

February 14, 2020

ONLINE SUBMISSION VIA wipo.int

World Intellectual Property Organization
34, chemin des Colombettes
CH-1211 Geneva 20, Switzerland

**Re: WIPO Conversation on Intellectual Property (IP) and Artificial Intelligence (AI)
Comments of Benjamin L. W. Sobel**

To Whom It May Concern:

Benjamin L. W. Sobel respectfully submits the following comments in response to the World Intellectual Property Organization's call for comments dated December 13, 2019.¹

Mr. Sobel is an Affiliate at Harvard University's Berkman Klein Center for Internet & Society and a J.D. student at Harvard Law School. His comments reflect his personal views only. Attached as an appendix to these comments, please find a copy of Mr. Sobel's article, *Artificial Intelligence's Fair Use Crisis*, as published in the *Columbia Journal of Law & the Arts* in 2017.

Comments

A. Issue 7

The Draft Issues Paper asks whether the use of the data subsisting in copyrighted works without authorization for machine learning may constitute an infringement of copyright. The answer is yes—but only with respect to a subset of AI applications, which I refer to as “market-encroaching” uses of copyrighted works. My comments divide relevant AI applications into two categories: “non-expressive” uses and “market-encroaching” uses. With a focus on United States policy and decisional law, these comments explain why the former do not constitute copyright infringement, while the latter may. Finally, these comments recommend an exceptions framework that considers whether or not a use is market-encroaching in purpose.

1. Non-Expressive Uses Do Not Infringe Copyright

Non-expressive uses reproduce copyrighted works in order to derive value from aspects of those works other than protected expression. *See generally* Matthew Sag, *Orphan Works as Grist for the Data Mill*, 27 Berkeley Tech. L.J. 1503 (2012); *see also* James Grimmelman, *Copyright for*

¹ These comments are adapted from the commenter's submission to the United States Patent and Trademark Office's request for comments on Intellectual Property Protection for Artificial Intelligence Innovation (PTO-C-2019-0038), available at <https://www.bensobel.org/files/misc/Sobel-AI-USPTO-Comment-PTO-C-2019-0038.pdf>.

Literate Robots, 101 Iowa L. Rev. 657, 665 (2015). Even though non-expressive uses often implicate large-scale, unauthorized reproductions of copyrighted works, courts consistently hold them to be fair uses. *Authors Guild v. Google*, a leading non-expressive fair use case, held it fair use for Google to create digital reproductions of copyrighted texts in order to provide the public with factual information *about* those texts—such as how many times a search term appears—without providing a substitute for the expression in those texts. 804 F.3d 202, 225 (2d Cir. 2015). Similarly, the United States Court of Appeals for the Fourth Circuit has held that fair use permits the reproduction of student papers for use in a plagiarism-detection tool in part because comparing works’ textual similarity “is . . . unrelated to any creative component.” *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630, 641-42 (4th Cir. 2009).

The premise underpinning cases like *Authors Guild* and *iParadigms* is that copyright protects authorial expression, not facts. *Feist Publications, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 348 (1991). This principle is easy to apply to many emerging AI applications. That today’s AI may be more technologically sophisticated than the methods in use a decade ago in *iParadigms* does not alter this legal rationale, provided that an AI application’s ultimate purpose remains non-expressive. For example, facial recognition algorithms may train on a database of copyrighted photographs that have been collected without authorization. See Benjamin L. W. Sobel, *Artificial Intelligence’s Fair Use Crisis*, 41 Colum. J.L. & Arts 45, 67 (2017). Assembling such a dataset and using it to train an algorithm may reproduce or modify these photographs in ways that constitute *prima facie* copyright infringement. But the information being extracted from these photographs is precisely the information that does not originate in any act of authorship: the facial geometry of the photographic subjects.²

2. Market-Encroaching Uses May Infringe Copyright

The second category of artificial intelligence is what I refer to as “market-encroaching” uses. Market-encroaching uses ingest copyrighted expression for a purpose that endangers the market *for that very expression*, rather than merely engaging with the non-expressive aspects of copyrighted works. Because market-encroaching uses are by definition *expressive* uses, rather than non-expressive uses, their legality is uncertain. In fact, I have argued that market-encroaching uses of copyrighted materials go well beyond what fair use caselaw has permitted, and that it is far from obvious that fair use should or will excuse such uses. See Sobel, *supra*, at 68-79.

As AI technologies improve at isolating and replicating the aesthetically appealing aspects of copyrighted media, market-encroaching AI will grow in legal and commercial significance. Already, however, commercial technology exists that could enable market-encroaching uses of

² This analysis is not meant to suggest that facial recognition technology should not be regulated at all. Rather, it argues only that *copyright law* would not be an appropriate regulatory mechanism because facial recognition technology uses copyrighted works only in order to analyze non-expressive information contained in those works.

copyrighted works.³ Deepart.io uses artificial intelligence to combine the artistic style of a user-uploaded image with the content of a user-uploaded photograph and sells prints of the composite images its software generates. *See* Deepart, <https://deepart.io/>. Several startups use machine learning to generate novel, musical sound recordings that they then assign or license to customers. One such startup, Jukedeck, was acquired earlier this year by the Chinese AI firm Bytedance. Brenda Goh, *China's ByteDance ventures into AI-generated music with Jukedeck deal*, Reuters (July 24, 2019), <https://www.reuters.com/article/us-china-bytedance/chinas-bytedance-ventures-into-ai-generated-music-with-jukedeck-deal-idUSKCN1UJ0NN>.

Market-encroaching uses differ from non-expressive uses with respect to the two most important statutory fair use factors: factor one, “the purpose and character of the use,” and factor four, “the effect of the use upon the potential market for or value of the copyrighted work.” 17 U.S.C. § 107. First, the purpose of market-encroaching uses is not transformative in the same manner that non-expressive uses are. Consider an AI model that ingests copyrighted musical compositions in order to generate novel musical compositions. To the extent this AI succeeds at its purpose, it has learned to identify and replicate expressive qualities of the copyrighted materials in its training corpus. Thus, to the extent that this AI derives value from its input data, it engages not with mere facts about copyrighted materials, but instead with the protected, expressive aspects of those materials. Such an expressive purpose does not make a use *per se* non-transformative. But it does make the rationale of non-expressive fair use unavailable. *See* Sobel, *supra*, at 72.

Second, because they are expressive in nature, market-encroaching uses threaten the potential market for the works used. A non-expressive AI analysis of a book might suggest with high certainty that it plagiarizes other texts. This revelation could, in turn, diminish the market for that book. But that market harm would not weigh against fair use because it results from promulgating information *about* a work, rather than usurping that work’s expressive value. In contrast, music-generating AI ingests musical works in order to learn how to generate aesthetically appealing musical works. If, as is likely, AI-generated music could be licensed at lower prices than conventionally-authored works, then this AI-generated music could diminish the demand for the human-authored works on which it trained. It could even force human composers out of some segments of the music market, like “stock” and background music. *See* Sobel, *supra*, at 79.

3. Exceptions to Copyright Should Contemplate a Use’s Market-Encroaching Potential

The Draft Issues Paper asks, “[i]f the use of the data subsisting in copyright works without authorization for machine learning is considered to constitute an infringement of copyright, should an exception be made for at least certain acts for limited purposes, such as the use in non-commercial user-generated works or the use for research?” Applying the rationale of market-encroaching and non-market-encroaching uses can simplify this question.

³ This comment should not be understood as a discussion of these companies’ particular data-use practices, which are unknown to the commenter. Rather, the comment refers to existing businesses only to illustrate that AI technologies capable of fulfilling market-encroaching purposes are in use today.

To the extent that uses of data for non-expressive purposes constitute *prima facie* infringements of copyright, an exception to copyright law should establish that non-expressive uses are non-infringing. Moreover, uses that are expressive in character are not necessarily market-encroaching. For example, academic research may entail training AI on copyrighted works without authorization, in order to generate expressive outputs. But so long as that research is conducted for academic purposes, rather than to usurp the market for its training data, it is not a market-encroaching use. Thus, any *prima facie* infringements that such academic work entails should fall within a copyright exception. The same is true for expressive uses that serve parodic, critical, or otherwise transformative purposes, and are therefore less likely to produce cognizable market harms. In contrast, machine learning that makes unauthorized use of expressive works, in order to create expressive works for commercial use in the same or similar markets that the training data occupy, is a market-encroaching use. An exception to copyright that also excused market-encroaching uses would be much more difficult to justify than an exception limited to non-market-encroaching uses.

An exception that does not give weight to a use's market-encroaching purpose risks being both over- and under-inclusive. For example, the European Union's Digital Single Market (DSM) Directive of 2019 requires member states to implement copyright exceptions related to text and data mining (TDM) by June 7, 2021. 2019 O.J. (L 130/92) 124. The Directive defines TDM as "any automated analytical technique aimed at analysing text and data in digital form in order to generate information which includes but is not limited to patterns, trends and correlations[.]" *Id.* at 112. It is not clear from this language alone whether the definition of TDM includes only non-expressive uses of copyrighted materials, or whether it would include market-encroaching uses as well. Article 3 of the DSM Directive permits certain reproductions of copyrighted materials by research and cultural heritage organizations for TDM research. *Id.* at 113. Article 4 extends Article 3's exception to *any entity* conducting TDM. *Id.* at 113-14. However, the Article 4 exception does not apply when rights holders have expressly reserved their TDM right. *Id.* Because it does not incorporate an assessment of the purpose of a use, Article 4's opt-out mechanism may be both over- and under-inclusive. A right to opt-out of TDM would be overbroad if it permitted rights holders to exclude others from making even non-expressive TDM uses of their works, which are rightfully outside copyright's monopoly over expression.

At the same time, the EU's opt-out mechanism may be too narrow, to the extent that it only permits rightsholders to withhold their works entirely or to extend a *gratis* license to TDM uses. In this regard, the United States fair use model is also ill-suited to respond to market-encroaching uses, because fair use either condemns such uses entirely or legalizes them outright. A more effective policy intervention would facilitate *licensed* market-encroaching uses, rather than giving a rights holder only a choice between extending a *gratis* license or asserting a right to exclude.

Conclusion

The commenter thanks WIPO for the opportunity to comment on these issues. The undersigned happily would provide further information in response to any questions WIPO may have.

*WIPO Conversation on Intellectual Property (IP)
and Artificial Intelligence (AI)
WIPO/IP/AI/2/GE/20/1
Comments of Benjamin L. W. Sobel*

Respectfully submitted,
Benjamin L. W. Sobel

Affiliate

Berkman Klein Center for Internet & Society at Harvard University
23 Everett Street, 2nd Floor
Cambridge, MA 02138

APPENDIX

Benjamin L. W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COLUM. J.L. & ARTS 45 (2017), available at https://www.bensobel.org/files/articles/41.1_Sobel-FINAL.pdf.

Artificial Intelligence's Fair Use Crisis

Benjamin L. W. Sobel*

ABSTRACT

As automation supplants more forms of labor, creative expression still seems like a distinctly human enterprise. This may someday change: by ingesting works of authorship as "training data," computer programs can teach themselves to write natural prose, compose music, and generate movies. Machine learning is an artificial intelligence ("AI") technology with immense potential and a commensurate appetite for copyrighted works. In the United States, the copyright law mechanism most likely to facilitate machine learning's uses of protected data is the fair use doctrine. However, current fair use doctrine threatens either to derail the progress of machine learning or to disenfranchise the human creators whose work makes it possible.

This Article addresses the problem in three Parts: using popular machine learning datasets and research as case studies, Part I describes how programs "learn" from corpora of copyrighted works and catalogs the legal risks of this practice. It concludes that fair use may not protect expressive machine learning applications, including the burgeoning field of natural language generation. Part II explains that applying today's fair use doctrine to expressive machine learning will yield one of two undesirable outcomes: if U.S. courts reject the fair use defense for machine learning, valuable innovation may move to another jurisdiction or halt entirely; alternatively, if courts find the technology to be fair use, sophisticated software may divert rightful earnings from the authors of input

* Affiliate, Berkman Klein Center for Internet & Society, Harvard University. The author thanks William Fisher and Jonathan Zittrain for their invaluable mentorship, and James Grimmelman for detailed comments. He also thanks Saptarishi Bandopadhyay, Christopher Bavitz, Andrei Ciupan, Jessica Fjeld, Priscilla Guo, Adam Holland, Kevin Hong, Lily Hu, Haijun Jin, Rachel Kalmar, Mason Kortz, David Lehr, Diana Liebenau, Fred von Lohmann, Rachel Orol, Alexander Reben, Sam Richman, Daniel Schönberger, Joel Sobel, Michael Veale, Jordi Weinstock, Jacques de Werra, and Kathryn Woolard for discussions and support during the writing of this Article. In addition, the author is grateful to the editors of this journal for their careful review of this manuscript, and to the staff and fellows of the Berkman Klein Center and its Ethics and Governance of Artificial Intelligence Effort. Drafts of this Article were presented at the 2017 Geneva Internet L@W Research Colloquium in Geneva, Switzerland, and the 2017 CopyrightX Summit in Cambridge, Massachusetts; the author thanks the organizers and attendees of these conferences.

© 2017 Sobel. This is an open access article distributed under the terms of the Creative Commons CC Attribution-NonCommercial-ShareAlike 4.0 license, which permits unrestricted non-commercial use, distribution, and reproduction, provided the original author and source are credited.

data. This dilemma shows that fair use may no longer serve its historical purpose. Traditionally, fair use is understood to benefit the public by fostering expressive activity. Today, the doctrine increasingly serves the economic interests of powerful firms at the expense of disempowered individual rights holders. Finally, in Part III, this Article contemplates changes in doctrine and policy that could address these problems. It concludes that the United States' interest in avoiding both prongs of AI's fair use dilemma offers a novel justification for redistributive measures that could promote social equity alongside technological progress.

Abstract	45
Introduction	47
I. Fair Use May Not Excuse Machine Learning	49
A. Mere Machinery	50
B. Non-Expressive Fair Use	51
C. A Technical Primer on Machine Learning	57
D. Potential Copyright Infringements in Machine Learning	61
1. Literal Reproduction in Datasets	61
2. Literal Reproduction as an Intermediate Step	62
3. Non-Literal Reproduction in Datasets and Models	63
4. Infringing Output	65
E. Machine Learning and Fair Use	66
1. Some Machine Learning is Non-Infringing	67
2. Google's Natural Language Generation	68
3. "The Next Rembrandt"	71
4. "Productive" or "Consumptive"?	72
5. Delineating the "Potential Market"	75
6. A New Threat of Market Substitution	77
II. AI's Fair Use Dilemma	79
A. One Bad Outcome: Training AI is Infringement	80
B. Another Bad Outcome: Training AI is Fair Use	81
C. Fair Use: Old and New Narratives	82
1. The Familiar Narrative: Big Content, Little Users	85
2. The Familiar Narrative Inverted	86
D. What Kind of Progress?	89
III. What Can Be Done?	90
A. Levies	91
B. Doctrinal Innovations	93
C. Moonshots	93
D. Copyright Is No Silver Bullet	94
1. Contracts	94
2. Algorithmic Bias	95
Conclusion	96

INTRODUCTION

“*Uno no es lo que es por lo que escribe, sino por lo que ha leído.*”¹

It is the year 2017, and artificial intelligence (“AI”) is in vogue. Technology that promises to boost productivity, fight wars, replace human labor, and perhaps even precipitate humankind’s demise has captivated engineers, policymakers, academics, artists, and the general public.² Copyright scholars have been eager to get in on the action. The fast-growing body of legal scholarship on AI and copyright tends to focus on one question: what happens when computers produce outputs that resemble conventional, expressive works of art by human creators, such as musical compositions, images, poetry, and prose—do these “works” have “authors?”³ “Can a computer be an author?” is an enticing query because it raises deep questions about the nature of creativity, autonomy, and human expression.

This question, though more popular than ever, is not as novel as it may seem.⁴ Over a century has passed since the Supreme Court first evaluated whether the outputs of a new creative technology, capable of operating with less human oversight than its predecessors, could manifest authorship to the degree intellectual property laws required.⁵ That technology was photography; courts asked the same questions of video games in the early 1980’s,⁶ and intellectual property scholars have been pondering for at least thirty years whether a computer program can be an

1. “One is not what he has written, but what he has read.” Jorge Luis Borges, *quoted in* Gideon Lewis-Kraus, *The Great A.I. Awakening*, N.Y. TIMES (Dec. 14, 2016), <https://perma.cc/SBQ9-K899>; *but see Talk: Jorge Luis Borges*, WIKIQUOTE, <https://perma.cc/2Z84-MQXD> (last visited July 29, 2017) (noting that the quotation is unsourced).

2. See EXEC. OFFICE OF THE PRESIDENT, *ARTIFICIAL INTELLIGENCE, AUTOMATION, AND THE ECONOMY* (2016), available at <https://perma.cc/5WN8-5UMD>; NICK BOSTROM, *SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES* 115 (2014); Stuart Russell, Daniel Dewey & Max Tegmark, *Research Priorities for Robust and Beneficial Artificial Intelligence*, 36 *AI MAG.* 105, 112–13 (2015).

3. See, e.g., Bruce E. Boyden, *Emergent Works*, 39 *COLUM. J.L. & ARTS* 377, 378 (2015); Annemarie Bridy, *Coding Creativity: Copyright and the Artificially Intelligent Author*, 5 *STAN. TECH. L. REV.* (2012); William Ralston, *Copyright in Computer-Composed Music: Hal Meets Handel*, 52 *J. COPYRIGHT SOC’Y USA* 281 (2005); Pamela Samuelson, *The Future of Software Protection: Allocating Ownership Rights in Computer-Generated Works*, 47 *U. PITT. L. REV.* 1185 (1986).

4. See James Grimmelmann, *There’s No Such Thing as a Computer-Authored Work - and It’s a Good Thing, Too*, 39 *COLUM. J.L. & ARTS* 403, 403 (2016).

5. See *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53 (1884); *see also infra* *Mere Machinery*.

6. See *Midway Mfg. Co. v. Artic Intern., Inc.*, 704 F.2d 1009, 1011 (7d Cir. 1983) (“There is a second difficulty that must be overcome if video games are to be classified as audiovisual works.... The question is whether the creative effort in playing a video game is enough like writing or painting to make each performance of a video game the work of the player and not the game’s inventor.”); *see also* *Atari Games Corp. v. Oman*, 888 F.2d 878, 879–80, 884 (D.C. Cir. 1989) (discussing, but ultimately rejecting, the theory that the audiovisual displays in a video game are uncopyrightable because they are created by the game’s player, not by its author).

“author” for the purposes of copyright.⁷ Moreover, even after all these years, it is not clear that the question of computer authorship is relevant to any real-world facts: James Grimmelmann noted in 2015 that “no one has ever exhibited even one work that could plausibly claim to have a computer for an ‘author’ in the sense that the Copyright Act uses the term.”⁸

It is not surprising that so much discussion of AI and copyright law focuses on the question of authorship. Indeed, it reflects the fascination with individual, mythic authors that is an entrenched—and rightly criticized—flaw of copyright jurisprudence more generally.⁹ That this emphasis is unsurprising makes it no less unfortunate. Just as collaboration and creative influence deserve more consideration in discussions of human authorship, so too do these concepts deserve more consideration in the analysis of artificial intelligence and copyright law.¹⁰ This is because today’s cutting-edge technology requires potential AI authors to first be readers.

While computers capable of authorship—if they exist—remain few and far between, artificial intelligence is now consuming copyrighted works at a tremendous pace. Much as human creators learn from the works of their human predecessors, a technology called “machine learning” allows today’s AI to emulate works of human authorship after being provided with many examples. Depending on the data on which it is trained, an AI could learn to generate prose, paintings, motion pictures, musical compositions, and so on. These “training data” often comprise thousands of unauthorized copies of copyrighted works, which are reduplicated and modified countless more times throughout the training process.

