

COPYRIGHT ACCELERATIONISM

BENJAMIN L.W. SOBEL *

<u>INTRODUCTION</u>	74
<u>I.WHY ACCELERATIONISM? COPYRIGHT INCREMENTAL-MINIMALISM’S GRADUAL DEMISE</u>	75
<u>A. A Taxonomy of Accelerationism (and Its Impersonators)</u>	77
1. <u>True Accelerationism Pursues Instability</u>	77
a) <u>Marx’s Accelerationism: “Accelerate the Bad Thing”</u>	77
b) <u>“Accelerate the Good Thing”</u>	78
2. <u>Ersatz Accelerationisms</u>	78
a) <u>Marc’s Accelerationism: Selfish Ersatz Accelerationism</u>	79
b) <u>Altruistic Ersatz Accelerationism</u>	81
<u>B. Copyright Accelerationism</u>	82
<u>II.A COPYRIGHT ACCELERATIONIST’S GUIDE TO GENERATIVE AI</u>	88
<u>A. Fair Use and Learning</u>	89
1. <u>Fair Use Does Not Privilege (Human) Learning</u>	89
2. <u>Incremental-Minimalist Approaches to Fair Use and AI</u>	91
3. <u>The Accelerationist Response</u>	94
4. <u>Pick Your Pony</u>	97
<u>B. The AI Industry’s Neo-Formalities Movement</u>	99
<u>CONCLUSION</u>	104

* Postdoctoral Fellow, Digital Life Initiative, Cornell Tech; Program Affiliate Scholar, New York University Information Law Institute. Thanks to James Grimmelman, Frank Pasquale, David Nimmer, Rebecca Wexler, Cornell Tech’s CTRL-ALT discussion group, the attendees of the 2024 WIPIP Conference at Santa Clara Law, and the participants in the *Chicago-Kent Law Review*’s AI Disrupting Law symposium for helpful feedback.

INTRODUCTION

[T]he free trade system is destructive. It breaks up old nationalities and pushes the antagonism of the proletariat and the bourgeoisie to the extreme point. In a word, the free trade system hastens the social revolution. It is in this revolutionary sense alone, gentlemen, that I vote in favor of free trade.

— Karl Marx, 1848¹

Humans doing the hard jobs on minimum wage while the robots write poetry and paint is not the future I wanted.

@KarlreMarks on X, 2023²

Modern copyright law seems determined to impede people's engagement with creative expression. The absence of formal prerequisites to copyright ownership locks up works with no commercial value, written by authors who may not wish to claim property rights. Copyright's originality requirement is so low that all but the most mindless emails and shopping lists satisfy it. Sweeping exclusive rights undermine salutary creative interchange. In tandem with these doctrines, an irrationally long copyright term ensures that nearly all recorded culture is encumbered not merely for years, but for generations.

Today, however, change is in the air—for all the wrong reasons. By historical accident, the same foundational properties of copyright law that have long undermined creators and audiences now happen to pose an existential threat to generative AI technology. Tech companies and their allies are pushing zealously to reform the very aspects of copyright law that impoverish traditional readership and authorship. But by and large, their proposals would change these doctrines only in ways that benefit the generative AI enterprise. AI could “learn” from pirated textbooks, but flesh-and-blood students would still pay full freight. Authors would have to affirmatively opt out of AI training datasets while other copyright entitlements continue to vest without formality. This is a world in which humans, as ever, are left high and dry.

1. Karl Marx, *On the Question of Free Trade*, (Jan. 9, 1848), <https://cooperative-individualism.org/marx-karl-on-the-question-of-free-trade-1848.htm> [perma.cc/VAK5-3BT7].

2. Karl Sharro (@KarlreMarks), X (May 15, 2023, 3:34 AM), <https://twitter.com/KarlreMarks/status/1658028017921261569> [perma.cc/JVS8-92K2].

This essay offers an alternative: copyright accelerationism. It is a proposal so staid that it sounds radical. Stop seeking exceptions to copyright law that, in purpose and effect, serve only the interests of AI firms. Stop opposing copyright's anti-reader doctrines on the grounds that they threaten the AI enterprise. Instead, take the law's apparent principles and commitments at face value and push them as far as they will go. Then, push them some more. Our copyright regime falls short of its constitutional mandate to promote the progress of knowledge. We will never correct course by letting the powerful exempt themselves from the copyright regime and leaving its intended beneficiaries to bear an inequitable burden. But we *will* correct course if we insist that copyright applies to everyone on equal terms—because it will make a regime that has long been untenable for some into a regime that is untenable for all.

I proceed in two parts. In Part I, I briefly define and contextualize the accelerationist ethos. I then explain what the politics of copyright accelerationism would look like and why I advocate it over the “incremental-minimalist” (or “copyright decelerationist”) position prevalent today. Part II focuses on two case studies: the AI industry's push to broaden the fair use doctrine to excuse AI training on unauthorized copies of copyrighted works and the AI industry's efforts to establish a *de facto* formalities regime that assumes permission-by-default for AI training. I contrast accelerationist approaches to these scenarios with the incremental-minimalist proposals that loom in academic and activist discourse. Copyright incremental-minimalists, I conclude, have lost the plot. They are so focused on fighting the last war—the war against the copyright-maximalist media behemoths that dominated the culture economy in the 1990s and early 2000s—that they view any assertion of authors' rights as jeopardizing their admirable dream of universal access to knowledge. Copyright incremental-minimalists appear to believe that if they align with the AI industry's interests today, they can still realize the permissionless remix culture they've longed for. They're not wrong, but there's a catch—that permissionless remix culture will be for AI alone. The better path is the one the incremental-minimalists reflexively resist: the self-destructive expansion of copyright itself.

I. WHY ACCELERATIONISM? COPYRIGHT INCREMENTAL-MINIMALISM'S GRADUAL DEMISE

Googling “accelerationism” yields a frightening panoply. Some of the first results are news articles that describe accelerationism as “the obscure

idea inspiring white supremacist killers around the world.”³ These headlines might lead you to believe that accelerationism is just for far-right reactionaries. That belief would be mistaken. Indeed, the critical theorist who coined the term, Benjamin Noys, did so to describe attitudes that he linked to the “ultra-left.”⁴ The theorists Noys associated with accelerationism—Deleuze and Guattari, Lyotard, and Baudrillard—sought to further Marxist goals by doubling down on the very systems Marxism opposed.⁵ Noys summarized their accelerationist thinking as follows: “if capitalism generates its own forces of dissolution then the necessity is to radicalise capitalism itself: the worse the better.”⁶

Accelerationism is a philosophy open to just about anyone who isn’t a moderate. One can hold accelerationist views about technological, social, and/or economic issues. Learning that someone is an accelerationist tells you little about her ideological commitments, since she could be anything from a left-wing, post-capitalist radical to a neo-Nazi.⁷ Because accelerationism can connote such disparate ideologies—and because so many people are now claiming to be accelerationists even when the label doesn’t accurately reflect their views—it’s worth taxonomizing the philosophy to explain copyright accelerationism’s place within it. Copyright accelerationism falls somewhere between the true accelerationisms and ersatz accelerationisms that the next sub-Part taxonomizes. Copyright accelerationism is true accelerationism in that it aspires, by heightening copyright law’s contradictions in the AI

3. Zack Beauchamp, *Accelerationism: The Obscure Idea Inspiring White Supremacist Killers Around the World*, VOX (Nov. 18, 2019, 8:15 AM), <https://www.vox.com/the-high-light/2019/11/11/20882005/accelerationism-white-supremacy-christchurch> [perma.cc/H5BR-FHHV]; accord *White Supremacists Embrace “Accelerationism,”* ANTI-DEFAMATION LEAGUE (Apr. 16, 2019), <https://www.adl.org/resources/blog/white-supremacists-embrace-accelerationism> [https://perma.cc/YLP7-9X7K]; Daveed Gartenstein-Ross et al., *The Growing Threat Posed by Accelerationism and Accelerationist Groups Worldwide*, FOREIGN POL’Y RSCH. INST. (Apr. 20, 2020), <https://www.fpri.org/article/2020/04/the-growing-threat-posed-by-accelerationism-and-accelerationist-groups-worldwide/> [https://perma.cc/TP9D-2NWR].

4. BENJAMIN NOYS, *THE PERSISTENCE OF THE NEGATIVE: A CRITIQUE OF CONTEMPORARY CONTINENTAL THEORY* 5 (2010), <http://www.cambridge.org/core/product/identifier/9780748643295/type/BOOK> [perma.cc/6TZA-WDKM]; *Accelerationism*, WIKIPEDIA (2023), <https://en.wikipedia.org/w/index.php?title=Accelerationism&oldid=1186854157> [perma.cc/WW6F-XT47] (stating that Noys coined the term “accelerationism.”).

5. NOYS, *supra* note 4, at 4-5.

6. *Id.* at 5.

7. Compare, e.g., Alex Williams & Nick Srnicek, *#Accelerate Manifesto for an Accelerationist Politics*, CRITICAL LEGAL THINKING, § 2(7) (May 14, 2013), <https://criticallegalthinking.com/2013/05/14/accelerate-manifesto-for-an-accelerationist-politics/> [perma.cc/WW6F-XT47] (“[T]he assessment of left politics as antithetical to technosocial acceleration is also, at least in part, a severe misrepresentation.”), with *Atomwaffen Division*, S. POVERTY L. CTR., <https://www.splcenter.org/fighting-hate/extremist-files/group/atomwaffen-division> [perma.cc/F3NE-VAET] (describing “Atomwaffen Division . . . a terroristic neo-Nazi organization” whose “members . . . can be fairly described as accelerationists”).

age, to destabilize modern copyright doctrines and thereby realize a better regime. But copyright accelerationism is more reformist than revolutionary. It's still legalism, after all; it assumes the stability of the larger legal structure that contains copyright, and in this regard, copyright accelerationism resembles ersatz accelerationisms that do not pursue instability.

A. A Taxonomy⁸ of Accelerationism (and Its Impersonators)

Accelerationism pushes a trend to extremes in order to unleash a runaway, destabilizing force. That force might be unthinkable powerful technology, the world-eating logic of unrestrained capitalism, or something else entirely. Accelerationism hopes that by provoking instability, it will transform or supersede the structures it destabilizes. This desire for instability is distinct from a desire merely for rapid development, which is what characterizes the ersatz accelerationism that is now popular among Silicon Valley elites.

1. True Accelerationism Pursues Instability

An accelerationist seeks to exaggerate destabilizing trends to create the conditions for a transformation more radical than could be realized by incremental change. Some accelerationists exaggerate trends they oppose in order to heighten their contradictions and force change. Others exaggerate trends they regard as beneficial in order to sow instability that they view as productive.

i. Marx's Accelerationism: "Accelerate the Bad Thing"

The Marx quote that begins this essay is paradigmatic accelerationism.⁹ In it, Marx explains that between the policies of free trade and economic protectionism, he prefers the unrestrained capitalism of free trade because its "destructive" power will "hasten[] the social revolution."¹⁰ Marx's stance here is the classic accelerationist posture: he argues for aggravating a dynamic that he disapproves of—unchecked capitalism—in order to destabilize the system that gave rise to that dynamic in the first place.¹¹ But this sort of

8. I thank James Grimmelman for suggesting the taxonomical categories I use here.

9. See Nick Land, *A Quick-and-Dirty Introduction to Accelerationism*, JACOBITE (May 25, 2017), <https://web.archive.org/web/20210829071454/https://jacobitemag.com/2017/05/25/a-quick-and-dirty-introduction-to-accelerationism/> [perma.cc/B48R-5CRA] (referring to the Marx quote as an "accelerationist fragment").

