U-Net architecture contains an example in which the first layer produces sixty-four filtered copies of the input image—all recognizable copies of the input image, with certain features highlighted.<sup>197</sup> Thus, in a training process, hundreds of copies of the input image are being made.

Of course, arguably, all these additional copies are not “copies” within the meaning for the Copyright Act. Although the Copyright Act leaves it to courts to determine what exactly counts as more than “transitory duration” under the definition of “fixed” in § 101,<sup>198</sup> it may be that none of these additional copies exists for more than what a court would find to be a transitory duration. However, they are still evidence of the consistency of how the machine learning algorithm processes each image. If similar consistency could be achieved by human beings only with the use of copies, then it is hard to argue that what the machine is doing is functionally equivalent to human reading.

That is the learning process. How about the resulting model? I argued above that the traces of the training works that are left in the generative model may in some sense be similar to what human beings remember. In both cases, what is remembered or modelled may be mostly patterns in the works that were experienced or processed, not the exact sequence of words in a particular textual work, or the exact graphic design of a particular image.<sup>199</sup> That suggests that human memory and the generative model may not be all that different. However, some have suggested that “memorization” of training works is inevitable in generative models, to a greater

image, and thus should count as a copy. See 17 U.S.C. § 101 (“copy”) (“Copies’ are material objects . . . from which the work can be perceived . . . either directly or with the aid of a machine or device.”) (emphasis added).

197. See Olaf Ronneberger, Philipp Fischer & Thomas Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation* 2, ARXIV (May 18, 2015), <http://arxiv.org/abs/1505.04597> [<https://perma.cc/9U5W-XHYJ>]

[<https://web.archive.org/web/20240922200846/https://arxiv.org/pdf/1505.04597>] (depicting a diagram showing a U-Net with a 64-channel first convolutional layer). For good visualizations of the output of a first convolutional layer (and subsequent layers) see Jay Wang et. al., CNN EXPLAINER (Aug. 29, 2020), <https://poloclub.github.io/cnn-explainer/> [<https://perma.cc/4MUQ-LYPA>]

[<https://web.archive.org/web/20240922202853/https://poloclub.github.io/cnn-explainer/>] (last visited Sept. 22, 2024); Adam Harley, An Interactive Node-Link Visualization of Convolutional Neural Networks, ISVC 865, 867–77 (2015), [https://adamharley.com/nn\\_vis/](https://adamharley.com/nn_vis/) [<https://perma.cc/UP6L-TU2E>] [[https://web.archive.org/web/20240922203823/https://adamharley.com/nn\\_vis/](https://web.archive.org/web/20240922203823/https://adamharley.com/nn_vis/)] (last visited Sept. 22, 2024).

198. 17 U.S.C. § 101 (“fixed”) (“A work is ‘fixed’ in a tangible medium of expression when its embodiment in a copy or phonorecord . . . is sufficiently permanent or stable to permit it to be perceived, reproduced, or otherwise communicated for a period of more than transitory duration.”).

199. See *supra* text accompanying notes 141–42.

degree than human beings would actually memorize verbatim text or images.<sup>200</sup> If that is the case, that is also evidence that models are doing something that human beings could not do without copies.<sup>201</sup>

Lastly, we consider ability or capacity: can a generative AI model accomplish more than a human author can? The answer has to be “no and yes.” As of Fall 2024, there are still many things human authors can do that generative AI tools cannot. Human authors can write longer coherent literary works than generative AI tools can. Human authors can fact-check better. Human artists can create works in physical media that generative AI tools cannot.<sup>202</sup>

On the other hand, generative AI tools can produce short textual works and medium-resolution digital images much faster and much less expensively than human authors can. That, is, of course, why companies are investing many billions of dollars in the technology. If those companies thought that their potential customers would be able to obtain text and images better and faster by employing human beings to do the work, they would not be making that investment. Thus, within the area of competence of generative AI tools, they can out-compete human beings, and they do so through analysis and modeling that human beings would not be able to do without making copies. If generative AI developers are actually making copies of training works, and if they are using computers to perform analysis and build models that human beings would need copies to equal, then it is hard to argue that the computers are “merely reading” the training works.