Massive, computerized reproduction of copyrighted works is not, in itself, a new legal issue. Google Images and Google Books are services powered by the unauthorized copying of protected expression, and both have been excused by the fair use doctrine. But applications like these have been found non-infringing largely because they do not purport to be expressive works in themselves and do not resemble copyright’s traditional subject matter.¹¹ The jurisprudence on these computerized copyists tends to treat them as mere processors of existing expression that assemble many individual works into non-expressive, factual “reference tool[s].”¹² New applications of machine learning, which use expressive works to teach AI expressive skills in order to generate new works, change this reasoning.

7. See REGISTER OF COPYRIGHTS, SIXTY-EIGHTH ANNUAL REPORT OF THE REGISTER OF COPYRIGHTS 5 (1966), available at <https://perma.cc/5VYM-RNZD>; see also Samuelson, *supra* note 3.

8. Grimmelmann, *supra* note 4, at 403.

9. See Peter Jaszi, *Toward a Theory of Copyright: The Metamorphoses of “Authorship”*, 1991 DUKE L.J. 455, 455 (1991); Peter Jaszi, *On the Author Effect: Contemporary Copyright and Collective Creativity*, 10 CARDOZO ARTS & ENT. L.J. 293, 294 (1992); see also Michel Foucault, *What Is an Author?*, THE NORTON ANTHOLOGY OF THEORY AND CRITICISM 1622 (Donald F. Bouchard & Sherry Simon trans., 2001) (discussing problems with overemphasizing an author-figure in literary criticism).

10. See, e.g., Giancarlo F. Frosio, *A History of Aesthetics from Homer to Digital Mash-ups: Cumulative Creativity and the Demise of Copyright Exclusivity*, 9 LAW HUMAN. 262, 263 (2015).

11. See James Grimmelmann, *Copyright for Literate Robots*, 101 IOWA L. REV. 657, 665 (2015).

12. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1165 (9th Cir. 2007).

Copyright law forces artificial intelligence into a binary: it is either a mystical author or a dumb machine. State-of-the-art machine learning is not exactly either. This is a lacuna in the scholarship, and the present Article fills it. Part I describes how machine learning works, and why it may not be able to rely on fair use to excuse the reproduction and analysis of copyrighted works that it entails. Part II describes how expressive machine learning¹³ poses a dilemma for the fair use doctrine: broadly speaking, neither of the two outcomes the doctrine can reach appears equitable or desirable. This dilemma shows that fair use may no longer serve its historical purpose. Traditionally, fair use is understood to benefit the public by fostering expressive activity. Today, the doctrine increasingly serves the economic interests of powerful firms at the expense of disempowered, individual rights holders. Finally, Part III contemplates how today's copyright doctrine might be used to promote distributive equity in the AI age.

I. FAIR USE MAY NOT EXCUSE MACHINE LEARNING

Technological innovation fosters new and valuable uses of copyrighted works: search engines index trillions of webpages, images, and videos; translation services use text from all over the internet to improve their fidelity; news aggregators direct readers to important stories written by third parties. These uses often take place without rights holders' explicit authorization, and, for this reason, depend on exceptions and limitations to copyright law. The most prominent of these limitations is the fair use doctrine, set forth in § 107 of the Copyright Act. Many technologies owe their existence to a successful fair use defense.¹⁴ Others owe their demise to a failed fair use defense.¹⁵ Fair use has yet to assess machine learning, and the doctrine could make or break the technology's future.

To adjudicate fair use, judges apply a four-factor standard that evaluates: "(1) the purpose and character of the use; (2) the nature of the copyrighted work; (3) the amount and substantiality of the portion used; (4) the effect of the use upon the potential market for or value of the copyrighted work."¹⁶ The doctrine is

13. "Expressive machine learning," an important term in this Article, is difficult to define precisely. This is largely because "expression" itself is difficult to define and delimit. Human expression takes infinitely many forms and can be embodied in countless distinct media; the United States copyright statute acknowledges that protectable expression may be "fixed in any tangible medium ... now known or later developed, from which [the work] can be perceived, reproduced, or otherwise communicated[.]" 17 U.S.C. § 102 (West 2017). In addition to the hazy nature of expression, the breadth of potential applications of machine learning makes "expressive machine learning" still harder to define. For the purposes of this Article, the term can be understood to refer to machine learning that trains on the expressive aspects of works, in order to fulfill an expressive purpose. Determining the exact legal boundaries of expressiveness in the artificial intelligence context will likely fall to courts in the coming years.

14. See *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007) (holding that an image search engine's unauthorized reproductions of copyrighted photographs are excused by fair use).

15. See *Capitol Records, LLC v. ReDIGI Inc.*, 934 F.Supp.2d 640 (S.D.N.Y. 2013) (holding that fair use does not excuse a secondhand marketplace for digital sound recordings from liability for infringement).

16. 17 U.S.C. § 107 (2017).

notoriously protean; today, the most important consideration is whether or not a use is “transformative.” In a 1990 article, Judge Pierre Leval outlined the canonical definition of transformativeness:

The use must be productive and must employ the quoted matter in a different manner or for a different purpose from the original. A quotation of copyrighted material that merely repackages or republishes the original is unlikely to pass the test; in Justice Story’s words, it would merely “supersede the objects” of the original. If, on the other hand, the secondary use adds value to the original — if the quoted matter is used as raw material, transformed in the creation of new information, new aesthetics, new insights and understandings — this is the very type of activity that the fair use doctrine intends to protect for the enrichment of society.¹⁷

The Supreme Court adopted Judge Leval’s transformativeness test four years later.¹⁸ Other considerations, such as a use’s effect on the potential market for the work used, still influence a fair use defense. But transformativeness is the weightiest factor, and “the more transformative the new work, the less will be the significance of other factors, like commercialism, that may weigh against a finding of fair use.”¹⁹

Transformative fair use protects some of machine learning’s precursor technologies, and many people doubtless assume it will shield machine learning, too. However, this Part argues that certain applications of machine learning may not be excused by fair use. First, it briefly documents copyright’s historical tendency to treat mechanical processes as separate from protectable expression. It then explains how this view of machinery has informed a doctrine of “non-expressive” fair use, which permits computers to reproduce and derive information from copyrighted works in ways that, if done by humans, would be infringement. After giving an account of the non-expressive use doctrine, this Part describes how machine learning technology operates and catalogs the copyright liabilities it entails. Finally, this Part explains why certain applications of machine learning challenge the doctrine of non-expressive use by recasting the analysis of the two most important factors of fair use: the purpose of the use, and its effect on the market for the works used.

A. MERE MACHINERY

When novel technologies emerge, society often doubts their ability to facilitate human expression, particularly when those technologies mediate between a human subject and the expressive output she creates.²⁰ Copyright jurisprudence, too, tends

17. Pierre N. Leval, *Toward a Fair Use Standard*, 103 HARV. L. REV. 1105, 1111 (1990).

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

19. *Id.* at 579.

20. For an over two thousand year-old criticism of written language as inferior to memorized language, see PLATO, PHAEDRUS (Benjamin Jowett trans., 1999), available at <https://perma.cc/TF8J-FXCY>. See also, Jordan Teicher, *When Photography Wasn’t Art*, JSTOR DAILY (Feb. 6, 2016), <https://perma.cc/P3DY-R26C> (surveying criticisms of photography); Ira Flatow, *Digital Music*

to look skeptically at machinery inserted into the expressive process. For example, when the Supreme Court first evaluated whether photographs ought to be considered copyrightable works of authorship in the 1884 case *Burrow-Giles Lithographic Co. v. Sarony*, the opinion observed in dicta that the “ordinary production of a photograph” may be a process that is “merely mechanical, with no place for novelty, invention or originality[.]”²¹ Only because the plaintiff in the case proved the “existence of those facts of originality, of intellectual production, of thought, and conception on the part of the author” of the photograph was the Court willing to grant it copyright protection.²² The *Burrow-Giles* analysis suggests that purely mechanical encodings or transcodings of observed phenomena are facts, not authorial expression.²³ More than a century of copyright jurisprudence has reinforced this attitude: to give one example, the most recent Supreme Court precedent on originality in copyright, *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, comments that the “mechanical or routine” composition of a work is insufficiently original to qualify for protection.²⁴

In general, reactions to new expressive technologies, in court or in culture at large, reveal a belief that machines cannot in themselves impart, apprehend, or evince authorial expression. Machinery must be used by or under the direction of an author who has “original intellectual conceptions” if it is to produce copyrightable subject matter.²⁵ Of course, as particular technologies are adopted, and observers become better attuned to their expressive affordances, this initial mistrust tends to lessen. The copyright statute itself accommodates this pattern by defining a copy as a work “fixed by any method now known or later developed” and a “device,” “machine,” or “process” as “one now known or later developed.”²⁶ In other words, while some form of skepticism towards undue mechanical intervention in creativity is more or less consistent over a century of technological innovation, notions of what exactly constitutes undue intervention vary with historical context.

B. NON-EXPRESSIVE FAIR USE

Just as copyright treats machines as too dumb to count as authors, it also treats machines as too dumb to count as readers.²⁷ If machines cannot create authorial expression by themselves, it makes sense to infer that machines cannot engage with

Sampling: Creativity Or Criminality?, NATIONAL PUBLIC RADIO (Jan. 28, 2011), <https://perma.cc/WQZ2-B9RZ>.

21. *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 59 (1884). The writings of early photographers and critics show a similar attitude: an 1865 criticism of photography notes, “[a]ll labor of love must have something beyond mere mechanism at the bottom of it.” *Art and Photography*, 2 NEW PATH 198, 199 (1865).

22. *Burrow-Giles*, 111 U.S. at 59–60.

23. *Id.* at 59.

24. *Feist Publications, Inc. v. Rural Telephone Service Co.*, 499 U.S. 340, 362 (1991); *see also* Bridy, *supra* note 3, at 6 n.29 (listing copyright decisions that disparage mechanical creation).

25. *Feist*, 499 U.S. at 346–47; *Burrow-Giles*, 111 U.S. at 58.

26. 17 U.S.C. § 101 (West 2017).

27. Grimmelmann, *supra* note 11, at 658.

or appreciate that expression, either. As computers have become more efficient at reproducing, storing, and analyzing vast amounts of copyrighted works, copyright law has distinguished this activity from human consumption and excused much of it as transformative fair use. The scholars who have identified this doctrine call it “non-expressive” use, and it has evolved considerably over the past quarter-century.²⁸

The Ninth Circuit’s *Sega v. Accolade* decision is one of the earliest judicial recognitions of non-expressive fair use.²⁹ *Sega* dealt with a dispute between defendant Accolade, a company that developed and marketed video games, and plaintiff Sega, the developer and marketer of the “Genesis” gaming console. Part of Sega’s business model was to license copyrighted software to independent video game developers, who in turn developed and sold video games compatible with the Genesis console. Accolade sought to develop Genesis-compatible games without entering into a licensing agreement with Sega. To do so, Accolade purchased a Genesis console and several Sega game cartridges and reverse-engineered them in order to copy the functional computer code that allowed game cartridges to operate with the console. Accolade’s reverse-engineering process necessarily generated verbatim reproductions of the source code for several Sega games, but only the functional code that pertained to the Genesis interface was ultimately included in Accolade’s games.³⁰

The Ninth Circuit held that Accolade’s “intermediate copying” of Sega games was fair use, because it was necessary to gain access to the “functional requirements for Genesis compatibility”—a functional element of Sega’s games ineligible for copyright protection.³¹ It did not rule on whether or not Accolade’s finished games infringed Sega’s copyright, but a claim of infringement could not be sustained solely by the inclusion of Sega’s unprotectable functional code in Accolade’s games, and the language of the opinion insinuates that the defendants’ final games did not infringe.³² This ruling set the precedent that the unauthorized reproduction of copyrighted works, if incidental to a non-expressive purpose, was non-infringing fair use.³³

The next major development in the doctrine of non-expressive use came with two cases evaluating unauthorized use of copyrighted images in search engines, *Kelly v. Arriba* and *Perfect10 v. Amazon*.³⁴ In both cases, the defendants were tech companies that operated image search engines, and the plaintiffs owned copyrights to images that had been reproduced in thumbnail form, stored on defendants’

28. See Grimmelmann, *supra* note 11, at 661; Matthew Sag, *Orphan Works as Grist for the Data Mill*, 27 BERKELEY TECH. L.J. 1503, 1503 (2012).

29. *Sega Enterprises Ltd. v. Accolade, Inc.*, 977 F.2d 1510 (9th Cir. 1992); Grimmelmann, *supra* note 11, at 661.

30. *Sega*, 977 F.2d at 1515.

31. *Id.* at 1518, 1526.

32. *Id.* at 1527–28; Grimmelmann, *supra* note 11, at 667 n.27.

33. Grimmelmann, *supra* note 11, at 662.

34. *Kelly v. Arriba Soft Corp.*, 336 F.3d 811 (9th Cir. 2003); *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007).

servers, and displayed to internet users by the defendants' services.³⁵ And in both cases, these activities were found to be transformative fair uses.

The image search cases added a wrinkle to *Sega's* rationale because, unlike *Accolade*, the defendants had duplicated expressive aspects of the plaintiffs' works and presented some of that expression to the public. In *Kelly*, the defendant, Arriba Soft Corp., operated a search engine that produced thumbnail-sized images in response to search queries. These thumbnails were collected using a "crawler," an automated computer program that traverses the web, visiting and indexing the pages it encounters.³⁶ When it came upon an image, Arriba's crawler would download full-size copies of that image to Arriba's servers. Arriba's software then reduced the images to thumbnail size, deleted the full size copies, and featured the thumbnails in its search results.³⁷ Google Image Search, the service at issue in *Perfect 10*, created, stored, and displayed thumbnails in more or less the same way.³⁸

To dispense with the argument that Arriba and Google unlawfully coopted protected expression, both decisions doubled down on the rhetoric of non-expressive machinery. The defendants' image search engines assembled photographs into "tool[s]"—mere machines—not vehicles for conveying expression.³⁹ The dispositive element of *Kelly's* fair use finding was Arriba's lack of artistic purpose in reproducing *Kelly's* images. While *Kelly's* photographs are "artistic works intended to inform and to engage the viewer in an aesthetic experience[.]" Arriba's thumbnails are merely instrumental: they are part of a "tool to help index and improve access to images[.]"⁴⁰ Arriba's use is not artistic expression, and, the court reasons, the thumbnails' low resolution makes it unlikely that any user would attempt to consume them for aesthetic, rather than referential, purposes: "The thumbnails do not stifle artistic creativity because they are not used for illustrative or artistic purposes and therefore do not supplant the need for the originals."⁴¹ This formulation of Arriba's purpose undergirded the court's conclusion that Arriba's use was transformative, which in turn allowed Arriba's fair use defense to prevail.⁴² In *Perfect 10*, too, Google Image Search repurposes

35. Both opinions also dealt with the legality of defendants' display of full-sized copies of plaintiffs' images that were stored on third party servers, through a practice called "inline linking" or "hotlinking." *Kelly* initially found such displays to be infringement, then withdrew its opinion for procedural reasons and issued an amended opinion that avoided the issue. *Kelly v. Arriba Soft Corp.*, 280 F.3d 934, 944-45 (9th Cir. 2002). The Ninth Circuit's ultimate resolution of this question in *Perfect 10* analyzed the display right to find that defendants did not in fact display the full-sized images. Because this dimension of the litigation does not implicate the fair use doctrine, this Article does not discuss it. *Perfect 10*, 508 F.3d at 1154-56; *Kelly*, 336 F.3d at 815.

36. *Kelly*, 336 F.3d at 815.

37. *Id.* at 815-16.

38. *Perfect 10*, 508 F.3d at 1155.

39. *Perfect 10*, 508 F.3d at 1165; *Kelly*, 336 F.3d at 818.

40. *Kelly*, 336 F.3d at 818.

41. *Id.* at 819-20.

42. *Id.* at 819-20.

images into “pointer[s] directing a user to a source of information” as part of an “electronic reference tool[,]” rather than aesthetic objects.⁴³

The capstone of non-expressive use jurisprudence is the Second Circuit’s recent decision in *Authors Guild v. Google Inc.*⁴⁴ Google Books, the service at issue in the case, adapted the doctrine to new subject matter: literature. In partnership with major libraries, Google scanned over twenty million books, some of which are copyright-protected, some of which are in the public domain, and most of which are out of print.⁴⁵ From these scans, Google assembled a corpus of machine-readable texts that powers its Google Books service. Google Books is a publicly-accessible search engine that performs several functions: first, it enables internet users to perform a keyword search on the Google Books corpus, which returns a list of all books in the corpus in which the queried terms appear, as well as the terms’ frequencies in each book.⁴⁶ These search results also include general information about the books returned by the search query, such as bibliographic information and frequent terms contained in the book, as well as links to sites where the book can be purchased, if available.⁴⁷ Second, Google Books sometimes enables users to view all or some of a book’s text. Books that are in the public domain are displayed in their entirety, as are books whose publishers have authorized Google to reproduce their text in full.⁴⁸ Other books are displayed in a “limited preview,” which displays a limited number of full-text pages, if the owners in those books’ copyright have consented to the display.⁴⁹ A third display option, “Snippet View,” shows key words and phrases in a book, as well as “a few snippets—a few sentences to display [a] search term in context.”⁵⁰ Third, Google Books offers an interface for researchers to “examine word frequencies, syntactic patterns, and thematic markers to consider how literary style has changed over time.”⁵¹ This information is conveyed chiefly through “n-grams,” phrases of up to five consecutive words that are matched with the frequency with which they appear in the Google Books corpus.⁵²

In *Authors Guild*, Judge Leval of the Second Circuit found Google Books’ unauthorized reproductions of copyrighted works a transformative fair use of the texts, largely because Google Books provides information “about” books, not the books’ expression.⁵³ Even though “snippet view” shows users the textual expression that surrounds a search term, it nevertheless furthers Google’s

43. *Perfect 10*, 508 F.3d at 1165.

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

45. *Id.* at 208.

46. *Id.* at 208–09.

47. *Id.* at 209.

48. *What you’ll see when you search on Google Books*, GOOGLE BOOKS, <https://perma.cc/3PJN-VNRF> (last visited Sept. 30, 2017).

49. *Id.*

50. *Id.*

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

52. *Authors Guild v. Google Inc.*, 804 F.3d 202, 209 (2d Cir. 2015); Alex Franz & Thorsten Brants, *All Our N-gram are Belong to You*, GOOGLE RESEARCH BLOG (Aug. 3, 2006), <https://perma.cc/EV9K-PBJU> (last visited May 23, 2017).

53. *Authors Guild v. Google Inc.*, 804 F.3d 202, 216–17 (2d Cir. 2015).

transformative purpose by contextualizing a term's usage within a book without revealing enough expression to "threaten the author's copyright interests[.]"⁵⁴

Authors Guild is notable because it deploys the logic of non-expressive use to circumscribe the "potential market" for a copyrighted work—among the most important factors in fair use analysis—in a way that the image search cases do not.⁵⁵ *Kelly* and *Perfect 10* do not spend much time dissecting the types of markets that may exist for a work and contemplating which of those markets a copyright owner is entitled to control. Instead, *Kelly* emphasized the low likelihood of meaningful market harms: if anything, Arriba's thumbnails would drive users towards the plaintiff's site, rather than detracting from his business.⁵⁶ Similarly, *Perfect 10* concluded that the factor favored neither party because evidence of market harms remained "hypothetical."⁵⁷

In contrast, *Authors Guild* explicitly notes that Google Books may well harm authors' markets, but such harms "will generally occur in relation to interests that are not protected by the copyright."⁵⁸ *Authors Guild* suggests that, for the purposes of fair use's fourth factor, the relevant potential market only encompasses consumers' interest "in the protected aspect of the author's work[.]"⁵⁹ A Google Books user interested in a single historical fact may encounter that fact in snippet view and, as a consequence, may decide not to procure an authorized copy from a bookstore or library. Google Books therefore might harm an author's market by deterring these purchasers. Nevertheless, this harm is immaterial to a fair use inquiry; Google does not implicate the owner's entitlements under copyright by furnishing a fact that appears amidst expression.⁶⁰ This reasoning is also used to rebut plaintiffs' claims that Google infringed their exclusive rights to prepare derivative works under § 106(2) of the Copyright Act.⁶¹ The opinion reiterates that plaintiffs' copyright interest does not extend to the information about their works that Google furnishes to the public, and, accordingly, it "does not include an exclusive derivative right to supply such information through query of a digitized copy."⁶²

Authors Guild's "protected aspect[s]" rationale adapts an approach to the fourth fair use factor that Judge Leval himself appears to have pioneered and popularized.⁶³ Taken at face value, it is a powerful limitation of the "potential

54. *Id.* at 218.