10. Marx, *supra* note 1.

11. Marx may not have been a thoroughgoing accelerationist. For example, he apparently wrote warmly of some reformist legislation that regulated the conditions of factory labor. See Ramesh Mishra, *Marx and Welfare*, 23 SOCIO. REV. 287, 293-95 (1975). But see KARL MARX, CAPITAL, § 2 (1867) <https://www.marxists.org/archive/marx/works/1867-c1/ch10.htm> [perma.cc/JXZ3-DPPG] ("[T]he

accelerationism is by no means limited to Marxists. Other “accelerate-the-bad-thing” accelerationists fixate not on technology or financial systems, but on disturbing visions of social relations. For example, some accelerationist white supremacists argue that their followers should support socially progressive policies that they oppose in substance in order to hasten reactionary backlash and racial polarization.¹²

ii. “Accelerate the Good Thing”

Accelerationists also pursue destabilizing change by accelerating trends they regard as good in themselves. For example, some contemporary anti- or post-capitalists favor the acceleration of technological progress because they believe that through technological “mastery over society and its environment,” humanity can supersede the capitalist social order.¹³ Nick Srnicek demands “a fully automated economy” in which “machines . . . produce all necessary goods and services, . . . releasing humanity from the effort of producing them.”¹⁴ Aaron Bastani calls for “fully automated luxury communism,” a future in which technological progress leads to society-wide abundance, and work becomes not a means of survival but a means of self-actualization.¹⁵ In this “accelerate-the-good-thing” paradigm, technological acceleration is both the *material* key to a lifestyle of universal superabundance and the destabilizing catalyst for the *political* change required to realize such a future.

2. Ersatz Accelerationisms

Ersatz accelerationisms are philosophies that resemble accelerationism—and sometimes purport to be accelerationist—but which assume the stable persistence of the structure in which they unfold. Two competing ersatz accelerationisms are popular in contemporary Silicon Valley, one that accelerates in the name of self-interest and another that accelerates in the name of altruism. Both resemble true accelerationism in their calls for rapid technological progress. In truth, however, they are ersatz accelerationisms, either because they envision that free-market capitalism will stably persist even as this rapid technological development unfolds or because they pursue

limiting of factory labour was dictated by the same necessity which spread guano over the English fields. The same blind eagerness for plunder that in the one case exhausted the soil, had, in the other, torn up by the roots the living force of the nation.”).

12. ANTI-DEFAMATION LEAGUE, *supra* note 3.

13. Williams & Srnicek, *supra* note 7, § 3(21).

14. NICK SRNICEK, INVENTING THE FUTURE: POSTCAPITALISM AND A WORLD WITHOUT WORK 109 (2015); *see also* Williams & Srnicek, *supra* note 7.

15. AARON BASTANI, FULLY AUTOMATED LUXURY COMMUNISM 50–56 (2019).

technological mastery in order to contain potential instability rather than to exacerbate it.

i. Marc's Accelerationism: Selfish Ersatz Accelerationism

Neoliberal capitalists are now styling themselves as “accelerationists.”¹⁶ Silicon Valley elites are embracing “effective accelerationism,” often shortened to “e/acc,” a philosophy that *Forbes* magazine recently described as “hasten[ing] the growth of technology and capitalism at the expense of nearly anything else.”¹⁷ The venture capitalist Marc Andreessen’s October 2023 “Techno-Optimist Manifesto,” a fervid panegyric to technology and free markets, sets out the e/acc vision. “[G]rowth,” Andreessen explains, “is progress – leading to vitality, expansion of life, increasing knowledge, higher well being,” and “the only perpetual source of growth is technology.”¹⁸ “[F]ree markets,” in turn, “are the most effective way to organize a technological economy . . . markets are an upward spiral.”¹⁹ Thus, the route to “material abundance for everyone” and “[l]iberat[ion] of human potential,” Andreessen argues, is runaway technological progress under a free-market-capitalist system whose participants pursue self-interest:

Combine technology and markets and you get . . . the techno-capital machine, the engine of perpetual material creation, growth, and abundance . . . We believe in accelerationism – the conscious and deliberate

16. See, e.g., Marc Andreessen, *The Techno-Optimist Manifesto*, ANDREESSEN HOROWITZ (Oct. 16, 2023), <https://a16z.com/the-techno-optimist-manifesto/> [perma.cc/LRL8-SZWB]; Ezra Klein, *The Chief Ideologist of the Silicon Valley Elite Has Some Strange Ideas*, N.Y. TIMES, (Oct. 22, 2023), <https://www.nytimes.com/2023/10/26/opinion/marc-andreessen-reactionary-futurism.html> [https://perma.cc/Z6SZ-GGC5].

17. Emily Baker-White, *Who Is @BasedBeffJezos, The Leader Of The Tech Elite's 'E/Acc' Movement?*, FORBES (Dec. 1, 2023), <https://www.forbes.com/sites/emilybaker-white/2023/12/01/who-is-basedbeffjezos-the-leader-of-effective-accelerationism-eacc/> [https://perma.cc/V7GJ-485G]. Effective accelerationism is itself a reference to “effective altruism,” a movement purportedly dedicated to “using evidence and careful reasoning to take actions that help others as much as possible,” see, e.g., *What is Effective Altruism?*, HARV. L. SCH., <https://orgs.law.harvard.edu/effectivealtruism/about-us/about-effective-altruism/> [perma.cc/9ZGQ-QVUZ], whose adherents have been preoccupied with the possibility of a superintelligent artificial intelligence eliminating humankind. See *The Sam Altman Drama Points to a Deeper Split in the Tech World*, ECONOMIST (Nov. 19, 2023), https://www.economist.com/business/2023/11/19/the-sam-altman-drama-points-to-a-deeper-split-in-the-tech-world?utm_medium=cpc.adword.pd&utm_source=google&ppccampaignID=17210591673&ppcadID=&utm_campaign=a.22brand.pmax&utm_content=conversion.direct-response.anonymous&gad_source=1&gbraid=0AAAAADBuq3ImUUBjZX3GQDgSLC-MJWqh4&gclid=Cj0KCQjwh_iBhCzARIsANimeoGmN7tKm7yp9pLOlaprG4f-qd3ADhb5Ksprj9x1-4jhDAdmc86l5zoaAve_EALw_wcB&gclid=aw.ds [https://perma.cc/92AY-5YPF].

18. Andreessen, *supra* note 16.

19. *Id.*

propulsion of technological development – to ensure . . . the techno-capital upward spiral continues forever.²⁰

Writing in the *New York Times*, Ezra Klein sums up Andreessen’s thesis: “Technology is good. Very good. Those who stand in its way are bad.”²¹

Andreessen claims to embrace “accelerationism,” but he and his fellow travelers misuse the term. Andreessen indeed desires *speed*; he wants to see a runaway trend—in this case, technological innovation—build upon itself in a positive feedback loop. But Andreessen’s ethic is not *accelerationism*: his vision is not that runaway technological progress will destabilize the status quo and thus catalyze radical change. Rather, his aspiration is that this upward spiral of technology will *entrench* the dominant economic paradigm of free-market capitalism. Writing a decade before Andreessen’s Techno-Optimist Manifesto, Nick Srnicek and Alex Williams recognized that an ethos like Andreessen’s is ersatz accelerationism: in their words, it “confuses speed with acceleration.”²² In this ersatz accelerationism, “[w]e may be moving fast, but only within a strictly defined set of capitalist parameters that themselves never waver.”²³

In other words, Andreessen’s e/acc is not accelerationism because it doesn’t seek to destabilize a stable hegemony.²⁴ While capitalism might itself be a state of instability and “creative destruction,” the e/acc aspiration is the stable persistence of free-market capitalism.²⁵ Andreessen’s brand of e/acc certainly opposes things, but the things it opposes do not represent the established, stable order so much as efforts to challenge that order. In a subsection titled “Enemies,” Andreessen’s manifesto explains that “[o]ur enemies are . . . bad ideas.”²⁶ As present-day manifestations of these bad ideas,

20. *Id.* Andreessen’s manifesto includes a citation to Nick Land, whom the *Guardian* in 2017 called a “central figure[] of accelerationism.” Andy Beckett, *Accelerationism: How a Fringe Philosophy Predicted the Future We Live In*, THE GUARDIAN (May 11, 2017, 12:30 AM), <https://www.theguardian.com/world/2017/may/11/accelerationism-how-a-fringe-philosophy-predicted-the-future-we-live-in> [<https://perma.cc/7DX9-5S8R>]. Land’s more recent pursuits include the dissemination of tracts denouncing democracy and espousing “neo-reactionary” politics. See Nick Land, *The Dark Enlightenment*, THE DARK ENLIGHTENMENT, <https://www.thedarkenlightenment.com/the-dark-enlightenment-by-nick-land/> [<https://perma.cc/2P4F-RVB2>].

21. Klein, *supra* note 16.

22. Williams & Srnicek, *supra* note 7, § 2(2).

23. *Id.*

24. Admittedly, at the margins, it can be difficult to discern whether e/acc adherents want stability or instability. Some of e/acc’s more esoteric apostles hope for a “technocapital singularity” that will create “unthinkable next-generation lifeforms and silicon-based awareness.” @zestular et al., *Effective Accelerationism — e/Acc*, E/ACC NEWSLETTER (Oct. 30, 2022), <https://effectiveaccelerationism.substack.com/p/repost-effective-accelerationism> [<https://perma.cc/7KXM-2H76>]. A commitment to hastening the arrival of “unthinkable next-generation lifeforms” is, I think, destabilizing enough to qualify as bona-fide accelerationism, even if it assumes the continued dominance of free-market capitalism. *Id.*

25. See Carol M. Kopp, *Creative Destruction: Out With the Old, in With the New*, INVESTOPEDIA (updated Feb. 20, 2023), <https://www.investopedia.com/terms/c/createdestruction.asp> [<https://perma.cc/7KXM-2H76>].

26. Andreessen, *supra* note 16.

Andreessen cites such concepts as “sustainability,” “social responsibility,” and “tech ethics.”²⁷ With the possible exception of “risk management,”²⁸ none of the “bad ideas” Andreessen cites is hegemonic today. In fact, Andreessen specifically characterizes his “enemies” as “zombie ideas, many derived from Communism, disastrous then and now – that have refused to die.”²⁹ These ideas can be quashed, but they can’t be *destabilized* in the sense that accelerationism contemplates, because they aren’t stably entrenched to begin with. Thus, because it does not pursue acceleration in order to sow instability, Andreessen’s e/acc is not accelerationism, as I use the term.³⁰

ii. *Altruistic Ersatz Accelerationism*

Another form of ersatz accelerationism is the avowedly altruistic futurism that was the founding mission of OpenAI. In a blog post from late 2015 titled “Introducing OpenAI,” OpenAI’s cofounders observed,

It’s hard to fathom how much human-level AI could benefit society, and it’s equally hard to imagine how much it could damage society if built or used incorrectly.

Because of AI’s surprising history, it’s hard to predict when human-level AI might come within reach. When it does, it’ll be important to have a leading research institution which can prioritize a good outcome for all over its own self-interest. We’re hoping to grow OpenAI into such an institution.³¹

OpenAI’s founding ethos was thus to lead the development of AI technology in order to ensure that potentially dangerous AI breakthroughs remained in conscientious hands. It was established as a non-profit in order “to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return.”³² In short, OpenAI was established ostensibly to constrain the out-of-control feedback

27. *Id.*

28. It’s unclear how risk management offends Andreessen’s worldview, as it would seem to be the stock-in-trade of free-market capitalism, under which investments are allocated based on expected returns. Cf. Will Kenton, *What is Risk Management in Finance, and Why is It Important?*, INVESTOPEDIA (July 13, 2024), <https://www.investopedia.com/terms/r/riskmanagement.asp> [<https://perma.cc/YJP3-WBVG>] (“Risk management essentially occurs when an investor . . . analyzes and attempts to quantify the potential for losses in an investment, . . . and then takes the appropriate action (or inaction) to meet their objectives and risk tolerance.”).