To be sure, some cases hold that computer processing of copyrighted works is noninfringing, even though copies are made of those works, using the doctrine of fair use. However, they do not do so by concluding that the computers are “just reading” the works. In *Authors Guild v. Google, Inc.*, for example, the Second Circuit does not hold that the Google Library Project is fair use because the computers that are building a full-text search tool for millions of books are “just reading” those books.<sup>203</sup> Rather, it considers a variety of factors, including the transformativeness of the use, and the effect on the use on the potential market for the copyrighted work. Full-text

---

200. See, e.g., Henderson et al., *supra* note 2, at 27 (noting that technical examinations of generative models “have consistently found that generative models memorize or plagiarize content,” and that “[t]he percentage of verbatim outputs varies depending on extraction strategy and the model.”).

201. Some will argue that “memorization” in a model is the making of a copy as that term is defined in copyright law, and therefore is a violation of the reproduction right in whatever work is memorized. That may be true, though there is the counterargument that if that copy is never perceived by anyone—either because it is unlikely to ever be generated given all of the possible generations that a model can produce, or because it will be filtered out by an effective closed-system filter—that it should not be counted as a copy. Here, I am interested, not in whether there is infringement of a particular training work, but what the presence of systematic memorization says about how generative models may exceed human limitations.

202. Pindar Van Arman is among the artists who have been building generative AI systems that are connected to robots capable of painting on physical media. See CLOUDPAINTER, <https://www.cloudpainter.com/> [<https://perma.cc/YJ8C-NK3G>] [<https://web.archive.org/web/20240923060113/https://www.cloudpainter.com/>] (last visited Sept. 23, 2024).

203. 804 F.3d 202 (2d Cir. 2015).

search indexing may be transformative, and may not impact the potential market for the copyrighted work. Generative AI training may not be transformative, in so far as the works in the training set are created at least in part to disseminate knowledge about expressive choices, and it may impact the potential market for the training work. In neither case is the analogy of human reading determinative.

## 2. Reading, Emotion, Memory, and Agency

Those who argue that generative AI training is fair use because it is a non-expressive use would readily agree that human experience of works—including human reading—is (felt) expressive. Human beings can be moved by what they read, while machine learning algorithms can't be. Here I want to add one more layer of distinction between human beings and generative AI training algorithms. Not only do we human beings have hedonic reactions to works that we experience, we also can remember those reactions, and we typically can and do take actions based on our remembered experience. That makes a big difference in the ecosystem of distribution and presentation of works, in a way that is relevant to the question of whether the algorithms can be said to be “reading” what they are processing.

When OpenAI argues that using copyrighted works to train AI models is fair use, it emphasizes that it is using “*publicly available* internet materials.”<sup>204</sup> The exact legal point it is making is not quite clear,<sup>205</sup> but the general atmospheric suggestion is that these are works that the copyright owners are willing to have anyone experience. The owners have just put them out there in plain view of any visitor, and therefore training use is not interfering with any of their expectations. Recall, also, that free public display and performance is an important way in which we human beings get to experience works and learn from them without making any direct payment for them.<sup>206</sup> Consider, however, *why* the owners (who are often also the authors) have decided to make their works available to the general public. Here are some reasons, expressed in the first person:

I have an opinion on a particular public issue, and I would like to make that opinion known, in the hope of convincing others and influencing the outcome of the debate.

I want people to listen to my recording of a song, in the hope that they will like it and decide to pay to hear my other recordings, or to pay to attend one of my live concerts.

I want people to see what the item I am selling looks like, in the hope that they like it, and that they then buy it.

