Artificial Infringement

MICHAEL P. GOODYEAR[†]

Generative artificial intelligence ("AI") is changing the way we do everything from legal research to artistic creation. These advances are possible through recent improvements in machine learning, but greater AI capacity comes with increasingly unpredictable outputs. AI systems often generate outputs the user and developer never considered, let alone intended. Sometimes, these unforeseen outputs can infringe others' copyrights. In the past two years, copyright law has become one of the leading legal and policy battlegrounds for generative AI. Yet the question of who should be liable when AI systems infringe remains unanswered.

By examining the historical doctrinal response of copyright law to new technologies, this Article offers a new analytical framework for determining liability for what it terms artificial infringement, where generative AI systems create infringing outputs. Time and again, new technologies have posed challenges to existing copyright law, straining its capacity to balance protecting authors' rights to incentivize new creative works and providing public access to those works. Courts and Congress have been able to maintain this balance by using a variety of doctrinal and regulatory tools, including fair use, compulsory licensing, and secondary liability. One undertheorized tool, however, is the refinement of the copyright infringement claim. Examining the evolution of copyright infringement reveals the emergence of the volition or causation requirement in response to increasingly complex machines.

This Article advances causation into the AI era and proposes that the AI system should be held directly liable for artificial infringement because it determines the expression in the infringing

---

[†] Associate Professor, New York Law School. J.D., University of Michigan Law School; B.A., University of Chicago. This Article was written and edited prior to the publication of the Supreme Court's opinion in *Cox Communications, Inc. v. Sony Music Entertainment*, 146 S. Ct. 959 (2026). The policy arguments set forth in most of the Article are largely unaffected by the decision and remain valid, but the decision will affect part of the proposal in Subpart.IV.B.2. For helpful discussions and comments on earlier drafts, I would like to thank BJ Ard, Oren Bracha, Ran Canetti, TJ Chiang, Natasha Chokhani, Tyler Rose Clemons, Kevin Collins, Courtney Cox, Jeannie Fromer, Jane Ginsburg, Patrick Goold, James Grimmelmann, Madeleine Gyory, Justin Hughes, Jyh-An Lee, Mark Lemley, Anat Lior, Elise Maizel, Sari Mazzurco, Jacob Noti-Victor, Lisa Larrimore Ouellette, Isaac Park, Zvi Rosen, Nathan Rouse, Jason Schultz, Omar Shehabi, David Schwartz, Chris Sprigman, Rebecca Tushnet, and Felix Wu, as well as the participants in the NYU Lawyering Scholarship Colloquium, Thirty-First Intellectual Property Law & Policy Conference at Fordham Law, NYU ACP Scholarship Clinic, the Twenty-Fourth Annual Intellectual Property Scholars Conference at Berkeley Law, Seventh Junior Faculty Forum for Law and STEM at the University of Pennsylvania Carey Law School, Digital Life Seminar at Cornell Tech, CS & Law 2025: ACM Symposium on Computer Science and Law in Munich, Germany, the Fifth Machine Lawyering Conference at the Chinese University of Hong Kong, and the "Dawn of AI, Dusk of WTO" symposium at Loyola Law School, Los Angeles.

output. Then, by making the AI system the direct infringer as a legal person or legal fiction, courts can remove copyright law from strict liability and instead utilize secondary liability doctrines to conduct a more nuanced, fault-based analysis of user and developer liability for AI-generated infringements. This conceptualization of the AI system as the direct infringer and users and developers as potentially secondarily liable helps resolve the seismic infringement battles between copyright owners and AI while maintaining balance between copyright's competing policy goals.

TABLE OF CONTENTS

INTRODUCTION 1156
I. AI'S THREAT TO COPYRIGHT'S BALANCE 1159
A. MACHINE LEARNING AND GENERATIVE AI 1160
B. THE INCOMPLETE FAIR USE DEFENSE 1166
C. THREATENING COPYRIGHT'S INCENTIVES-ACCESS BALANCE 1171
II. THE EVOLUTION OF COPYRIGHT DIRECT LIABILITY 1175
A. CORPOREAL INFRINGEMENT 1176
B. MECHANICAL INFRINGEMENT 1178
III. THE CAUSATION POLESTAR 1182
IV. ARTIFICIAL INFRINGEMENT 1187
A. DIRECT LIABILITY FOR AI 1187
1. Limited Instances of Potential User and Developer Causation 1189
2. The Promise of AI System Causation 1192
3. Normative Detriments of De Facto Liability for Developers and Users 1195
4. Parallel Problems with a Principal-Agent Relationship 1199
B. SECONDARY LIABILITY FOR HUMANS 1202
1. Vicarious Liability 1204
2. Contributory Liability and Inducement 1205
CONCLUSION 1211

INTRODUCTION

Generative AI offers tantalizing creative possibilities. Take a twenty-one-year-old college student, Shane, hoping to write a song for his boyfriend. Having difficulty stringing together lyrics, Shane decides to use Anthropic's AI chat system, Claude. He asks Claude to write him a love song. In response, Claude produces melodic lyrics. Shane, happy with the result, adds notes, rhythm, tempo, and dynamics. He sings the song, and it is a hit with his boyfriend. He even decides to post a recording to YouTube, where it garners 100,000 views.

However, unbeknownst to Shane, this song's lyrics are rather similar to those in Taylor Swift's hit song "Love Story." Copyright infringement risks abound. If his song is found to be infringing, he would be held liable as the direct infringer for posting it on YouTube.[1] YouTube could also be secondarily liable for infringement if it does not remove the infringing song.[2]

But who is liable for the initial creation of the infringing lyrics? Shane only requested a love song. Anthropic, the developer of Claude, did not want the infringing song and, in fact, specifically tried to avoid infringing outputs.[3] Courts and scholars have barely contemplated how to answer this question.

In the past few years, copyright law has been at the vanguard of legal issues surrounding generative AI. Over four dozen pending lawsuits and a growing body of literature have examined whether training a generative AI model on copyrighted works is infringement.[4] Several scholars have plausibly suggested that including copyrighted works in AI systems' training data is lawful, either as an exception to copyright infringement under the fair use doctrine or because the training only utilizes factual, nonprotected parts of the training data.[5] Two early

---

1. Bus. Casual Holdings, LLC v. YouTube, LLC, No. 21-CV-3610, 2022 WL 17177970, at *6 (S.D.N.Y. Nov. 22, 2022).

2. Bus. Casual Holdings, LLC v. YouTube, LLC, No. 22-CV-3007, 2023 WL 6842449, at *2 (2d Cir. Oct. 17, 2023).

3. This is especially true given that Anthropic is involved in a copyright infringement lawsuit filed by Universal Music Group, which serves as Swift's record label. *See* Complaint & Demand for Jury Trial at 3, Concord Music Grp., Inc. v. Anthropic PBC, No. 23-CV-01092 (M.D. Tenn. Oct. 18, 2023).

4. *See, e.g.*, Complaint at 4–5, Getty Images (US), Inc. v. Stability AI, Ltd., No. 25-CV-06891 (N.D. Cal. Aug. 14, 2025); Complaint at 2, Authors Guild v. OpenAI Inc., No. 23-CV-08292 (S.D.N.Y. Sep. 19, 2023); Complaint at 2, N.Y. Times Co. v. Microsoft Corp., No. 23-CV-11195 (S.D.N.Y. Dec. 27, 2023). Edward Lee maintains a helpful website called "ChatGPT Is Eating the World," where he tracks the progress of these lawsuits. *See Master List of Lawsuits v. AI, ChatGPT, OpenAI, Microsoft, Meta, Midjourney & Other AI Cos.*, CHAT GPT IS EATING THE WORLD (Feb. 7, 2026), https://chatgptiseatingtheworld.com/2024/08/27/master-list-of-lawsuits-v-ai-chatgpt-openai-microsoft-meta-midjourney-other-ai-cos [hereinafter *Master List of Lawsuits*]. While these cases are slowly moving toward resolution, the question addressed in this Article is unlikely to be answered any time soon.

5. *See* Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 770–79 (2021) (describing training an AI system on copyrighted data as relying on the unprotectable parts of copyrighted works); Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence's Implicit Bias Problem*, 93 WASH. L. REV. 579, 619–28 (2018) (describing AI training as being highly transformative and only using the factual content of copyrighted works); Matthew Sag, *Copyright Safety for Generative AI*, 61 HOUS. L. REV. 295, 340–41 (2023) ("Using copyrighted works as training data for generative AI is likely to be fair use if appropriate precautions are taken."); Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley & Percy Liang, *Foundation Models and Fair Use*, 24 J. MACH. LEARNING RSCH., Sep. 2023, at 1, 6 ("For generative models

court decisions seem to have endorsed this view, at least under certain circumstances.[6]

Yet even if training an AI system on copyrighted works is lawful, there is still the largely unexamined question of who should be liable under copyright law when a generative AI system creates an infringing output. Even the staunchest advocates of fair use have admitted that at least some outputs would likely infringe copyright because they are substantially similar to or derivatives of works in the training data and compete in the same market as the original copyrighted work.[7] Examples of likely infringing AI-generated outputs are starting to appear in litigation, including in that brought by Disney and Warner Brothers against Midjourney for reproducing copyrighted characters such as Buzz Lightyear and Batman in outputs.[8]

The generative nature and unpredictability of AI systems' outputs pose novel questions for copyright law. Overbroad liability could risk stifling the generative AI industry, including its ability to promote access to information and creative works. Yet depriving authors of their full rights could disincentivize the creativity copyright law seeks to foster. As others have recognized in different areas of the law, the unprecedented nature of AI may necessitate a new legal paradigm.[9]

---

like DALL-E or GPT that produce creative outputs, the situation is less likely to be problematic if the outputs do not copy a substantial portion of any existing work but instead transform the input into totally different outputs, in line with the fair use doctrine."); Pamela Samuelson, *Fair Use Defenses in Disruptive Technology Cases*, 71 UCLA L. REV. 1484, 1557–59 (2024) (considering the ingestion of training data by AI systems to be copyright fair use because it is a non-expressive use); Katherine Lee, A. Feder Cooper & James Grimmelmann, *Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain*, 72 J. COPYRIGHT SOC'Y 251, 367–69 (2025) (concluding that training a generative AI system and AI-generated outputs could be fair use, or not, depending on the nature of the generated outputs); Oren Bracha, *The Work of Copyright in the Age of Machine Production*, 38 HARV. J.L. & TECH. 171, 202 (2024) (arguing, instead of fair use, that training copies are not infringing because they do not fall within copyright's scope). *But see* David W. Opderbeck, *Copyright in AI Training Data: A Human-Centered Approach*, 76 OKLA. L. REV. 951, 953–57 (2024) (questioning the fair use defense for AI training); Benjamin L.W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COLUM. J.L. & ARTS 45, 67 (2017) (arguing that fair use may not apply).

6. Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007, 1033 (N.D. Cal. 2025); Kadrey v. Meta Platforms, Inc., 788 F. Supp. 3d 1026, 1060 (N.D. Cal. 2025).

7. *See, e.g.*, Henderson et al., *supra* note 5 ("When the downstream product based on such a model is *not* transformative (e.g., model outputs similar content to copyrighted training data, or the application's market is similar to original data markets), courts may decide that the generated content, the model deployment, and even potentially the model parameters themselves are not covered by fair use."); Sag, *supra* note 5, at 309–10 ("The critical difference is that, although [large language models] do not generally produce pseudo-expressive works that mimic their training data, they may do so under specific circumstances, particularly in the context of copyrightable characters and analogous situations.").

8. Complaint & Demand for Jury Trial at 40, Disney Enters., Inc. v. Midjourney, Inc., No. 25-5275 (C.D. Cal. June 11, 2025); Complaint & Demand for Jury Trial at 56, Warner Bros. Ent. Inc. v. Midjourney, Inc., No. 25-8376 (C.D. Cal. Sep. 4, 2025).

9. *See, e.g.*, Kenneth S. Abraham & Robert L. Rabin, *Automated Vehicles and Manufacturer Responsibility for Accidents: A New Legal Regime for a New Era*, 105 VA. L. REV. 127, 128 (2019) ("We are on the verge of another new era, requiring yet another revision to the legal regime."); Anna Beckers & Gunther Teubner, *Responsibility for Algorithmic Misconduct: Unity or Fragmentation of Liability Regimes?*, 25 YALE J.L. & TECH. (SPECIAL ISSUE) 76, 78–80 (2023) (identifying liability gaps for automated investment machine failure, inaccurate human-machine allegations of misconduct, and multiple algorithms acting correctly at an

While approaches to AI-generated harms have started to ossify in the broader tort context, discussions of what this Article terms "artificial infringement"—infringing outputs created through generative AI systems—are still nascent. A rare exception is a proposal by Mala Chatterjee and Jeanne Fromer that AI systems might, under the right circumstances, be sufficiently volitional within the meaning of copyright law for the developer to be held liable because the doctrine focuses on function rather than consciousness.[10] They argue that while a copy machine would not have functional volition, "a more sophisticated computer could plausibly be designed to 'choose' whether to make a copy."[11]

This Article takes a different approach, relying not on volition but a different aspect of the copyright infringement claim that is often imprecisely entangled with volition: causation. It proposes that AI systems themselves should be directly liable for infringing outputs, whether as a legal person or a legal fiction, because they determine the expression, causing the infringement to occur. To make this contribution, I offer a new paradigm for examining artificial infringement liability by looking to the historical evolution of the copyright infringement claim. Courts and Congress have used a variety of tools to accommodate new technologies while maintaining a balance between copyright's goals of protecting copyright owners' rights to incentivize new creative works and providing public access to those works.[12] However, refinements to the (direct) copyright infringement claim have been underexamined as a tool. This Article identifies the historical transition from liability for what it terms corporeal infringement to mechanical infringement through the causation requirement, which helped preserve copyright's incentives-access balance in the face of new, complex machines.[13]

Situating AI within this historical narrative, this Article's primary contribution is proposing that the AI system itself could be directly liable, which would shift the determination of developer and user liability to a more intent-focused secondary liability analysis. Holding the developer or user strictly liable for all artificial infringements would seem to punish them for offering or using a semiautonomous generative AI system even though they have limited control

individual level but "caus[ing] damage due to their mutual reinforcement"); Andrew D. Selbst, *Negligence and AI's Human Users*, 100 B.U. L. REV. 1315, 1375 (2020) ("[W]e will likely need a new regulatory paradigm to compensate the victims of AI's use."); Mihailis E. Diamantis, *Reasonable AI: A Negligence Standard*, 78 VAND. L. REV. 573, 603 (2025) (proposing an AI negligence standard); W. Nicholson Price II & I. Glenn Cohen, *Locating Liability for Medical AI*, 73 DEPAUL L. REV. 339, 341–42 (2024) (arguing that existing tort doctrines are inappropriate for medical AI-based liability).

10. Mala Chatterjee & Jeanne C. Fromer, *Minds, Machines, and the Law: The Case of Volition in Copyright Law*, 119 COLUM. L. REV. 1887, 1914 (2019).

11. *Id.* at 1910.

12. *See infra* Subpart.I.C.

13. *See infra* Part.II and Part.III.

over the contours of the specific output.[14] Holding the AI system directly liable is more accurate because the infringing quality of the output is attributable to it. Instead of an overbroad strict liability rule, recognizing the AI system as a direct infringer allows the analysis to shift from strict liability to the more flexible, fault-based contributory liability and inducement standards for determining when to hold developers or users liable for their roles in the resulting infringement.[15] Eschewing relying on a single doctrine, this Article's approach uses two tools in tandem—refinements to direct and secondary liability—to maintain copyright's balance between its competing goals of incentives and access in the face of new challenges posed by artificial infringement.

This Article proceeds in four parts. First, in Part I, it explains how generative AI systems operate and the novel issues they pose to the copyright infringement liability status quo. It then briefly summarizes existing fair use and related noninfringement arguments, finding that AI systems will generate infringing outputs in at least some cases. Part I concludes by showing how liability for AI-generated infringing outputs risks overturning the delicate balance between incentives and access that copyright law seeks to preserve. Part II turns to how courts have refined copyright infringement doctrine to accommodate new technologies by analyzing two distinct phases in its evolution: corporeal and mechanical infringement. Part III elucidates and defends the causation polestar courts used for mechanical infringement. The remainder of the Article, Part IV, then applies causation to the issue at hand: artificial infringement. To answer who should be liable for those outputs, this Article applies causation in the direct liability context to conclude that, in most cases, the AI system should be the direct infringer in artificial infringement cases, whereas courts should examine the prospective liability of developers and users under the more nuanced secondary liability doctrines of contributory liability and inducement.

I. AI'S THREAT TO COPYRIGHT'S BALANCE

New technologies have long posed significant challenges to copyright law. One of the latest innovations, generative AI, is no different. Scholars and litigants have already debated whether training generative AI models on copyrighted materials should be lawful.[16] However, another copyright liability question remains largely unexamined: When an AI system generates an infringing copy, who should be liable? The unpredictability of AI-generated outputs makes this issue a case of first impression for copyright law. Courts must proceed carefully because an improper liability framework could upset the delicate balance copyright law seeks to maintain between its competing goals of

14. *See infra* Subpart.IV.A; Mark A. Lemley & Bryan Casey, *Remedies for Robots*, 86 U. CHI. L. REV. 1311, 1327 (2019) (describing the unpredictable nature of AI systems' outputs).

15. *See infra* Part.IV.

16. *See supra* notes 4–5.

protecting copyright owners' rights to incentivize new creative works and providing public access to those works.[17]

This Part first provides an overview of machine learning and generative AI technology, which raise the novel liability questions that copyright law will need to address. Next, it discusses the state of fair use and related arguments, the current focus of copyright and generative AI, and explains how liability will remain uncertain for AI-generated infringing outputs. Finally, this Part explains how uncertainty around liability for infringing outputs threatens the balance between incentives and access at the core of copyright law.

A. MACHINE LEARNING AND GENERATIVE AI

There is no singular definition of AI, but it can generally be defined as a machine-based system that makes predictions and acts in ways that mimic human intelligence.[18] While futurists predict a world where AI will replace human decision-makers,[19] the reality of AI's capabilities today is more limited. Although machines can emulate humans, a true thinking machine—the so-called Artificial General Intelligence ("AGI") or "Strong AI"—is still elusive.[20] AI systems are able to produce useful, intelligent results primarily through heuristics: they detect statistically relevant patterns in the data on which they are trained and produce results accordingly.[21] Therefore, AI systems are generating results based on computation, not cognition.[22] In fact, all algorithmic commands are either AND, OR, or NOT.[23]

While generative AI is a recent development, AI more generally has been around for years. For most of that time, AI systems were entirely hand-coded, with explicit algorithms governing all of the system's conduct—an approach called Good Old-Fashioned AI ("GOFAI").[24] GOFAI used common algorithms and code to perform very specific tasks that, while seemingly intelligent, were based on the designer's explicit instructions for each scenario.[25] This approach achieved remarkable results, including the IBM supercomputer Deep Blue beating chess world champion Garry Kasparov in a 1997 match.[26]

---

17. *See infra* Subpart.I.C.

18. *See, e.g.*, National Artificial Intelligence Initiative Act of 2020, 15 U.S.C. § 9401(3) ("The term 'artificial intelligence' means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.").

19. Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1, 2–3, 46 (2019).

20. Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 GA. ST. U.L. REV. 1305, 1308–09 (2019).

21. *Id.* at 1308.

22. *Id.*

23. Lemley & Casey, *supra* note 14, at 1322.

24. *Id.*

25. Omri Rachum-Twaig, *Whose Robot Is It Anyway?: Liability for Artificial-Intelligence-Based Robots*, 2020 U. ILL. L. REV. 1141, 1146.

26. Deborah Yao, *25 Years Ago Today: How Deep Blue vs. Kasparov Changed AI Forever*, AI BUS. (May 10, 2022), https://aibusiness.com/ml/25-years-ago-today-how-deep-blue-vs-kasparov-changed-ai-forever.