29. Andreessen, *supra* note 16.

30. Andreessen also asserts in his manifesto that “Techno-Optimism is . . . *not* a political philosophy,” and techno-optimists are neither “necessarily left wing” nor “necessarily right wing.” *Id.* I’ll leave it to the reader to evaluate the plausibility of that assertion.

31. Greg Brockman & Ilya Sutskever, *Introducing OpenAI*, OPENAI BLOG (Dec. 11, 2015), <https://openai.com/blog/introducing-openai> [<https://perma.cc/L4H4-RR6X>].

32. *Id.*

loop that accelerated AI might precipitate, not to trigger it. Like *e/acc*, then, OpenAI's founding mandate venerates technological acceleration—but, like *e/acc*, it does so in service of stability, not instability. It is, therefore, not accelerationism.

After a recent shake-up in internal governance, OpenAI has aligned with selfish ersatz accelerationism and shunned the altruistic ersatz accelerationism it initially embraced.³³ As a result, some associated with altruistic ersatz accelerationism are now being labeled “decelerationists.”³⁴ Decelerationism is an ideology unto itself, and like accelerationism, it's a big tent. Some decelerationists are radical leftist Luddites who view unconstrained technological development as a tool for entrenching capitalist hegemony.³⁵ Other self-identified decelerationists might be Silicon Valley executives — like Emmett Shear, who served as OpenAI's interim CEO for 72 hours during the internal reshuffling that solidified the company's capitalist turn, and who identifies as “basically *e/acc* on literally everything except the attempt to build a human level [] AI”³⁶—who embrace the capitalist social order but view superintelligent AI as a unique threat that warrants deceleration.

B. Copyright Accelerationism

But what does *copyright* accelerationism look like, and why am I presenting it for your consideration? I'll begin with a premise, and then a definition. Premise: rather than encouraging meaningful creativity and expressive engagement, modern copyright law impedes these desiderata, and we'd be better served by a dramatically different regime. Definition: copyright accelerationism demands that we insist on the rigid application of contemporary copyright doctrine and thereby heighten its contradictions. Doing so will achieve two things. First, it will make it undeniable that the copyright system is unworkable in the present day. Second, *it will actually make the copyright*

33. Nitish Pahwa, *Sam's Town*, SLATE (Nov. 22, 2023, 1:59 PM), <https://slate.com/technology/2023/11/sam-altman-rehired-openai-effective-altruism.html> [https://perma.cc/UYQ6-LP4D] (“OpenAI's current board has been purged of effective-altruist influence.”).

34. See Saqib Shah, *What Is a Decel? The Term That Encapsulates the Dueling Ideologies in AI*, EVENING STANDARD (Nov. 21, 2023), <https://www.standard.co.uk/news/tech/what-is-a-decel-dueling-ideologies-in-ai-artificial-intelligence-b1121591.html> [https://perma.cc/C55Q-5QBY].

35. See, e.g., Gavin Mueller, *Decelerate Now*, LOGIC(S) MAG. (Dec. 20, 2020), <https://logicmag.io/commons/decelerate-now/> [https://perma.cc/TE7W-MSAY].

36. Emmett Shear (@eshear), X (Aug. 25, 2023, 3:06 PM), <https://twitter.com/eshear/status/1695165810325832040> [https://perma.cc/A9VC-MKJS]; Matthew Loh, *Emmett Shear, Who Ran Openai for 72 Hours Before Sam Altman's Return Was Confirmed, Says He's "Deeply Pleased" About the News*, YAHOO FIN. (Nov. 22, 2023), <https://finance.yahoo.com/news/emmett-shear-ran-openai-72-081856736.html> [https://perma.cc/L8HJ-TW5Q].

system unworkable and, in turn, catalyze systemic reforms that are long overdue.

Let me indict our copyright system with an example: as I was writing this paragraph, I picked up my phone, looked down, and took the photo shown in Figure 1. That I now own a copyright in this photograph—a copyright my heirs will hold for seven decades after I die—illustrates the infirmity of our copyright regime. It should *astonish* us that I or my successors in interest will be able to obtain injunctive relief and damages against those who reproduce this photograph without my authorization and without privilege, from now until (knock on wood) roughly the year 2150.



Figure 1: “Table with Pant Leg” (© Ben Sobel, 2023)

Table with Pant Leg (Sobel, 2023) isn’t just a depiction of a table, my leg, and the seat of a chair—it’s an illustration of modern copyright’s shortcomings. First, it shows that the fixation requirement does little to limit what can be propertized. Images that in the past would’ve required laborious sketching or a chemistry lab are now trivially easy to record. Second, it shows how feeble copyright’s originality requirement is. I thought very little about this image; it contains just a slight amount of anything I’d call my “expression.” Third, it illustrates how unconditional copyright withholds incalculable amounts of information from the public domain. All I had to do for my copyright to vest was tap my phone’s screen a single time. I didn’t have to fill out any forms or pay any administrative fees. I didn’t even have

to think about whether I wanted to own the photograph. It's unclear that I could do anything to place it truly in the public domain before the expiration of its term, although, with a bit of work, I could attach a license that approximates public-domain status or file a notice of abandonment with the Copyright Office for a \$125 fee.³⁷ In short, we are inundated with low-value information that someone owns for just short of two lifetimes; there's often no good way to figure out who owns it; and we risk financial consequences if we err in our determinations.³⁸ If a historian years from now wanted to reproduce "Table with Pant Leg" in an academic analysis of office furnishings in the early 21st century, or perhaps in an encyclopedia of my sartorial choices, a risk-averse publisher might refuse to reproduce the image or force the historian to sign an indemnity agreement.

So, to put my cards on the table, I tend to think that copyright law should reimplement formalities,³⁹ impose higher originality standards,⁴⁰ and dramatically shorten the term of protection.⁴¹ I tend to think that no sensible copyright regime would protect mindless dross like "Table with Pant Leg" in the manner ours does, and I doubt that the architects of our present system of copyright envisioned that it would do so. Finally, I tend to think that these reforms would foster a culture that better empowers citizens to lead fulfilling lives of creating and engaging with expression while still "secur[ing] a fair return for . . . creative labor" and thereby "stimulat[ing] artistic creativity for the general public good."⁴²

Why, then, am I insisting that the proper application of some of copyright's most regressive features presents an ineluctable impediment to generative AI? The first reason is that, as a doctrinal matter, I'm correct—but

37. See *Valancourt Books, LLC v. Garland*, 82 F.4th 1222, 1235-36 (D.C. Cir. 2023) (observing that "there is no indication that" "a known and costless [copyright] abandonment option . . . generally exists"); Matthew P. Gelfand, *A Perfect (Copyright) Union: Uniting Registration and License Designation*, 25 HARV. J.L. & TECH. 711, 714 n.18 (2012) ("The legal status of public domain declarations is unsettled."); CC0, CREATIVE COMMONS, <https://creativecommons.org/public-domain/cc0/> [<https://perma.cc/C82S-DQLP>]. I thank David Nimmer for referring me to the *Valancourt Books* case.

38. See David R. Hansen et al., *Solving the Orphan Works Problem for the United States*, 37 COLUM. J.L. & ARTS 1, 3-11 (2013). In some cases, of course, the cost of defending or settling a claim may be a more meaningful financial deterrent than the relief a copyright owner could actually recover in court.

39. See generally, e.g., Christopher Sprigman, *Reform(aliz)ng Copyright*, 57 STAN. L. REV. 485 (2004).

40. See generally, e.g., William W. Fisher III, *Recalibrating Originality*, 54 HOUS. L. REV. 437 (2016).

41. Justice Breyer's dissenting opinion in *Eldred v. Ashcroft* provides cogent reasons to doubt the utility of long copyright terms. See 537 U.S. 186, 248-55 (2003) (Breyer, J., dissenting).

42. *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 156 (1975).

I've argued that elsewhere, and I don't intend to relitigate it here.⁴³ The second reason, which is most salient for present purposes, is that faithfully and uncompromisingly applying copyright's most baleful doctrines to generative AI might accomplish some radical good in the world.

Copyright accelerationism has features in common with true accelerationism and ersatz accelerationism. Copyright accelerationism resembles Marxist "accelerate-the-bad-thing" accelerationism in that it seeks to heighten the deleterious contradictions in today's copyright regime and thereby catalyze change. Copyright accelerationism is genuine accelerationism in that it aspires to sow instability. But the instability it aspires to sow is quite modest: it aims only to destabilize the copyright regime.⁴⁴ Thus, like Andreessen's ersatz accelerationism, my vision of copyright accelerationism assumes stability in the dominant social order; it assumes that it will spark, at most, a controlled burn. Although copyright accelerationism envisions a significant reconstituting of copyright law, it assumes the persistence of American legalism, just as e/acc assumes the persistence of free-market capitalism.

I offer copyright accelerationism as an alternative to what I call "copyright incremental-minimalism," an approach that advocates for limiting copyright's scope by opposing perceived expansions of copyright, and which might also be called "copyright decelerationism." Today's incremental-minimalism is a product of the copyright wars of the 1990s, when "rightsholders" tended to be industry behemoths like record labels and film studios, and "users" tended to be ordinary people trying to consume or create media on networked devices.⁴⁵ According to the incremental-minimalists, their opponents are the "copyright maximalists," who seek to control and monetize every use of copyrighted works.⁴⁶ The incremental-minimalists' project, as they conceive of it, is to defend the traditionally narrow scope of copyright

43. See generally Benjamin L. W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COLUM. J.L. & ARTS 45 (2017); Benjamin L. W. Sobel, *Elements of Style: Copyright, Similarity, and Generative AI*, 38 HARV. J.L. & TECH. 49 (2025).

44. Indeed, copyright accelerationism aspires to provoke such a slight amount of instability that it might draw the same criticism that Williams and Srnicek make of neoliberal capitalist faux-accelerationism: that it is not imaginative enough to be true accelerationism, and is instead "only the increasing speed of a local horizon, a simple brain-dead onrush rather than an acceleration which is also navigational, an experimental process of discovery within a universal space of possibility." Williams & Srnicek, *supra* note 7, § 2(2).

45. For a discussion of the copyright wars, see generally Jessica D. Litman, *The Copyright Wars*, in DIGITAL COPYRIGHT 151 (2006), <https://repository.law.umich.edu/books/1> [<https://perma.cc/PSH2-GKUA>]. For a discussion of how the AI business modifies and complicates the narrative about copyright that present-day incremental-minimalists have inherited from the copyright wars, see Sobel, *Artificial Intelligence's Fair Use Crisis*, *supra* note 43, at 82–89.

46. See, e.g., Pamela Samuelson, *The Copyright Grab*, WIRED (Jan. 1, 1996, 12:88 PM), <https://www.wired.com/1996/01/white-paper/> [<https://perma.cc/PSH2-GKUA>].

from copyright maximalists' efforts to expand it in furtherance of digital-age business interests—in other words, to decelerate what they perceive as an expansive power grab.⁴⁷

Like copyright accelerationism, incremental-minimalism strives for a copyright regime that is more limited in scope than today's.⁴⁸ And like copyright accelerationists, copyright incremental-minimalists are reformers rather than revolutionaries, although copyright accelerationism tolerates a bit more upheaval on its route to reform.⁴⁹ For years, the incremental-minimalists have fought a noble fight. Their hard-won achievements include exemptions from the Digital Millennium Copyright Act's anticircumvention provisions and a judicial holding that copyright holders must consider fair use before submitting a takedown notice.⁵⁰ However, the incremental-minimalists have responded to generative AI by losing the plot.

A good distillation of the incremental-minimalist response to AI is a recent social media post by Luis Villa, a lawyer, computer programmer, and board member of Creative Commons.⁵¹ On November 6, 2023, Villa wrote,

Reminder: using other people's work without consent or permission is something we as humans do all the time, and can be ethical in many contexts. Please don't buy into the copyright maximalist position that every use must require permission.