I want people to visit my webpage, because I display advertisements on that page, and I earn money from those advertisements.

---

204. *Open AI and Journalism*, *supra* note 17; see OpenAI's MTD, *supra* note 17, at 12.

205. The “public availability” of the works may eliminate other legal grounds for objection, such as trade secrecy or terms of service agreements, or may be a reference to publication as a factor in fair use analysis (although it is rarely the determining factor, particularly after a 1992 amendment to § 107).

206. See *supra* text accompanying note 52.

I want people to experience my work, in the hope that many people will like it, because I want to be liked and respected, apart from any potential economic gain.

I want my family and friends to know what I am doing, and I am making work available to the public mainly as a byproduct of reaching those friends and family.

Notice that all of these reasons require a human audience—human beings who can emotionally react, and remember those reactions, and act on them. Thus, our traditional concept of human experience of works—of reading, viewing, and listening—is a concept in which the possibility of emotion, memory, and action plays a central role. True, not every audience member will react the way an author hopes, but human audience members have the basic capacity to feel, to remember the feeling, to act.

Generative AI development all but destroys any connection between what is extracted from a work during training, on the one hand, and emotion, memory, and action, on the other. Indeed, that is the selling point of those who argue that training use is a “non-expressive use”—computer processing to train a model avoids producing any emotion in anyone, let alone an emotion that could be remembered and motivate action. Of course, the eventual products of the generative models are, and are meant to be, experienced by feeling and acting human beings. That means there is a small possibility that there could be some connection between the reasons people freely publish their works and the generative AI output. For example, someone might hope that their expression of opinion on a topic could influence the attitudes expressed in the output of a generative AI model. However, that influence is greatly attenuated, and most generative AI developers dismiss the effect of any one work on the model as infinitesimal.

That leads us to revisit the question whether generative model training can be justified, not as a non (felt) expressive use, but as an activity that is just like human reading. When answering that question, generative AI’s putative strengths become serious weaknesses. The more complete the break between extracting constitutive expression from works and any possibility of emotional reaction and resulting action, the less that generative AI training is anything like human reading, viewing or listening. That should at least give us serious pause before we accept the argument that because there is no infringement liability for human reading, there should not be infringement liability for the processing that occurs during generative AI training either.

To be sure, there is some Ninth Circuit precedent which suggests that, in some contexts, the use of technical means to avoid the conditions under which the copyright owner made a work publicly available does not weigh against fair use or in favor of infringement. In the 2014 case of *Fox Broadcasting Co., Inc. v. Dish Network LLC*,<sup>207</sup> for example, Dish Network offered its customers a service it called “AutoHop.” AutoHop enabled Dish Network customers to skip the commercials on some television shows, including some shows broadcast by Fox. The customers recorded the shows on DVRs provided by Dish Network. Simultaneously, Dish Network

---

207. 747 F.3d 1060 (9th Cir. 2014).

employees watched the shows and electronically marked where the commercials began and ended. Dish Network then sent the file with the electronic markings to its customers. The customers' DVRs could use the files with the markings to skip the ads on the shows they had recorded.

Is it fair use for a customer to copy a show while using technology to skip all of the ads? The Supreme Court precedent obviously looming over this question is *Sony Corp. of America v. Universal City Studios, Inc.*<sup>208</sup> The *Sony* Court famously held that private copying of broadcast television for purposes of time-shifting was a fair use. Yet it is hardly a foregone conclusion that it would also have held copying for purposes of commercial-skipping to be a fair use. After all, the *Sony* majority noted how hard it was to skip commercials given the videotape technology of the time, and it concluded that commercial skipping would not substantially affect the market value of the Universal's works, since a survey showed that "92% of the programs were recorded with commercials and only 25% of the owners fast-forward through them."<sup>209</sup>