55. 4 MELVILLE B. NIMMER & DAVID NIMMER, *NIMMER ON COPYRIGHT* § 13.05 (2017).

56. *Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 821 (9th Cir. 2003).

57. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1168 (9th Cir. 2007).

58. *Authors Guild v. Google Inc.*, 804 F.3d 202, 224 (2d Cir. 2015).

59. *Id.* at 224.

60. *Id.* at 224.

61. *Id.* at 225.

62. *Id.* at 225.

63. Nimmer's copyright treatise explains, "Only the impact of the use in defendant's work of material that is protected by plaintiff's copyright need be considered under this factor. Thus, a court need not take into account the adverse impact on the potential market for plaintiff's work by reason of defendant having copied from plaintiff noncopyrightable factual material." 4 *NIMMER ON COPYRIGHT*, *supra* note 55, at § 13.05 (2017). To support this contention, the treatise cites two cases; the first is the Second Circuit's opinion in *Harper & Row, Publishers, Inc. v. Nation Enterprises*, which discounted the

market” factor. However, the extent of its power is not entirely clear. If a use does not appropriate works’ protectable aspects for human consumption, is this dispositive of the fourth factor? Or does factor four demand a broader inquiry, of which the protectable aspects analysis is simply a component? The opinion sometimes appears to limit the fourth factor to substitution of works’ “protected aspects;” at other points, it suggests that a fair use inquiry must assess whether the use “provides a meaningful substitute for the original,” with no mention of whether or not doing so must implicate authors’ protected expression.⁶⁴

If a work’s potential market is truly limited to consumers’ interest in its expression, then the narrower formulation of the authors’ market would be the correct one. However, in some contexts, Judge Leval seems hesitant to embrace this conclusion fully. For instance, he notes it a virtue of Google Books that it “does not provide snippet view for types of books, such as dictionaries and cookbooks, for which viewing a small segment is likely to satisfy the searcher’s need[,]” thereby avoiding revealing something that “... could usefully serve as a competing substitute for the original.”⁶⁵ If copyright owners’ interests in their works extend only to the works’ protected aspects, why is it pertinent that Google restricts snippet view for these especially factual categories of works? The “protected aspect[s]” view of the potential market suggests that Google could furnish a great many snippets from cookbooks without infringing the authors’ copyrights, notwithstanding the impact of that activity on the market for those works. After all, just as some readers purchase biographies to ascertain facts, many people surely purchase cookbooks and dictionaries solely to consult the factual information they comprise.

relevance of market harms due to the copying of “uncopyrightable information in which Harper & Row and Reader’s Digest may not claim ownership” and ultimately found fair use. *Harper & Row, Publishers, Inc. v. Nation Enterprises*, 723 F.2d 195, 208 (2d Cir.1983), *rev’d*, 471 U.S. 539 (1985). While Nimmer notes that the decision was “*rev’d* on other grounds” than its interpretation of factor four, it bears noting that the Supreme Court’s reversal in *Harper & Row* in fact rejected the Second Circuit’s analysis of market harms. *Harper & Row*, 471 U.S. 539 at 567. The second case Nimmer cites is *Salinger v. Random House*, in which then-District Judge Leval held that an unauthorized biography of J.D. Salinger, which quoted from unpublished letters authored by Salinger, did not harm the market for those letters “as contemplated by the fair use statute” because “readership may indeed be affected, but only by [defendant]’s report of unprotected, historical facts, not by his use of copyright protected expression.” *Salinger v. Random House, Inc.*, 650 F. Supp. 413, 425 (S.D.N.Y. 1986), *rev’d*, 811 F.2d 90 (2d Cir. 1987). On appeal, the Second Circuit reversed Judge Leval’s order and rejected his analysis that the market for Salinger’s letters would be unimpaired by their use in the biography, citing the new guidance it had received from *Harper & Row* throughout its opinion. *Salinger v. Random House, Inc.*, 811 F.2d 90, 99 (2d Cir. 1987). In sum, Judge Leval appears more willing to limit a copyright holder’s market entitlements under factor four than other jurists have been when faced with the same or similar facts. The “protected aspect[s]” interpretation of fair use’s fourth factor, while now venerable enough to appear in a copyright treatise, owes much of its prominence—and perhaps even its existence—to Judge Leval.

64. *Authors Guild v. Google Inc.*, 804 F.3d 202, 220, 225 (2d. Cir. 2015).

65. *Id.* at 222.

Computation enables novel forms of value extraction from copyrighted materials, and with it, new markets for the information in those works.⁶⁶ Because mechanical analysis of text is presumptively non-expressive, the “protected aspect[s]” rationale may dramatically reduce the fraction of the market over which copyright grants exclusive rights. Indeed, inventive entrepreneurs are already testing the bounds of rightsholders’ potential markets beyond the circumstances *Authors Guild* evaluated. While Google forbore from displaying the cookbooks in its book corpus, there is clearly immense value in such activity: Yummly, a startup that aggregates the non-protectable elements of online recipes, was valued at one hundred million dollars in 2015 and was acquired by Whirlpool in May 2017.⁶⁷

C. A TECHNICAL PRIMER ON MACHINE LEARNING

The doctrine of non-expressive fair use relies on two core premises to excuse massive, unauthorized copying undertaken by computers. The first is that machinery cannot, by itself, consume copyrighted expression in an infringing manner. Accordingly, the mechanical ingestion of works is a non-expressive purpose, provided it is not to facilitate human engagement with the works’ expression. The second premise is that these uses do not affect works’ potential markets in a way that is material to copyright law, because copyright owners’ entitlements do not encompass the non-expressive components of their works—the very components with which computerized analysis engages, and from which it can derive value.

Emerging applications of machine learning challenge both these premises of non-expressive use. First, machine 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; as a result, these uses of machine learning may no longer qualify as non-expressive in character. Second, machine learning technology could present a new type of threat to markets for authorial expression: rather than merely supplanting the market for individual works, expressive machine learning could also supersede human authors by replacing them with cheaper, more efficient automata.

This sub-Part describes the technology of machine learning to show how it can extract value from authorial expression. The discussion below distinguishes machine learning from other forms of AI, which may not pose similar copyright liabilities; it also delineates uses of machine learning that likely do qualify as non-expressive fair use from uses that do not.

66. For a discussion of some forms of value extraction from machine-readable works, see Maurizio Borghi & Stavroula Karapapa, *Non-Display Uses of Copyright Works: Google Books and Beyond*, 1 QUEEN MARY J. INTELL. PROP. 21, 32–37 (2011) (“Certainly, these new uses generate a ‘value’ over which many parties have a legitimate claim. It is not clear on which grounds all the value should be appropriated by one party only, namely the ‘user’!”).

67. Ingrid Lunden, *Whirlpool Acquires Yummly, the Recipe Search Engine Last Valued at \$100M*, TECHCRUNCH (May 4, 2017), <https://perma.cc/FBG4-DJR2>.

AI is a broad and nebulous field that encompasses far more methodologies than just machine learning. From a copyright scholar's point of view, what makes machine learning distinct is that it learns from vast corpora of input data without the guidance of human-programmed rules, rather than by applying narrower sets of rules and facts that are predetermined by its programmers. Indeed, well before machine learning was as prominent as it is today, computer programs could generate content that would resemble—but perhaps not rival—human expression. These types of programs tend to rely on rule based representations of the world that have been programmed into them by their designers. A form of AI called an “expert system” combines a knowledge base of facts, and rules derived from those facts, with an inference engine that reaches conclusions.⁶⁸ For example, early attempts to generate literature using AI often worked by codifying literary conventions as abstract “story grammars.”⁶⁹ These grammars could be deployed in tandem with databases of facts about a story—including information about the author, the universe in which it takes place, and so on—in order to generate a narrative.⁷⁰ Because this type of AI relies on uncopyrightable facts and procedures, it is less likely to entail copyright liabilities, although a copyright may of course subsist in the program as a whole. Furthermore, any copyrightable information that ends up incorporated into this AI is likely to be imparted by its developers or their collaborators, rather than by nonconsenting third parties.

Whereas the expert system approach described above starts with a small, high-quality compilation of human knowledge and builds a system around that knowledge, machine learning analyzes troves of data to discern valuable information without human intervention.⁷¹ Unlike traditional approaches to statistics, some machine learning techniques do not require researchers to make assumptions about the distribution of, or relationships within, the data they seek to analyze.⁷² Thus, rather than analyzing “hard-coded” knowledge baked in by human designers, machine learning ascertains patterns from “raw” data.⁷³ These methodologies encourage and reward the acquisition of large amounts of data.

Machine learning is chiefly a predictive technology, and many of its tasks fit into two general categories: classification and regression. Classification tasks associate input data with labels; an example is optical character recognition, a process that aims to identify written or printed characters as letters of the

68. SIMON KENDAL & MALCOLM CREEN, AN INTRODUCTION TO KNOWLEDGE ENGINEERING 29-32 (2007).

69. Ira Goldstein & Seymour Papert, *Artificial Intelligence, Language, and the Study of Knowledge*, 1 COGN. SCI. 84-123, 94-96 (1977); see also Lyn Pemberton, *A Modular Approach to Story Generation*, in PROCEEDINGS OF THE FOURTH CONFERENCE ON EUROPEAN CHAPTER OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS 217 (1989).

70. Pemberton, *supra* note 69, at 218.

71. Ali S. Hadi, *Expert Systems*, in INTERNATIONAL ENCYCLOPEDIA OF STATISTICAL SCIENCE 480 (2011); KEVIN P. MURPHY, MACHINE LEARNING: A PROBABILISTIC PERSPECTIVE 1-2 (2007).

72. MURPHY, *supra* note 71, at 18-19; Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1156-57 (2017).

73. IAN GOODFELLOW ET AL., DEEP LEARNING 2-3 (2016), available at <https://perma.cc/2SXZ-ZLVE>.

alphabet.⁷⁴ Regression tasks attempt to predict a continuous variable given a set of data that may influence that variable.⁷⁵ These tasks belong to a category of processes called “supervised learning.”⁷⁶ Supervised learning requires labeled data, that is, data that have been associated with facts about them.⁷⁷ An example of labeled data is a set of photographs of human faces, tagged with the identities of the persons pictured. Machine learning can also be conducted “unsupervised.” Unsupervised learning, by contrast, apprehends patterns in data without being prompted with a particular kind of label to predict; it just uncovers “interesting structure.”⁷⁸ A solution to a given machine learning task may, at various stages, use both supervised and unsupervised learning techniques.

By uncovering “interesting structure” in data rather than emphasizing features that humans have predetermined to be salient, machine learning techniques can extract and mimic features that people might find ineffable or difficult to discern, but that nevertheless encapsulate some unifying quality of a particular group of data. When trained on samples of an individual’s handwriting or recordings of a human voice, for example, machine learning models can mimic scrawls and drawls with uncanny accuracy.⁷⁹ The constellations of features that machine learning can appropriate might well be called “personality.” This capability complicates a prevailing assumption in copyright law: that traces of an author’s “personality” uniquely individuate works of authorship, and that those traces ought to be copyrightable.⁸⁰

The touchstone case on this subject, *Bleistein v. Donaldson Lithographic Co.*, dealt with the copyrightability of three chromolithographed advertisements for a circus. The Supreme Court, proclaiming that courts should not arbitrate artworks’ merit, held that the posters’ lack of artistic pretension did not preclude their copyrightability.⁸¹ Justice Holmes’s majority opinion reads:

The copy is the personal reaction of an individual upon nature. Personality always contains something unique. It expresses its singularity even in handwriting, and a

74. MURPHY, *supra* note 71, at 3.

75. *Id.* at 8–9.

76. *Id.* at 1–9.

77. *Id.* at 2.

78. MURPHY, *supra* note 71, at 9; *see also* Andrew Ng, *MachineLearning-Lecture01*, <https://perma.cc/3ZRV-RP5S>.

79. *See* Tom S. F. Haines et al., *My Text in Your Handwriting*, 35 ACM TRANS. GRAPH. 26:1 (2016) (describing a machine learning model that mimics handwriting); Natasha Lomas, *Lyrebird is a Voice Mimic for the Fake News Era*, TECHCRUNCH (Apr. 25, 2017), <https://perma.cc/ZQ3X-RCL7> (describing machine learning-based software that can speak text in famous humans’ voices); Alexander Reben, *Artificially Intelligent Impressions*, YOUTUBE (Mar. 13, 2017), <https://www.youtube.com/watch?v=CWLqIXCu3OM> (demonstrating imitations of famous vocal cadences generated by machine learning).

80. Benjamin L. W. Sobel, *The Hermeneutics of Authorship in Copyright Law—and their Subversion in Computer-Generated Works of Art 21–22* (Feb. 28, 2015) (unpublished manuscript) (on file with the author).

81. *Id.* at 20.

very modest grade of art has in it something irreducible, which is one man's alone. That something he may copyright unless there is a restriction in the words of the act.⁸²

Bleistein, a hugely influential judgment, enfeebled copyright's originality requirement.⁸³ Instead of the proof of "original intellectual conceptions" that courts had previously demanded, the *Bleistein* rationale gives any work wrought by a human author a presumed trace of copyrightable personality.⁸⁴ Subsequent case law diminished the standard even further: in *Alfred Bell & Co. Ltd. v. Catalda Fine Arts, Inc.*, Judge Frank observed in dicta that "A copyist's bad eyesight or defective musculature, or a shock caused by a clap of thunder, may yield sufficiently distinguishable variations. Having hit upon such a variation unintentionally, the 'author' may adopt it as his and copyright it."⁸⁵

A personality-centric copyright jurisprudence may never have been philosophically coherent or normatively desirable, but it nevertheless holds an intuitive appeal. The factors that individuate humans seem ineffable and inimitable, yet are nevertheless unmistakable. If these qualities are indeed intrinsic, affording property rights in them could strengthen autonomy, and if they are indeed unique to each person, individuals' rights would not rival one another. The uniqueness of every human face, for instance, was a premise that Francis Hargrave, an eighteenth-century advocate of copyright, used to justify proprietary, personal authorship.⁸⁶

However, recent advances in machine learning may refute the idea that unique personality subsists only in the individual human gesture, or at least undermine it as a justification of intellectual property rights. Machine learning techniques enable machines to identify and mimic the features that distinguish sensory data, even when those features are not qualities that humans can easily express or represent.⁸⁷ Today's technology can isolate the characteristics that individuate a human face and human handwriting, the very attributes that Hargrave and Justice Holmes used long ago as metonyms for protectable authorial personality.⁸⁸

82. *Bleistein v. Donaldson Lithographing Co.*, 188 U.S. 239, 250 (1903).

83. Diane Leenheer Zimmerman, *It's an Original!(!): In Pursuit of Copyright's Elusive Essence*, 28 COLUM. J.L. & ARTS 187, 203–207 (2015); Barton Beebe, *Bleistein, the Problem of Aesthetic Progress, and the Making of American Copyright Law*, 117 COLUM. L. REV. 319, 329 (2017).

84. *Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53, 58 (1884).

85. *Alfred Bell & Co. v. Catalda Fine Arts*, 191 F.2d 99 (2d Cir. 1951); Zimmerman, *supra* note 83, at 204.

86. FRANCIS HARGRAVE, AN ARGUMENT IN DEFENSE OF LITERARY PROPERTY 7 (1774); *see also* Mark Rose, *The Author as Proprietor: Donaldson v. Becket and the Genealogy of Modern Authorship*, REPRESENTATIONS 51, 72–73 (1988).

87. GOODFELLOW ET AL., *supra* note 73, at 1–8.

88. These applications of machine learning implicate other areas of law, such as the right of publicity, in addition to raising alarming ethical questions. *See* Reben, *supra* note 79; Lomas, *supra* note 79; *see also* Yaniv Taigman et al., *DeepFace: Closing the Gap to Human-Level Performance in Face Verification*, in 2014 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION 1701 (2014).

D. POTENTIAL COPYRIGHT INFRINGEMENTS IN MACHINE LEARNING

Training AI on copyrighted works can raise a number of copyright liabilities. The clearest potential infringement takes place when training data are reproduced in order to be incorporated into a dataset. Other, more esoteric infringement claims may arise when the data are analyzed, or when a trained AI produces output similar to the data on which it was trained. This sub-Part offers a technical description of the infringements that machine learning may engender. Its analysis is diagnostic, not prescriptive.

1. Literal Reproduction in Datasets

The clearest copyright liability in the machine learning process is assembling input datasets, which typically requires making digital copies of the data. If those input data contain copyrighted materials that the engineers are not authorized to copy, then reproducing them is a prima facie infringement of § 106(1) of the Copyright Act. If the data are modified in preprocessing, this may give rise to an additional claim under § 106(2) for creating derivative works. In addition to copyright interests in the individual works within a dataset, there may be a copyright interest in the dataset as a whole.⁸⁹

Some input datasets are assembled by digitizing physical media. Others are made by copying and processing born-digital data like ebooks and news photographs.⁹⁰ Many such datasets appear to have been assembled without the authorization of the copyright owners whose work they contain. There are several explanations for this practice: computer engineers may not be informed about or concerned with intellectual property law, and therefore may not realize that assembling datasets through the mass reproduction of copyrighted materials could infringe copyright. Some may anticipate that fair use will excuse their behavior. Others may not believe they have any effective way of licensing the data they use and simply hope not to get caught. However reasonable their justifications, researchers who reproduce and/or distribute unauthorized copies of copyrighted materials as input data typically depend on the fair use exception to do so.

Of course, some datasets are compiled through unambiguously legal means. Datasets may exclusively comprise public domain works⁹¹ or use only copyrighted material that has been specifically licensed for data mining and machine learning

89. See, e.g., *Data Sets*, COMPUTER VISION RESEARCH LABORATORY, <https://perma.cc/6ALB-YLBL> (last visited June 13, 2017) (making available input data for facial recognition algorithms, claiming copyright in these datasets and conditioning use on licensing agreements that restrict users' ability to reproduce, modify, distribute, or make commercial uses of the data).

90. Yukun Zhu et al., *Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books*, ARXIV (June 22, 2015), <https://perma.cc/3587-SVVK>; Erik Learned-Miller et al., *Labeled Faces in the Wild: A Survey*, in *ADVANCES IN FACE DETECTION AND FACIAL IMAGE ANALYSIS* 189 (Michal Kawulok, M. Emre Celebi, & Bogdan Smolka eds., 2016).

91. See, e.g., *THE NEXT REMBRANDT*, <https://perma.cc/7A34-8EMT> (last visited June 6, 2017).

purposes.⁹² Moreover, a great deal of copyrighted machine learning data are obtained through private contracting. Large internet platforms' terms of service typically grant the platforms broad licenses to any intellectual property that users upload to the sites. If written well, these licenses almost certainly entitle companies to train machine learning algorithms using content uploaded to their services by their users.⁹³

Liability stemming from literal reproduction in datasets may be less likely to arise in the future, as technological advances make it possible to conduct more and more machine learning without centralized datasets. In February of 2017, Google announced "a completely new, lightweight, machine learning architecture" that enables Android wearable devices to generate predictive text using users' locally stored data, without having to copy those data to cloud servers.⁹⁴ Two months later, the company unveiled a technique called "federated learning," which allows for these local models to train and be updated by a shared model stored in the cloud.⁹⁵ In federated learning, all training data are stored on users' devices, and only small updates are transmitted to the cloud. This, performed across many devices, updates a shared model that can in turn be downloaded back onto users' devices.⁹⁶

Federated learning obviates the need to construct a centralized database of user data for certain applications of machine learning; Google primarily touts this as a boon to user privacy and data security. Just as importantly, this technology and others like it may reduce the copyright liabilities associated with machine learning, because they can eliminate some of the copying typically required to compile and consolidate training data. However, even learning processes that avoid initial reproductions of training data may lead to infringements at later stages, as the following sub-Parts explain.