And yes, this is about AI, and no, I am unsure whether AI scraping is ethical. But we cannot let distaste for AI chip away at the small space we've painstakingly carved out for fair, ethical, permissionless use.⁵²

With respect, I believe that Villa misapprehends the stakes. Villa is posting from a defensive crouch: he touts the "small space" for users' rights

47. See *id.*; Neil Weinstock Netanel, *Copyright and a Democratic Civil Society*, 106 YALE L.J. 283, 372 (1996) ("Minimalist critics, on the other hand, insist that the 'free use zone' of the hard copy world. . . must be maintained in cyberspace."); see generally, e.g., Jessica Litman, *Lawful Personal Use*, 85 TEX. L. REV. 1871 (2007).

48. See Xiyin Tang, *Copyright's Techno-Pessimist Creep*, 90 FORDHAM L. REV. 1151, 1154 (2021) (describing copyright minimalism as "roughly corresponding to broader users' rights like fair use or broad liability shields like safe harbor laws").

49. In differentiating between reformers and revolutionaries, I am once again cribbing from Marxist history. Compare, e.g., EDUARD BERNSTEIN, BERNSTEIN: THE PRECONDITIONS OF SOCIALISM 191-92 (Henry Tudor ed., 1993) (reform socialism), with ROSA LUXEMBURG, REFORM OR REVOLUTION (1900) (Marxist Archive 1999), <https://www.marxists.org/archive/luxemburg/1900/reform-revolution/intro.htm> [<https://perma.cc/3GZW-SU8W>].

50. *Lenz v. Universal Music Corp.*, 815 F.3d 1145, 1153 (9th Cir. 2016); *DMCA Rulemaking*, ELEC. FRONTIER FOUND., <https://www.eff.org/issues/dmca-rulemaking> [perma.cc/PK6D-5LN7]. I cite the accomplishments of the Electronic Frontier Foundation (EFF) here because the EFF is emblematic of old-guard copyright minimalism; it represented Lenz and petitioned for numerous DMCA exceptions, among many other things. See *Lenz*, 815 F.3d at 1147.

51. Luis Villa, *About Me*, LUIS VILLA: OPEN TECH & POL'Y, <https://lu.is/> [<https://perma.cc/4RB6-2RL5>].

52. Luis Villa (@luis_in_brief@social.coop), MASTODON (Nov. 6, 2023, 10:16 AM), https://social.coop/@luis_in_brief/111364496878745261 [<https://perma.cc/5YT5-F7NF>].

that the minimalists have built in the hostile wilderness of copyright, and he suggests that subjecting AI-related practices to copyright liability will threaten that space.

I'd frame the problem differently. Instead of the little oasis Villa identifies, I'd focus on the inhospitable expanse that surrounds it. I'd ask the incremental-minimalists: sunk costs aside, how much do you really have to lose if a policy response to AI does, in fact, "chip away at" your "small space"? The little oasis has been drying up since long before machine learning took off. Jessica Litman, for example, has cataloged "increasingly forceful encroachment" on personal use of copyrighted media, an activity that she argues "copyright law rarely concerned itself with" in the mid-20th century; Litman argued in 2007 that, "over the past 30 or even 50 years," "the size of rights granted by the copyright law" has "expanded, extraordinarily."⁵³ The available evidence suggests that the oasis will only continue to shrink. For instance, just last term, the Supreme Court narrowed the availability of the fair use defense (and, of course, it did so in a case that concerned copying by a human artist).⁵⁴

In sum, copyright's incremental-minimalist decelerationists fear that treating copyright as a barrier to generative AI will jeopardize users' rights across the board. The incremental-minimalists' paranoia is in some measure justified—their oasis has really been drying up—but this particular concern is misplaced. It seems implausible that hostility to AI will erode the limited rights of users that exist today. After all, my thesis—and that of many critiques of AI-friendly copyright policy—is that AI should be subject to the same copyright strictures as human readers, not that all users should enjoy fewer rights. In other words, an ebbing tide for AI would not lower all boats, as the incremental-minimalists seem to fear.

More importantly, a rising tide for AI would not raise all boats! If an influential court of appeals holds that training commercial, generative AI on copyrighted works is fair use, nobody is going to cite that decision as overruling pre-generative-AI caselaw denying fair use for artistic pastiche,⁵⁵

53. Litman, *supra* note 45, at 1873–93; Jessica Litman, *Billowing White Goo*, 31 COLUM. J.L. & ARTS 587, 587 (2007).

54. *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508 (2023).

55. *See, e.g., Dr. Seuss Enters., L.P. v. ComicMix LLC*, 983 F.3d 443, 448–51 (9th Cir. 2020); *Dr. Seuss Enters., L.P. v. Penguin Books USA, Inc.*, 109 F.3d 1394 (9th Cir. 1997).

reproductions of personal letters in a biography,⁵⁶ fan fiction,⁵⁷ and so on. In context, a win for AI would be a win for AI—nothing more and nothing less. The incremental-minimalists seem to think that carrying the AI industry’s water will help them achieve their dream of a permissionless “remix culture.”⁵⁸ That’s partially true. But the most likely future is one in which this permissionless remix culture is an oasis that excludes human beings: it will be reserved for the generative AI enterprise alone. Unless!

II. A COPYRIGHT ACCELERATIONIST’S GUIDE TO GENERATIVE AI

We are hurtling towards a future in which all of copyright’s anti-reader provisions remain in place for traditional human audiences and creators, while the generative AI enterprise gets *carte blanche* to consume copyrighted works without authorization. This, in my view, is not a good world. Copyrighted works are “training data” for humans just as much as for artificial intelligence. A copyright regime that makes these works available *gratis* to train AI but inhibits their circulation among human readers is a regime that misconceives “progress” by defining it not as human engagement with expressive works but as a narrow sort of technological development.⁵⁹ It is also a legal framework that compromises rule-of-law values: one should be reflexively suspicious of a legal regime that limits human engagement with expressive works (to the primary benefit of a few rightsholders)⁶⁰ and simultaneously gives AI unfettered access to expressive works (to the primary benefit of a few powerful technology companies).⁶¹

56. See, e.g., *Salinger v. Random House, Inc.*, 811 F.2d 90, 96 (2d Cir. 1987) (granting preliminary injunction); *Folsom v. Marsh*, 9 F. Cas. 342, 349 (C.C.D. Mass. 1841) (Story, J.). See also *Harper & Row Publrs, Inc. v. Nation Enters.*, 471 U.S. 539, 569 (1985) (rejecting the fair use defense for the unauthorized reproduction of former president Gerald Ford’s memoirs in a political periodical).

57. See, e.g., *Salinger v. Colting*, 607 F.3d 68, 83 (2d Cir. 2010) (considering preliminary injunction). Follow-on works that subvert the original work have strong claims to fair use, however. Compare, e.g., *Suntrust Bank v. Houghton Mifflin Co.*, 268 F.3d 1257, 1270-71 (11th Cir. 2001) (considering preliminary injunction), with *Warner Bros. Entm’t Inc. v. RDR Books*, 575 F. Supp. 2d 513, 551 (S.D.N.Y. 2008).

58. See generally LAWRENCE LESSIG, *REMIX: MAKING ART AND COMMERCE THRIVE IN THE HYBRID ECONOMY* (2008).

59. Cf. Barton Beebe, *Bleistein, the Problem of Aesthetic Progress, and the Making of American Copyright Law*, 117 COLUM. L. REV. 319, 346-47 (2017) (describing a “pragmatist” theory of aesthetic progress that measures progress “by the extent of popular, democratic participation in aesthetic practice”).

60. See, e.g., *Eldred v. Ashcroft*, 537 U.S. 186, 248-49 (2003) (Breyer, J., dissenting) (observing that the Copyright Term Extension Act of 1998 represents a “transfer of several billion extra royalty dollars to holders of” a relatively small fraction of old copyrights).

61. Cf. generally David Gray Widder et al., *Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI 2* (Aug. 16, 2023) (available at SSRN) <https://papers.ssrn.com/abstract=4543807> [<https://perma.cc/LZ76-435W>] (arguing that “open” AI may redound to the benefit of a few powerful corporate incumbents). But see Andreessen Horowitz (a16z), *Comment on Artificial Intelligence and Copyright 8* (Oct. 30, 2023) <https://www.regulations.gov/comment/COLC-2023-0006-9057>

This section focuses on two flashpoint issues in copyright and generative AI and describes how a copyright accelerationist would approach them. First, I consider the availability of the fair use defense for training generative AI on unauthorized copies of copyrighted works. Second, I examine AI firms' attempt to bend copyright's exclusive rights into a *de facto* permission-by-default model that presumes authorization for AI-related uses of copyrighted works.

A. Fair Use and Learning

1. Fair Use Does Not Privilege (Human) Learning

One might argue—as many have⁶²—that training AI on unauthorized copies of copyrighted works so that it can learn to produce similar expression is fair use because the AI is only “learning” uncopyrightable facts, abstract patterns, and the like. There are excellent reasons to doubt a key premise of this argument, which is the seemingly categorical assertion that AI traffics only in facts and not in expression.⁶³ But even if we accept that premise, this fair use argument is one that founders when human copyists make it. In *American Geophysical Union v. Texaco*, the Second Circuit held that it was not fair use for a scientist in Texaco's research division to photocopy scientific journal articles in order to “facilitate . . . current or future professional research.”⁶⁴ The court noted explicitly that the copied journal articles were “essentially factual in nature,” and, moreover, that the evidence indicated that a representative Texaco researcher “was interested exclusively in the facts, ideas, concepts, or principles contained within the articles.”⁶⁵ The court elaborated: “[t]hough scientists surely employ creativity and originality to develop ideas and obtain facts and thereafter to convey the ideas and

[<https://perma.cc/W99H-DDXK>] (arguing that treating AI training as copyright infringement would reduce competition and entrench the power of dominant technology firms).

62. As illustrative examples, see Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 307-08 (2023); Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 772-74 (2020); Cory Doctorow, *Copyright Won't Solve Creators' Generative AI Problem*, MEDIUM (Feb. 9, 2023), <https://doctorow.medium.com/copyright-wont-solve-creators-generative-ai-problem-92d7adbcc6e6> [<https://perma.cc/EHQ5-JNVF>] (“[M]achine learning systems ingest a lot of works, analyze them, find statistical correlations between them, and then use those to make new works. It's a math-heavy version of what every creator does We should not create a new right to decide who is allowed to think hard about your creative works and learn from them.”).

63. See Sobel, *Artificial Intelligence's Fair Use Crisis*, *supra* note 43, at 68-72; see generally Sobel, *Elements of Style*, *supra* note 43.

64. *Am. Geophysical Union v. Texaco Inc.*, 60 F.3d 913, 915, 931 (2d Cir. 1994). The litigation in *Texaco* concerned the practices of hundreds of researchers at Texaco, but the parties agreed to choose one at random as a representative example. *Id.* at 915. I therefore discuss Texaco's representative researcher in the singular throughout.

65. *Id.* at 925, 925 n.11.

facts in scholarly articles, it is primarily the ideas and facts themselves that are of value to other scientists in their research.”⁶⁶ Nevertheless, the court rejected the fair use defense, primarily because the purpose of the copying was to create more copies for which Texaco could have purchased a photocopying license from the rightsholders.⁶⁷ Thus, *Texaco* expressly rejected the argument that a purpose of extracting unprotectable material categorically privileges unauthorized copying as fair use.