The GOFAI model worked well for narrow, rule-based objectives. For example, Deep Blue was an expert system that incorporated chess experts' knowledge and replicated it well because chess is a rule-based game in which all the stimuli are known.[27] Other systems included medical AI to diagnose diseases based on the medical knowledge upon which they were trained, or tax preparation software coded with specific rules based on the guidance of a tax attorney or accountant.[28] We have come to rely on these systems for highly rule-based tasks. Tax preparation software such as TurboTax, personalized legal forms provided by LegalZoom, and a parking ticket appeal app called DoNotPay have all replicated, to varying degrees of success, the more procedural functions of human accountants and attorneys.[29]

But the GOFAI model inherently limited AI to coded knowledge since the developer must explicitly provide the AI system with all its operating and decision-making rules in advance.[30] There was no intuitive learning by the AI.[31] GOFAI accordingly struggled with more amorphous tasks such as image generation.[32] For these kinds of tasks, there is no single correct result. They therefore require more complex algorithms. For example, hard coding an understanding of "dog" as only a German Shepherd or a Poodle may miss other types of dogs, such as Scottish Terriers and Shetland Sheepdogs, whereas an organic learning process could lead to an AI system recognizing dogs more holistically.[33]

Given this limitation, researchers started to explore other options for approximating human cognition in machines, leading to machine learning emerging as the dominant model for AI innovation. Machine learning is the process by which computers use algorithms to learn from and improve through data.[34] This type of AI involves two types of code: programmer-created and machine-created code.[35] First, the programmer writes the code for the model. Second, the program then consumes the (programmer-selected) training data and generates its own code that executes outputs in response to specific user

27. Clive Thompson, *What the History of AI Tells Us About Its Future*, MIT TECH. REV. (Feb. 18, 2022), https://www.technologyreview.com/2022/02/18/1044709/ibm-deep-blue-ai-history.

28. Surden, *supra* note 20, at 1316–17.

29. Pasquale, *supra* note 19, at 8–17.

30. Surden, *supra* note 20, at 1317.

31. PEDRO DOMINGOS, THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD 89–90 (2015) (describing the GOFAI approach as subject to a "knowledge acquisition bottleneck").

32. Surden, *supra* note 20, at 1324–25.

33. *See* Deven R. Desai & Joshua A. Kroll, *Trust but Verify: A Guide to Algorithms and the Law*, 31 HARV. J.L. & TECH. 1, 28 (2017) (explaining that adding randomness into an AI system can further help it gain a better understanding of the task and address any challenges it encounters).

34. Andres Guadamuz, *A Scanner Darkly: Copyright Liability and Exceptions in Artificial Intelligence Inputs and Outputs*, 73 GRUR INT'L 111, 113 (2024); *see also* 15 U.S.C. § 9401(11) (defining machine learning).

35. *See* Jatinder Singh, Ian Walden, Jon Crowcroft & Jean Bacon, Responsibility & Machine Learning: Part of a Process 9 (Oct. 27, 2016) (unpublished manuscript), https://ssrn.com/abstract=2860048.

inquiries.[36] Developers provide raw components in the form of the AI system's goal and a set of training materials, but the system "learns" from the training data to compute its own algorithms to best achieve its goal.[37] Unlike GOFAI, "the machine essentially programs itself."[38] This allows AI systems to become increasingly self-reliant or autonomous.[39] They can program themselves to achieve goals in ways that are not necessarily predictable by the developer or user.[40] Although machine learning predates recent advances, faster computers and larger quantities of data facilitated the recent explosion in machine learning capabilities.[41]

Email spam filters are an example of machine learning: Existing email datasets train spam filters, which then use this training to move likely matches into the spam folder. The spam filter is continuously improved by users, as each time users mark an email as spam (or not), they are continuing the filter's training. With more data, the algorithm should be able to incrementally improve its performance.[42] While AI systems are not intelligent in the way humans are, with enough data, they may reach the same result.[43]

Machine learning-based AI does not just mimic human intelligence; it also discovers solutions no human would have uncovered by themselves.[44] This is especially true with a subset of machine learning called deep learning, where an AI system analyzes statistical patterns using layers of algorithms to mimic a neural network to recognize complex patterns in pictures, text, and other data, sometimes even surpassing human comprehension.[45] For example, the AI may be able to compute complex chains of data that would be too difficult for a human to process.[46] Machine learning also allows the AI system to continuously improve based on the information and feedback it processes.[47] Many AI systems will be connected to the cloud, allowing them to be continually updated based on new sources of information and even other AI systems.[48] These benefits allow AI systems to significantly improve efficiency and possibilities across

36. *See id.* at 7.

37. Lemley & Casey, *supra* note 14, at 1324; Rachum-Twaig, *supra* note 25, at 1146–47.

38. Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), https://www.technologyreview.com/2017/04/11/5113/the-dark-secret-at-the-heart-of-ai.

39. Andrew Tutt, *An FDA for Algorithms*, 69 ADMIN. L. REV. 83, 100 (2017).

40. Knight, *supra* note 38; Curtis E.A. Karnow, *The Application of Traditional Tort Theory to Embodied Machine Intelligence*, *in* ROBOT LAW 51, 52 (Ryan Calo, A. Michael Froomkin & Ian Kerr eds., 2016); Surden, *supra* note 20, at 1314.

41. Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399, 405 (2017).

42. Surden, *supra* note 20, at 1314.

43. *Id.* at 1314–15.

44. *Id.* at 1318; Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513, 538 (2015).

45. *What Is Deep Learning?*, AMAZON WEB SERVS., https://aws.amazon.com/what-is/deep-learning (last visited Mar. 21, 2026).

46. Surden, *supra* note 20, at 1318.

47. Calo, *supra* note 44, at 539; Emiliano Marchisio, *In Support of "No-Fault" Civil Liability Rules for Artificial Intelligence*, SN SOC. SCIS., Jan. 11, 2021, at 1, 6.

48. Jack M. Balkin, *The Path of Robotics Law*, 6 CALIF. L. REV. CIR. 45, 54 (2015).

many industries, leading to experts terming machine learning-based AI one of the few great "general-purpose technologies" ("GPTs") in world history, alongside the internal combustion engine, electricity, and the Internet.[49]

Due to these advantages, today's AI systems primarily rely on machine learning or a hybrid to perform their tasks.[50] One of the most significant advances in machine learning-based AI is generative AI. Generative models "learn" how to create content by extracting statistical patterns from existing media.[51] As with AI more broadly, the developers of the generative models choose what training material to use.[52] In response to user prompts, the AI system then generates an output based on its machine learning.[53]

Generative AI systems are primarily based on diffusion models and large language models ("LLMs") to produce images and prose.[54] These models are trained on vast datasets and allow AI systems to execute complex behaviors including language processing, object recognition, and content generation.[55] Diffusion models, including Imagen, DALL-E, Stable Diffusion, and Midjourney, corrupt large amounts of training data by adding noise and then teach the AI system to reform the data by reversing the corruption.[56] LLMs such as OpenAI's GPT, Meta's Llama, or Anthropic's Claude, instead use a natural language-processing model to digest millions or billions of pieces of data to elucidate patterns from them.[57] Nonetheless, LLMs are still "stochastic parrots" that generate the statistically most likely response given their training material and prompt.[58]

Through techniques like diffusion and LLMs, generative AI has advanced to a level where AI systems can have a deeper understanding of the meaning and correlations of data.[59] Generative AI systems have achieved remarkable results

49. Lemley & Casey, *supra* note 14, at 1325.

50. Surden, *supra* note 20, at 1310; Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson & Harlan Yu, *Accountable Algorithms*, 165 U. PA. L. REV. 633, 679 (2017).

51. Ziv Epstein, Aaron Hertzmann & Investigators of Human Creativity, *Art and the Science of Generative AI*, 380 SCI. 1110, 1110 (2023).

52. *See* Tarleton Gillespie, *The Relevance of Algorithms*, *in* MEDIA TECHS.: ESSAYS ON COMMC'N, MATERIALITY, & SOC'Y 167, 169 (Tarleton Gillespie, Pablo J. Boczkowski & Kirsten A. Foot eds., 2014).

53. Guadamuz, *supra* note 34.

54. Epstein et al., *supra* note 51.

55. Lemley & Casey, *supra* note 14, at 1322.

56. Mengchun Zhang, Maryam Qamar, Taegoo Kang, Yuna Jung, Chenshuang Zhang, Sung-Ho Bae & Chaoning Zhang, *A Survey on Graph Diffusion Models: Generative AI in Science for Molecule, Protein and Material*, ARXIV Aug. 2015, at 1, 3; Guadamuz, *supra* note 34, at 114.

57. Jaron Lanier, *There Is No A.I.*, NEW YORKER (Apr. 20, 2023), https://www.newyorker.com/science/annals-of-artificial-intelligence/there-is-no-ai; Sag, *supra* note 5, at 300.

58. Emily M. Bender, Timnit Gebru, Angelina McMillan-Major & Shmargaret Shmitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, *in* FACCT '21: PROCS. OF THE 2021 ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 610, 616–17 (2021).

59. Guadamuz, *supra* note 34, at 114.

in the last two years. They have created vivid award-winning art,[60] excelled at the bar exam,[61] and written new songs in the voice of deceased stars.[62]

These advancements are based on how the models are trained. Like all AI, the developer of diffusion and LLM models selects the material on which the AI system will be trained.[63] After the training material and goal are set, the level of involvement by human developers and users in machine learning can vary: there is, theoretically at least, a spectrum from autonomous to human-influenced decision-making. Yet, in practice, due to the sheer scale required for these models, most data collection is done automatically rather than manually.[64] Often times, this automatic gathering is achieved through scraping the Internet for content.[65] Other AI systems rely on large datasets that have already been collected, such as the LAION or WebVid datasets.[66] Similarly, preparing the dataset—transforming raw data into analysis-ready formats and assigning data labels—is usually achieved through automation.[67] Developers could engage in more heavily supervised learning where they pre-label the training materials to influence the training process and help the system understand the data.[68] But given the size of datasets, the vast majority (80 percent or more) of all data is left unstructured, without any human-defined parameters to organize the content therein.[69]

Along with the increasing promise of these AI systems comes greater unpredictability around their specific outputs. General AI tools are non-deterministic, meaning that the same inputs can produce different outputs.[70] The exact path the AI system follows can be complex, messy, unexpected, and even unexplainable.[71] While the general functions are desired and predictable, the

60. Kevin Roose, *An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy.*, N.Y. TIMES (Sep. 2, 2022), https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html.