2. Literal Reproduction as an Intermediate Step

Once an input dataset has been compiled, it may be copied, emulated, and re-copied thousands of times during the learning process.⁹⁷ Some of these copies may exist for such transitory durations that they may not meet the statutory definition of

92. See, e.g., David D. Lewis, *Reuters-21578 Text Categorization Test Collection Distribution 1.0 README File* (Sep. 26, 1997), <https://perma.cc/V7JJ-CNVW>.

93. See *infra*, Part I.E.5 (discussing the market for training data licensed from internet platforms' end users).

94. Sujith Ravi, *On-Device Machine Intelligence*, GOOGLE RESEARCH BLOG (Feb. 9, 2017), <https://perma.cc/WQ8L-WS5D>.

95. Brendan McMahan & Daniel Ramage, *Federated Learning: Collaborative Machine Learning without Centralized Training Data*, GOOGLE RESEARCH BLOG (Apr. 6, 2017), <https://perma.cc/XVA2-J96J>. For a more technical discussion of federated learning, see Brendan McMahan et al., *Communication-Efficient Learning of Deep Networks from Decentralized Data*, ARXIV (Feb. 28, 2017), <https://perma.cc/P7GX-LXB9>.

96. *Id.*

97. See, e.g., MURPHY, *supra* note 71, at 1003–05 (discussing deep auto-encoders, which repeatedly encode and decode training data or elements thereof).

a “copy.”⁹⁸ If, in a literal sense, these steps do create infringing copies or derivative works from input data, this intermediate copying is not of great concern. For one, the spirit of the copyright statute seems to exempt this type of copying. The law is explicit that copies of computer programs created by their owners “as an essential step in the utilization of the computer program in conjunction with a machine and ... used in no other manner” are not infringements.⁹⁹ This portion of the statute does not, of course, excuse intermediate copying in machine learning outright: input data are not always a “computer program,” and engineers training a model are often not the “owners” of the input data. But the reduplication of data that takes place during training is much like an “essential-step” replication of software in memory, in that it yields ephemeral copies not for further consumption.

Furthermore, an additional infringement claim during the training stage would mean little in practice. If the initial use of input data is excused by fair use, then this intermediate copying will certainly be excused as well; alternatively, if the initial use of input data is unfair, the intermediate reproduction of data during training would be unlikely to further prejudice the defendant’s case. Intermediate copies are not likely to inflict distinct, calculable harms on a plaintiff, and statutory damages are assessed per work infringed, not per instance of infringement.¹⁰⁰

3. Non-Literal Reproduction in Datasets and Models

Because so much machine learning takes place on unauthorized copies of training data, the most pressing concern is infringement at the input stage, as described in Part I.D.1. However, it is possible that infringement could occur even when learning is performed using authorized copies of works as training data, because machine learning could yield a probabilistic model of the copyrightable aspects of a work that falls under the copyright statute’s capacious definition of a “copy” or a “derivative work.”¹⁰¹ Take one example: in 2016, a master’s student trained a model to reconstruct the 1982 Ridley Scott film *Blade Runner*. This model’s training set consisted of every frame of the movie.¹⁰² The model’s objective was to learn how to maximize the similarity between frames that it generated and frames copied directly from the film.¹⁰³ After teaching itself to

98. *Cartoon Network v. CSC Holdings, Inc.* held that data stored in a buffer for 1.2 seconds did not meet copyright’s fixation requirement, and therefore did not constitute a copy for the purposes of the Act. *Cartoon Network LP, LLLP v. CSC Holdings, Inc.*, 536 F.3d 121, 129–30 (2d Cir. 2008). For further discussion of this issue with respect to AI, see Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence’s Implicit Bias Problem*, WASH. L. REV. 15 (forthcoming 2017) (filed with <https://perma.cc/E3TR-7DDV>); see also Aaron Perzanowski, *Fixing RAM Copies*, 104 NW. U. L. REV. 1067 (2010); but see *MAI Systems Corp. v. Peak Computer, Inc.*, 991 F.2d 511, 518 (9th Cir. 1993) (holding that “‘copying’ for purposes of copyright law occurs when a computer program is transferred from a permanent storage device to a computer’s RAM.”).

99. 17 U.S.C. § 117(a)(1) (West).

100. 17 U.S.C. § 504 (West).

101. 17 U.S.C. § 101 (West).

102. Terence Broad, *Autoencoding Video Frames 26–27* (May 12, 2016) (unpublished Master’s dissertation, Goldsmiths, University of London) (filed with <https://perma.cc/7PC9-ZDUG>).

103. *Id.* at 25.

recreate *Blade Runner*, the model outputted a version of the film that resembles a rough bootleg.¹⁰⁴ The generated version of *Blade Runner* lacks an audio track and, during busy scenes, can have such poor visual fidelity that it appears only to be swirling, impressionistic blobs of color. However, it sometimes contains recognizable cinematography, and, most importantly, was derived exclusively by copying and recopying *Blade Runner*, albeit in an innovative way.

In a situation like this, when all the input data can be assimilated into a single “work” for the purposes of copyright law, it seems plausible to deem the model derived from these data a derivative work.¹⁰⁵ After all, the *Blade Runner* model is purpose-built to create ersatz versions of *Blade Runner*. If a trained model always ends up replicating its input data, it would be sensible to call the model itself a copy or a derivative work—just as a VHS copy of a film always ends up replicating the underlying film when inserted into a VCR.

Even if a model was not intentionally built to mimic a copyrighted work, it could still end up doing so to an infringing degree. Machine learning models sometimes reconstruct idiosyncrasies of input data instead of reflecting underlying trends about those data. In the technical parlance, these models are “overfitted;” they are undesirable in a predictive context because they capture “noise” rather than “signal.”¹⁰⁶ A number of machine learning models analyze sound recordings and/or musical compositions in order to learn how to generate music.¹⁰⁷ If these models had been given a homogeneous enough training set, or if they had been overfitted to that training set, they might only be capable of producing output that was substantially similar to, or outright duplicative of, their input data.

Another remote concern is that some machine readable representations of copyrighted works, despite comprising metadata rather than full, human readable reproductions, could qualify as “copies.”¹⁰⁸ The Million Song Dataset, for example, contains factual information about one million songs, as well as information about their key, pitch, timbre, tempo, and so on.¹⁰⁹ Google has promulgated metadata about the Google Books corpus in the form of n-grams, which represent the frequency of words and phrases within the corpus.¹¹⁰ The

104. Whitney Museum of American Art, *Terence Broad, Blade Runner – Autoencoded*, VIMEO (Oct. 20, 2016), <https://vimeo.com/188157459>. Interestingly, Warner Brothers sent a Digital Millennium Copyright Act takedown notice to Vimeo, requesting that the autoencoded film be removed. In response to press coverage, Warner Brothers permitted the video to be restored. Aja Romano, *A Guy Trained a Machine to “Watch” Blade Runner. Then Things Got Seriously Sci-fi*, VOX (June 1, 2016), <https://perma.cc/FX6U-4XDY>.

105. See *Castle Rock Entertainment, Inc. v. Carol Publishing Group*, 150 F.3d 132, 138 (2d Cir. 1998) (discussing precedential treatment of TV series as individual episodes versus as aggregate works).

106. MURPHY, *supra* note 71, at 22, 593.

107. Aäron van den Oord et al., *Wavenet: A Generative Model for Raw Audio*, ARXIV 6-8 (Sept. 19, 2016), <https://perma.cc/H7L2-9MLL> [hereinafter, *Wavenet*]; see also *WaveNet: A Generative Model for Raw Audio*, DEEPMIND, <https://perma.cc/B8WK-B66H> (last visited Mar 7, 2017) (providing snippets of the generated audio); George Dvorsky, *This Artificially Intelligent Robot Composes and Performs Its Own Music*, GIZMODO (June 14, 2017), <https://perma.cc/S4AS-R9XR>.

108. 17 U.S.C. § 101 (West).

109. MILLION SONG DATASET, <https://perma.cc/2RD7-5RUS> (last visited Jun 6, 2017).

110. Franz and Brants, *supra* note 52.

encodings of works in these datasets are not detailed enough to enable the works to be “perceived, reproduced, or otherwise communicated,” as the Act requires of “copies,” but a sufficiently granular representation isn’t difficult to imagine. Unless and until machine learning takes place without using unauthorized copies of copyrighted works as training data, these more esoteric reproductions will be of lesser consequence in practical copyright litigation. However, the possibility that machine learning could produce derivative works suggests that an enterprise would still risk liability even if it used authorized copies of copyrighted works as training data.

4. Infringing Output

Legal concerns may persist after the development of a model. After all, protectable input data are commonly used to train models to generate similar output. If that similarity is “substantial,” then that output may infringe copyright in the pre-existing work or works to which it is similar—or, at least, it could be found infringing if it were rendered by a human. While this is an intriguing question, this Article’s primary focus is on possible copyright infringement in the inputs to artificial intelligence rather than the outputs of artificial intelligence. This Article focuses on inputs rather than outputs for two principal reasons: first, there has already been more legal scholarship about ownership of—or liability for—the output of AI. Second, a practical concern: while machine learning is developing quickly, its outputs have not yet supplanted works of human authorship. Works generated by AI are fascinating and entertaining, but today they remain novelties rather than mainstream sources of entertainment or compelling substitutes for human expression.

Nonetheless, the outputs of machine learning models raise several questions that are worth cataloging. How might one assess output for the properties of originality and minimal creativity that are required for copyright to subsist?¹¹¹ If these qualities are present, does there exist a copyright over the output, and in whom does it vest? Most interesting, what would happen if an author alleged that a machine learning model’s output infringed her right of reproduction without copying her work verbatim? In conventional copyright litigation, such an allegation of non-literal copying requires the plaintiff to demonstrate: (1) that the defendant copied her protected expression; (2) that the allegedly infringing work constitutes a “copy”; and (3) that the defendant’s work evinces “substantial similarity” to the plaintiff’s.¹¹² Against human defendants, courts can satisfy the first requirement with direct evidence of copying, likelihood of access and probative similarity between the two works, and “striking similarity” that could only have arisen through copying.¹¹³ Determining access has, historically, been a murky inquiry. Humans simply do not recall all the cultural products they encounter throughout their lives, and they certainly do not store their recollections

111. See, e.g., *Feist Publications, Inc. v. Rural Telephone Service Co.*, 499 U.S. 340, 345 (1991).

112. 4 MELVILLE B. NIMMER & DAVID NIMMER, *NIMMER ON COPYRIGHT* § 13.03 (2017).

113. *Id.*

in an orderly manner. The vagaries of the human brain have led to some awkward rulings: George Harrison was found to have “subconsciously” copied a song by the Chiffons, and a jury found that the singer Michael Bolton had access to an obscure Isley Brothers song because it was played on radio and television in the areas where Bolton lived as an adolescent, nearly three decades before he released an allegedly infringing song.¹¹⁴

When a putative infringer is a machine learning model, or the legal persons that are its stewards, some aspects of this analysis are simpler. While it is impossible to exhaustively list the copyrighted works that a human being has encountered throughout her life, doing so for some machine learning models would simply entail examining their training corpora. If the allegedly copied work is present in input data, this would demonstrate access to the plaintiff’s work.

The situation becomes more complicated, however, when an allegedly copied work is absent in input data. This phenomenon could have several causes: perhaps the model was programmed to find input data by crawling the web on its own, and its activities were not logged. Or, most intriguing, perhaps the model was given a known, finite set of input data that excluded the plaintiff’s work, and nevertheless created substantially similar output. This is not as farfetched as it sounds: consider a machine learning model instructed to generate funk music, trained on a corpus of all funk recordings except the work of James Brown, the definitive originator of the genre. All funk songs bear a debt to James Brown, so it would not be altogether surprising if the funky AI generated output that was substantially similar to his work.

Because creativity is cumulative, rather than *ex nihilo*, a work can bear the mark of works that its author has not encountered firsthand. This indirect chain of influence may be unsurprising to anyone familiar with the creative process, but it is not something the judiciary is well equipped to handle. Precise fact-finding on human authorial influence is impossible; to cope with that difficulty, copyright doctrine has developed shortcuts that transpose very poorly onto creative situations in which direct influence could actually be proven or refuted, like machine learning.¹¹⁵ It remains to be seen how the doctrine for non-literal copying would address full-scale works of authorship generated by machine learning programs or other artificial intelligences.

These provocations with respect to infringing AI output will remain provocations. Issues of copyright in input data have more practical urgency and, to date, have received less academic scrutiny.

E. MACHINE LEARNING AND FAIR USE

The previous sub-Part listed the various steps of machine learning that could constitute *prima facie* copyright infringement. If performed on unauthorized copies

114. *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).

115. See Jessica Litman, *The Public Domain*, 39 EMORY L.J. 965, 1000–01 (1990) (discussing the legal fictions surrounding copying in the human context).

of input data, these learning processes must rely on fair use. Novel applications of machine learning push the boundaries of non-expressive use in several ways, and may do so significantly enough to alter the fair use calculus. While the non-expressive use doctrine readily excuses some applications of machine learning, other applications have a much more tenuous fair use defense. Some cutting-edge machine learning changes fair use analysis in two major ways: first, sophisticated machine learning may no longer be “non-expressive,” and therefore may not be sufficiently transformative to satisfy fair use’s first factor; second, expressive machine learning presents a new threat of market substitution that alters the analysis of the fourth fair use factor.

1. Some Machine Learning is Non-Infringing

Many forms of machine learning use copyrighted input data for purposes that, under current doctrine, appear to constitute fair use. Facial recognition is a good example. Training a model to identify human faces requires many photographs of the individuals one hopes to identify, and these photographs are likely to be copyright protected. Consider, for instance, the Labeled Faces in the Wild (“LFW”) dataset, a popular benchmark for measuring the performance of facial recognition algorithms. LFW is a database of 3,233 images of 5,749 people, derived from captioned news pictures featured on Yahoo News in 2002 and 2003.¹¹⁶ Because they correspond to fifteen year old news stories, it is safe to assume that most, if not all, of these images were created recently enough that they remain copyright protected.¹¹⁷

Absent some sort of licensing arrangement, anyone who reproduces the LFW dataset infringes, *prima facie*, the reproduction rights of the copyright owners of the photographs. In addition, the academics who promulgate the dataset, and others like it, risk liability for distributing the data. However, several considerations strongly suggest this activity is non-infringing. The first is that the photos in “Faces in the Wild” and “Labeled Faces in the Wild” are not complete reproductions of their source photographs. Only the portions of the photographs that show the subjects’ faces are reproduced, and they are edited in such a way that elides most copyrightable expression in the photographs. Each is compressed to a low resolution and cropped tightly around the subject’s face, seldom including even the full head and shoulders. This processing would eliminate most original expression in the photograph, leaving only the physical likeness of the subject. Because so little copyrightable content remains in the dataset, a court might well

116. *README, LABELED FACES IN THE WILD*, <https://perma.cc/AB8J-XXDJ> (last visited June 13, 2017); Learned-Miller et al., *supra* note 90, at 5; T. L. Berg et al., *Names and Faces in the News*, PROCEEDINGS OF THE 2004 IEEE COMPUTER SOCIETY CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION 2 (2004), <https://perma.cc/S57H-NC7T>.

117. Indeed, a spot-check of the LFW database confirms that at least one image is copyrighted: “Faces in the Wild” includes an image that Google’s “Search by Image” feature identifies as an October 7, 2002 file photo of United States General Tommy Franks, credited to Scott Martin and the Associated Press. Learned-Miller et al., *supra* note 90, at 5.

find no “improper appropriation,” and therefore that no § 106(1) violation could have occurred.

Use and distribution of the LFW database may not even require a fair use defense, but if they did, the *Authors Guild* precedent suggests that the defense would prevail. Training facial recognition algorithms on copyrighted photographs does not implicate the works’ protectable aspects. The use analyzes factual information—the unique physical features of a subject’s face—in the photographs, rather than photographers’ expressive choices. Moreover, the copied portions of the photographs would only receive thin protection, and the amounts taken are small. Finally, the use of these images to perform facial recognition does not disturb the market for file photos of notable people to be displayed alongside news stories.¹¹⁸

2. Google’s Natural Language Generation

While facial recognition may qualify as non-expressive fair use, certain cutting-edge applications of machine learning may not. Put simply, they are much more expressive than anything courts have evaluated in the past. In late 2015, Google added a feature called “Smart Reply” to its Inbox email service, which is an application distinct from but compatible with its Gmail email service.¹¹⁹ Smart Reply uses machine learning to automatically generate up to three responses to the emails that Inbox users receive, which users can select instead of composing replies themselves.¹²⁰ In its first iteration, the Smart Reply algorithm was trained on a corpus of 238 million email messages, presumably sampled from Gmail accounts.¹²¹ Using a combination of statistical analysis and human review, the Smart Reply research team manipulated “the most frequent anonymized sentences” in its dataset to train an AI engine that could express the same intention in different words, while avoiding redundant suggestions.¹²²

Unsurprisingly, reading millions of emails may not have taught Smart Reply to write sparkling prose. Google has since sought to improve Smart Reply’s writing style by giving the software novels to read. In a research paper from May 2016, six Google employees describe using the BookCorpus dataset, a massive collection of ebooks popular in academic research, to train a program to “generate coherent novel sentences” that could make Smart Reply more conversational.¹²³ At least

118. There may be liability for wholesale copying of pre-existing databases, rather than assembling one’s own database of copyrighted works.

119. Bálint Miklós, *Computer, respond to this email: Introducing Smart Reply in Inbox by Gmail*, THE KEYWORD (Nov. 3, 2015), <https://perma.cc/2AJM-KLEY>; *How Inbox works with Gmail*, INBOX BY GMAIL HELP, <https://perma.cc/2YCD-M9WX> (last visited May 31, 2017).

120. Anjali Kannan et al., *Smart Reply: Automated Response Suggestion for Email*, PROCEEDINGS OF THE 22ND ACM SIGKDD INTERNATIONAL CONFERENCE ON KNOWLEDGE DISCOVERY AND DATA MINING 955–64 (2016), <https://perma.cc/M797-7QNC> (last visited May 31, 2017).

121. Google’s paper refers to its data sources only as “accounts.” *Id.* at 962.

122. *Id.* at 959, 961.

123. Two of these six researchers were affiliated with academic institutions at the time of the article’s publication; the article notes that the research was conducted “when all authors were at Google, Inc.” Samuel R. Bowman et al., *Generating Sentences from a Continuous Space*, ARXIV 1 (May 12,

some of the novels in BookCorpus were unauthorized copies of copyrighted works.¹²⁴ Their authors were not notified, credited, or compensated for Google's use of their works.¹²⁵ Many of the books in BookCorpus appear to have been copied from the ebook distributor Smashwords.com and contain the platform's standard license, which includes the stipulation, "This ebook is licensed for your personal enjoyment only."¹²⁶

Duplicating these novels and using them in machine learning is presumptively copyright infringement, but Google is confident that fair use excuses its conduct. A Google spokesman wrote that the use "doesn't harm the authors and is done for a very different purpose from the authors', so it's fair use under US law."¹²⁷ But what exactly is Google's "very different purpose"? The same spokesman explained that romance novels made good input data because they "frequently repeated the same ideas, so the model could learn many ways to say the same thing – the language, phrasing and grammar in fiction books tends [sic] to be much more varied and rich than in most nonfiction books."¹²⁸ In other words, Google sought to make use of authors' varied and rich expression of ideas. This is the essence of copyrightable subject matter. Google's use cannot be called non-expressive; no longer is the company merely providing facts about books or furnishing a reference tool.

Daniel Schönberger¹²⁹ suggests that Google's BookCorpus research merely analyzes "the basic building blocks and patterns of human language," which are "entirely within the public domain."¹³⁰ He continues, "the ideas of the authors were not in scope of the use and the translation of the model into natural language text does not transpose any of the expressive elements in the training material."¹³¹ This characterization seems at odds with claims in Google's own research paper that the model can "explicitly model holistic properties of sentences such as style, topic, and high-level syntactic features[.]"¹³² features far more expressive than mere linguistic building blocks. Moreover, while it is true that machine learning recognizes patterns and is rooted in statistical analysis of data, the fact that something is a mechanically recognizable pattern does not necessarily place it in

2016), <https://perma.cc/9XQF-NEGD>; Alex Kantrowitz, *Google Is Feeding Romance Novels To Its Artificial Intelligence Engine To Make Its Products More Conversational*, BUZZFEED (May 5, 2016), <https://perma.cc/BX79-PCBG>.