Texaco also makes clear that unauthorized copying in service of a generalized project of “learning” or “research” is not *per se* fair use. On this point, the Second Circuit was unmistakable:

Texaco cannot gain fair use insulation for [its researcher’s] archival photocopying of articles (or books) simply because such copying is done by a company doing research. It would be equally extravagant for a newspaper to contend that because its business is “news reporting” it may line the shelves of its reporters with photocopies of books on journalism or that schools engaged in “teaching” may supply its [sic] faculty members with personal photocopies of books on educational techniques or substantive fields. Whatever benefit copying and reading such books might contribute to the process of “teaching” would not for that reason satisfy the test of a “teaching” purpose.⁶⁸

Justice Stevens’s majority opinion in *Sony v. Universal* made a similar observation in a footnote: “the notion of social ‘productivity’ cannot be a complete answer to [the fair use] 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.”⁶⁹ Even in the 1841 case that birthed the American fair use doctrine—which held that the defendants had infringed copyright by reproducing some of George Washington’s letters in a biography of the first president—Justice Story concluded by remarking that he was “not without some regret, that [the disposition of the case] may interfere, in some measure, with the very meritorious labors of the defendants, in their great undertaking of a series of works adapted to school libraries. But a judge is entitled . . . only to know and to act upon his duty.”⁷⁰

The reasoning of *Texaco* is uncontroversial as far as human learners are concerned. Could you *imagine* a defendant having the temerity to argue that she only torrented a Beatles album because she wanted to learn to write songs in the style of Lennon and McCartney? She’d be laughed out of court,

66. *Id.*

67. *Id.* at 924-25, 930-31.

68. *Id.* at 924.

69. *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 455 n.40 (1984).

70. *Folsom v. Marsh*, 9 F. Cas. 342, 349 (C.C.D. Mass. 1841) (Story, J.).

irrespective of whether she ultimately wrote songs that were substantially similar to Beatles tunes.⁷¹ The law is clear: the fact that a use is edifying does not make it fair. Proposals that training AI should be fair use, either because it implicates only factual material or because “learning is good,” extend a largesse to AI that copyright denies to humans. Why should AI get to learn on more favorable terms than human beings do?

By describing fair use’s hostility to learning, I don’t mean to endorse it. It’s tragic that copyright fetters learning! On this point, I’m more or less in agreement with the outspoken copyright skeptic Cory Doctorow, who has commented that “[t]he universal access to all human knowledge” that digital media technologies facilitate “is the realization of one of the most important dreams of humanity, and . . . complaining about it is morally indefensible.”⁷² Conventional copyright minimalists and I agree that the present copyright regime does not strike a good balance between incentives and access. We disagree on how to respond.

2. Incremental-Minimalist Approaches to Fair Use and AI

Mainline copyright incremental-minimalists generally argue that training expressive AI is fair use because “[f]air use protects . . . forms of analysis that create new knowledge about works or bodies of works,” and “it’s no more illegal for [an AI] model to learn a style from existing work than for human artists to do the same.”⁷³ A weakness in this line of reasoning is that copyright often makes it difficult and costly for humans to obtain lawful access to existing works, even if their purpose is merely to learn from them. At the same time, the AI industry has allegedly trained its models on vast “shadow libraries” of copyrighted works.⁷⁴ Some incremental-minimalist analyses of fair use and AI, however, do recognize the ways that copyright impedes human learning. In *Fair Learning*, Mark Lemley and Bryan Casey propound a theory of fair use that would, as a general matter, permit training

71. I use this example in a recent op-ed. See Ben Sobel, *Don’t Give AI Free Access to Work Denied to Humans, Argues a Legal Scholar*, THE ECONOMIST (Feb. 16, 2024), <https://www.economist.com/by-invitation/2024/02/16/dont-give-ai-free-access-to-work-denied-to-humans-argues-a-legal-scholar> [<https://perma.cc/V7RF-8WTJ>].

72. Diane Coutu, *Predicting the Present*, HARV. BUS. REV. (2009), <https://hbr.org/2009/07/predicting-the-present> [<https://perma.cc/EGY3-HF9Y>].

73. Kit Walsh, *How We Think About Copyright and AI Art*, ELEC. FRONTIER FOUND. (Apr. 3, 2023), <https://www EFF.org/deeplinks/2023/04/how-we-think-about-copyright-and-ai-art-0> [<https://perma.cc/6BVT-JJ28>]. See also, e.g., Creative Commons, Comments on United States Copyright Office Notice of Inquiry on Artificial Intelligence & Copyright (Oct. 30, 2023), at 3-4 https://downloads.regulations.gov/COLC-2023-0006-8735/attachment_1.docx [<https://perma.cc/D6JE-SELE>].

74. See Complaint at ¶ 35, *Silverman v. OpenAI, Inc.*, (No. 3:23-cv-03416), 2023 WL 4448006, at 5 (N.D. Cal. July 7, 2023); Complaint at ¶¶ 23-30 *Kadrey v. Meta*, (No. 3:23-cv-03417), 2023 WL 4463445, at 4 (N.D. Cal. July 7, 2023),

AI on copyrighted works without authorization.⁷⁵ Their article proposes that, “[i]f the purpose of the AI’s use is not to obtain or incorporate the copyrightable elements of a work but to access, learn, and use the unprotectable parts of the work, that use should be presumptively”⁷⁶ favored under the first factor and, save perhaps for some cases of expressive substitution, be treated as fair use overall.⁷⁷ To Lemley and Casey’s great credit, their analysis doesn’t start and end with AI. Instead, they argue that their fair learning proposal “may also help courts do a better job of identifying and protecting fair learning by humans too.”⁷⁸ They even suggest that their approach to fair use would have flipped the outcome for the human learners in *Texaco*.⁷⁹ Lemley and Casey’s fair learning proposal is thus an illustrative incremental-minimalist foil to copyright accelerationism. It proposes to tweak the well-established doctrine of fair use in ways that will save the generative AI industry and, they argue, enhance copyright’s esteem for human expressive engagement as well.

Another minimalist analysis comes from Cory Doctorow, whose views on copyright reform are more radical and less incremental than the proposal in *Fair Learning*.⁸⁰ In an essay titled “Copyright Won’t Solve Creators’ Generative AI Problem,” Doctorow comes painfully close to making the argument for copyright accelerationism. But he can’t shake his copyleft priors, and he ends up advocating the familiar, incremental-minimalist copyright-decelerationist response to generative AI.⁸¹ Doctorow’s essay argues that treating unauthorized training of AI on copyrighted works as copyright infringement would not benefit authors and artists.⁸² This, he argues, is because creators’ compensation depends not on the scope of their copyright entitlements but instead on “the structure of the creative market.”⁸³ Because monopolies dominate publishing, music, and ad-tech, “giving a creator more

75. Lemley & Casey, *supra* note 62, at 748.

76. *Id.* at 776.

77. *Id.* at 776-79.

78. *Id.* at 779.

79. *Id.* at 779-80.

80. See, e.g., Benjamin Aleksandr Franz, *Cory Doctorow (1971–)*, in *FIFTY KEY FIGURES IN CYBERPUNK CULTURE* (Anna McFarlane et al. eds., 2022) (“Doctorow has spoken in opposition to intellectual property, which he sees as a destructive and pervasive trend.”).

81. Doctorow’s essay recites the typical arguments that training AI on copyrighted works without authorization is not copyright infringement: “machine learning systems ingest a lot of works, analyze them, find statistical correlations between them, and then use those to make new works. It’s a math-heavy version of what every creator does . . . We should not create a new right to decide who is allowed to think hard about your creative works and learn from them . . .” Doctorow, *supra* note 62. For a debunking of these arguments, see Part II.A.1, *supra*.

82. Doctorow, *supra* note 62.

83. *Id.* (emphasis omitted).

copyright is like giving a bullied schoolkid extra lunch money”: the “monopolists who control the creative industries” will just strip that right away, too, via extractive, take-it-or-leave-it contracts.⁸⁴

Up until this point, Doctorow is making a standard-enough antimonopoly argument. Suddenly, however, his head peeks out of the cave. He writes, “creative workers can’t afford to let corporations create [an exclusive right over ML training] — and not just because they will use it against us. These corporations also have a track record of creating new exclusive rights that bite *them* in the ass.”⁸⁵ As an example, Doctorow cites the broadening scope of the reproduction right, which he argues has exposed major recording artists to spurious infringement suits and threatened major labels’ own bottom lines. Of these major labels, Doctorow writes,

They are completely wedded to the idea that every part of music should be converted to property, so that they can expropriate it from creators and add it to their own bulging portfolios. Like a monkey trapped because it has reached through a hole into a hollow log to grab a banana that won’t fit back through the hole, the labels can’t bring themselves to let go.

That’s the curse of the monkey’s paw: the entertainment giants argued for everything to be converted to a tradeable exclusive right—and now the industry is being threatened by trolls and ML creeps who are bent on acquiring their own vast troves of pseudo-property.⁸⁶

This is a stunning realization: unchecked copyright maximalism undermines the very structures that copyright maximalism builds! Indeed, this insight is the very premise of copyright accelerationism.

But alas, this glimpse of sunlight blinds Doctorow, and he retreats. He ends his essay back in the cave, implying that we will win more equitable treatment for creators by *rejecting* the impulse to impede AI with copyright.⁸⁷ “Turning every part of the creative process into ‘IP’ hasn’t made creators better off. All that’s [sic] it’s accomplished is to make it harder to create without taking terms from a giant corporation, whose terms inevitably include forcing you to trade all your IP away to them.”⁸⁸

To recapitulate, then, Doctorow makes the following points: (1) artists’ inequitable compensation is caused not by inadequate copyright protection for their work, but by the structure of the content industry; (2) the content industry treats artists unfairly by using its outsize bargaining power to extract

84. *Id.*

85. *Id.*

86. *Id.* (citation omitted).

87. *Id.*

88. *Id.*

artists' IP rights in inequitable deals; (3) the content industry is so obsessed with IP that it sometimes creates rights that undermine its own business model; and (4) artists should not invoke copyright to exclude unauthorized uses of their work to train AI, because using IP rights in this manner will just reinforce the content industry's inequitable business model. As I hope is obvious, Doctorow's jump from (3) to (4) is unsupported. He explicitly contemplates that applying copyright to AI training could undermine major companies' business models—which he insists are what cause artists' poor compensation—and then apparently forgets about the possibility.⁸⁹ Ironically, the shortcomings of Doctorow's argument are best explained by reflecting a mirror image of his own criticism of the content industry at him: Doctorow is “completely wedded to the idea that [no] part of [expression] should be converted to property . . . Like a monkey trapped because it has reached through a hole into a hollow log to grab a banana that won't fit back through the hole, [Doctorow] can't bring [himself] to let go” of his copyleft priors—even when doing so would further the vision of culture that he prizes.⁹⁰

3. The Accelerationist Response

By contrast, the accelerationist response to artificial intelligence's fair use crisis is to acknowledge, embrace, and exploit fair use's historical hostility to learning. This is a response addressed not only to copyright's incremental-minimalists, but also to the e/acc industrialists. Consider a representative e/acc perspective on fair use: in recent comments to the United States Copyright Office, Marc Andreessen's venture capital firm, Andreessen Horowitz, asserted that “[i]mposing infringement liability for the use of copyrighted works in AI model training . . . would upset at least a decade's worth of investment-backed expectations that were premised on” an understanding that such uses were fair use—a figure that Andreessen Horowitz pegged at “billions and billions of dollars.”⁹¹ It continued: “imposing the cost of actual or potential copyright liability on the creators of AI models will either kill or significantly hamper their development.”⁹²

The copyright accelerationist's reply to this e/acc angst is, “[s]ounds like you might've made a bad bet. What would you be willing to do to

89. *Id.* (arguing that “creative workers can't afford to let corporations create [an exclusive right over ML training]” because “corporations . . . have a track record of creating new exclusive rights that bite *them* in the ass”).