In *Fox Broadcasting*, however, the Ninth Circuit held that the Dish Network customers who used the AutoHop feature to skip commercials on shows that they had recorded were still engaged in fair use. According to its logic, no analysis of how commercial skipping might affect Fox's revenues was needed, because Fox did not own copyright in the ads: "If recording an entire copyrighted program is a fair use, the fact that viewers do not watch the ads not copyrighted by Fox cannot transform the recording into a copyright violation."<sup>210</sup> Although the court rests its reasoning on Fox's lack of ownership of copyright in the advertisements, it is hard to imagine that that fact makes any difference. Suppose that Fox did own copyright in the ads, or that customers used AutoHop to skip over parts of the show in which Fox did own copyright, like the introduction, or scenes that were too violent or too boring. Would the Ninth Circuit then hold that home copying was no longer fair use, because of what the home copiers decided *not* to watch?<sup>211</sup> That seems very unlikely. The opinion

---

208. 464 U.S. 417 (1984).

209. 464 U.S. at 453 n.36 (quoting *Universal City Studios, Inc. v. Sony Corp. of Am.*, 480 F. Supp. 429, 468 (C.D. Cal. 1979)). The District Court noted that there were probably actually fewer commercials skipped than those statistics suggested. "[T]o omit commercials, [as 8% of recordings did,] Betamax owners must view the program, including the commercials, while recording." *Universal City Studios, Inc.*, 480 F. Supp. at 468. "To avoid commercials during playback [as 25% of viewers did,] the viewer must fast-forward and, for the most part, guess as to when the commercial has passed"—and a wrong guess would result in viewing part of the commercial, or missing part of the show and then possibly rewinding and viewing part of the commercial. *Id.*

210. *Fox Broadcasting Co.*, 747 F.3d at 1068–69.

211. A parallel question is whether private performances of a work that omit some of it infringe the adaptation right of 17 U.S.C. § 106(2). Although some of the legislative history of the 1976 Act suggested that the adaptation right could be infringed without fixation, courts have tended to shy away from so holding. See, e.g., *Lewis Galoob Toys, Inc. v. Nintendo of Am., Inc.*, 964 F.2d 965, 967 (9th Cir. 1992) (holding that "[a] derivative work must incorporate a protected work in some concrete or permanent 'form'"). The Copyright Office took the position that specially programmed DVD players which muted or skipped objectionable portions of movies did not violate the adaptation right. See *Family Movie Act of 2004: Hearing on H.R. 4586 Before the Subcomm. on Cts., the Internet, and Intell. Prop. of the H. Comm. on the Judiciary*, 108th Cong. 12–13 (2004) (statement of Marybeth Peters, Register of Copyrights, Copyright Office of the

is better interpreted as enlarging the power and the liberty of home copiers, and protecting the interests of the tech companies that assist them, at the expense of the copyright owners and program originators. Under that interpretation, once a copyright owner makes its work publicly available, it cannot prevent private persons using technical means to avoid the advertisements that finance the public availability of the work, at least when that use is incidental to the time-shifting purpose sanctioned in *Sony*.

In *Perfect 10, Inc. v. Amazon. com, Inc.*, an earlier Ninth Circuit panel articulates a similar limit on copyright owner control.<sup>212</sup> When an internet browser requests a web page from a server, the server, using traditional web page technology, first sends the text and code of the web page, but not whatever images might be included on the page. The images are stored on a server at separate URLs, which are included in the page's code. The browser uses those URLs to request each image from the server where it is stored, and it then incorporates them into the display of the requested web page. Because the images are stored separately, they are technically available for display on web pages that are completely different than the page intended by the person who placed the images on the server.

Suppose that I take beautiful photographs of national park scenery and display them on a web page on which I also display advertisements. I gain some income from those advertisements to support myself as a photographer. Along comes someone else—call them Opportunist—who decides to create a different web page that features my photographs. Opportunist incorporates my photographs by copying the image URLs contained in the code of my web page, so that the browser of a visitor to Opportunist's web page requests my photographs from the server on which I placed them, and that server responds to those requests. Opportunist also makes money by displaying advertisements and charging for them. The appearance of an image file on a new page has been called "framing," since the image has been provided with a new context.<sup>213</sup> Do I have any legal recourse in this situation, or is it the case that once I have made my photographs publicly available in this way, anyone can "frame" them, and I have no legal control?