124. Zhu et al., *supra* note 90; Richard Lea, *Google swallows 11,000 novels to improve AI's conversation*, THE GUARDIAN (Sept. 28, 2016), <https://perma.cc/LG94-ZXZA>.

125. Lea, *supra* note 124.

126. Zhu et al., *supra* note 90; Lea, *supra* note 124; Mark Coker, *Are Copyright Statements Copyrightable?*, SMASHWORDS (May 21, 2010), <https://perma.cc/YD5H-GYS9>.

127. Lea, *supra* note 124.

128. *Id.*

129. Schönberger is employed as Google's Head of Legal for Switzerland and Austria, but his article on this topic represents Schönberger's personal opinions, and not necessarily those of his employer.

130. Daniel Schönberger, *Deep Copyright: Up- and Downstream - Questions Related to Artificial Intelligence (AI) and Machine Learning (ML)*, in DROIT D'AUTEUR 4.0 / COPYRIGHT 4.0 (Jacques de Werra ed., forthcoming 2017).

131. *Id.*

132. Bowman et al., *supra* note 123.

the public domain. Zechariah Chafee famously opined that copyright protects the “pattern” of a work; in the context of plays, this meant “the sequence of events and the development of the interplay of the characters[.]”¹³³ Tellingly, another research paper in machine learning uses BookCorpus and other data to learn precisely these patterns, by training a model to align books’ text with corresponding scenes in film adaptations.¹³⁴ Finally, Schönberger’s claim that the text output of Google’s model does not transpose any expressive elements in training data, if true, may not suffice to avoid copyright liability. Reproducing copyrighted books for human consumption without authorization is presumptively infringement, irrespective of whether those readers go on to write text that resembles the works’ expression. Similarly, computerized consumption of authorial expression might also constitute infringement if that consumption implicates the expressive value in those works.¹³⁵

Google’s use of BookCorpus appropriated authors’ copyrighted expression to a greater degree than anything the doctrine of non-expressive use has previously justified. The company’s conduct might fall under fair use, but it is not obvious that it would. This particular research is unlikely to provoke litigation, but in principle, it exposes Google to significant damages. BookCorpus contains thousands of copyrighted works, and some of those copyrights were registered years before Google published its research. If Google’s actions were infringing, authors with registered copyrights would be able to recoup statutory damages of \$750 to \$150,000 per work infringed.¹³⁶

For a company like Google, BookCorpus is a triflingly small dataset. What other data might the company use for similar purposes, and with what consequences? Google has access to a great deal of web content that the company may not be explicitly licensed to use. Consider the email messages that non-Gmail users send to Gmail’s more than one billion monthly active users.¹³⁷ Gmail users license their copyrighted email messages to Google as a condition of using the service, so Google would be within its rights to conduct expressive machine learning using their emails as training data.¹³⁸ But what about non-Gmail users who correspond with Gmail users; can Google extract expressive value from their work? Google’s Terms of Service do treat “send[ing] . . . content to or through our services” as a grant of license, but it is not clear that non-Gmail users accept these terms simply by sending a message to a Gmail user.¹³⁹ Indeed, in a 2016 order in an ongoing class action wiretapping suit against Google, Judge Koh of the Northern

133. Zechariah Chafee, Jr., *Reflections on the Law of Copyright: I*, 4 COLUM. L. REV. 503, 513-14. (1945).

134. Zhu et al., *supra* note 90.

135. See *infra* Part I.E.4 (discussing the distinction between “productive” and “consumptive” uses of copyrighted works).

136. 17 U.S.C. § 504 (West).

137. Frederic Lardinois, *Gmail Now Has More Than 1B Monthly Active Users*, TECHCRUNCH (Feb. 1, 2016), <https://perma.cc/ZDJ7-89FG>.

138. *Google Terms of Service*, GOOGLE (last modified Apr. 14, 2014), <https://perma.cc/3F6T-BH8R>. See also *Matera v. Google, Inc.*, No. 15-CV-04062-LHK, 2016 WL 5339806, at *17 (N.D. Cal. Sept. 23, 2016) (finding that Gmail users accept the 2014 version of the terms).

139. *Google Terms of Service*, *supra* note 138.

District of California found that non-Gmail users do not consent to Google's scanning of their messages simply by sending emails to Gmail users.¹⁴⁰

Finally, Google has a goldmine of expressive training data unlike any other: the full Google Books database contains twenty-five million machine-readable copies of print books.¹⁴¹ Even before Google Books was found to be fair use, the company had been improving its search results and other services by making “non-display uses” of these books.¹⁴² The proposed Google Books settlement would have affirmed Google's right to make non-display uses, which it defined as “uses that do not display Expression from Digital Copies of Books or Inserts to the public[.]” including “internal research and development using Digital Copies[.]”¹⁴³ In 2005, a Google employee told the science historian George Dyson, “We are not scanning all those books to be read by people . . . We are scanning them to be read by an AI.”¹⁴⁴ *Authors Guild* authorized Google to construct this library for a host of non-expressive purposes. Could Google train expressive, commercial AI on the entire Google Books library? Doing so could train brilliant artificial intelligence, but this purpose goes far beyond the uses of the scans that *Authors Guild* held to be fair.

3. “The Next Rembrandt”

Expressive machine learning has found purchase in media other than literature, too. In 2014, Microsoft and the bank ING began work on a project entitled “The Next Rembrandt.” A team of engineers collected data about works in Rembrandt van Rijn's oeuvre, from the demography of the subject of the paintings to their three-dimensional topography.¹⁴⁵ This wealth of data—“150 gigabytes of digitally rendered graphics” derived from 346 paintings¹⁴⁶—was first used to ascertain that a novel Rembrandt painting would likely have been a portrait of a Caucasian male with facial hair; age thirty to forty; wearing dark clothing, a collar, and a hat; and

140. *Matera*, 2016 WL 5339806, at *18. See also *In re Google, Inc.*, No. 13–MD–02430–LHK, WL 5423918, at *14 (N.D. Cal. Sept. 26, 2013) (“Accepting Google's theory of implied consent—that by merely sending emails to or receiving emails from a Gmail user, a non-Gmail user has consented to Google's interception of such emails for any purposes—would eviscerate the rule against interception The Court does not find that non-Gmail users who are not subject to Google's Privacy Policies or Terms of Service have impliedly consented to Google's interception of their emails to Gmail users.”).

141. James Somers, *Torching the Modern-Day Library of Alexandria*, THE ATLANTIC (Apr. 20, 2017), <https://perma.cc/99FV-VKW8>.

142. Pamela Samuelson, *The Google Book Settlement as Copyright Reform*, 2011 WIS. L. REV. 479, 516 (2011); Supplemental Memorandum of Amicus Curiae Open Book Alliance in Opposition to the Proposed Settlement Between the Authors Guild, Inc., Association of American Publishers, Inc., et al., and Google Inc. at 17–18, *Authors Guild, Inc. v. Google, Inc.* 954 F. Supp. 2d 282 (S.D.N.Y. 2013) (No. 05 Civ. 8136 (DC)).

143. Amended Settlement Agreement, § 1.94, § 2.2, *Authors Guild, Inc. v. Google, Inc.*, 954 F. Supp. 2d 282 (S.D.N.Y. 2013) (No. 05 Civ. 8136 (DC)), <https://perma.cc/U28J-5NUL>.

144. George Dyson, *Turing's Cathedral*, EDGE (Oct. 23, 2005), <https://perma.cc/EHM4-TLWY>. See also GOOGLE AND THE WORLD BRAIN (2013) (adumbrating the possibility of Google using the Google Books corpus to train AI).

145. THE NEXT REMBRANDT, *supra* note 91.

146. *Id.*

facing to the right. Once this subject was determined, the data were used to train a machine learning model that mimicked Rembrandt's use of geometry, shading, and light.¹⁴⁷ This model generated a novel work, which was fixed as a physical painting by a 3-D printer using many layers of paint-based ink.¹⁴⁸ A recent copyright article somewhat paradoxically describes the output of "The Next Rembrandt" as a "new, creative, independent, and original work of art that mimicked, but was entirely different from, a genuine Rembrandt."¹⁴⁹ This author finds it easier to call the painting a "derivative work," perhaps even one with the same "aesthetic appeal" as one or more of the works it copied for input data.¹⁵⁰ Of course, Rembrandt has been dead for almost 350 years.¹⁵¹ His paintings are in the public domain. But if this project had been conducted without authorization on the oeuvre of a living painter, it is doubtful that the doctrine of non-expressive fair use would have excused it.

4. "Productive" or "Consumptive"?

The previous sub-Part argued that emerging applications of machine learning can no longer be described as non-expressive in character. This jeopardizes their most promising route to a finding of transformative fair use, but it does not, in itself, guarantee that the use is unfair. After all, a great deal of transformative fair uses, like parodies, critiques, and collages, are expressive in character. These uses are non-infringing because they are, in Judge Leval's words, "productive" uses.¹⁵² Rather than being mere consumers or usurpers of others' expression, the critic, the parodist, the collagist are themselves the authors of new expression.

The dichotomy between productive and consumptive use is a familiar one in copyright jurisprudence. It emerges most clearly in the dialogue between Justice Stevens's majority opinion and Justice Blackmun's dissenting opinion in *Sony Corporation of America v. Universal City Studios, Inc.* In *Sony*, the Supreme Court evaluated whether the use of a videocassette recorder ("VCR") to tape copyrighted television broadcasts infringed the copyright holders' exclusive rights, and, in turn, whether electronics manufacturer Sony was secondarily liable for infringement facilitated by its VCRs.¹⁵³ The *Sony* majority held that taping copyrighted programs for the purpose of "time-shifting"—that is, one-time viewing at a later point in time—was a fair use.¹⁵⁴ The Court reached this conclusion

147. *Id.*

148. *Id.*

149. Shlomit Yanisky-Ravid & Samuel Moorhead, *Generating Rembrandt: Artificial Intelligence, Accountability and Copyright*, 2017 MICH. ST. L. REV. (forthcoming 2017).

150. See *Peter Pan Fabrics, Inc. v. Martin Weiner Corp.*, 274 F.2d 487, 489 (2d Cir. 1960) (applying substantial similarity analysis to a cloth design and an alleged infringement and concluding, "[T]he ordinary observer, unless he set out to detect the disparities, would be disposed to overlook them, and regard their aesthetic appeal as the same. That is enough[.]").

151. *Rembrandt van Rijn*, ENCYCLOPÆDIA BRITANNICA, <https://perma.cc/4DVG-NFSH>.

152. Leval, *supra* note 17, at 1111.

153. *Sony Corporation of America v. Universal City Studios*, 464 U.S. 417, 419–20 (1984).

154. *Id.* at 419.

largely on economic grounds, with little consideration of whether home viewing was a productive or transformative activity: “The statutory language does not identify any dichotomy between productive and nonproductive time-shifting, but does require consideration of the economic consequences of copying.”¹⁵⁵

Justice Blackmun’s dissent reached an opposite conclusion by distinguishing between “productive” uses of copyrighted material and “ordinary,” consumptive uses.¹⁵⁶ Productive uses “[result] in some added benefit to the public beyond that produced by the first author’s work.”¹⁵⁷ While productive uses are not always fair uses, paradigmatic fair uses, such as those enumerated in section 107’s suggestive list, are productive uses.¹⁵⁸ In Justice Blackmun’s assessment, reproducing television programs for home viewing was “purely consumptive,” not productive.¹⁵⁹ Accordingly, it was inappropriate for the Court to excuse it as fair use: “There is no indication that the fair use doctrine has any application for purely personal consumption on the scale involved in this case, and the Court’s application of it here deprives fair use of the major cohesive force that has guided evolution of the doctrine in the past.”¹⁶⁰

Justice Blackmun’s dissent has proven to be at least as influential as Justice Stevens’s majority opinion.¹⁶¹ Productivity is a cornerstone of transformativeness, today’s most salient criterion. It so behooves defendants to demonstrate a productive rather than a consumptive purpose that “non-expressive use” is sometimes referred to as “non-consumptive use.”¹⁶² Accordingly, whether a court deems expressive machine learning productive or consumptive may determine the fate of its fair use defense.

On one hand, valorizing productivity suggests that Google’s BookCorpus research and the “Next Rembrandt” project would be fair use because they constitute technological progress. Making gigabytes upon gigabytes of copies of copyrighted art, in order to teach a machine to mimic that art, is indeed a remarkable technological achievement. An artificially intelligent painter or writer may yield social benefits and enrich the lives of many beholders and users. However, this view of productivity is overbroad. No human can rebut an infringement claim 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. Justice Stevens’s *Sony* opinion articulates this problem:

The distinction between “productive” and “unproductive” uses may be helpful in calibrating the balance, but . . . the notion of social “productivity” cannot be a

155. *Id.* at 455 n.40.

156. *Id.* at 478–79 (Blackmun, J., dissenting).

157. *Id.* at 478–79 (Blackmun, J., dissenting).

158. *Id.* at 478–79 (Blackmun, J., dissenting).

159. *Id.* at 496 (Blackmun, J., dissenting).

160. *Id.* at 495 (Blackmun, J., dissenting).

161. See, e.g., William W. Fisher, *Reconstructing the Fair Use Doctrine*, 101 HARV. L. REV. 1659, 1672 (1988); Frank Pasquale, *Breaking the Vicious Circularity: Sony’s Contribution to the Fair Use Doctrine*, 55 CASE W. RES. L. REV. 777, 778 (2005).

162. *Sag*, *supra* note 28, at 1505–06.

complete answer to this analysis. A teacher who copies to prepare lecture notes is clearly productive. But so is a teacher who copies for the sake of broadening his personal understanding of his specialty. Or a legislator who copies for the sake of broadening her understanding of what her constituents are watching; or a constituent who copies a news program to help make a decision on how to vote.¹⁶³

The teacher who copies to broaden his personal understanding is a productive consumer, but he nonetheless must pay for the works he consumes. If the teacher's consumption of copyrighted works inspires him to create new scholarship, so much the better, but his subsequent productivity does not entitle him to a refund for the works that influenced him. In much the same way, machine learning makes consumptive use of copyrighted materials in order to facilitate future productivity. If future productivity is no defense for unauthorized human consumption, it should not excuse robotic consumption, either.

Copyright's compulsion to differentiate human authorship from mechanical outputs obfuscates this simple proposition. Believing that machines differ fundamentally from human authors could imply that expressive machine learning always transforms the meaning of the works it appropriates. In a sense, this is true. The "meaning" of a work does depend on its author and its reader. A word-for-word duplicate of *Don Quixote*, written by a nineteenth-century Frenchman, is not the same work of literature as Cervantes's version.¹⁶⁴ Text purportedly written by software elicits a different reaction than the same text presented as the work of a human.¹⁶⁵ In theory, copyright acknowledges this proposition: Judge Learned Hand wrote that an independently created version of Keats's "Ode on a Grecian Urn" would be a distinct, protectable work.¹⁶⁶ In practice, however, the law disregards the idea, because it threatens to turn the doctrine to unenforceable mush.¹⁶⁷ Every quotation reshapes meaning,¹⁶⁸ but this does not turn every act of

163. Sony Corporation of America v. Universal City Studios, 464 U.S. 417, 455 n.40 (1984).

164. See Jorge Luis Borges, *Pierre Menard, Author of the Quixote*, in BORGES, A READER: A SELECTION FROM THE WRITINGS OF JORGE LUIS BORGES 96 (Emir Rodriguez Monegal & Alastair Reid trans., 1st ed. 1981); see also ABRAHAM DRASSINOWER, WHAT'S WRONG WITH COPYING?, 59–60, 63 (2015) (discussing Borges's short story in the context of copyright law).

165. See, e.g., Rob Dubbin, *The End of Horsebooks Is Hardly the End of Anything*, THE NEW YORK REVIEW OF BOTS (Sept. 24, 2013), <https://perma.cc/G5TF-6EPU> (describing the public perception of a popular account on the microblogging site Twitter that purported to be a robot gone haywire, but was later revealed to be the work of a human writer); see also Adrian Chen, *How I Found the Human Being Behind Horse_ebooks, The Internet's Favorite Spambot*, GAWKER (Feb. 23, 2012), <https://perma.cc/25NL-2QBW>.

166. Sheldon et al. v. Metro-Goldwyn Pictures Corp. et al., 81 F.2d 49, 54 (2d Cir. 1936).

167. Rebecca Tushnet, *Copy This Essay: How Fair Use Doctrine Harms Free Speech and How Copying Serves It*, 114 YALE L.J. 535, 566–67 (2004); Litman, *supra* note 115, at 1000–03; see also Laura A. Heymann, *Everything Is Transformative: Fair Use and Reader Response*, 31 COLUM. J.L. & ARTS 445 (2008).

168. The literary critic Leo Spitzer observed, "When we reproduce in our own speech a portion of our partner's utterance, then by virtue of the very change in speakers a change in tone inevitably occurs: *the words of 'the other person' always sound on our lips like something alien to us, and often have an intonation of ridicule, exaggeration, or mockery,*" quoted in M. M. BAKHTIN, PROBLEMS OF DOSTOEVSKY'S POETICS 194 (1984).

copying into transformative fair use; copying undertaken by artificial intelligence should be regarded with no less skepticism.

5. Delineating the “Potential Market”

The previous sub-Parts have argued that the first fair use factor, “the purpose and character of the use,” may weigh against some emerging uses of machine learning. While damaging, this in itself would not vitiate a defense. Fair use also considers the effect a use has on the potential market for the works used. It is difficult to anticipate how courts might assess this factor in the context of expressive machine learning. However, there are two key reasons to think that this factor may not point in a defendant’s favor. The first is that an established market for training data already exists. The second is that expressive machine learning presents a novel danger to works’ potential markets by threatening to usurp the position of authors themselves, rather than supplanting individual works.

An instructive precedent is *American Geophysical Union v. Texaco Inc.*, in which the Second Circuit held that the photocopying of articles in a scientific journal by Texaco’s research scientists, largely to facilitate researchers’ access to the articles, was not a fair use. The facts of this case—the bulk copying of copyrighted works by the research arm of a for-profit company—resemble those of Google’s BookCorpus research. The Texaco court did not hold that plaintiffs are entitled to compensation whenever their works are used commercially; what is relevant is the use’s impact on “traditional, reasonable, or likely to be developed markets.”¹⁶⁹ The Second Circuit found that Texaco had adversely impacted such a market, noting that the publishers had established a clearance center through which interested parties could obtain a license to photocopy articles.¹⁷⁰

Does training data for machine learning constitute a market that is traditional, reasonable, or likely to develop? Surprisingly, it often does. It is tempting to view machine learning as an alchemical process that spins value out of valueless data and creates a market where none previously existed. Considered individually, the bits of expression on which a machine learning model is trained are of infinitesimal value in comparison to the resulting model. The expression in a single email exchange from the Smart Reply training corpus is, optimistically, of interest only to the interlocutors; Smart Reply, on the other hand, could save time and effort for hordes of Inbox users. The idea of a legitimate “market” for emails—at least, for emails that would be voluntarily published by their authors—seems preposterous; we don’t buy volumes of emails at the bookstore.¹⁷¹ Even when input data comprise conventional “works” like the BookCorpus dataset does, it still seems ridiculous to compare those works’ value to that of a machine learning model that powers an innovative web service.

169. *American Geophysical Union v. Texaco*, 60 F.3d 913, 929–30 (2d Cir. 1994).

170. *Id.* at 931.

171. Of course, plenty of unscrupulous people would sell hacked emails that the individuals who own the copyright in those emails would not sell.