90. *See id.*

91. Andreessen Horowitz, *supra* note 61, at 6–7.

92. *Id.* at 8.

salvage it?”⁹³ It isn’t anything new for copyright to pose an existential threat to innovation; indeed, copyright “kill[ing] or significantly hamper[ing]” progress is business as usual. In the 1990s, case law and business practices adverse to hip-hop artists erected practically insurmountable obstacles for sample-heavy music production.⁹⁴ An entire musical genre was forced to change course.⁹⁵ The rapper Chuck D of Public Enemy reflected that once copyright enforcement began targeting sample-based hip-hop, “we had to change our whole style.”⁹⁶ Unlike the generative AI industry, Public Enemy didn’t have “billions and billions of dollars,” and its representatives weren’t invited to congressional hearings about the legality of sampling.⁹⁷ But could you imagine the legal changes that Public Enemy and their colleagues might have achieved if they had those resources?

A copyright accelerationist would call Andreessen Horowitz’s bet and raise the stakes. She would refuse to cede an inch on AI-specific exemptions from copyright liability and demand instead that the AI industrialists yoke their interests to those of the broader community of human readers and authors—or else risk fulfilling Andreessen Horowitz’s warning of “copyright liability . . . kill[ing]” AI.⁹⁸ In other words, in order to shield themselves from our runaway copyright regime, the AI industrialists would have to save *everyone* from our runaway copyright regime.

There are any number of possibilities for that sort of reform, many of which are compatible with one another. Enact legislation that restores the copyright terms of the Copyright Act of 1790, which provided for an initial 14-year term of protection, followed by a 14-year renewal period.⁹⁹ Jack up

93. And my response would of course be that I warned about this in 2017. See generally Sobel, *Artificial Intelligence’s Fair Use Crisis*, *supra* note 43.

94. See KEMBREW MCLEOD, FREEDOM OF EXPRESSION®: OVERZEALOUS COPYRIGHT BOZOS AND OTHER ENEMIES OF CREATIVITY 87 (1st ed. 2005) (“The way copyright law was enforced helped turn on its head hip-hop’s original creative method: reworking prerecorded sounds.”); Kembrew McLeod, *How Copyright Law Changed Hip Hop*, ALTERNET (June 1, 2004), https://www.alternet.org/2004/06/how_copyright_law_changed_hip_hop [<https://perma.cc/MN6X-HBB5>] (observing that the law of sampling shifted hip-hop songwriting towards “one primary sample, instead of a collage of different sounds”).

95. See MCLEOD, *supra* note 94, at 87–88 (discussing rappers’ pivot to recreating sounds with live musicians instead of sampling); McLeod, *How Copyright Law Changed Hip Hop*, *supra* note 94 (quoting Hank Shocklee of Public Enemy: “So those things [about live instruments] change your mood, the feeling you can get off of a record” versus sampling).

96. *Id.*

97. Although Chuck D did testify to Congress in support of peer-to-peer filesharing technology in 2003. *Chuck D Testifies*, C-SPAN (Sept. 30, 2003), <https://www.c-span.org/clip/senate-committee/user-clip-chuck-d-testifies/4738659>.

98. Andreessen Horowitz, *supra* note 61, at 8.

99. An Act for the Encouragement of Learning by Securing the Copies of Maps, Charts, and Books to the Authors and Proprietors of Such Copies, During the Times Therein Mentioned, Ch. 15, 1 Cong., Pub. L. 1-15, 1 Stat 124 (1790). A similar initiative, the proposed Copyright Clause Restoration Act of

originality requirements.¹⁰⁰ Reimplement formalities.¹⁰¹ Modify the fair use statute to provide, in so many words, that “any use that promotes learning is fair use” and provide government subsidies to creators. Establish a “governmentally administered reward system” that provides universal access to creative works and also compensates creators.¹⁰² Abolish copyright entirely and enact the no-strings-attached regulation that Doctorow has proposed: “If you call yourself an artist, the government will pay you [~\$74,000] a year until you stop calling yourself an artist.”¹⁰³ These are all profoundly different options, and some are far more appealing than others. Each of them, however, plausibly improves upon our current system. Tens or hundreds of millions of people have copyright entitlements that, if accorded the solemnity and breadth we’ve historically afforded to such rights, could be wielded to destroy the generative AI enterprise.¹⁰⁴ By and large, these small-time

2023, would reduce copyright terms to 28 years, with a 28-year renewal period, and would apply retroactively to some copyrights. See H. R. 576, 118th Cong. (1st Sess. 2023); S. 4178, 117th Cong., 2d Sess. (2022). This particular copyright reform initiative appears to have originated as an effort by Senator Josh Hawley of Missouri to punish the Walt Disney Company for taking political positions that Senator Hawley disfavours. See Sarah Jeong, *Josh Hawley Wants to Punish Disney by Taking Copyright Law Back to 1909 and That Sucks*, VERGE (May 10, 2022, 5:12 PM), <https://www.theverge.com/2022/5/10/23066030/hawley-copyright-disney-mickey-mouse> [https://perma.cc/T3ZG-VNTA]. There are, of course, serious questions as to the constitutionality of such a reform, but at least one commentator has concluded that “shortening the copyright term is probably constitutional” under the Takings Clause if term-shortening legislation contains “a minimum period that every existing copyright will last after the reform legislation is implemented.” Note, *Copyright Reform and the Takings Clause*, 128 HARV. L. REV. 973, 975, 990-91 (2015).

100. See generally Fisher, *supra* note 40.

101. See generally Sprigman, *supra* note 39.

102. See WILLIAM W. FISHER III, PROMISES TO KEEP: TECHNOLOGY, LAW, AND THE FUTURE OF ENTERTAINMENT 202 (2004). Andreessen in 2014 expressed some openness towards universal basic income, but more recently has deployed the infelicitously worded metaphor that “a Universal Basic Income would turn people into zoo animals to be farmed by the state.” Compare Kevin Roose, *Marc Andreessen on Why Optimism Is Always the Safest Bet*, INTELLIGENCER (Oct. 19, 2014, 9:03 PM), <http://nymag.com/daily/intelligencer/2014/10/marc-andreessen-in-conversation.html> [https://perma.cc/GHK2-B7GL] (“It is a very interesting idea.”) with Andreessen, *supra* note 16. A less dramatic option would be to follow Oren Bracha’s proposal to enact a global opt-out rule through the fair use doctrine and/or a statutory safe harbor. See Oren Bracha, *Standing Copyright Law on Its Head? The Googleization of Everything and the Many Faces of Property*, 85 TEX. L. REV. 1799, 1855-68 (2006); see also *infra*, Part II.B (discussing opt-out versus opt-in).

103. Cory Doctorow, *What Do We Want Copyright to Do?*, THE GUARDIAN (Nov. 23, 2010), <https://www.theguardian.com/technology/2010/nov/23/copyright-digital-rights-cory-doctorow> [https://perma.cc/2WSY-4XRG]. Doctorow’s figure was “£40,000,” which I adjusted for inflation using a calculator on the Bank of England’s website, converted into dollars using the exchange rate Google provided on November 22, 2023, and rounded to the nearest thousand. Inflation Calculator, <https://www.bankofengland.co.uk/monetary-policy/inflation/inflation-calculator> (last visited Nov. 22, 2023); GBP/USD Currency Exchange Rate & News, GOOGLE FIN., <https://www.google.com/finance/quote/GBP-USD> (last visited Nov. 22, 2023).

104. Cf. Katherine Lee et al., *Talkin’ ‘Bout AI Generation: Copyright and the Generative-AI Supply Chain*, 71 J. COPYRIGHT SOC’Y U.S.A. (forthcoming 2025) (manuscript at 129) (discussing equitable remedy of destruction).

rightsholders have little to lose and something to gain from embracing copyright accelerationism.

4. Pick Your Pony

If you—like Lemley and Casey, like Doctorow, and like me—would like to see a copyright regime that better fosters human learning, you should choose between incremental-minimalist and accelerationist approaches based on how well you think a given approach will effectuate that goal. For both doctrinal and pragmatic reasons, accelerationism is a better bet, although I acknowledge that it's far from a sure shot.

First, as a doctrinal matter, the federal courts' common-law approach to fair use stacks the deck against the incremental-minimalists from the get-go. Fair use is effectively a common-law doctrine; the statute provides almost no guidance and, according to the legislative history, was "intended to restate the present judicial doctrine of fair use, not to change, narrow, or enlarge it in any way."¹⁰⁵ Because fair use is an accretion of hundreds of case-specific outcomes, any given decision about AI is unlikely to disturb established doctrine in other sectors. The reason for this is obvious: any decision involving AI can and will be distinguished from past decisions *on the ground that it involves AI*. Historically, fair use has differentiated between human and non-human readers.¹⁰⁶ As a practical matter, this trend has engendered what James Grimmelman calls "a two-tracked copyright law: one for human readers and one for robots. Uses involving human readers receive close and exacting scrutiny to make sure that no market belonging to the copyright owner is being preempted. Uses involving robotic readers are fast-tracked for fair use."¹⁰⁷ What this means for the future is that a "fair learning" victory for AI would not entail that fair learning would suddenly shield human learners, too. Much more likely is that fair learning's opponents would, in a subsequent case involving training data for humans, distinguish the AI-friendly precedent on the ground that it involves AI and recapitulate the durable human-robot distinction that Grimmelman has documented. So, I repeat: there are strong doctrinal reasons to doubt that a rising tide for AI would raise all boats.¹⁰⁸

105. H.R. REP. NO. 1476, at 66 (1976).

106. James Grimmelman has chronicled the law's tendency to presume that human readers engage with expression in a way copyright entitles rightsholders to control, and that robot readers don't. *See generally* James Grimmelman, *Copyright for Literate Robots*, 101 IOWA L. REV. 657 (2015).

107. *Id.* at 667.

108. *See supra* Part I.B.

The next knock on the incremental-minimalists is pragmatic. Imagine a watershed court of appeals decision upholds the fair use defense for training expressive AI, and the AI business chugs on apace. Now imagine the next big fair use case rolls around, this time involving reproductions of copy-righted materials for human consumption. The incremental-minimalists will surely do their best to advocate for an outcome friendly to human readers. But will the AI industry, now assured that its business model is lawful, be as steadfast an ally to the incremental-minimalists as the incremental-minimalists are to today's AI industry? I can only speculate, but I doubt it. If anything, if and when the big players in AI establish the footing in copyright law that they desire, they will advocate for strengthening at least some forms of intellectual property protection, especially trade secret.¹⁰⁹ I don't mean to suggest here that the incremental-minimalists defend AI training as fair use because they expect reciprocal favors down the line; I only mean to suggest that, as a matter of political economy, a copyright victory for AI is not likely to entail a broader victory for human learning of all sorts.

Of course, the progress of AI itself is, or can be, the progress of human learning and cultural engagement. As a general matter, it is aesthetically, politically, and economically salutary to ensure citizens' freedom to develop new expressive technologies and to use those technologies to express themselves. The creation or use of AI models is just as legitimate a vehicle for human self-development as reading books the old-fashioned way is. *But that's exactly my point.* Yes, the copyright accelerationist is willing to use the future of AI as a bargaining chip to force systemic reforms. That doesn't, however, mean that the copyright accelerationist opposes AI. The copyright accelerationist is agnostic towards AI; the future she seeks to realize is one that is *equally as accommodating* of human self-development through old-fashioned reading as it is of human self-development through the creation and use of expressive AI.