61. Debra Cassens Weiss, *Latest Version of ChatGPT Aces Bar Exam with Score Nearing 90th Percentile*, ABA J. (Mar. 16, 2023, at 13:59 CT), https://www.abajournal.com/web/article/latest-version-of-chatgpt-aces-the-bar-exam-with-score-in-90th-percentile.

62. Steve Brachmann, *Senate IP Subcommittee Mulls Federal Right of Publicity at AI and Copyright Hearing*, IPWATCHDOG (July 13, 2023, at 15:15 PT), https://ipwatchdog.com/2023/07/13/senate-ip-subcommittee-mulls-federal-right-publicity-ai-copyright-hearing.

63. *See* Gillespie, *supra* note 52.

64. Guadamuz, *supra* note 34, at 116.

65. *Id.* at 115.

66. *See generally* Christoph Schuhmann et al., *LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models*, ARXIV 2 (Oct. 16, 2022), https://arxiv.org/pdf/2210.08402 (LAION-5B is a dataset with over 5 billion image-text pairs); Max Bain (@m-bain), *WebVid*, GITHUB, https://github.com/m-bain/webvid (last visited Mar. 21, 2026) (WebVid is a dataset of 10 million scraped video-text pairs).

67. Guadamuz, *supra* note 34.

68. Rachum-Twaig, *supra* note 25, at 1147; Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 678 n.24 (2016).

69. Eric Avidon, *Generative AI Hype Evolving into Reality in Data, Analytics*, TECHTARGET (July 6, 2023), https://www.techtarget.com/searchbusinessanalytics/feature/Generative-AI-hype-evolving-into-reality-in-data-analytics.

70. A. Feder Cooper, Jonathan Frankle & Christopher De Sa, *Non-Determinism and the Lawlessness of Machine Learning Code*, *in* CSLAW '22: PROCS. OF THE 2022 SYMP. ON COMPUT. SCI. & L. 1, 2 (2022).

71. Desai & Kroll, *supra* note 33, at 26.

non-deterministic nature of generative AI systems means that the contours of specific outputs often cannot be predicted in advance, especially with larger generative models.[72] Even so-called explainable AI is a red herring in that it offers a plausible post-hoc explanation rather than prediction of or the actual reasoning for a decision.[73] In addition, the more complex or ill-defined the goal, the less stable, predictable, and explainable the AI system's actions will be.[74] AI models can misinterpret prompts, over-rely on specific training material, or suffer from any other manner of analysis disconnect that could lead to a very different output from what the user intended.[75] Machine learning often incorporates randomness, which can further contribute to accountability and reproducibility shortcomings.[76] Iterative learning by AI systems further exacerbates the problem of humans failing to predict outputs in advance.[77] AI systems' incorporation of randomness results in the same inquiries creating varied outcomes each time—even if they are only seconds apart.[78] Those outputs are often unexpected, let alone intended.[79] The unpredictability of an AI system's outputs means that even developers in charge of creating the AI model may not be able to precisely foresee the outputs.[80] Their opacity can make machine learning-based AI systems "seemingly magical technologies."[81]

Yet the unpredictability of an AI system's outputs is in some ways a feature, not a fault.[82] This output unpredictability is what enables AI systems to go beyond human capabilities and offer novel solutions.[83] The unpredictability of outputs can also be a positive for users, who may wish to choose from a variety of options and use the AI system as a tool rather than have answers dictated to them.[84] For example, generating a dog in the pop art style does not have just one right answer. As shown in Figure 1, that prompt could result in a

72. Deep Ganguli et al., *Predictability and Surprise in Large Generative Models*, *in* FACCT '22: PROCS. OF THE 2022 ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 1747, 1748 (2022).

73. Boris Babic & I. Glenn Cohen, *The Algorithmic Explainability "Bait and Switch,"* 108 MINN. L. REV. 857, 864 (2023).

74. Rachum-Twaig, *supra* note 25, at 1147.

75. *See id.*; *see also* Barocas & Selbst, *supra* note 68, at 681 (explaining that algorithms may unintentionally perpetuate discrimination).

76. Kroll et al., *supra* note 50, at 684.

77. *Id.* at 669.

78. *See* Lanier, *supra* note 57; James Grimmelmann, *There's No Such Thing as a Computer-Authored Work—and It's a Good Thing, Too*, 39 COLUM. J.L. & ARTS 403, 412 (2016).

79. Calo, *supra* note 44, at 554–55; *see also* Selbst, *supra* note 9, at 1332–33 (describing how traditionally strict liability requires foreseeability yet AI injuries are often unforeseeable).

80. P. Bernt Hugenholtz & João Pedro Quintais, *Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output?*, 52 IIC – INT'L REV. INTELL. PROP. & COMPETITION L. 1190, 1212 (2021); Anat Lior, *AI Entities as AI Agents: Artificial Intelligence Liability and the AI Respondeat Superior Analogy*, 46 MITCHELL HAMLINE L. REV. 1043, 1057 (2020).

81. DOMINGOS, *supra* note 31, at xv.

82. Mindy Nunez Duffourc, *Malpractice by the Autonomous AI Physician*, 2023 U. ILL. J.L., TECH. & POL'Y 1, 19 (2023).

83. *See id.* (discussing the benefits of unpredictable outputs in the health AI context).

84. Lanier, *supra* note 57.

happy, bejeweled Scottie or a cool, neon-haired Poodle. The user also cannot predict specific expressive choices such as color, clothing, hair style, and background. Therefore, machine learning poses a significant tradeoff between achieving the richest predictive results and being able to foresee them.[85]

FIGURE 1: FIREFLY-GENERATED POP ART DOGS

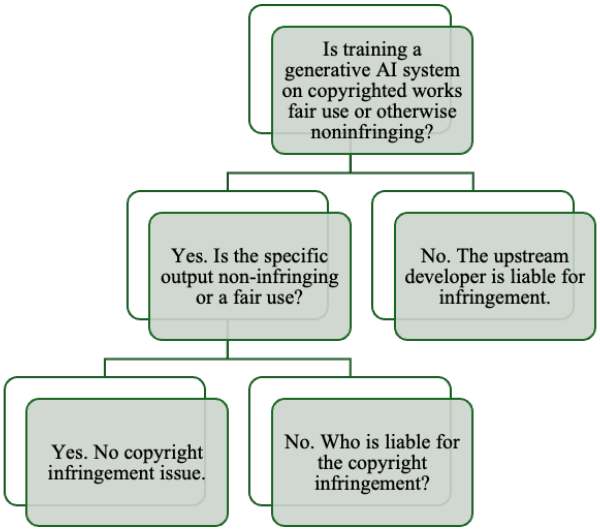


B. THE INCOMPLETE FAIR USE DEFENSE

The unprecedented and complex nature of generative AI poses weighty questions for copyright law. As shown in Figure 2, there are three sequential liability questions that courts must address in copyright infringement litigation involving generative AI. First, is the training of a generative AI system on copyrighted material fair use or otherwise noninfringing? Second, if so, is a specific output from a generative AI system fair use or otherwise noninfringing? Finally, if the output is infringing, who is held liable and under what circumstances?

85. Bryce Goodman & Seth Flaxman, *European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation,"* ARXIV 6 (Aug. 31, 2016), https://arxiv.org/pdf/1606.08813.pdf.

FIGURE 2: GENERATIVE AI INFRINGEMENT QUESTIONS



So far, scholars and litigants have largely restricted themselves to examining the first two questions. At the heart of most of these discussions is copyright's fair use doctrine. Fair use is a statutory exception to copyright infringement for certain societally desirable uses such as criticism, education, and other transformative purposes.[86] Courts weigh four non-exclusive factors as a sort of proxy for differentiating between uses we wish to encourage and those we would rather prohibit in favor of users' rights.[87] In particular, courts have tended to emphasize the transformativeness of the use and the effect of the use on the copyrighted work's market.[88]

There are at least two opportunities for a copyright infringement claim and fair use defense with generative AI: at the input (training) stage and at the output stage. The first question of the analysis is input-side fair use, whether training a generative AI model on copyrighted works infringes the reproduction or derivative works right, or is a lawful fair use. During the training process, temporary or incidental copies of the data may be created, which could result in a prima facie infringement claim.[89] While some of the millions of works on which these models are trained are in the public domain or otherwise free to use,

86. 17 U.S.C. § 107.

87. *Id.*

88. Barton Beebe, *An Empirical Study of U.S. Copyright Fair Use Opinions Updated, 1978–2019*, 10 N.Y.U. J. INTELL. PROP. & ENT. L. 1, 4, 33 (2020); Michael P. Goodyear, *Culture and Fair Use*, 32 FORDHAM INTELL. PROP., MEDIA & ENT. L.J. 334, 368–69 (2022).

89. *See* Pamela Samuelson, *Generative AI Meets Copyright*, 381 SCI. 158, 159 (2023).

many are also copyrightable.[90] This is the main question with which courts are grappling in the plethora of generative AI-related lawsuits across the country.[91]

Several scholars have already weighed in on this question of input infringement, offering robust fair use or other non-infringement arguments. Many argue that training an AI system on copyrighted content is likely transformative.[92] These arguments place generative AI within a long line of functionally transformative products that courts have deemed to be fair uses, including search engines, plagiarism checkers, and Google Books.[93] Alternatively, at least one scholar has argued that AI training relies on the uncopyrightable parts of copyrighted works, such as their underlying facts or grammatical structure, and therefore should not implicate copyright law.[94]

If these input-side fair use and non-infringement arguments are successful, we would still have to address the second question: whether a particular output is infringing. An AI-generated output can be a copy, and therefore implicate copyright law, because it is fixed in a computer's random-access memory ("RAM") for "more than a transitory duration."[95] Unlike the input analysis, whether an output is infringing will likely require individualized examinations of each output. In many cases, the output will be noninfringing, or fair use will likely apply.[96]

AI systems likely engage in actual copying, a requirement for copyright infringement, at least some of the time. Copyright infringement requires copying in fact and improper appropriation.[97] While substantial similarity would be

90. Sag, *supra* note 5, at 300; *see also* Guadamuz, *supra* note 34, at 116 (describing how training data can include uncopyrightable items such as facts, numbers, statistics, and open source material, but also copyrightable works such as text, photographs, music, and art).