This issue has been litigated, not as a matter of fair use, but as a matter of whether the person who "frames" or "embeds" an image is making a new public display of that image, thus implicating the public display right of § 106(4) of the Copyright Act.<sup>214</sup> The *Perfect 10* court held that framers do not publicly display images, because they did not place the copy of the image on the server from which it is communicated to the viewer of the framing web page.<sup>215</sup> According to *Perfect 10*, only the person who has placed the image on the server is publicly displaying it. That holding, like the holding in *Fox Broadcasting*, limits copyright owner control. Once the owner makes an image

---

United States, the Library of Congress). In spite of that position, Congress passed the Family Movie Act, codified at 17 U.S.C. § 110(11), which specifically immunizes such technology from copyright liability.

212. 487 F. 3d 701 (9th Cir. 2007).

213. See *id.* at 712.

214. 17 U.S.C. § 106(4).

215. See *Perfect 10 Inc.*, 487 F.3d at 716–18.



publicly available, they cannot legally prevent the image from being placed in a new context that may undercut their original goals in displaying the image to the public.

In the case of framing and the public display right, the Ninth Circuit's holding has not gone unchallenged. Three judges in the Southern District of New York and one in the Northern District of Texas have come out the other way, holding that framing is public display, in spite of the Ninth Circuit precedent.<sup>216</sup> However, none of those cases made to an appellate court, and the Ninth Circuit recently reaffirmed and broadened its commitment to the "server test" in *Hunley v. Instagram*.<sup>217</sup> The *Fox Broadcasting* issue of ad-skipping has not produced a judgment in any other circuit.

The Ninth Circuit's precedent on ad-skipping and framing may or may not hold up in the long run. Even if it does, however, it does not support the position that machine extraction of constitutive expression from copyrighted works is equivalent to human reading, viewing or listening. When human beings view a TV show without advertisements, or view a photograph framed by a different web page than the one intended by the photographer, they still have an aesthetic reaction to those works, and can remember and act on that aesthetic reaction. In other words, the filmmaker and the photographer are still building reputations with viewers. Indeed, if a provider of ad-skipping technology also edited out TV show credits, or if a framer also obscured photograph credits, either visible or in metadata, they would be liable for removing copyright management information, and the Ninth Circuit decisions would not save them.<sup>218</sup> By contrast, the generative machine learning algorithm has no capacity for aesthetic reaction, memory, or ability to act. It is extracting constitutive expression without taking part in any of the social, cultural, or economic world within which works of authorship are produced and appreciated, and within which authors gain compensation, social recognition, or whatever else they desire. Thus, generative machine learning algorithms do not "read" in the sense in which humans do, and their processing of works should not count as immunized "reading" for copyright purposes.

### C. REMEDIES

Suppose that a court concludes that generative AI training is not a fair use of the copyrighted works that are in the training set, and that the generative AI developer has therefore infringed copyright in the works that it did not get permission to use. What remedies should be available to the plaintiffs? Here I will say something that won't be quite as favorable to them as the arguments I have made against fair use. I argued above that even if a person learns something very valuable from an infringing copy of a work, and incorporates what they learned into a new work, actual damages

---

216. See *McGucken v. Newsweek LLC*, 2022 WL 836786 (S.D.N.Y. Mar. 21, 2022) (per Judge Failla); *Nicklen v. Sinclair Broad. Grp., Inc.*, 551 F. Supp.3d 188 (S.D.N.Y. 2021) (per Judge Rakoff); *Goldman v. Breitbart News Network, LLC*, 302 F. Supp.3d 585 (S.D.N.Y. 2018) (per Judge Forrest); *Leader's Institute, LLC v. Jackson*, 2017 WL 5629514 (N.D. Tex. Nov. 22, 2017) (per Judge Boyle).