Appealing as it may seem, the alchemical view of machine learning is a misconception. Computation alone cannot impart value to worthless input: “garbage in, garbage out” is a concept that dates back over a century and is a pillar of computer science.¹⁷² If a machine learning model has value, each training datum contributes in some miniscule way to that value. Of course, a work’s value as input data may not have anything to do with the value that copyright allows authors to control. Facebook’s immense collection of images, “tagged” with the identities of the people pictured, is valuable largely because it matches one set of facts (individuals’ identities) with another set of facts (individuals’ facial geometries). These works may contain copyrighted expression, but a facial recognition algorithm doesn’t implicate that expression in the least. Because training a facial recognition model does not engage with copyright-protected aspects of training images, any market for images qua facial recognition input data is unlikely to be a market over which copyright affords rights holders a monopoly. As *Authors Guild* states, copyright owners’ interests only extend to the “protected aspect” of their works.

However, as the previous sub-Parts explain, some applications of machine learning do capitalize on the protectable aspects of input data. In these cases, analysis of market harms must be more nuanced. Does expressive machine learning encroach upon a market in which authors were likely to participate, or does it create a new market into which the authors’ legal monopoly should not extend? The appealing answer is that the market created by expressive machine learning is a distinct market in which authors are unlikely to participate. After all, wouldn’t these romance novels and emails be practically worthless but for their innovative uses in machine learning technology?

This answer is wrong: the market for training data has already developed. A paradigmatic business model for technology platforms is to acquire user data in exchange for providing gratis services to users. This user data can then enable a firm or its partners to serve users with precisely targeted advertisements, calibrated to their particular habits or needs.¹⁷³ User data, in aggregate, is immensely valuable.¹⁷⁴ Of course, some of these data are uncopyrightable: a user’s browsing history, the time she spends on particular pages, or the items she adds to and removes from her online shopping cart are all facts. But some data, like private messages, emails, wall posts, and the like, may well be copyrightable, and emerging applications of machine learning could implicate the expression the data contain.

Thus, there is already a thriving market for the data that fuel expressive machine learning. When users accept platforms’ terms of use, they take part in a transaction. Oftentimes, that transaction involves an explicit grant of intellectual

172. See *Garbage in, garbage out*, WIKIPEDIA, <https://perma.cc/B2XU-F65Y> (last visited June 4, 2017).

173. Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES (Feb. 16, 2012), <https://perma.cc/P29W-UD85>.

174. *The World’s Most Valuable Resource Is No Longer Oil, But Data*, THE ECONOMIST (May 6, 2017), <https://perma.cc/JZN3-YXHJ>.

property rights. In Facebook's latest terms of use, one of the first grants the document makes is:

For content that is covered by intellectual property rights, like photos and videos (IP content), you specifically give us the following permission, subject to your privacy and application settings: you grant us a non-exclusive, transferable, sub-licensable, royalty-free, worldwide license to use any IP content that you post on or in connection with Facebook (IP License).¹⁷⁵

Google's terms grant the company a similar license over user-generated content.¹⁷⁶ So do Amazon's.¹⁷⁷ It is terms of service like these that authorize Google researchers to do things like generating novel piano music using hours of piano videos uploaded by YouTube users and distributing the results, with no fair use defense necessary.¹⁷⁸ This market for training data is not restricted to platforms' terms of service, either. Google attempted to license works for "non-display" purposes in the rejected Google Books settlement agreement.¹⁷⁹ Not only did the Google Books settlement show Google's willingness to negotiate non-display uses in the market, it also offered the plaintiffs a chance to assert that such uses presumptively required licensing.¹⁸⁰

The broad IP licenses common to terms of use show a clear market for the copyrightable aspects of user-generated content, even if that information on its own seems pedestrian and worthless. There is a great deal of debate about whether this market as currently structured is efficient or inefficient, laudable or lamentable, but it is certainly a market of some kind.¹⁸¹ If fair use categorically protected companies that sought to conduct expressive machine learning for a commercial purpose, it would bypass this market and presumably harm the individuals who own the rights to training data.¹⁸²

6. A New Threat of Market Substitution

Expressive machine learning not only jeopardizes the market for the works on which it is trained, it also threatens to marginalize authors completely. Another case study elucidates these risks: Jukedeck is a company that uses artificial intelligence, trained via machine learning, to compose and adapt music.¹⁸³

175. *Terms of Service*, FACEBOOK (Jan. 30, 2015), <https://perma.cc/LAH3-2VEE>.

176. *Google Terms of Service*, *supra* note 138.

177. *Conditions of Use*, AMAZON (May 30, 2017), <https://perma.cc/PU8N-4SV8>.

178. Wavenet, *supra* note 107, at 8.

179. Amended Settlement Agreement, *supra* note 143, § 2.2 at 25.

180. Samuelson, *supra* note 142, at 515.

181. See Kenneth C. Laudon, *Markets and Privacy*, 39 COMMUN. ACM 92–104, 93 (1996); Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: J.L. & POL'Y FOR INFO. SOC'Y 543 (2009); see also Nathan J. Davis, *Presumed Assent: The Judicial Acceptance of Clickwrap*, 22 BERKELEY TECH. L.J. 577, 577–80 (2007) (surveying advantages of standard-form contracts in cyberspace).

182. *Cf.* Sony Corporation of America v. Universal City Studios, 464 U.S. 417, 451 (1984) (discussing presumed harms of unauthorized commercial use of copyrighted works).

183. *About*, JUCEDECK, <https://perma.cc/NK4C-UPJK> (last visited May 20, 2017).

Jukedeck users can create tracks by selecting from a range of parameters, including duration, tempo, mood (“uplifting” or “melancholic”), genre (“piano,” “folk,” “rock,” etc.), and instrumentation; once the user sets the parameters, Jukedeck’s AI will generate a track that fits the specifications.¹⁸⁴ Jukedeck offers three licensing options to govern use of the songs its AI generates: (1) “Individuals, Small Businesses or Non-Profits”; (2) “Large Businesses”; and (3) “Buy The Copyright.”¹⁸⁵ The first two licensing options allow for royalty-free commercial and non-commercial uses, but prohibit the user from reselling a track or making it available for others to use; the third option is self-explanatory.¹⁸⁶

Jukedeck has not fully disclosed what data it uses to train its AI, although input data that Jukedeck has released has been in the public domain.¹⁸⁷ It is conceivable that the Jukedeck training corpus contains exclusively public domain music. However, many of the genres that Jukedeck allows users to select came into existence no more than forty or so years ago, such as “Drum ‘N’ Bass” and “Synthpop.”¹⁸⁸ It is reasonable to assume that the algorithm’s input data contains music from these genres. After all, without such exemplars, it would be difficult for Jukedeck’s algorithm to convincingly generate music from those genres. It is therefore probable that at least some of the works in Jukedeck’s training data are copyright protected in the United States, since copyright in post-1978 works subsists for seventy years past the death of the author.¹⁸⁹

Because Jukedeck’s input data and methods are kept secret, it is unclear whether it needs to rely on the fair use defense. However, imagining how that defense would play out elucidates a new complication of the market impact factor. There is no doubt that AI generated, royalty-free sound recordings would jeopardize the market for recordings that are composed and performed by humans in a traditional fashion. Jukedeck’s rates are lower than the cost of licensing a conventional sound recording.¹⁹⁰ Its output is not limited by the constraints human composers or recording artists face. Moreover, the market impact analysis does not merely scrutinize the effects of Jukedeck’s use: rather, it “poses the issue of whether unrestricted and widespread conduct of the sort engaged in by the defendant (whether in fact engaged in by the defendant or by others) would result in a substantially adverse impact on the potential market for, or value of, the plaintiff’s

184. *Create unique, royalty-free soundtracks for your videos*, JUKEDeck, <https://perma.cc/L6EW-9K4L> (last visited May 20, 2017).

185. *Licensing*, JUKEDeck, <https://perma.cc/9AP9-3PEP> (last visited May 20, 2017).

186. *Id.* Jukedeck’s terms do not clarify some issues that a transfer of copyright might raise, such as whether or not Jukedeck would exercise its termination rights under § 203 of the copyright statute (if copyright protection subsists at all in Jukedeck’s works under US law).

187. See Jukedeck R&D Team, *Releasing a Cleaned Version of the Nottingham Dataset*, JUKEDeck RESEARCH (Mar. 7, 2017), <https://perma.cc/3SHH-XF7Q>.

188. *Jungle*, ENCYCLOPEDIA OF POPULAR MUSIC (4th ed. 2009), available at <https://perma.cc/DJ36-34EJ> (last visited Oct. 23, 2017); COLETTE SIMONET, *Synthpop*, GROVE MUSIC ONLINE, available at <https://perma.cc/WTF7-63CG> (last visited Oct. 2, 2017).

189. 17 U.S.C. § 302(a) (West).

190. Alex Marshall, *From Jingles to Pop Hits, A.I. Is Music to Some Ears*, N.Y. TIMES, Jan. 22, 2017, <https://nyti.ms/2jOdbWj>.

present work.”¹⁹¹ If use of copyrighted recordings as input data for AI composers were unrestricted and widespread, there is no doubt that some section of the market for compositions would become completely robotized.

It should now be clear that some expressive uses of machine learning would, if unrestricted, deprive authors of markets they currently exploit. The affected subject matter is by no means limited to sound recordings, either. Open-source software provides a massive training corpus for artificial intelligence,¹⁹² which has already learned to repair software glitches and even generate novel programs after being trained on existing codebases.¹⁹³ Writers, too, might find their jobs in jeopardy, particularly those in data oriented fields like sports journalism or financial reporting. For instance, Kristian Hammond, a founder of Narrative Science, a company whose algorithms generate news articles, estimated that “[m]ore than 90 percent” of news journalism will be computer generated by 2027.¹⁹⁴ Expressive machine learning shifts the balance of fair use’s fourth factor because it could substitute for the individual works on which it trains and for the authors of those works.

II. AI'S FAIR USE DILEMMA

Part I discussed the operation of machine learning and listed the potential risks of this practice. It then analyzed the current state of the fair use doctrine as applied to computer uses of copyrighted materials and suggested that some expressive machine learning cannot rely on present fair use doctrine. This Part presents normative reasons why the fair use doctrine is ill equipped to deal with expressive machine learning. To contextualize artificial intelligence’s fair use dilemma, this Part shows how today’s digital economy challenges accepted narratives about the nature of digital intellectual property, and, as a consequence, inverts traditional paradigms of fair use. Ensuring that fair use continues to serve the public benefit in this new ecosystem will require reassessing the scope and availability of the doctrine.

191. 4 NIMMER ON COPYRIGHT § 13.05.

192. Every open-source project on the code sharing website Github has been indexed in Google’s BigQuery database tool; while at present the tool only appears to facilitate factual analysis, one can imagine it as fodder for more sophisticated machine learning. See Felipe Hoffa, *All the Open Source Code in GitHub Now Shared Within BigQuery: Analyze All the Code!*, MEDIUM (June 29, 2016), <https://perma.cc/W2YT-QM7G>.

193. See, e.g., Fan Long & Martin Rinard, *Automatic Patch Generation by Learning Correct Code*, 51 ACM SIGPLAN NOTICES 298, 298–309 (2016); Matt Reynolds, *AI Learns to Write Its Own Code by Stealing from Other Programs*, NEW SCIENTIST (Feb. 22 2017), <https://perma.cc/S7AD-QQGB>; Matej Balog et al., *DeepCoder: Learning to Write Programs*, ARXIV (Mar. 8, 2017), <https://perma.cc/R7D5-7JE4>.

194. Quoted in Steven Levy, *Can an Algorithm Write a Better News Story Than a Human Reporter?*, WIRED (Apr. 24, 2012 4:46 PM), <https://perma.cc/FF4S-F7RD>.

A. ONE BAD OUTCOME: TRAINING AI IS INFRINGEMENT

The preceding analysis suggests several reasons why current formulations of the fair use doctrine may not excuse cutting edge applications of expressive machine learning. If machine learning is indeed not categorically fair use, then an author would have a plausible infringement claim against an engineering team for reproducing her work in input data without authorization. While this Article criticizes business practices associated with expressive machine learning, it recognizes the unique value of the technology. Make no mistake: a categorical rejection of fair use for expressive machine learning would have disastrous ramifications.

This outcome would be devastating because the remedies that copyright law offers are mismatched with the harms an author would suffer from inclusion in input data. If the work in question were registered prior to the infringement, the author could claim statutory damages of at least \$750 per infringed work, and up to \$150,000 per work if the infringement were deemed willful.¹⁹⁵ Because machine learning datasets can contain hundreds of thousands or millions of works, an award of statutory damages could cripple even a powerful company. Conceivably, a plaintiff could enjoin a defendant from proceeding with a machine learning operation, though it is unlikely that a court would offer such a drastic equitable remedy in a case involving input data.

Most worrisome, however, is the chilling effect that a single adverse ruling could have on the development of machine learning. Machine learning promises to streamline legal drudgery and bring a form of advocacy to parties who otherwise would not be able to access it.¹⁹⁶ Writers who were employed to perform formulaic composition might be able to devote their energies to more creative forms of self-expression once machines supplant them. Professors might find more time to do research if an artificial intelligence engine automatically emailed with students for them. Copyright law can make or break emerging technologies, and it would be a shame if machine learning's immense potential benefits foundered on copyright concerns.

Limits on machine learning could also impoverish research efforts. For example, in the European Union, member states typically rely on enumerated exceptions to copyright law, rather than the more protean, and potentially more capacious, fair use doctrine. The United Kingdom's copyright statute contains an exception for "text and data analysis for non-commercial research," but in general, legal hurdles to computational analyses of copyrighted works remain higher in Europe than in the United States.¹⁹⁷ Some empirical analysis suggests that these

195. 17 U.S.C. § 504(e) (West).

196. See Cara McGoogan, *19-Year-Old's "Robot Lawyer" Overtakes 160,000 Parking Tickets*, THE TELEGRAPH (June 29, 2016), <https://perma.cc/9GLR-G37U>; Hugh Son, *This Software Does in Seconds What Took Lawyers 360,000 Hours*, THE INDEPENDENT (London) (Feb. 28, 2017), <https://perma.cc/8JX7-NY32>.

197. Copyright, Designs and Patents Act 1988 § 29A. The European Parliament, noting that "there is a risk that the Union's competitive position as a research area will suffer unless steps are taken to address the legal uncertainty for text and data mining," has proposed an exception for non-profit or

restrictions have hampered relevant academic research, in comparison to countries with broader exceptions to copyright protection.¹⁹⁸ What's more, restrictions on unauthorized uses of copyrighted input data might afford powerful or deep-pocketed incumbents singular access to high-quality data and thereby entrench them as the only innovators in the field. Still more worrisome, machine learning models trained only on readily available or freely licensed data might codify pernicious biases present in those data.¹⁹⁹

Kneecapping cutting edge machine learning with an adverse fair use judgment could jeopardize the technology's social value, or drive innovation to a foreign jurisdiction with relaxed copyright constraints. That said, a finding of fair use for expressive machine learning is hardly an appealing alternative.

B. ANOTHER BAD OUTCOME: TRAINING AI IS FAIR USE

Imagine that an influential precedent, handed down in early 2018, resolves the doctrinal ambiguities discussed in this Article, declares broadly that expressive machine learning is fair use, and is understood by engineers and general counsels as giving carte blanche to unauthorized uses of copyrighted input data in AI development. Call it an *Authors Guild* for expressive uses of copyrighted material. This world, in all likelihood, would not be all that different from the world we currently inhabit. After a few years of unrestricted expressive machine learning, however, it could be drastically worse.

It is well known that, unless counterbalanced by major policy changes, advances in artificial intelligence threaten to exacerbate inequality. The Obama White House's recent report on "Artificial Intelligence, Automation, and the Economy" forecasts a shift in income from labor to capital as automata begin to supplant human laborers.²⁰⁰ In response, among other things, the report recommends bolstering the social safety net to support the millions of people whose livelihoods will be disrupted or eliminated by artificial intelligence.²⁰¹ These and other attempts to equalize the distribution of wealth could be funded in part by a modernized tax regime, the report suggests.²⁰² Elon Musk advocates a universal basic income that could leave individuals free to pursue meaningful work while robots take over the drudgery, while Bill Gates recommends a redistributive tax on "the robot that takes your job."²⁰³

public interest text and data mining as part of its Digital Single Market initiative. Proposal for a Directive of the European Parliament and of the Council on Copyright in the Digital Single Market, COM/2016/0593 Final 14, 23 (Sept. 2016).

198. See, e.g., Christian Handke et al., *Is Europe Falling Behind in Data Mining? Copyright's Impact on Data Mining in Academic Research* (June 9, 2015) (unpublished manuscript) (on file with <https://perma.cc/EQX4-D33D>)

199. See Amanda Levendowski, *supra* note 98; see also *infra* Part III.D.2.

200. EXEC. OFFICE OF THE PRESIDENT, *supra* note 2, at 41.

201. *Id.* at 36–37.

202. *Id.* at 41.

203. Catherine Clifford, *Elon Musk: Robots Will Take Your Jobs, Government Will Have to Pay Your Wage*, CNBC (Nov. 4, 2016), <https://perma.cc/PS7B-QRED>; Kevin J. Delaney, *The Robot That*

It would indeed be wonderful if advances in AI and progressive policy changes provided ordinary humans with greater economic stability and more leisure time to devote to meaningful work. However, there is little to guarantee that such a safety net will ever be in place for the individuals displaced by AI, let alone in the near future. This political reality makes it difficult to countenance a fair use paradigm that allows sophisticated actors to extract expressive value from copyrighted works, without compensation to authors, in the service of technology that may well deprive those authors of a livelihood. Furthermore, encouraging robotic forms of reading may impoverish the human expressive faculties we presumably cherish—and place perverse restrictions on the kinds of expressive activities humans may use their newfound leisure time to explore. Why should a digital humanities scholar devour millions of texts without compensating their authors, while a more conventional literary hermeneut—or an ordinary reader—must pay for the copyrighted works she interprets?²⁰⁴

Of course, no copyright policy could fully address the concerns that the Obama administration or Elon Musk or Bill Gates have articulated. Copyright will not restore to truck drivers the livelihoods that many of them will soon lose to automation.²⁰⁵ Nor will copyright law secure workers whose every move in the workplace is recorded and used to train their robotic replacements.²⁰⁶ But the need for broader social reforms should not lead us to overlook existing doctrine that could be used to promote greater distributive justice.

In addition to its potential to worsen economic inequality, a final, extreme concern is that a permissive fair use doctrine for expressive machine learning could precipitate an existential risk to humanity. As Grimmelman has suggested, progress in AI, enabled by broad fair use exceptions, may not correspond to progress for the human species. Unchecked robotic reading could hasten the rise of superintelligent machines, technology that could easily outsmart and outgun its human handlers. Unless programmed to adhere to some human view of the good—on which, of course, consensus would be elusive—a superintelligent computer could wipe out humanity if doing so were instrumentally useful to it.²⁰⁷

C. FAIR USE: OLD AND NEW NARRATIVES

Taken together, the preceding two sub-Parts contend that the fair use doctrine, as currently construed, cannot equitably resolve some cases involving expressive machine learning. This is because each outcome appears to contradict at least one of the underlying objectives of the fair use doctrine, and of copyright law more generally.²⁰⁸ An expansive fair use defense for machine learning would promote

Takes Your Job Should Pay Taxes, Says Bill Gates, QUARTZ (Feb. 17, 2017), <https://perma.cc/W6TK-DVC2>.

204. Grimmelman, *supra* note 11, at 675.

205. EXEC. OFFICE OF THE PRESIDENT, *supra* note 2, at 13.

206. Daisuke Wakabayashi, *Meet the People Who Train the Robots (to Do Their Own Jobs)*, N.Y. TIMES (Apr. 28, 2017), <https://nyti.ms/2qdGWkk>.

207. Grimmelman, *supra* note 11, at 676–78.

208. Fisher, *supra* note 161, at 1686–91.

technological progress, but it would deprive “contributors to the store of knowledge” of “a fair return for their labors.”²⁰⁹ In turn, a finding of infringement would justly affirm authors’ exclusive rights to exploit their protected expression, but it could hamstring a new, promising technology.