91. *See, e.g.*, Complaint at 3, Andersen v. Stability AI Ltd., No. 23-CV-00201 (N.D. Cal. Jan. 13, 2023); Complaint at 2, Getty Images (US), Inc. v. Stability AI, Ltd., No. 25-CV-06891 (N.D. Cal. Aug. 14, 2025). For a running list of generative AI-related cases, see *Master List of Lawsuits*, *supra* note 4.

92. Lemley & Casey, *supra* note 5, at 776–77; Levendowski, *supra* note 5, at 623–24; Sag, *supra* note 5, at 309; Samuelson, *supra* note 5, at 1558; *see also* Lee et al., *supra* note 5, at 369 (concluding that training a generative AI system and AI-generated outputs could be fair uses depending on their nature).

93. *See, e.g.*, Authors Guild v. Google, Inc., 804 F.3d 202, 225 (2d Cir. 2015) (holding that Google Books' search and snippet function was a fair use); Perfect 10, Inc. v. Amazon.com, Inc., 508 F.3d 1146, 1168 (9th Cir. 2007) (holding that Google Images' use of thumbnail images was likely a fair use); Kelly v. Arriba Soft Corp., 336 F.3d 811, 822 (9th Cir. 2003) (holding that thumbnail use in a searchable database was fair use); A.V. *ex rel.* Vanderhye v. iParadigms, LLC, 562 F.3d 630, 645 (4th Cir. 2009) (holding that copying student papers to incorporate into a plagiarism checker was fair use); *see also* Matthew Sag, *Copyright and Copy-Reliant Technology*, 103 NW. U. L. REV. 1607, 1618–24 (2009) (describing four non-expressive fair use cases).

94. Bracha, *supra* note 5, at 198. *But see* U.S. COPYRIGHT OFF., COPYRIGHT AND ARTIFICIAL INTELLIGENCE PART 3: GENERATIVE AI TRAINING 47–48 (2025) ("Where the resulting model is used to generate expressive content, or potentially reproduce copyrighted expression, the training use cannot be fairly characterized as 'non-expressive.'").

95. MAI Sys. Corp. v. Peak Comput., Inc., 991 F.2d 511, 518 (9th Cir. 1993).

96. *See* Henderson et al., *supra* note 5 (positing that AI-generated outputs are likely fair uses so long as they do not copy a substantial portion of a copyrighted work); Sag, *supra* note 5, at 326 (suggesting that the attenuated link between training data and outputs should generally make generative AI a non-expressive use entitled to fair use).

97. Rentmeester v. Nike, Inc., 883 F.3d 1111, 1117 (9th Cir. 2018).

determined on a case-by-case basis, training the AI system on the original work could support copying in fact. Computer science literature suggests that if an AI model "memorizes its training data, these training data are *in the model*."[98] Regurgitation of training data implies memorization.[99] At least for some literary works, users are able to extract significant amounts of copyrighted text from several prominent AI tools.[100] A recent decision from the Munich I Regional Court held that lyrics that were reproduced were fixed within the model parameters.[101] Even if we were to reject this copying in fact finding, exposure to a work can be sufficient because even subconscious copying counts. For example, in *Three Boys Music Corp. v. Bolton*, the court found that the defendants may have listened to the plaintiff's song, "Love Is a Wonderful Thing," twenty years prior as teenagers and subconsciously copied it from the depths of their memories.[102] Similarly, if the AI system was trained on a particular work, it is possible the system is copying in fact and that a substantially similar output infringes the work.

Infringing AI-generated outputs are certain to occur in at least some circumstances. Although the creators of AI systems can—and should—do their best to restrict infringement, the unpredictability of AI outputs mean that infringement is bound to occur.[103] Many scholars recognize that at least a portion of outputs may be substantially similar to copyrighted works but do not fall within the confines of the fair use doctrine.[104] For example, one study found that, of the 175 million images they generated using the most-duplicated examples in Stable Diffusion training data, 109 outputs were near copies of training data.[105] More alarming, a report by Patronus AI, the developer of a copyright-detection application programming interface ("API") for AI, found that OpenAI produced

98. A. Feder Cooper & James Grimmelmann, *The Files Are in the Computer: On Copyright, Memorization, and Generative AI*, 100 CHI.-KENT L. REV. 141, 146 (2025).

99. *Id.* at 165.

100. *See, e.g.*, A. Feder Cooper, Aaron Gokaslan, Ahmed Ahmed, Amy B. Cyphert, Christopher De Sa, Mark A. Lemley, Daniel E. Ho & Percy Liang, *Extracting Memorized Pieces of (Copyrighted) Books from Open-Weight Language Models*, ARXIV 5 (Nov. 30, 2025), https://arxiv.org/pdf/2505.12546; Ahmed Ahmed, A. Feder Cooper, Sanmi Koyejo & Percy Liang, *Extracting Books from Production Language Models*, ARXIV 19 (Jan. 6, 2026), https://arxiv.org/pdf/2601.02671.

101. Landgericht München I [LG] [Munich I Regional Court] Nov. 11, 2025, 42 O 14139/24 ¶ 52 (Ger.).

102. 212 F.3d 477, 484 (9th Cir. 2000).

103. *See* Marchisio, *supra* note 47, at 2 (describing how even the best intentions of developers cannot stop this inevitability).

104. *See, e.g.*, Lemley & Casey, *supra* note 5, at 766–67 (noting that machine learning systems are capable of replicating works); Sag, *supra* note 5, at 309 ("[T]hese machine learning models still qualify as nonexpressive use so long as the outputs are not substantially similar to any particular original expression in the training data."); Henderson et al., *supra* note 5, at 3 ("[T]he risk of infringement is real, and fair use will not cover every scenario where a foundation model is created or used."); Lee et al., *supra* note 5, at 370 ("Some outputs from these models will incorporate copyrighted material that will be seen by humans—indeed, some generations will infringe."); Samuelson, *supra* note 5, at 1555 ("It is, of course, possible for outputs of generative AI to [infringe] derivative work rights.").

105. Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramer, Borja Balle, Daphne Ippolito & Eric Wallace, *Extracting Training Data from Diffusion Models*, ARXIV 5–6 (Jan. 30, 2023), https://arxiv.org/pdf/2301.13188.pdf.

copyrighted content in response to 44% of the prompts, Meta's Llama 2 on 10% of them, and Anthropic's Claude 2 on 8% of them.[106] Studies have also shown that significant extraction of at least some literary works is possible.[107] This shows that while infringements are potentially rare, they do exist. Infringement is especially likely with copyrighted works that are more abstractly protected, such as characters, which led Matthew Sag to dub this the "Snoopy problem."[108] Recent lawsuits by Disney and Warner Brothers pounce on this problem, identifying outputs that likely infringe on characters such as Darth Vader, Buzz Lightyear, Tweety Bird, Batman, and others.[109] A recent decision by the Munich I Regional Court relied on the substantial reproduction of song lyrics by ChatGPT.[110]

There is also the issue of prompt injection or adversarial machine learning, through which individuals can make machine learning systems ignore instructions and act in unintended ways.[111] Sag, for example, successfully provoked Midjourney and Stable Diffusion into generating likely infringing images of Peanuts' Snoopy, Disney's Mickey Mouse, and Banksy's *Girl with Balloon*.[112] Others have similarly made Stable Diffusion, GPT-4, and DALL-E 2 generate images that are analogous to Nintendo's cartoon Italian plumber Mario and comic book hero Spider-Man.[113]

While these scholars intended to make the AI systems generate infringing outputs, infringements may also occur where they are not intended. For example, in Sag's work, one of the prompts that generated an output likely infringing on Banksy's *Girl with Balloon* was "black and white stencil little girl reaching for heart-shaped red balloon."[114] It is foreseeable that others may use a similar prompt without thinking of or even being aware of Banksy's work. Nonetheless,

106. *Introducing CopyrightCatcher, the First Copyright Detection API for LLMs*, PATRONUS AI, https://www.patronus.ai/blog/introducing-copyright-catcher (last visited Apr. 12, 2026) [hereinafter, *Introducing CopyrightCatcher*].

107. *See* Cooper et al. *supra* note 100, at 17; Ahmed et al., *supra* note 100, at 3.

108. Sag, *supra* note 5, at 327–28.

109. Complaint & Demand for Jury Trial at 3, Disney Enters., Inc. v. Midjourney, Inc., No. 25-5275 (C.D. Cal. June 11, 2025); Complaint & Demand for Jury Trial at 2, Warner Bros. Ent. Inc. v. Midjourney, Inc., No. 25-8376 (C.D. Cal. Sep. 4, 2025).

110. Landgericht München I [LG] [Munich I Regional Court] Nov. 11, 2025, 42 O 14139/24 ¶ 51 (Ger.).

111. Ryan Calo, Ivan Evtimov, Earlence Fernandes, Tadayoshi Kohno & David O'Hair, *Is Tricking a Robot Hacking?*, 6–7 (Univ. Wash. Sch. L. Legal Stud. Rsch. Paper, Paper No. 2018-05), https://ssrn.com/abstract=3150530; Simon Willison, *Prompt Injection Attacks Against GPT-3*, SIMON WILLISON'S WEBLOG (Sep. 12, 2022, at 22:20 PT), https://simonwillison.net/2022/Sep/12/prompt-injection; *see also* Jane Bambauer & Tal Zarsky, *The Algorithm Game*, 94 NOTRE DAME L. REV. 1, 5–22 (2018) (describing algorithms as gameable).

112. Sag, *supra* note 5, at 330–35; *see also* Guadamuz, *supra* note 34, at 125 (remarking that AI tools such as Midjourney and Stable Diffusion often do a good job of replicating famous characters).

113. Timothy B. Lee & James Grimmelmann, *Why the New York Times Might Win Its Copyright Lawsuit Against OpenAI*, ARS TECHNICA (Feb. 20, 2024, at 09:05 PT), https://arstechnica.com/tech-policy/2024/02/why-the-new-york-times-might-win-its-copyright-lawsuit-against-openai; BJ Ard, *Copyright's Latent Space: Generative AI and the Limits of Fair Use*, 110 CORN. L. REV. 509, 543–44 (2025).

114. Sag, *supra* note 5, at 335.

the AI system may generate an infringement like the one Sag identified, and someone will be liable for it.