217. 73 F.4th 1060 (9th Cir. 2023).

218. See 17 U.S.C. § 1202(b).

should not take into consideration profits earned from that new work, so long as that person does not incorporate what I have called “actionable expression” from the infringing copy into the new work.<sup>219</sup> I also argued that so long as the new work contains no actionable expression from the infringed work, a court is not justified in ordering the destruction of the new work.<sup>220</sup>

Those limitations on remedies should also apply in the generative AI training context, with some elaboration. The value of particular outputs or generations from generative AI models should not be taken into account in assessing actual damages, unless those outputs are substantially similar to one or more training works. However, the best measure of damages is not the price that a human being would pay for an authorized copy of an infringed work. A deployed generative AI model like OpenAI’s ChatGPT may have 10 million or more paying subscribers.<sup>221</sup> That represents economic activity on a completely different scale than that of an individual human author. There is a licensing market developing for AI training use of works, and ideally damages should be keyed to a reasonable licensing fee. Although several of the plaintiffs in the generative AI training infringement lawsuits have asked the court to order destruction of the models,<sup>222</sup> that seems warranted, if at all, only if the models have “memorized” substantial numbers of training works. If not—if the models do not themselves contain infringing copies—and they should not be subject to destruction orders. That may give the plaintiffs’ attorneys less leverage than they would like, but it avoids a dramatic expansion of remedies that would not be available if a human author had learned from infringing copies. To this extent, the limitations in copyright remedies should be respected. Of course, the attorneys who are representing plaintiffs who have registered their works before their use in generative training are salivating over statutory damages and attorneys fees, and those can be awarded under current copyright law, though there is room for legislative reform in that area.

#### IV. CONCLUSION

Copyright law does not grant human authors blanket fair use immunity from infringement liability when they are learning from copyrighted works. Should there be such immunity when people are using computers to train a generative AI model from copyrighted works? This Article concludes that there most likely should not be.

Generative AI model training does use the constitutive expressive choices embodied in the training works. While the computers processing the works do not have any hedonic reactions to those works, and they thus make non (felt)expressive

---

219. See *supra* text accompanying notes 54–57.

220. See *supra* text accompanying note 59.

221. ChatGPT has an estimated 9.88 million paying subscribers. See *ChatGPT Has More Paid Subscribers than What Streaming Service?*, GOV’T TECH., (July 12, 2024), <https://www.govtech.com/question-of-the-day/chatgpt-has-more-paid-subscribers-than-what-streaming-service> [https://perma.cc/BB8T-YH2Q].

222. See, e.g., Complaint at 68, *New York Times Co. v. Microsoft Corp.*, No. 23-cv-11195 (S.D.N.Y. Dec. 27, 2023) (“ask[ing] the court to “[o]rde[r] destruction under 17 U.S.C. § 503(b) of all GPT or other LLM models and training sets that incorporate Times Works”).

use of those works, that should not make a legal difference. Copyright law protects the learning value of works as well as their hedonic value.

Neither should generative model training be treated as functionally equivalent to human reading, viewing, or listening that does not implicate any of the exclusive rights. Generative model training transcends the human limitations that underlie the structure of the exclusive rights, and it therefore cannot take advantage of that structure. Moreover, the very inability of computers to have any hedonic reactions to the works they are processing, and their inability to act on those reactions, makes computer processing fundamentally different than human experience of works.

There are serious policy concerns about the effect of copyright law on generative AI. Those include concerns about performance, concentration, bias, and international competitiveness. At this early point in the history of generative AI, however, none of those concerns are so clear that courts should step in and grant blanket immunity for generative AI training. The current case that generative AI training is a fair use is weak.