Expressive machine learning presents an especially difficult dilemma because it undermines accepted narratives about fair use in the digital economy. In the late twentieth century, the proliferation of media technology pitted incumbent rights holders against individual end users. New technologies simultaneously facilitated widespread copying of copyrighted works and afforded rights holders more granular control over how their works could be consumed. Unauthorized consumption by individual users became a threat to rights holders’ bottom lines, and in response, copyright owners sought and secured greater powers to control uses of their works.²¹⁰ In turn, scholars and activists worried that rights holders’ greater control over works would chill expression and turn every act of consumption into an opportunity for rent seeking.²¹¹

The theoretical backdrop to this debate is the market failure theory of fair use, which Wendy Gordon articulates in a seminal article from 1983. Gordon, criticizing the Ninth Circuit’s judgment against Sony in *Universal City Studios, Inc. v. Sony Corp. of America*—which would be reversed by the Supreme Court the following year²¹²—contends that uses should be considered fair if they: (1) occur under conditions of market failure; (2) serve the public interest; and (3) do not present such a substantial injury to a plaintiff that incentives would be impaired.²¹³ Gordon uses “market failure” to refer to scenarios in which a socially desirable use of intellectual property cannot be effectuated by market forces.²¹⁴ As examples of possible failed markets, Gordon cites uses of copyrighted materials that are rendered inefficient by high transaction costs; uses that engender nonmonetizable positive externalities, like scholarship; and uses that implicate difficult-to-price values such as wellbeing and free expression.²¹⁵ Market failure, Gordon writes, is a precondition for an economically justified finding of fair use.²¹⁶

As digital technologies reduced the costs of licensing transactions, some scholars adapted market failure theory to justify extensive “fared use” regimes in which essentially all uses of copyrighted materials—perhaps even parody—would

209. *Harper & Row, Publishers v. Nation Enter.*, 471 U.S. 539, 546 (1985).

210. PETER BALDWIN, *THE COPYRIGHT WARS: THREE CENTURIES OF TRANS-ATLANTIC BATTLE* 263–65, 273–88 (2014).

211. “[I]n an era of exploding information, we also have to guarantee access to libraries and also university researchers, to make sure we do not enter a new era of pay per view, where the use of a library card always carries a fee and where the flow of information comes with a meter that rings up a charge every time the Internet is accessed.” 144 Cong. Rec. 7074, 7100 (1998) (statement of Rep. Klug on the Digital Millennium Copyright Act); BALDWIN, *supra* note 210, at 284, 318–19.

212. *Sony Corp. v. Universal City Studios*, 464 U.S. 417 (1984).

213. Wendy J. Gordon, *Fair Use as Market Failure: A Structural and Economic Analysis of the Betamax Case and Its Predecessors*, 30 J. COPYRIGHT SOC’Y U.S.A. 253, 254–55, 271–76 (1983).

214. *Id.* at 272.

215. *Id.* at 287–98.

216. *Id.* at 272.

be paid for, rather than excused by fair use.²¹⁷ In response, other scholars developed distributive and normative critiques of this expansion of market failure theory.²¹⁸ These writers' concern was that a "fared use" model would advance the interests of established rights holders while inhibiting participatory semiosis and downstream uses of copyrighted works.²¹⁹ Matthew Africa notes that market failure analysis encourages licensing creep, whereby "borderline uses"—uses that may well be noninfringing fair uses—are licensed by risk averse parties with the means to do so.²²⁰ This threatens less powerful creators, who cannot afford to pay customary licensing fees for the building blocks of their expression, even though that expression may in fact be non-infringing.

As an alternative to ubiquitous licensing, Robert Merges proposes reemphasizing the redistributive aspects of copyright law.²²¹ By recognizing that copyright is itself a subsidy to creators, fair use can in turn be understood as a tax on rights holders that facilitates certain categories of users, like critics, parodists, and educators.²²² Fair use should not subsist only in areas where the market fails; rather, Merges advocates focusing on the categories of uses that should be effectively subsidized with permission to circumvent the market through fair use.²²³

The contemporary perception of fair use emerged from this factual and theoretical context. On one side, commentators projected that fair use would obsolesce as technological advances facilitated ever more granular licensing arrangements that allowed individual users to pay for the works they consumed.²²⁴ On the other side, commentators characterized fair use as a public entitlement and a safety valve for expressive freedom, and suggested that the stringent controls afforded to rights holders in the 1990's often amounted to anti-competitive

217. Tom W. Bell, *Fair Use vs. Fared Use: The Impact of Automated Rights Management on Copyright's Fair Use Doctrine*, 76 N.C. L. REV. 557, 595 (1998).

218. See Robert P. Merges, *The End of Friction? Property Rights and Contract in the "Newtonian" World of On-line Commerce*, 12 BERKELEY TECH. L.J. 115, 131–32 (1997) (describing the narrow applicability of a fair use defense if market failures are a necessary predicate); Matthew Africa, *The Misuse of Licensing Evidence in Fair Use Analysis: New Technologies, New Markets, and the Courts*, 88 CAL. L. REV. 1145, 1164–75 (2000) (critiquing market failure theory as an oversimplification of fair use).

219. Negativland, *Two Relationships to a Cultural Public Domain*, 66 LAW & CONTEMP. PROBS. 239, 257–58 (2003).

220. Africa, *supra* note 218, at 1172.

221. Merges, *supra* note 218, at 134.

222. *Id.*, at 134–35. Copyright law contains provisions like termination rights, which could be understood as measures to ensure some distributive justice for authors who alienate lucrative intellectual property through ignorance, folly, or lack of bargaining power. See, e.g., *M. Witmark & Sons v. Fred Fisher Music Co.*, 125 F.2d 949, 955 (2d Cir. 1942), *aff'd*, 318 U.S. 643 (1943) (Frank, J., dissenting) (rejecting the holding that authors of copyrighted works may assign renewal interests in their copyrights).

223. Merges, *supra* note 218, at 135.

224. PAUL GOLDSTEIN, *COPYRIGHT'S HIGHWAY: FROM GUTENBERG TO THE CELESTIAL JUKEBOX* 207 (revised ed. 2003); see also INFORMATION INFRASTRUCTURE TASK FORCE, *INTELLECTUAL PROPERTY AND THE NATIONAL INFORMATION INFRASTRUCTURE* 82–84 (1995).

corporate overreach.²²⁵ The vision of fair use that this discourse produced was, depending on one's perspective, either a tax on rights holders or a bulwark against monopolistic enclosure of intellectual property. Whether one agreed or disagreed with how it operated, fair use was characterized as a redistributive mechanism that subsidized public pursuits at major content owners' expense.

Today's digital economy upends this narrative. Today's ordinary end users are not passive consumers of others' intellectual property. Rather, they create troves of text, images, video, and other data that they license to large companies in exchange for gratis services. Powerful technology companies are now users of copyrighted material, and the companies' end users are the rights holders. This pivot in market dynamics should prompt a corresponding shift in attitudes towards fair use. The doctrine no longer redistributes wealth from incumbents to the public; it shifts wealth in the other direction, from the public to powerful companies.

1. The Familiar Narrative: Big Content, Little Users

Copyright rhetoric in the digital age tends to reflect two worries, both due to digital media's unprecedented enabling of the duplication, recombination, and exchange of information.²²⁶ The first worry—palpable in the debate over the Sony VCR in the 1980's²²⁷ to peer-to-peer file sharing in the early 2000's²²⁸ to the present-day paradigm of unauthorized copies on file locker or streaming sites²²⁹—was that unauthorized personal consumption by individual users would kill rights holders by a thousand cuts. Individual users' ability to harm copyright owners through personal use is, in Niva Elkin-Koren's words, a “fundamental feature of cyberspace.”²³⁰ The more straightforward and more pervasive copying became, the more that unauthorized personal uses threatened familiar business models.

225. JAMES BOYLE, *THE PUBLIC DOMAIN* 58–82 (2008); Fred Von Lohmann, *Unintended Consequences: Twelve Years under the DMCA*, ELECTRONIC FRONTIER FOUNDATION (Feb. 2010), <https://perma.cc/R5F5-4DUD>.

226. See, e.g., Timothy Wu, *Copyright's Communications Policy*, 103 MICH. L. REV. 278, 278 (2004) (“There is something for everyone to dislike about early twenty-first century copyright. Owners of content say that newer and better technologies have made it too easy to be a pirate. Easy copying, they say, threatens the basic incentive to create new works; new rights and remedies are needed to restore the balance. Academic critics instead complain that a growing copyright gives content owners dangerous levels of control over expressive works.”).

227. See, e.g., *Home Recording of Copyrighted Works: Hearing on H.R. 4783, H.R. 4794, H.R. 4808, H.R. 5250, H.R. 5488, and H.R. 5705 Before the Subcomm. on Courts, Civil Liberties, & the Admin. of Justice of the H. Comm. on the Judiciary*, 97th Cong. (1982) (statement of Jack Valenti, President, Motion Picture Association of America) (“I say to you that the VCR is to the American film producer and the American public as the Boston strangler is to the woman home alone.”).

228. See, e.g., *Music on the Internet: Is There an Upside to Downloading? Hearing Before the S. Comm. on the Judiciary*, 106th Cong. (2000) (containing testimony from music industry executives, technology developers, and others on the effects of peer-to-peer file sharing on the sound recording business).

229. See, e.g., Comments of Universal Music Group, United States Copyright Office Section 512 Study (2016) (No. 2015-7), available at <https://perma.cc/Z4HM-4ZYJ> (describing difficulties of policing online platforms for unauthorized reproductions of works owned by Universal).

230. Niva Elkin-Koren, *Cyberlaw and Social Change: A Democratic Approach to Copyright Law in Cyberspace*, 14 CARDOZO ARTS & ENT. L.J. 215, 285 (1996).

As this narrative of individual users threatening copyright owners developed, a reciprocal narrative emerged: that overzealous copyright enforcement by rights holders would chill individual users' self-expression. In addition to facilitating unauthorized private copying, networking technologies gave rights holders new tools to monitor and regulate uses of their works and create licensing markets where none previously existed. Incumbents successfully lobbied to strengthen their copyrights in at least two ways.²³¹ First, rights holders gained statutory entitlements through the Digital Millennium Copyright Act, which permitted them to curtail the scope of fair use through technological self-help strategies.²³² Second, rights holders extended the duration of copyright protections, which chiefly benefited the companies and heirs who retained the rights to a small percentage of lucrative works.²³³ As corporations clamored for longer and stricter enforcement of copyright, scholars worried that a pivot to stringently regulating individual uses of copyrighted works would hinder individual freedom to participate in culture, to learn, and to exchange information.²³⁴ Accordingly, scholars like Merges emphasized fair use's redistributive origins and advocated viewing the doctrine as an express subsidy for certain privileged activities.²³⁵

Though they may be opposing narratives, the two strains of copyright rhetoric in the digital age share two general themes: (1) that widespread infringement threatens established, incumbent businesses; and (2) that aggressive copyright enforcement by these incumbents could jeopardize the expressive freedoms of entities with less economic or legal power.²³⁶ Together, they form a discourse in which Little Users infringe the economic rights of Big Content and Big Content tramples the expression of Little Users.

2. The Familiar Narrative Inverted

Undergirding the familiar narrative of digital copyright is the gospel (or dogma) that digital media fosters participatory culture. Networking, more than earlier technologies, enables ordinary individuals with minimal equipment to create,

231. If it is tendentious to characterize the Digital Millennium Copyright Act and the Sonny Bono Copyright Term Extension Act (CTEA) as the consequences of corporate lobbying, that tendentiousness is widely shared. In his official opposition to the CTEA, Senator Hank Brown wrote, "This bill reduces competition and primarily rewards those who lobby Congress rather than those who create." S. REP. NO. 104-315, at 36.

232. 17 U.S.C. § 1201 (West 2008) (prohibiting the circumvention of "a technological measure that effectively controls access to a work" protected by the Copyright Act).

233. *Eldred v. Ashcroft*, 537 U.S. 186, 254 (Breyer, J., dissenting).

234. Elkin-Koren, *supra* note 230, at 288–89.

235. Merges, *supra* note 218, at 134–35.

236. Of course, threats to established business models don't just harm kingpins. Content industry rhetoric rightly emphasizes the diversity of the artists and businesspeople it sustains. In Senate testimony, the musician Lars Ulrich described the recorded music business as "an industry with many jobs, a few glamorous ones like ours, and lots more covering all levels of the pay scale and providing wages which support families and contribute to our economy." See *Music on the Internet: Is There an Upside to Downloading? Hearing Before the S. Comm. on the Judiciary*, 106th Cong. (2000) (statement of Lars Ulrich, Member and Co-Founder, Metallica Musical Group).

exchange, and recombine media. Culture's "recipients," once networked, become "participants."²³⁷ Sometimes, more participation is deleterious: individuals who might have passively watched broadcast television can now actively share infringing copies of copyrighted media and threaten creators in the process. Other times, that participation is salutary: the internet enables innovations from the backwater to reach the mainstream, and, better yet, to commingle with other memes into something entirely new.²³⁸ In the headiest days of early user generated content, *Time* commemorated "You"—every reader who held that issue of the magazine and saw his or her reflection in its mirrored cover, a synecdoche for the millions of bloggers, video uploaders, and social networkers that the World Wide Web empowered—as its 2006 person of the year.²³⁹

The Internet's affordances for fixing and disseminating expression turned ordinary end users into "authors" in the legal sense. Of course, this is not to say that people expressed themselves more in the digital age than in any other, or that they did so more creatively or originally. Their expression simply became easier to fix "in a tangible medium" and share with large numbers of people.²⁴⁰ User generated content, unless it improperly appropriates some other expression, is user owned content. Media platforms like Facebook, YouTube, and Twitter do not own the copyrights to much of the content they furnish to users; rather, the platforms are licensees of content owned and uploaded by their users. This business model—a few Big Users monetizing lots and lots of Little Content—upends the premises of 1990's copyright rhetoric.

The metamorphosis of users into authors appears to have democratized semiotic power, for better and for worse.²⁴¹ In all likelihood, it has also enlarged and diversified the group of people that owns the rights to popular media. But these changes do not appear to have triggered a commensurate disaggregation of economic power. As Internet users took advantage of networking technology to create, share, and consume expression, companies emerged to monetize these expressive acts. Today, it is a handful of monolithic corporations that facilitate the web's participatory meaning making.

Ultimately, user generated data are what drive artificial intelligence. Cutting edge AI needs large amounts of high quality data to perform well; companies with the best and broadest access will be able to harvest the best data. For example,

237. This characterization flattens the nuances of user agency in digital and predigital media, but it is faithful to a general trend towards individuals' more participatory engagement with culture in the digital age (and to the rhetoric that accompanied that trend). See José van Dijck, *Users Like You? Theorizing Agency in User-Generated Content*, 31 MEDIA CULT. SOC. 41, 42–43 (2009); see also Jack M. Balkin, *Digital Speech and Democratic Culture: A Theory of Freedom of Expression for the Information Society*, 79 N.Y.U. L. REV. 1, 22 (2004).

238. Balkin, *supra* note 237, at 33–34.

239. Lev Grossman, *You — Yes, You — Are TIME's Person of the Year*, TIME (Dec. 25, 2006).

240. 17 U.S.C. § 101 (West).

241. See, e.g., Zeynep Tufekci, *Adventures in the Trump Twittersphere*, N.Y. TIMES (March 31, 2016), <https://nyti.ms/2k0WR1E> ("For decades, journalists at major media organizations acted as gatekeepers who passed judgment on what ideas could be publicly discussed, and what was considered too radical. This is sometimes called the 'Overton window.' . . . For worse, and sometimes for better, the Overton window is broken. We are in an era of rapidly weakening gatekeepers.").

Google's online translation service may behave as if it were pure machinery, but it in fact comprises thousands of insights gleaned from real human translators whose work appears online, or who furnish improved translations while using the service. Jaron Lanier writes:

A giant act of statistics is made practically free because of Moore's Law, but at core the act of translation is based on the real work of people. Alas, the human translators are anonymous and off the books. The act of cloud-based translation shrinks the economy by pretending the translators who provided the examples don't exist. With each so-called automatic translation, the humans who were the sources of the data are inched away from the world of compensation and employment.²⁴²

Because better data breed better services, scholars have observed that technology platforms can follow a "winner-take-all" economic distribution.²⁴³ The more user data a company can collect, the more it can improve data driven services like machine learning. These improved offerings will attract more users, and thus more data. This positive feedback loop enables so-called "super-platforms" to consolidate market power.²⁴⁴

In today's platform economy, value emerges not from the ownership of intellectual property rights in data, but from the ability to make licensed use of large amounts of data. Big Users are hegemonic, not Big Content. This reshuffling of owners and users demands a corresponding readjustment in intuitions about fair use.

Fair use limits rights holders' economic entitlements to subsidize certain classes of users.²⁴⁵ This subsidy is often justified by the particular "public benefit" that these uses provide.²⁴⁶ A fair user is motivated by a "humanitarian impulse" beyond personal profit.²⁴⁷ The uses that the copyright statute enumerates as paradigmatically fair—"criticism, comment, news reporting, teaching, ... scholarship, or research"—are characterized by a dedication to a purpose distinct from profit.²⁴⁸ They are, in Justice Blackmun's words, "activities the primary benefit of which accrues to others", and they stand in implicit contrast to purely commercial uses.²⁴⁹

Moreover, fair use is characterized as a safety valve to prevent the powerful from smothering the expressive rights of the less powerful. The Supreme Court has repeatedly described fair use as one of the "speech-protective purposes and safeguards" built into the law to offset rights holders' monopolies.²⁵⁰ Fair use

242. JARON LANIER, WHO OWNS THE FUTURE? 16 (2013).

243. PHILIP E. AUERSWALD, THE CODE ECONOMY: A FORTY-THOUSAND-YEAR HISTORY 198 (2017).

244. Ariel Ezrachi & Maurice E. Stucke, *Virtual Competition*, 7 J. EUR. COMPET. LAW & PRACT. 585, 585–86 (2016).

245. Merges, *supra* note 218, at 134–35.

246. *Sony Corp. v. Universal City Studios*, 464 U.S. 417, 480 (1984) (Blackmun, J., dissenting).

247. *Id.* at 496.

248. 17 U.S.C. § 107 (West).

249. *Sony*, 464 U.S. at 496; Fisher, *supra* note 161, at 1673.

250. *Golan v. Holder*, 132 S. Ct. 873, 876 (2012); *see also Eldred v. Ashcroft*, 537 U.S. 186, 218–19 (2003).

redistributes economic and expressive power. It curtails an otherwise outsized legal and economic entitlement so that “the public” can undertake certain socially beneficial activities. If the doctrine develops to give carte blanche to expressive machine learning, it will redistribute in the opposite direction: it will serve the economic interests of incumbent firms at the expense of disempowered rights holders.

Commercial machine learning, trained on expressive media, promises tremendous social value. But it is not the sort of value that fair use exists to foster. Unlike the benefits realized by, say, scholarship, the value of advanced machine learning services is internalized by the large firms that furnish those services. The companies with the best machine learning technology can use this to consolidate their market power, which in turn gives these firms the best access to user data that will further improve their machine learning technology.

The historical narrative of copyright and technology is one of powerful rights holders and marginal users. Today's tech business turns this structure on its head. Accordingly, scholars and jurists ought to recalibrate their intuitions about what fair use is and does. A progressive interpretation of copyright does not, in this circumstance, entail a broad construction of fair use. Indeed, upholding copyright's redistributive roots may require a return to the market based reasoning that, at the time, seemed to move against redistribution.

D. WHAT KIND OF PROGRESS?

The Constitutional mandate of copyright law is “to promote the Progress of Science and useful Arts[.]”²⁵¹ Expressive AI's fair use dilemma puts technological development at odds with certain forms of human authorship. Which kind of progress should the law privilege? Permissive fair use for machine learning would undeniably foster progress in the scientific field of artificial intelligence. It might also foster a certain kind of artistic progress. Unencumbered by copyright, AI could learn from all the greatest books, movies, and music. Perhaps this erudite AI would become so adept at making art as to supersede human creativity. The public, enamored of masterful robot art, might come to view human art as a primitive novelty, like paintings done by elephants. Human creators, in turn, might not derive any incentives from copyright law if robotic rivals undercut their earning potential. If robotic creators gave the public access to more, and better, works of art than any human artistic establishment could deliver—and, in so doing, marginalized the human artistic establishment—would that be the progress copyright law exists to promote?