Therefore, generative AI raises an important but largely unexamined third question: who is liable when generative AI inevitably creates substantially similar, infringing outputs? Scholars have also posed the question of liability for AI actions before.[115] But answers have been lacking, especially in copyright.

The unpredictability of AI-generated outputs makes output liability a novel question for copyright law. The increasing size, complexity, and reliance on machine learning of AI systems is making AI outputs increasingly hard to predict.[116] A group of computer scientists and law professors suggested that machine learning researchers should ensure that AI systems create transformative outputs that would likely be fair use.[117] But even when various precautions are implemented in advance (as they should be), developers still cannot always prevent undesired outcomes.[118] It is difficult or even impossible to establish foolproof guardrails to prevent AI systems from ever generating infringing outputs.[119]

C. THREATENING COPYRIGHT'S INCENTIVES-ACCESS BALANCE

Determining who should be held liable for infringing AI-generated outputs, and under what circumstances, threatens to unmoor copyright law from its goals. Even before the founding of the United States, copyright sought to encourage the creation of new expressive works and the dissemination thereof.[120] The Framers later enshrined these goals in the United States Constitution.[121] Today, copyright gives authors of original, fixed expressive works a set of exclusive rights—including the right to copy, make derivatives, and publicly display and

115. *See, e.g.*, Balkin, *supra* note 48 ("[F]rom the standpoint of law—as opposed to the standpoint of engineering—the problem posed by emergence is the problem of assigning responsibility for the unpredictable behavior of robots and AI systems."); Beckers & Teubner, *supra* note 9, at 78 ("[I]f humans cannot understand the algorithm's decisions, who is liable if something goes wrong?"); Lemley & Casey, *supra* note 14, at 1353 ("[I]f the AI is self-learning, we may really never know who is to blame.").

116. Samuel R. Bowman, *Eight Things to Know About Large Language Models*, ARXIV (Apr. 2, 2023), https://arxiv.org/pdf/2304.00612.pdf; Ganguli et al., *supra* note 72.

117. Henderson et al., *supra* note 5, at 3.

118. Lemley & Casey, *supra* note 14, at 1335.

119. *See* Bowman, *supra* note 116 (describing how approaches such as the constitutional AI technique can "dramatically reduce *explicit* examples of [the disfavored norms or output]," but that such outputs can still emerge due to the unpredictability of what LLMs will generate); *see also* Simon Willison, *I Don't Know How to Solve Prompt Injection*, SIMON WILLISON'S WEBLOG (Sep. 16, 2022, at 16:28 PT), https://simonwillison.net/2022/Sep/16/prompt-injection-solutions ("A big problem here is provability. Language models like GPT-3 are the ultimate black boxes. It doesn't matter how many automated tests I write, I can never be 100% certain that a user won't come up with some grammatical construct I hadn't predicted that will subvert my defenses.").

120. *See* JULIE E. COHEN, LYDIA PALLAS LOREN, RUTH L. OKEDIJI & MAUREEN A. O'ROURKE, COPYRIGHT IN A GLOBAL INFORMATION ECONOMY 24 (4th ed. 2015) ("[T]he [eighteenth century] Statute of Anne was expressly meant to be, as its title stated, '[a]n act for the encouragement of learning.'").

121. U.S. CONST. art. I, § 8, cl. 8 ("Congress shall have Power . . . [t]o promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.").

perform those works—to incentivize them to create.[122] This goal is not about remuneration or rewarding "sweat of the brow" labor, but generating new expressive works.[123]

But the exclusive rights of the copyright owner have their limits. Granting unfettered rights would harm copyright's other goal, the "dissemination of [said] creative works."[124] As the Supreme Court noted fifty years ago in *Twentieth Century Music Corp. v. Aiken*, "private motivation must ultimately serve the cause of promoting broad public availability of literature, music, and the other arts."[125] The benefits of creation would be nullified if no one could access them.[126] Therefore, as the Supreme Court later noted in *Sony Corp. of America v. Universal City Studios, Inc.*, copyright law seeks "a balance between a copyright holder's legitimate demand for effective—not merely symbolic—protection of the statutory monopoly, and the rights of others freely to engage in substantially unrelated areas of commerce [including new information technologies]."[127]

Scholars across the ideological spectrum have long strived to maintain a balance between copyright's competing goals of protecting rights to incentivize new creative works and providing public access to those works. Scholars favoring access emphasized this importance when they concluded that "well-functioning copyright law balances the interests of the public and of copyright owners."[128] Those favoring rights have likewise noted how the law should "strike a happier balance between the copyright owner's, the intermediary's, and the end-user's interests (or greeds) . . . to continue to afford a hospitable environment for the creation and dissemination of works of authorship, to the ultimate enrichment of the public."[129] Maintaining this balance between

122. 17 U.S.C. §§ 102, 106; *see also* Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 477 (1984) (Blackmun, J., dissenting) (explaining that by granting copyrights, authors "are given an incentive to create").

123. *See* Feist Publ'ns, Inc. v. Rural Tel. Serv. Co., 499 U.S. 340, 349 (1991) ("The primary objective of copyright is not to reward the labor of authors, but '[t]o promote the Progress of Science and useful Arts.'"); MELVILLE B. NIMMER & DAVID NIMMER, NIMMER ON COPYRIGHT § 1.03[A] (finding that the "primary purpose of copyright [i]s not to reward authors," but to promote creation).

124. Stewart v. Abend, 495 U.S. 207, 228 (1990).

125. 422 U.S. 151, 156 (1975).

126. *See* Malla Pollack, *What Is Congress Supposed to Promote?: Defining "Progress" in Article I, Section 8, Clause 8 of the United States Constitution, or Introducing the Progress Clause*, 80 NEB. L. REV. 754, 809 (2001) (finding that "progress" in the Intellectual Property Clause of the Constitution refers to spreading knowledge and technology).

127. 464 U.S. at 442; *see also* Eleonora Rosati, *Infringing AI: Liability for AI-Generated Outputs Under International, EU, and UK Copyright Law*, 1 EUR. J. RISK REGUL. 603, 608 (2025) (noting that international copyright law also seeks a fair balance between, on the one hand, incentives and rights, and, on the other, public access).

128. Pamela Samuelson & Members of the CPP, *The Copyright Principles Project: Directions for Reform*, 25 BERKELEY TECH. L.J. 1175, 1181 (2010).

129. Jane C. Ginsburg, *How Copyright Got a Bad Name for Itself*, 26 COLUM. J.L. & ARTS 61, 64 (2002).

incentives and access is essential to achieving the Constitution's mandate to "promote the progress of knowledge and learning."[130]

New technologies pose a particularly thorny challenge to the incentives-access balance, placing an onus on courts and Congress to maintain this balance. As U.S. District Judge for the Southern District of New York Lewis Kaplan put it, new technologies lead to a "push and pull . . . over the breadth or narrowness" of copyright law.[131] Although a balance between incentives and access is universally desired, the optimal balance is difficult to ascertain.

The novel features of generative AI challenge copyright law's balance through artificial infringement. Generative AI has the almost unprecedented potential to democratize access to information and production of creative outputs.[132] However, an overbroad direct liability rule for artificial infringement could discourage robust, communication-enhancing generative AI innovation due to the resulting liability risk of up to $150,000 per infringed work.[133] But too weak a rule could allow infringements to proliferate, discouraging creation of new works. These concerns are salient. A recent U.S. Copyright Office report on the economic implications of AI and copyright policy recognized that infringement liability for outputs could challenge copyright's incentives-access balance.[134] Furthermore, lawmakers and scholars have struggled to craft new liability rules for AI-caused harms in other areas of law, including autonomous vehicle accidents and medical AI malpractice.[135]

---

130. *See* Lydia Pallas Loren, *Untangling the Web of Music Copyrights*, 53 CASE W. RSRV. L. REV. 673, 675 (2003) ("The fundamental purpose of copyright law is to promote the progress of knowledge and learning. Thus, examining the reasons to provide a preference for the author over the public becomes critical.").

131. Lewis A. Kaplan, Walter A. Effross & Joyce E. Taber, *Keynote Address: Resolving Tensions Between Copyright and the Internet*, 50 AM. U.L. REV. 409, 419 (2000).

132. *See, e.g.*, Dara Tafazoli, *Exploring the Potential of Generative AI in Democratizing English Language Educatio*n, COMPUTS. & EDUC.: A.I., Dec. 2024, at 1, 9 (finding that generative AI can help democratize English language learning by mitigating resource constraints).

133. 17 U.S.C. § 504(c).

134. U.S. COPYRIGHT OFF., IDENTIFYING THE ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE FOR COPYRIGHT POLICY: CONTEXT AND DIRECTION FOR ECONOMIC RESEARCH 16 (Brent Lutes ed., 2025). The report did not address *who* should be held liable; instead, it discussed market value vis-à-vis AI-created infringements. *Id.* at 17–22.

135. *See, e.g.*, Abraham & Rabin, *supra* note 9, at 147–60 (proposing a manufacturer-financed strict liability compensation regime referred to as "Manufacturer Enterprise Responsibility" in the context of automated vehicle accidents); David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 148–49 (2014) (proposing common enterprise liability in which the manufacturer of the vehicle and makers of the component parts would be jointly and severally liable); F. Patrick Hubbard, *"Sophisticated Robots": Balancing Liability, Regulation, and Innovation*, 66 FLA. L. REV. 1803, 1866–67 (2014) (describing the creation of a compensation fund by automobile distributors); Price II & Cohen, *supra* note 9 (arguing that existing tort doctrines are inappropriate for medical AI-based liability and advocating for enterprise liability to hold hospitals liable when they are able to adapt and monitor, and developers in all other cases); Sara Gerke, *Health AI for Good Rather Than Evil? The Need for a New Regulatory Framework for AI-Based Medical Devices*, 20 YALE J. HEALTH POL'Y, L., & ETHICS 433, 480 (2021) (proposing a new FDA regulatory framework for AI-driven medical devices, including more stringent pre-market approvals, black box-specific requirements, and monitoring and improvement obligations); Charlotte A. Tschider, *Medical Device Artificial Intelligence: The New Tort Frontier*, 46 BYU L. REV. 1551, 1615 (2021) (doubting the FDA's ability

Copyright law uses a toolkit of imperfect proxies and polestars to tease out what threatens and fosters the incentives-access balance in order to maintain it. A prominent example is fair use, which is often considered the most important exception to copyright infringement.[136] The literature is replete with analyses of how fair use balances copyright incentives and access.[137] As explained above, fair use will have an important role in balancing copyright law at the training stage (and possibly output stage) for generative AI.[138] However, fair use cannot solve all the risks related to liability for infringing AI-generated outputs.