In other words, the realistic best-case scenario for the incremental-minimalist approach is a regime that denigrates old-fashioned expressive engagement and venerates the use and development of AI. This is precisely the future that the copyright accelerationists are willing to hold hostage in order

109. For example, OpenAI refused to disclose information about the architecture and training data of its GPT-4 model, citing "the competitive landscape." *GPT-4 Technical Report*, OPENAI, 2 (2023), <https://cdn.openai.com/papers/gpt-4.pdf> [<https://perma.cc/2CZ9-8J96>]. See also Isha Marathe, *In OpenAI Copyright Lawsuits, Discovery Complications Likely to Take Center Stage*, LEGALTECH NEWS (July 21, 2023, 12:57 PM), <https://www.law.com/legaltechnews/2023/07/21/in-openai-copyright-lawsuits-discovery-complications-likely-to-take-center-stage/> [<https://perma.cc/X9BZ-GLG3>] (stating that, in a pending copyright lawsuit, OpenAI "is likely to request the court grant it a protective order that would make [information about its models] only visible to attorneys in the case, by invoking trade secret protections").

to demand something better. By contrast, the future the copyright accelerationists can achieve is one that demolishes the same barriers to learning that the incremental-minimalist future does—but it demolishes those barriers across the board, not just for the AI enterprise. By and large, the only people who stand to lose by choosing copyright accelerationism are also the only people who stand to gain anything from siding with the incremental-minimalists: they are the people who care only about the development and use of AI and not about traditional modes of engagement with expression. Everyone else might as well try the accelerationist gambit.

B. *The AI Industry's Neo-Formalities Movement*

I now briefly turn to a second case study: the AI industry's campaign to treat authors as having opted in by default to AI training on their works. This campaign illustrates the industry's ambitions for *de facto* if not *de jure* copyright reform, and the debate it has elicited reveals how much the copyright minimalists' position has eroded in just a few years. Less than twenty years ago, when Google Books was nascent, copyright minimalists were championing legal reforms that would have required authors to opt out of the universal public display of their works. Now, the mainstream opt-in/opt-out debate doesn't contemplate human access to works: it concerns opt-out of AI training alone. This change in the terms of the discourse illustrates just how much ground the minimalists have lost.

At the center of AI's neo-formalities movement is Spawning.ai. Spawning is the work of the academic Mat Dryhurst and the musician Holly Herndon, in collaboration with Jordan Meyer and Patrick Hoepner, the founders of a studio that develops AI software for artists.¹¹⁰ Last May, Spawning raised \$3 million in venture capital to develop what Meyer calls "IP standards for the AI era."¹¹¹ Spawning provides tools that help artists opt out of AI training and that help AI enterprises avoid using opted-out works. Its *Have I Been Trained?* database lets artists search the data used for training AI models to learn whether their works have been used to train generative AI.¹¹² In late 2022, Stability AI partnered with Spawning to enable artists to

110. Chris Stokel-Walker, *This Couple is Launching an Organization to Protect Artists in the AI Era*, INPUT (Sept. 14, 2022), <https://www.inverse.com/input/culture/mat-dryhurst-holly-herndon-artists-ai-spawning-source-dall-e-midjourney> [https://perma.cc/5RVB-QY5M]; WOLF BEAR STUDIO, <https://web.archive.org/web/20221206010744/http://wolfbearstudio.com/> (last visited Nov. 3, 2022).

111. Kyle Wiggers, *Spawning Lays out Plans for Letting Creators Opt out of Generative AI Training*, TECHCRUNCH (May 3, 2023, 6:00 AM), <https://techcrunch.com/2023/05/03/spawning-lays-out-its-plans-for-letting-creators-opt-out-of-generative-ai-training/> [https://perma.cc/PA2T-K7SM].

112. *Have I Been Trained?*, <https://haveibeentrained.com/> [https://perma.cc/9D4C-MFW].

opt out of having their images included in StableDiffusion's training data.¹¹³ A report from March 2023 indicated that 80 million images have been opted out of training StableDiffusion's next iteration, StableDiffusion 3.¹¹⁴ In written testimony to the Senate Judiciary Committee, Stability AI's Head of Public Policy touted its opt-out system as a "[b]est practice[] in training."¹¹⁵ OpenAI has also instituted an opt-out process.¹¹⁶

Spawning's tools for AI enterprises include an API that "enables [its users] to validate the data [they] collect, ensuring that [they're] not using opted-out work that respects the original copyright owners' consent preferences."¹¹⁷ The API functions in tandem with ai.txt, a standard that Spawning introduced in May 2023. Ai.txt is

a file placed at the root of a website, which selectively restricts or permits access to the site's content and media—mirroring the widely adopted robots.txt standard With ai.txt, website owners can control whether or not their work is used to train new AI models and can continue to use robots.txt to manage permissions for popular search engines.¹¹⁸

Spawning's API and its ai.txt standard work together to enable content owners to opt out of AI training, at least as conducted by entities that use Spawning's API and respect preferences expressed in ai.txt files.

Notably, Spawning has presented itself as hostile to copyright law. In its telling, instead of "copyright," Spawning would prefer that artists have "tools."¹¹⁹ An earlier version of Spawning's "About" page explained,

Copyright is an outdated system that is a bad fit for the AI era.

113. Emad (@EMostaque), X (Dec. 14, 2022, 3:59 PM), <https://x.com/EMostaque/status/1603147709229170695> [https://perma.cc/PF85-GGEZ]. "Opt-in as well of course, about 50:50 both ways. Technically this is tags for LAION and coordinated around that. It's actually quite difficult due to size (eg [sic] what if your image is on a news site?) Exploring other mechanisms for attribution etc, welcome constructive input." *Id.*; Spawning (@spawning_), X, (Mar. 23, 2023, 4:43 PM) https://x.com/spawning_/status/1639020027989962752 [https://perma.cc/S247-H4GC].

114. Heather Tal Murphy, *A.I. Is Sucking the Entire Internet In. What If You Could Yank Some Back Out?*, SLATE (Mar. 27, 2023, 5:50 AM), <https://slate.com/technology/2023/03/how-holly-herndon-and-mathew-dryhurst-brokered-an-a-i-deal-with-stable-diffusion.html> [https://perma.cc/3VXE-DGNJ].

115. Artificial Intelligence and Intellectual Property – Part II: Copyright, Hearing Before the S. Comm. on Intellectual Property (July 12, 2023), (Written Statement of Ben Brooks, Head of Public Policy, Stability AI), at 9 [hereinafter *Brooks Statement*], available at <https://www.judiciary.senate.gov/download/2023-07-12-pm-testimony-brooks> [https://perma.cc/6WYN-AFES].

116. Kali Hays, *OpenAI Offers a Way for Creators to Opt out of AI Training Data. It's so Onerous That One Artist Called It 'Enraging.'* BUS. INSIDER, <https://www.businessinsider.com/openai-dalle-opt-out-process-artists-enraging-2023-9> [https://perma.cc/HV8G-T9WV].

117. *SpawningAPI*, SPAWNING AI, <https://api.spawning.ai/spawning-api> [https://perma.cc/XA6P-H79F].

118. Cullen Miller, *Ai.Txt: A New Way for Websites to Set Permissions for AI*, SPAWNING BLOG (May 30, 2023), <https://spawning.substack.com/p/aitxt-a-new-way-for-websites-to-set> [https://perma.cc/6WXN-X633].

119. *About - Spawning*, <https://spawning.ai/About> [https://perma.cc/3RQ4-QX8Y].

A new era offers us the opportunity to reconfigure how we treat IP! We believe that the best path forward is to offer individual artists tools to manage their style and likenesses, and determine their own comfort level with a changing technological landscape.

We are not focused on chasing down individuals for experimenting with the work of others. Our concern is less with artists having fun, rather with industrial scale usage of artist training data.¹²⁰

Elsewhere, Dryhurst has stated that Spawning's project is "not . . . to build tools for DMCA takedowns and copyright hell That's not what we're going for, and I don't even think that would work."¹²¹

Spawning, Stability AI, and OpenAI appear to be treating opt-out more like a public-relations device than a legal one.¹²² Stability and OpenAI maintain that their unauthorized use of copyrighted training data is protected by fair use.¹²³ If that legal assertion proves to be correct, then these companies have no obligation to permit opt-out: they can train its AI on whatever they want.¹²⁴ And if the companies' legal assertion is wrong, then permitting opt-out is immaterial; all the rightsholders whose silence the companies have treated as consent could pursue infringement claims. In other words, existing law is clear: an opt-in regime is required if training expressive AI is not fair use. If training expressive AI is fair use, there's no duty to honor opt-outs.

120. *Id.*

121. Stokel-Walker, *supra* note 110. It's worth noting that, notwithstanding Dryhurst's stated intention of avoiding "copyright hell," ai.txt seems likely to consign AI-related intellectual property disputes to "contract hell." Compliance with ai.txt is not compulsory. If a defendant disregarded instructions in an ai.txt file, a website proprietor might assert claims for breach of contract and common-law trespass to chattels, or perhaps for violation of the Computer Fraud and Abuse Act. *See* Benjamin L. W. Sobel, *A New Common Law of Web Scraping*, 25 LEWIS & CLARK L. REV. 147, 173–76, 182–83 (2021).

122. I make this observation with reference to United States law. Article 4 of the European Union's 2019 Directive on the Digital Single Market permits "[r]eproductions and extractions of lawfully accessible works and other subject matter for the purposes of text and data mining" so long as rightsholders have not "expressly reserved" their rights "in an appropriate manner, such as machine-readable means in the case of content made publicly available online." 2019 O.J. (L 130/92) 113–14.

123. In Senate testimony, Stability's representative asserted that "training AI models is an acceptable, transformative, and socially-beneficial use of existing content that is protected by the fair use doctrine." *Brooks Statement*, *supra* note 115, at 8. *See also* Comment of OpenAI, LP to the United States Patent and Trademark Office, Request for Comments on Intellectual Property Protection for Artificial Intelligence Innovation, https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf [<https://perma.cc/3A4X-MX8J>], at 1 (asserting that "[u]nder current law, training AI systems constitutes fair use").

124. Matthew Sag has suggested that an AI firm's attentiveness to opt-outs might bear on the fair use analysis. *See* Matthew Sag, *Fairness and Fair Use in Generative AI*, 92 FORDHAM L. REV. 1887, 1920 (2024) ("The unfairness of systematic indirect expressive substitution seems particularly pronounced if that extraction is done by breaching paywalls, violating terms of service, or disregarding bot exclusion headers. It seems quite plausible that a court might extend the fourth factor to consider whether, in scraping material from the Internet, the defendant ignored robot.txt files indicating a desire to opt out of search engine indexing and similar activities."). *But see* Pierre N. Leval, *Toward a Fair Use Standard*, 103 HARV. L. REV. 1105, 1126 (1990) ("The [fair use] inquiry should focus not on the morality of the secondary user, but on whether her creation claiming the benefits of the doctrine is of the type that should receive those benefits No justification exists for adding a morality test.").

Yet, in response to a senator's question—"[w]hat steps, if any, does Stability AI take to ensure that your training data does not include copyrighted material?"—Stability's representative touted the company's adherence to opt-out requests.¹²⁵

There's a name for the policies that Spawning, Stability AI, and OpenAI are insinuating: formalities. The U.S. copyright system used to rely heavily on formalities. In order to keep their works out of the public domain, authors had to take affirmative steps like affixing copyright notice to their published works and renewing their copyrights after the expiration of an initial term of protection.¹²⁶ Under current law, however, copyright is "unconditional:" it subsists from the moment an original work is fixed in tangible form, and formalities like notice and renewal are not required.¹²⁷ The opt-out standard that the AI industry is trying to establish is, in practical effect, a formality.¹²⁸ In order to exercise a right to exclude AI training, authors would have to take the affirmative step of opting out—otherwise, their acquiescence is presumed.

There are a lot of reasons to like formalities. They create useful data about copyright ownership and keep the public domain populated with works not being exploited commercially. But the AI industry isn't pushing for formalities to realize these sorts of broad benefits. Rather, all the enterprise wants is a very narrow carveout to unconditional copyright. The protection-by-default regime would persist for nearly everyone. The biographer who wants to publish old correspondence, the publisher who wants to re-release an out-of-print book, the producer who wants to revive a forgotten musical—they'd all have to track down a copyright holder and obtain explicit permission. But if the unauthorized use of copyrighted works is to train generative AI, the rule would flip to permission-by-default.