Some formulations of copyright's purpose suggest that this is the progress the law prizes. Courts often describe copyright as an incentive system, designed with the ultimate goal of facilitating the creation of works for public consumption.²⁵² If the goal is to ensure a supply of works for the public, it may not matter whether

251. U.S. CONST. art. I, § 8, cl. 8.

252. Fox Film Corp. v. Doyal, 286 U.S. 123, 127–28 (1932); *Sony*, 464 U.S. at 431–32.

those works were created by humans or robots. If literate AI can furnish works the public wants to consume at a lower price point than human competitors, is this not the precise societal benefit that copyright exists to foster?

Maybe, maybe not. There is reason to think that copyright's progress mandates something more than the accrual of works. Barton Beebe distinguishes between "accumulationist" accounts of progress and a "pragmatist aesthetics" of progress. Beebe's aesthetic progress "focuses not on the stockpiling over time of fixed, archivable works but rather on the quality of ephemeral aesthetic experience in the present. . . . [P]ragmatist aesthetics measures aesthetic progress (or regress) largely by the extent of popular, democratic participation in aesthetic practice."²⁵³ In other words, the existence or the consumption of many cultural products may not suffice as progress; progress also entails meaningful participation in the creation of those cultural products. The progress of science is not, in Jessica Litman's words, "a giant warehouse filled with works of authorship."²⁵⁴ The value in human authorship flourishes still further when it is consumed, appreciated, and transformed by other humans. This cycle of creation and engagement is what the law clumsily tries to protect and propagate. Indeed, copyright places special value on human creativity and human reading; it "protects humans writing for humans."²⁵⁵ Therefore, doctrine that diminishes humans' capacities to create and engage with expression should be viewed warily. Fair use may be moving in this direction: Grimmelmann cautions that regulation of human readers and leniency towards robotic readers "discourages the personal engagement with a work that [copyright law] claims to value. Copyright's expressive message here—robots good, humans bad—is the exact opposite of the one it means to convey."²⁵⁶ If human authorship and human readership are indeed as special as copyright law suggests they are, it would be perverse for the law to marginalize human authors and readers in favor of erudite, eloquent machines. For this reason, a fair use regime that privileges the progress of robotic authors over incentivizing human authors warrants caution.

III. WHAT CAN BE DONE?

If given the opportunity to rebuild copyright from the ground up, a sensible person motivated by the public interest would probably not reconstruct it exactly as it exists today. But the current chimerical doctrine offers an extraordinary opportunity: a faithful interpretation of the doctrine suggests that progressive intervention is warranted. Parts I and II of this Article argued, respectively, that it is neither doctrinally correct nor normatively desirable for fair use to privilege unauthorized, expressive uses of copyrighted works in commercial machine learning. This Part contemplates several ways out of artificial intelligence's fair use dilemma. The purpose of this Part is not to identify a single panacea for AI's

253. Beebe, *supra* note 83, at 346.

254. Jessica Litman, *Lawful Personal Use*, 85 TEX. LAW REV. 1871, 1880 (2007).

255. Grimmelmann, *supra* note 11, at 660.

256. *Id.* at 675.

fair use crisis, let alone for the larger societal issues that AI raises. Rather, it is to demonstrate that there is no inevitable course the doctrine must take, and that some potential outcomes are more desirable than others.

A. LEVIES

Past technologies have presented copyright dilemmas not unlike the one described in Part II. On at least one occasion, Congress was willing to enact levies to address the problem. In the late 1980's, the advent of digital audio tape (DAT) technology allowed consumers to create "perfect" copies of recorded audio. The United States music recording industry, worried that widespread adoption of DAT technology would endanger their business model, threatened to sue the manufacturers and distributors of DAT technology.²⁵⁷ To mediate, Congress passed the Audio Home Recording Act ("AHRA"). The AHRA addressed copyright owners' concerns by attaching levies to the distribution of "any digital audio recording device or digital audio recording medium," which were allocated to two funds defined in the statute: the Sound Recordings Fund and the Musical Works Fund.²⁵⁸ Interested parties were further subdivided by the statute and disbursed royalties.²⁵⁹ In return, the Act shielded hardware manufacturers and distributors from liability by prohibiting infringement actions related to the distribution or noncommercial use of digital audio recording devices or media.²⁶⁰

The AHRA was no panacea.²⁶¹ In particular, the Act's narrow definitions of digital audio and media meant it applied only to outmoded technologies.²⁶² In the leading AHRA case, *Recording Industry Association of America v. Diamond Multimedia Systems*, a recording industry association sued the manufacturers of an MP3 player for noncompliance with the AHRA. The defendants maintained their device was not a "digital audio recording device" for the purposes of the Act, an argument that the district court concluded "would effectively eviscerate the AHRA."²⁶³ On appeal, the Ninth Circuit found for the defendants as a matter of first impression, concluding that "the Act seems to have been expressly designed to create this loophole."²⁶⁴

As ineffectual as it was, the AHRA might illuminate a solution the machine learning dilemma. The AHRA addressed a juridical logjam not unlike the problem

257. H.R. REP. No. 102-873, pt. 2.

258. 17 U.S.C. § 1003(b) (West).

259. 17 U.S.C §§ 1003-1004 (West).

260. 17 U.S.C. § 1008 (West).

261. DAVID NIMMER, *COPYRIGHT ILLUMINATED: REFOCUSING THE DIFFUSE US STATUTE* 103 (2008); Monica Zhang, Note, "Fair Compensation" in the Digital Age: Realigning the Audio Home Recording Act, 38 HASTINGS COMM. & ENT. L.J. 145, 158 (2016); see also Benton Gaffney, *Copyright Statutes that Regulate Technology: A Comparative Analysis of the Audio Home Recording Act and the Digital Millennium Copyright Act*, 75 WASH. L. REV. 611, 631-33 (2000).

262. *Recording Indus. Assoc. of Am. v. Diamond Multimedia Sys.*, 29 F.Supp.2d 624, 630 (C.D. Cal. 1998); Zhang, *supra* note 261, at 158.

263. *Diamond Multimedia Sys.*, 29 F.Supp.2d at 630.

264. *Recording Indus. Assoc. of Am. v. Diamond Multimedia Sys.*, 180 F.3d 1072, 1077 (9th Cir. 1999).

of expressive machine learning. As is the case with machine learning, pre-AHRA copyright law could either have buried a new technology by finding DAT distributors secondarily liable for infringement, or undermined rights holders' economic interests by allowing the technology to proliferate without restriction. Congress addressed the legal uncertainty with a compromise that enabled DAT technology to enter the market and provided some compensation for the music industry's foregone revenue.

Of course, home audio taping is not machine learning. The AHRA regulated technologies that might be used to make unauthorized copies of works. In machine learning, the clearest copyright infringements happen before, or during, a model's public deployment.²⁶⁵ Machine learning corpora comprise far more works by far more individuals, and the individual harms it causes are far smaller. The number and variety of the works implicated in expressive machine learning could make for an accounting regime even more convoluted than that of the AHRA.

Still, this approach—levying one's way out of a doctrinal and economic dilemma—could bear fruit, and the AHRA shows that the United States legislature was willing to support something like it the past. It also resembles the more general solutions that are already being proposed to allay AI-related fears. In a draft report to the European Parliament's Commission on Civil Law Rules on Robotics, rapporteur Mady Delvaux highlighted “the possible need to introduce corporate reporting requirements on the extent and proportion of the contribution of robotics and AI to the economic results of a company for the purpose of taxation and social security contributions[.]”²⁶⁶ The Obama White House report on “Artificial Intelligence, Automation, and the Economy” notes that “[a]dvanced AI systems could reinforce trends of national income shifting from labor to capital . . . [and] significantly exacerbate the rise in income inequality seen over the past few decades, absent an appropriate policy response[.]” such as a progressive taxation regime with higher taxes on capital.²⁶⁷

Moreover, while comprehensive tax reform might seem redistributionist and unpalatable to a large segment of the United States, AI's fair use dilemma gives progressive policymakers greater leverage. Because an honest interpretation of fair use threatens the future of expressive AI, levies on machine learning could be seen as a bargain rather than a taking. In exchange for a levy or taxation regime, the legislature could prohibit infringement actions related to machine learning and thereby secure the future of the technology.

Calculating the appropriate levy, and prescribing its disbursement, is a more ambitious task than this Article can fulfill. But potential models already exist. One would be to establish AHRA-style funds for authors of training data and outline royalty rates by statute; in the AHRA's case, royalties were set at a percentage of the transfer price bounded by statutory maxima and minima that could be revised

265. See *supra* Part I.D.

266. Draft Report with Recommendations to the Commission on Civil Law Rules on Robotics, EUR. PARL. DOC. 2015/2103 (INL) (2016).

267. EXEC. OFFICE OF THE PRESIDENT, *supra* note 2, at 41.

by Copyright Royalty Judges.²⁶⁸ A similar example is the rejected Google Books Settlement, which proposed a Book Rights Registry to collect proceeds of Google Books sales and distribute them to appropriate rights holders.²⁶⁹ Alternatively, because individual uses are difficult to value and to track, machine learning levies could bypass the accounting stage entirely. They might be remitted to appropriate artistic charities, as was prescribed for unclaimed funds in the Google Books settlement.²⁷⁰ They might reinforce a general safety net for all citizens. Or they might be put towards government programs specifically designed to support individuals who seek to make a living from creative, intellectual work.

B. DOCTRINAL INNOVATIONS

AI's fair use crisis is a result of a doctrinal logjam, so doctrine will be unlikely to solve it. Nonetheless, judges could take some imaginative actions to enforce compromise. One way to do so would be taking one of Lanier's admonishments to heart: "Big data has to be understood as repackaging of human labor. It's not stand-alone automation, it's human-driven automation[.]"²⁷¹ As a practical matter, the progress of science does not appear *ex nihilo*; even the most original authors are influenced by others' works. This influence is devilishly hard to quantify, but the law has tried to do so in the past by apportioning works' proceeds to rights holders with competing interests, proportional to their contributions to a work.²⁷² Theoretically, the judiciary could offer plaintiffs compensation on a case by case basis without affording them the power to destroy expressive machine learning. In general, copyright is structured to avoid overlapping entitlements.²⁷³ However, a judge could find a machine learning model to be a derivative work based on its training data, and use this finding to mandate apportionment of profits.

C. MOONSHOTS

The levy system described above is a modest compromise—a step, not a leap, in the right direction, and perhaps a foot in the door to justify broader reform. It will not, by itself, ensure distributive equity or shepherd the middle class through the AI age. Far bolder proposals exist. Some are technological: the abortive Project Xanadu has sought for decades to supplant the internet's HyperText Markup Language ("HTML") standard with a system of two way linking that affords

268. 17 U.S.C. § 1004 (West).

269. Amended Settlement Agreement, *supra* note 143, § 3.2(d)(i), at 29.

270. *Id.* § 6(i)(3), at 81.

271. *Quoted in* Auerswald, *supra* note 243, at 197.

272. In the context of an infringement action, Judge Learned Hand famously called the problem "insoluble." *Sheldon v. Metro-Goldwyn Pictures Corp.*, 106 F.2d 45, 48 (2d Cir. 1939).

273. Joint authorship, for example, permits two or more individuals to be the authors of a single work, but the restrictions it places on the types of works that may qualify are far too stringent to include machine learning models. A joint work is "a work prepared by two or more authors with the intention that their contributions be merged into inseparable or interdependent parts of a unitary whole." 17 U.S.C. § 101 (West).

content owners greater control over the use of their works.²⁷⁴ At the very bottom of its bare bones homepage, Project Xanadu offers the gnomic disclaimer, “DIFFERENT COPYRIGHT SYSTEM REQUIRED. Our matching copyright system for indirect content: the transcopyright license, permitting remix in any context and any quantity, and with automatic attribution of authorship.”²⁷⁵ Xanadu’s two way linking scheme, which would ensure that original works maintain a connection to quotations of them that appear elsewhere, undergirds Lanier’s idea for “humanistic computing,” an information infrastructure that tracks and values individuals’ contributions to networked resources, and disburses “nanopayments” accordingly.²⁷⁶ Other proposals use doctrinal and political reforms to shape a more efficient information ecosystem. William Fisher propounds the idea of a government funded reward system for entertainment media to facilitate public access to cultural products, which would compensate creators through revenue drawn from increased taxes, in lieu of exclusive economic rights.²⁷⁷ This Article cannot list, let alone evaluate, all such proposals, but they are a helpful contrast to its more modest suggestions.

D. COPYRIGHT IS NO SILVER BULLET

AI threatens to increase economic and social inequality in many ways. Copyright law, deployed judiciously, could mitigate some of these harms—but only a narrow subset of them. While many expressive uses of AI may promote economic inequality, only some of them raise genuine copyright issues. Negative social consequences of AI stem from a constellation of social and legal factors, and changes to copyright alone will not remedy all, or even most, of these consequences.

1. Contracts

A great deal of expressive machine learning can take place without copyright liability. This is because companies often license training data from end users through broad terms of service. These arrangements give a few powerful platforms access to immense stockpiles of data that can train and refine various forms of artificial intelligence. If Google wanted to train the Smart Reply engine to write like David Foster Wallace, the company might have to license his oeuvre, but if Google were content with Smart Reply writing like every Gmail user, no additional licensing would be needed.

274. Lanier, *supra* note 242, at 221; *Project Xanadu*®, PROJECT XANADU, <https://perma.cc/W29K-R342> (last visited Oct. 23, 2017); *The Curse of Xanadu*, WIRED (June 1, 1995 12 PM), <https://perma.cc/UA98-T6G6>.

275. Project Xanadu®, *supra* note 274; *The Transcopyright License*, PROJECT XANADU, <https://perma.cc/G8L6-MRSJ> (last visited Oct. 23, 2017).

276. Lanier, *supra* note 242, at 235, 242–43, 258–62.

277. WILLIAM W. FISHER, PROMISES TO KEEP: TECHNOLOGY, LAW, AND THE FUTURE OF ENTERTAINMENT 199 (2004).

Expressive AI trained on data licensed by end user assent may present similar economic issues to AI trained on unauthorized copies of copyrighted works, but contracting conventions in cyberspace place these issues outside the scope of copyright law. Bartering one's personal information for gratis services may already be an unfair transaction for end users. If it is the case that this AI will threaten individuals' livelihoods more in the future than it does in the present moment, this transaction may become more unfair.²⁷⁸ In any event, as data driven AI flourishes, the standard form contract model for data collection may not be able to adjust to changes in the balance of the transaction. Research suggests, for example, that too few users read terms of service to sustain an "informed minority" that polices platforms' behavior to promote an efficient marketplace.²⁷⁹

2. Algorithmic Bias

The decision-making capabilities of a machine learning algorithm depend on its training data. Biased data could lead an algorithm to perpetuate and reinforce those biases. Some facial recognition algorithms perform better on whites than on minorities; a Nikon camera insists that its Asian-American owner is blinking due to the shape of her eyes; some algorithms used to predict risk of criminality may disproportionately disadvantage blacks.²⁸⁰ Easy access to less biased training data could mitigate these problems. Insofar as copyright fetters this access, it does the public a disservice. This rationale might suggest that limiting the availability of fair use for machine learning could have undesirable consequences for social justice, and therefore that this Article's recommendations are unwise. In fact, the relative inaccessibility of low-bias training data is a problem for which copyright is, at most, only partially responsible.

First, although this Article's analysis of fair use suggests that the doctrine should not excuse expressive, commercial machine learning, it would permit many uses motivated by social justice. Issues of disparate impact in AI often pertain to the analysis of factual data and/or the output of purported facts—for example, the analysis of factual data about detainees or criminals and the output of purportedly factual risk scores, or the analysis of photographs and the output of purportedly factual identity matches.²⁸¹ These uses do not implicate protected expression in source data. Auditing data to determine the nature and extent of bias would be similarly non-expressive, because it matches factual outcomes with factual demographics. Some biased applications of machine learning may be more expressive in character. For example, machine translation can associate gender

278. Of course, the powerful AI services of the future could offer such dramatic benefits as to outweigh increased economic threats to users.

279. See Yannis Bakos et al., *Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts*, 43 J. LEGAL STUD. 1, 2–3 (2014).

280. Levendowski, *supra* note 98, at 6, 19.

281. Julia Angwin et al., *Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks*, PROPUBLICA (May 23, 2016), <https://perma.cc/LG4Q-MCBQ>; Claire Garvie et al., *The Perpetual Line-Up*, GEORGETOWN LAW CENTER ON PRIVACY & TECHNOLOGY (Oct. 18, 2016), <https://perma.cc/9RZH-D7F5>.

with certain professions or traits as a result of biases in training data. If engineers made unauthorized use of copyrighted data for the sole purpose of debiasing an expressive program, this Article's formulation of fair use would excuse it. But even if fair use were unavailable, altruistic engineers would not be out of luck. Bias issues of this kind may be ameliorated by mathematically transforming problematic models, rather than acquiring new data.²⁸²

Second, there is little reason to believe that copyright liability deters the use of particular data in machine learning. Recall a Google spokesperson's response to its use of the BookCorpus dataset: "The machine learning community has long published open research with these kinds of datasets, including many academic researchers with this set of free ebooks—it doesn't harm the authors and is done for a very different purpose from the authors', so it's fair use under US law."²⁸³ Little about the conduct of machine learning researchers or companies suggests an awareness of copyright liability, let alone any forbearance because of it.

Third, even if suitably unbiased training datasets in fact exist, it is unlikely that an ungenerous fair use doctrine is the sole force keeping them out of circulation. Imagine, for argument's sake, that the stock photography company Getty Images has a copyright protected repository of images that could be used to train less biased facial detection software, and that low resolution samples of these images are posted on Getty's website for public viewing. Even with an ironclad fair use defense, other statutes may outlaw the use of these photos to train AI. Assembling the dataset could violate the anticircumvention provisions of the Digital Millennium Copyright Act,²⁸⁴ as well as the Computer Fraud and Abuse Act, which restricts unauthorized access to computers and has been described as a "para-copyright tool to secure exclusivity to otherwise publicly accessible data."²⁸⁵

Finally, the most obvious obstacle to egalitarian machine learning is that the highest quality datasets are inaccessible not because of copyright law, but because of secrecy. Data titans have little incentive to license their user data *en masse* to interested parties. This is not surprising: these companies' troves of data are singularly valuable assets, and exclusive possession confers a distinct competitive advantage that in turn allows them access to more high quality data.²⁸⁶ Moreover, a company's willingness to license user data to third parties might discourage privacy conscious individuals from using the service.

CONCLUSION

Advances in artificial intelligence and changes in the digital information economy have placed the fair use doctrine in crisis. Today, economic dominance

282. Tolga Bolukbasi et al., *Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings*, ARXIV 11 (July 21, 2016), <https://perma.cc/2CNS-BFFZ>.

283. Lea, *supra* note 124.

284. Von Lohmann, *supra* note 225, at 9.

285. Nicholas A. Wolfe, *Using the Computer Fraud and Abuse Act to Secure Public Data Exclusivity*, 13 NW. J. TECH. & INTELL. PROP. 301, 303 (2015).

286. *See supra* Part II.C.2.

belongs not to incumbent rights holders or exclusive licensees, but to the large internet platforms that use others' data. Machine learning technology empowers these companies to extract value from authors' protected expression without authorization, and to use that value for commercial purposes that may someday jeopardize the livelihoods of human creators. Construing fair use to protect this activity will place the doctrine at odds with the public interest and potentially exacerbate the social inequalities that AI threatens. At the same time, finding that expressive machine learning is not fair use would frustrate the progress of the promising technology.

The numerous challenges AI poses for the fair use doctrine are not, in themselves, reasons to despair. Machine learning will realize immense social and financial benefits. Its potency derives in large part from the creative work of real human beings. The fair use crisis is a crisis precisely because copyright's exclusive rights may now afford these human beings leverage that they otherwise would lack. The fair use dilemma is a genuine dilemma, but it offers an opportunity to promote social equity by reasserting the purpose of copyright law: to foster the creation and dissemination of human expression by securing, to authors, the rights to the expressive value in their works.