While fair use has been called the "*de facto* regulator of new media technologies,"[139] it is not the lone tool courts can use. Fair use cannot maintain the delicate incentives-access balance in copyright law by itself.[140] Other avenues are needed to permit innovation and maintain the incentives-access balance.

Congress and courts have maintained copyright's incentives-access balance in new technology cases primarily through the use of one of five doctrines: (1) fair use, (2) bespoke exceptions, (3) licensing, (4) direct liability, and (5) secondary liability. Recent scholarship has sought to explore these other approaches. For example, prior literature has excavated the role of licensing in facilitating modern music business models.[141] In my prior work, I sought to explain how secondary liability refinements guided by considerations of intent facilitated (and could continue facilitating) new information technologies while maintaining copyright's incentives-access balance.[142]

The largely missing approach to balancing these competing goals in copyright discourse is direct liability. Herein lies the answer to liability for AI-generated copyright infringement. While courts have continued to refine aspects of the prima facie copyright infringement claim,[143] the primary concern has not

---

to effectively regulate black-box health AI systems); Duffourc, *supra* note 82, at 20 (proposing a control-centered framework and limited legal personhood for Autonomous AI Physicians).

136. 17 U.S.C. § 107.

137. *See, e.g.*, Edward Lee, *Technological Fair Use*, 83 S. CAL. L. REV. 797, 800–02 (2010) (describing how the question of technological fair use has a significant impact on technological development); Jacob Victor, *Utility-Expanding Fair Use*, 105 MINN. L. REV. 1887, 1900–01 (2021) (emphasizing the importance of fair use for permitting utility-expanding technologies such as search engines, plagiarism checkers, and book digitization); Pamela Samuelson, *Unbundling Fair Uses*, 77 FORDHAM L. REV. 2537, 2540 (2009) (explaining that one of the primary benefits of the fair use doctrine is balancing exclusive rights and access and follow-on uses of copyrighted works); Edward Lee, *Fair Use Avoidance in Music Cases*, 59 B.C. L. REV. 1873, 1878 (2018) (describing fair use as one of the "safety valve[s]" for maintaining copyright's balance); Sag, *supra* note 93, at 1615.

138. *See supra* Subpart.I.B.

139. Ard, *supra* note 113, at 513.

140. Gideon Parchomovsky & Philip J. Weiser, *Beyond Fair Use*, 96 CORN. L. REV. 91, 95 (2010).

141. *See* Jacob Victor, *Reconceptualizing Compulsory Copyright Licenses*, 72 STAN. L. REV. 915, 921–22 (2020) (describing the development of compulsory music licenses in response to new technologies).

142. Michael P. Goodyear, *Infringing Information Architectures*, 58 U.C. DAVIS L. REV. 1959, 1975–79 (2025).

143. *See, e.g.*, Skidmore v. Led Zeppelin, 952 F.3d 1051, 1064 (9th Cir. 2020) (clarifying the Ninth Circuit's extrinsic-intrinsic test for determining substantial similarity).

necessarily been optimizing balance in the face of technological change. A prominent exception to that rule is the introduction of the concept of "volition or causation" into direct liability, as will be discussed in depth below.[144] While several scholars have analyzed this requirement, they do not discuss how this could fit into an overall approach for accommodating new technologies while also maintaining copyright's incentives-access balance.[145] This Article seeks to fill this lacuna by first understanding how direct copyright infringement liability responded to technology, and then using that historical approach to propose a balance-maintaining liability framework for AI-generated infringing outputs.

II. THE EVOLUTION OF COPYRIGHT DIRECT LIABILITY

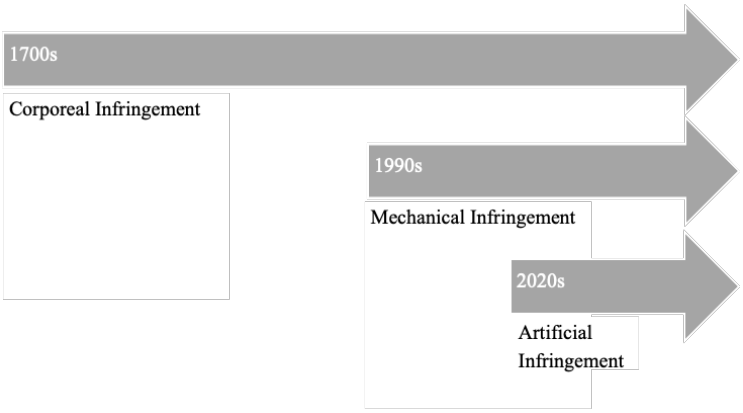
This Part examines two distinct stages in the evolution of direct copyright infringement claims, which this Article terms corporeal infringement and mechanical infringement, as shown in Figure 3 below. For much of U.S. history, no one was all that concerned with who would be held directly liable for copyright infringement if it occurred. The direct infringer—if there was one—was the corporeal individual who physically carried out the infringement. But the advent of more complex technologies with automatic functions and multiple actors complicated the corporeal infringement standard, leading to the development of a new paradigm, premised on causation, for mechanical infringement. In response to these new challenges to the incentives-access balance and innovation, courts applied a bespoke analytical framework to hold liable only those who had acted with "volition or causation." In copyright law, this requirement has primarily focused on causation rather than willfulness-like volition.

Generative AI presents distinct questions from these mechanical tools of the past few decades. Nonetheless, the history of mechanical infringement is important for elucidating how direct copyright infringement liability has been refined to accommodate new technologies and could similarly evolve in the future. Building on this history, Part III will examine the lessons from this evolution, and Part IV will apply them to the emerging third stage of liability involving generative AI.

144. *See infra* Subpart.II.B.

145. *See, e.g.*, Robert C. Denicola, *Volition and Copyright Infringement*, 37 CARDOZO L. REV. 1259, 1276 (2016) (describing the absence of volition when a service provider could only prevent an infringement by policing its customers' conduct); David Nimmer, *Volition in Violation of Copyright*, 43 COLUM. J.L. & ARTS 1, 6 (2019) (describing volition as causation, where liability does not lie if the chain of causation is too attenuated); Chatterjee & Fromer, *supra* note 10, at 1889 (challenging the notion that machines do not have sufficient mental states to achieve volition).

FIGURE 3: THE EVOLUTION OF DIRECT INFRINGEMENT



A. CORPOREAL INFRINGEMENT

This Article coins the term corporeal infringement to describe the original, quintessential type of infringement that is carried out by the human hand. While this term was never used by the courts or scholars, it was the reality for much of U.S. copyright history prior to the development of more complex machines. An individual reproducing, making a derivative of, distributing, publicly or digitally performing, or publicly displaying another's copyrighted work by their own hand, without an exception or defense, is liable for copyright infringement.[146] For example, a person uploading an infringing song to YouTube would be corporeal infringement.

To establish a prima facie claim for direct copyright infringement, a copyright owner must show copying in fact and improper appropriation.[147] Copying in fact requires that the defendant's work is substantially similar and the defendant had access to the plaintiff's work.[148] Improper appropriation requires that sufficient protected expression was copied from the original.[149] Corporeal infringement decisions were not concerned with *who* copied because that was self-evident.

All cases until the last three decades (and even after that, the vast majority) involved corporeal infringement. I provide but a few examples of well-known cases here. First, let us consider *Nichols v. Universal Pictures Corp.*, which is often taught as a paradigmatic example of the infringement analysis and how to

146. 17 U.S.C. § 501.

147. Rentmeester v. Nike, Inc., 883 F.3d 1111, 1117 (9th Cir. 2018).

148. *See* Boisson v. Banian, Ltd., 273 F.3d 262, 267–68 (2d Cir. 2001) (explaining how access and "similarities that are probative of copying" can show copyright infringement).

149. *Rentmeester*, 883 F.3d at 1117.

dissect non-protectable elements such as ideas and commonplace themes. In that case, the plaintiff playwright accused the defendant of making a motion picture that infringed his copyrighted play, "Abie's Irish Rose."[150] The case turned on whether the similarities between the two works—the children of staunchly Jewish and Irish Catholic families falling in love, their marriage, and the birth of grandchildren and subsequent reconciliation of the families—were copyrightable expression or merely unprotectable ideas and scènes à faire.[151] But the court was never in doubt that if anyone had infringed, it was the defendant, because the defendant was the one that produced the film.[152]

Similarly, in *Three Boys Music Corp. v. Bolton*, often the paradigmatic example of subconscious copying, the question was not the identity of the direct infringer, but whether he had access and whether his song was substantially similar.[153] The court determined that there was sufficient evidence that singer-songwriter Michael Bolton had subconsciously copied the Isley Brothers' 1960s hit "Love Is a Wonderful Thing" in his own song of the same name.[154]

Finally, in *Steinberg v. Columbia Pictures Industries, Inc.*, the court faced the question of whether the producers, promoters, distributors, and advertisers of the movie "Moscow on the Hudson" infringed the plaintiff's New Yorker magazine cartoons and illustrations by creating a promotional poster for their film that was similar in style.[155] The court focused on whether the font and respective illustrations of Manhattan in the foreground with the rest of the world and Moscow in the background were substantially similar.[156] The fact that there were many defendants involved in the copying did not affect the analysis because they had each unambiguously acted to create or distribute the poster.[157] The court held that the works were substantially similar and that the fair use and laches defenses did not apply, so the defendants had infringed.[158] Again, it was never in doubt that the defendants would be the direct infringers.