The evolution of the formalities discourse illustrates just how much ground the copyright minimalists have lost. Today's opt-in/opt-out debate in AI training recalls a far more utopian project from the not-so-distant past: the push to establish an opt-out framework for *universal* access to copyrighted works. In late 2004, Google announced the book-digitization

125. Artificial Intelligence and Intellectual Property – Part II: Copyright, Hearing Before the Subcomm. on Intellectual Property (Response to Questions from Senator Tillis for Mr. Ben Brooks, Head of Public Policy, Stability AI) at 8, *available at* https://www.judiciary.senate.gov/imo/media/doc/2023-07-12_pm_-_qfr_responses_-_brooks.pdf [<https://perma.cc/Z7C9-YYBX>].

126. See Sprigman, *supra* note 39, at 487–88.

127. See *id.*

128. See Benjamin Sobel, *A Taxonomy of Training Data: Disentangling the Mismatched Rights, Remedies, and Rationales for Restricting Machine Learning*, in *ARTIFICIAL INTELLIGENCE AND INTELLECTUAL PROPERTY* 221, 238 (Jyh-An Lee et al. eds., 2021) (discussing the opt-out text and data mining regime in the European Union).

initiative that would come to be known as Google Books.¹²⁹ Google touted an option that would permit copyright holders to “opt out” of having text from their books displayed to Google users.¹³⁰ In a 2005 letter to Google, the Association of American University Presses objected to the opt-out framework:

Google’s response to publishers’ objections to Google Print for Libraries that they may “opt out” of the program seems . . . legally irrelevant Among other reasons, it is irrelevant because all a publisher can do under this option is assert its control over the right of display by Google after the infringing copies have been made. It ignores the fundamental exclusive right of copyright owners to make copies in the first place, and it ignores the exclusive right of distribution, since a copy or copies will have already been given to the participating libraries.¹³¹

In a 2006 article titled *Standing Copyright Law on its Head*, Oren Bracha argued in favor of a general opt-out regime for commercial and noncommercial digital libraries accessible to the public.¹³² Other commentators followed suit.¹³³ The proposals didn’t catch on, and for that reason (and a number of others), Google Books fell short of Google’s soaring ambitions. In 2011, a federal district court refused to approve a proposed class-action settlement concerning the project, which, among other things, would have authorized Google to display digitized books unless the relevant rightsholders opted out.¹³⁴ The court observed, “many of the concerns raised in the objections [to the proposed amended settlement agreement] would be ameliorated if [it] were converted from an ‘opt-out’ settlement to an ‘opt-in’ settlement.”¹³⁵

Nearly two decades have passed since Bracha’s article and the initial Google Books opt-in/opt-out debate. In that time, what’s up for debate has shifted significantly. Bracha’s article is written broadly enough that it appears to advocate an opt-out regime even for websites that publish

129. *The Google Print Project is Announced*, HIST. INFO., <https://www.historyofinformation.com/detail.php?id=1160> [https://perma.cc/DSC5-NRTJ].

130. Stefanie Olsen, *Publishers Balk at Google Book Copy Plan*, CNET NEWS.COM, (May 24, 2005, 5:50 PM), https://web.archive.org/web/20050723084408/http://news.com.com/Publishers+balk+at+Google+book+copy+plan/2100-1025_3-5719156.html [https://perma.cc/UGK9-48ZF]; see also Bracha, *supra* note 102, at 1802.

131. *The University Press Assn.’s Objections*, BLOOMBERG (May 23, 2005), <https://www.bloomberglaw.com/product/blaw/bloombergterminalnews/bloomberg-terminal-news/O1T2R16JTSE9> [https://perma.cc/26LL-ARVB].

132. Bracha, *supra* note 102, at 1821–22, 1854–55.

133. See, e.g., Joy Su, *Google Book Search and Opt-Out Procedures*, 56 J. COPYRIGHT SOC’Y U.S.A. 947, 982 (2008); Giancarlo F. Frosio, *Google Books Rejected: Taking the Orphans to the Digital Public Library of Alexandria*, 28 SANTA CLARA COMPUT. & HIGH TECH. L.J. 81, 96–97 (2011); Robert Darnton, *Digitize, Democratize: Libraries and the Future of Books*, 36 COLUM. J.L. & ARTS 1, 17 (2012).

134. See *Authors Guild v. Google, Inc.*, 770 F. Supp. 2d 666, 686 (S.D.N.Y. 2011).

135. *Id.*

unauthorized copies of copyrighted works in full.¹³⁶ Today, such institutions are generally referred to as “shadow libraries”—although a more felicitous name might be “training datasets for human beings.”¹³⁷ Present-day opt-out discourse does not revolve around publication in shadow libraries. Rather, it is limited to the far narrower issue of AI training. Commentators seem to presume that an opt-out mechanism won’t be used to promote *human* access to copyrighted works; what’s on the table is merely whether an opt-out mechanism can be used to facilitate AI training on those works. Even in the AI-training realm, the shadow libraries are a toxic brand. Matthew Sag, a scholar supportive of extending the fair use defense to permit the unauthorized training of AI on copyrighted works, recently suggested that training AI on shadow libraries might not be fair use.¹³⁸

To put it bluntly, we are now far, far removed from the vision of opt-out that copyright minimalists apparently thought was attainable eighteen years ago. Opt-out has gone from a proposal to promote “universal access to cultural materials” to a device that would facilitate AI training alone.¹³⁹ A truly minimal copyright regime is the proverbial frog in the pot. The water has been heating up for decades and is now at a rolling boil. Today’s copyright-decelerationist incremental-minimalism, at least as far as AI is concerned, is an attempt to negotiate the temperature down to a simmer. That’s a bargain that would do the frog no good.

CONCLUSION

Sam Altman, past and present CEO of OpenAI, has characterized AI as “unstoppable.”¹⁴⁰ Fair enough. By embracing copyright accelerationism, we can pit the unstoppable force of AI against an object that we thought was immovable: bad copyright doctrine.

There are, of course, many reasons to be skeptical of what I’ve described. For one, copyright accelerationism could sputter out and end up enriching major rightsholders by requiring AI firms to license training data. This is the “licensing world,” and it is the future the incremental-minimalists

136. Bracha, *supra* note 102, at 1817, 1817 n.80 (defining “digital libraries” as “organized collections of informational items in digital format, accessible through computers” and stipulating that “[t]he definition [of ‘digital libraries’] supplied in the text is purposefully loose in order to encompass many of the relevant variations”).

137. For a recent lawsuit against shadow libraries, see Complaint, Cengage Learning, Inc. et al. v. Does 1-50, No. 1:23-cv-08136 (S.D.N.Y. Sept. 14, 2023).

138. See Sag, *supra* note 124, at 1917-18.

139. Bracha, *supra* note 102, at 1802.

140. Sam Altman, *Moore’s Law for Everything*, (Mar. 16, 2021), <https://moores.samaltman.com/> [<https://perma.cc/V5Z3-UD29>].

fear—in part, I suspect, because of its superficial resemblance to the defeat at the hands of big content that they dreaded in the 1990s. But for most people, the licensing world is roughly comparable to a future in which fair use protects the AI enterprise. The main difference is how big businesses divide the spoils of the AI boom. Instead of AI firms winning unconditional victory, major rightsholders will get to extract additional rents. For the reasons I’ve suggested, however, such a regime doesn’t leave human users’ rights much worse off, and it at least would direct some royalty revenue to small-time rightsholders.

Another potential downside, or perhaps upside, is that copyright accelerationism could impede the runaway progress of AI. Of course, copyright accelerationism doesn’t aspire to thwart AI. Rather, its goal is to use copyright’s coercive power to capture some of the AI industry’s inertia and direct it towards systemic copyright reforms. But if copyright accelerationism sputters and leaves us in the licensing world, the cost, availability, and potency of AI would probably suffer to some degree. At the extreme, copyright accelerationism might even “kill” the progress of AI.¹⁴¹ It’s hard to assess just how bad, or how good, this range of outcomes is. Some people believe that breakthroughs in AI can solve humanity’s problems and enrich us beyond our wildest dreams; others view AI as an existential threat. I can’t evaluate the probability that AI will save humanity or doom it, and copyright accelerationism is agnostic on the question. I’m just going to call this issue a wash.

A more serious shortcoming is that the coalition that copyright accelerationism envisions is a weak and implausible one. Instead of siding with either the behemoth content or AI industries, my proposal would require small-time rightsholders to band together and rankle both behemoths. Asking an amorphous group to organize and defy powerful interest groups is probably not the path to political victory—even if, in theory, copyright gives that diffuse group of small-time rightsholders extraordinary leverage at this particular moment.¹⁴² And if copyright accelerationism does succeed in destabilizing copyright, it may be naïve to think that this instability can be harnessed to enact a better regime.¹⁴³ Perhaps a controlled burn would fail, and a contagion of undesirable instability would infect other areas of law and society.

141. See Andreessen Horowitz, *supra* note 61, at 8.

142. Indeed, this sort of collective action is not all that dissimilar from Spawning’s vision of the future, in which “[l]esser-known artists could band together, in I.P. unions of sorts: a group of architectural photographers might create a data set large enough to appeal to corporate entities, and split the proceeds.” Anna Wiener, *Holly Herndon’s Infinite Art*, THE NEW YORKER (Nov. 13, 2023), <https://www.newyorker.com/magazine/2023/11/20/holly-herndons-infinite-art> [https://perma.cc/CF35-JJVE].

143. See *supra* Part I.B (discussing the assumption that copyright accelerationism would precipitate at most a controlled burn).

Or perhaps today's bad copyright system would, if destabilized, be replaced by something even worse.

At a minimum, however, copyright accelerationism is worth taking seriously as a challenge to the progressive intellectual property commentariat's received wisdom. Far too many copyright scholars and activists still believe they're fighting the copyright wars of the 1990s, in spite of the fact that today's AI revolution has drastically reconfigured copyright's political economy. The AI enterprise, perhaps today's most powerful business lobby and an industrial-scale "user" of copyrighted works, is championing exceptions and limitations to copyright—but ostensibly progressive copyright organizations are still framing their remit as the defense of small-time users' rights against the overreach of a powerful cabal of copyright maximalists.¹⁴⁴

I end with a warning for copyright incremental-minimalists and copyright decelerations of all stripes: if you reject copyright accelerationism, what you will probably get instead is *accelerationist copyright*. You'll get a regime that privileges technical development at all costs, even as it belittles the creation and consumption of expression by and for human beings. Something better is possible. Strong copyright might not be the weapon you want, but it's the one you have.

144. See, e.g., Kit Walsh, *EFF to Copyright Office: Copyright Is Indeed a Hammer, But Don't Be Too Hasty to Nail Generative AI*, ELEC. FRONTIER FOUND. (Oct. 31, 2023), <https://www.eff.org/deeplinks/2023/10/eff-copyright-office-copyright-indeed-hammer-dont-be-too-hasty-nail-generative-ai> [<https://perma.cc/C7HJ-DJXX>] (referring to "the imbalance in bargaining power between creators and . . . publishing gatekeepers" and "the desire of content industry players to monopolize any expression that is reminiscent of or stylistically similar to the work of an artist whose rights they own"); cf., e.g., Creative Commons, *supra* note 73, at 3 ("We are concerned about new restrictions on the ability to make use of a copyrighted work in order to extract data from it to be processed by a machine. Treating copying to train AI as per se infringing copyright would in effect shrink the commons and impede others' creativity in an over-broad way. It would expand copyright to give certain creators a monopoly over ideas, genres, and other concepts not limited to a specific creative expression, as well as over new tools for creativity.").