New technology was also not necessarily a bar for a straightforward application of corporeal infringement. In *Computer Associates International, Inc. v. Altai, Inc.*, the court considered whether the elements Altai copied from the plaintiff's computer program into its own OSCAR 3.5 program were copyrightable, ultimately determining that they were not.[159] Despite the new technology, there was no dispute that Altai employees had written the OSCAR 3.5 code and would be the direct infringers if infringement occurred.[160]

150. Nichols v. Universal Pictures Corp., 45 F.2d 119, 120 (2d Cir. 1930).
151. *Id.* at 121–22.
152. *Id.* at 120.
153. 212 F.3d 477, 481 (9th Cir. 2000).
154. *Id.* at 482–86.
155. 663 F. Supp. 706, 708 (S.D.N.Y. 1987).
156. *Id.* at 711.
157. *Id.* at 708.
158. *Id.* at 713–16.
159. 982 F.2d 693, 702–03, 715 (2d Cir. 1992).
160. *See id.* at 699–700.

B. MECHANICAL INFRINGEMENT

The rules of corporeal infringement existed, largely unencumbered, for decades, but technological developments put them under pressure. The advent of more complicated machines caused uncertainty to emerge in the corporeal infringement test. Copying occurred, but the identity of the copier was less clear. When many actors were involved in the operation of a facially neutral machine that produced infringement, who really caused the infringement to occur?

To determine the correct answer and avoid stymying innovation, courts elucidated an additional requirement in the doctrine: "volition or causation," which courts often condense into simply "volition."[161] The requirement was muddled in the initial decision that introduced it, although that case's analysis focused on actions rather than mental state.[162] Scholars and courts have subsequently suggested the term means willfulness, causation, or both.[163] However, causation seems to be driving the analysis for the increasingly complex machines of this second era of direct liability, which I term "mechanical" infringement. This Article therefore focuses on the causation half of the "volition or causation" requirement.

The "volition or causation" requirement is absent from the Copyright Act but emerged in one of the early cases involving infringement liability over the Internet.[164] In *Religious Technology Center v. Netcom On-Line Communication Services, Inc.*, the plaintiffs sued Dennis Erlich, a former Scientology minister, for publishing the writings of Scientology founder L. Ron Hubbard, along with Erlich's critiques, in an online discussion forum.[165] The plaintiffs sued not just Erlich, but also the operator of the online bulletin board that provided Erlich with access to the forum, and Netcom, which provided Internet access to the bulletin board.[166]

---

161. A recent district court decision explicitly noted this relationship, saying that causation is "also referred to as 'volitional conduct.'" Robinson v. Binello, 771 F. Supp. 3d 1114, 1122 (N.D. Cal. 2025).

162. Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1370 (N.D. Cal. 1995).

163. Chatterjee & Fromer, *supra* note 10, at 1890–91, 1901, 1902 n.81; Hunley v. Instagram, LLC, 73 F.4th 1060, 1074–75 (9th Cir. 2023); Denicola, *supra* note 145, at 1268–69; CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544, 550 (4th Cir. 2004) ("There must be actual infringing conduct with a nexus sufficiently close and causal to the illegal copying that one could conclude that the machine owner himself trespassed on the exclusive domain of the copyright owner."); Dmitry Karshtedt, *Causal Responsibility and Patent Infringement*, 70 VAND. L. REV. 565, 642 (2017) ("Nonetheless, *Netcom* specifically alluded to the idea that some threshold causal contribution, which the court also characterized as 'volition,' must be made before a defendant could be held directly liable."); Patrick R. Goold, *Is Copyright Infringement a Strict Liability Tort?*, 30 BERKELEY TECH. L.J. 305, 315 (2015) (describing volition as acting voluntarily, i.e., with a requisite mental state); *see also* NIMMER & NIMMER, *supra* note 123, § 13E.09[B][1] (describing volition as an alternative to causation); BWP Media USA Inc. v. Polyvore, Inc., 922 F.3d 42, 62 (2d Cir. 2019) (Newman, J., concurring) ("'Volition' in Judge Whyte's phrase is best understood to mean a concept essentially reflecting tort law causation."); *BWP Media USA*, 922 F.3d at 47 (Walker, J., concurring) (explaining that volition requires "some affirmative . . . step to infringe").

164. *See* NIMMER & NIMMER, *supra* note 123, § 13E.09.

165. *Netcom*, 907 F. Supp. at 1365–66.

166. *Id.*

Judge Whyte acknowledged that (direct) copyright infringement was largely strict liability, limiting his options, but surmised there must be a limit to who could be held liable in order to maintain copyright's incentives-access balance.[167] Otherwise, copyright infringement claims would disincentivize innovation by aggrandizing the bounds of copyright to hold liable tangentially related entities such as Internet service providers that undertook no conduct other than initial setup for all users.[168] He thought that "there should still be some element of *volition or causation* which is lacking where a defendant's system is merely used to create a copy by a third party . . . without any human intervention."[169] He used this concept of "volition or causation" to evaluate the connection between the defendant's actions and the infringement, not the harm.[170] He held that Netcom was not directly liable because its actions were only "automatic and indiscriminate"—the users were the ones who had caused the infringing copies to be made.[171] Netcom could still potentially be held liable under secondary liability, but that required a separate legal analysis.[172]

The volition or causation requirement serves as a rule that creation, supply, or ownership of a machine, by itself, does not make one directly liable for subsequent infringements. "[S]omething more"—volition or causation—is necessary for direct liability.[173] This requirement makes intuitive sense, as it would be undesirable to hold technology providers liable for all infringements resulting from their technologies that are capable of non-infringing, value-adding uses.[174] In the contributory copyright infringement context, the Supreme Court's *Sony* rule performs a similar role by shielding technology providers from per se liability where their technologies are capable of substantial non-infringing uses.[175]

Since *Netcom*, "volition or causation" has become an established feature of copyright law that has helped maintain the incentives-access balance. However, due to how Judge Whyte phrased it, subsequent courts seem to have used the term "volition" to refer to both volition *and* causation. This led to several courts referring to "volition" when they seem to have meant causation. The terms should be understood as separate concepts.[176]

---

167. *Id.* at 1369–70.

168. *See id.*

169. *Id.* (emphasis added).

170. *Id.* at 1370–72.

171. *Id.* at 1372.

172. *Id.* at 1373, 1375.

173. CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544, 550 (4th Cir. 2004).

174. While providers of anti-circumvention technologies can be held liable for infringements, that is because the technology is primarily for that illegal purpose rather than being a general purpose technology with substantial non-infringing uses. 17 U.S.C. § 1201(a)(2).

175. Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984). For a discussion on how the *Sony* rule developed, see Goodyear, *supra* note 142, at 1979–83.

176. *See* BWP Media USA Inc. v. Polyvore, Inc., 922 F.3d 42, 52 (2d Cir. 2019) (Walker, J., concurring) (referring to *Netcom* and concluding that "volition and causation are different concepts").

For mechanical infringement, courts such as that in *Netcom* had to carefully parse causation where users acted and systems automatically responded, ultimately holding the users, not the technology companies, directly liable.[177] For example, in *CoStar Group, Inc. v. LoopNet, Inc.*, the Fourth Circuit, heavily citing *Netcom*, compared an automatic online function to a copy machine, finding that "an [Internet service provider] who owns an electronic facility that responds automatically to users' input is not a direct infringer."[178] In *Parker v. Google, Inc.*, the Third Circuit, relying on *CoStar,* upheld a finding that "Google's automatic archiving of USENET postings and excerpting of websites in its results to users' search queries do not include the necessary volitional [or causal] element to constitute direct copyright infringement."[179] Similarly, in *Field v. Google Inc.,* the District of Nevada held that "when a user requests a Web page contained in the Google cache by clicking on a 'Cached' link, it is the user, not Google, who creates and downloads a copy of the cached Web page."[180] Therefore, Google's "automated, non-volitional [(i.e., non-causational)] conduct . . . in response to a user's request does not constitute direct infringement."[181] In *CoStar*, *Parker*, and *Field*, mere automatic functions were not enough—the direct infringer had to press the button that caused the creation of a copy.

Although not every case after *Netcom* addressed causation, the canonical cases of *Cablevision* and *Aereo* further developed that requirement. In *Cartoon Network, LP, LLLP v. CSC Holdings, Inc.* ("*Cablevision*"), Cartoon Network sued the defendant cable company for allegedly exceeding its broadcast license and directly infringing copyrighted works by offering its viewers a remote-storage digital video recorder.[182] Judge Walker reasoned that while Cablevision had designed and maintained the system, the user of the system selects and directs the actual reproduction of the specific programming.[183] Therefore, the court held that the direct infringer was "the person who actually presses the button to make the recording."[184] Like *Netcom*, *Cablevision* protected the defendant who offered a facially neutral product from being held liable for direct infringement. Courts used causation to distinguish between man and machine and balance rightsholders' rights against new technologies.

Six years later, "volition or causation"—the core of the mechanical infringement paradigm—was put at risk before the Supreme Court in *American Broadcasting Cos. v. Aereo, Inc.* The Supreme Court—without mentioning volition or causation—held that the company operating a series of streaming

177. Denicola, *supra* note 145, at 1272.
178. 373 F.3d at 550.
179. 422 F. Supp. 2d 492, 497 (E.D. Pa. 2006), *aff'd*, 242 F. App'x 833, 836 (3d Cir. 2007).
180. 412 F. Supp. 2d 1106, 1115 (D. Nev. 2006).
181. *Id.*
182. 536 F.3d 121, 124 (2d Cir. 2008).
183. *Id.* at 131–32.
184. *Id.* at 131.

antennae performed the broadcasts, even though the subscribers chose which broadcasts to stream.[185] Below, the Second Circuit had decided *Aereo* based on *Cablevision*, holding that the direct infringers were the users who pushed the button that caused the content to play.[186] But the Supreme Court disagreed. Justice Breyer wrote that "both the broadcaster and the viewer of a television program 'perform,' because they both show the program's images and make audible the program's sounds."[187] He likened Aereo to cable television, which Congress explicitly addressed in the Copyright Act of 1976 (well before mechanical infringement emerged).[188] Rather than rely on volition or causation, Justice Breyer determined that the blatant co-opting of copyrighted content by Aereo—like cable television—needed to be contained so as to not unduly support access to copyrighted works over rightsholders' rights.[189]

In his dissent, Justice Scalia criticized Justice Breyer for not considering volition (or, implicitly, causation) as a prerequisite for infringement.[190] He, like the *Cablevision* court, compared the service at hand to a photocopier in a copy shop, where the customer selects the content and presses "copy" on the machine.[191] Therefore, in Justice Scalia's opinion, Aereo should not have been directly liable (although it could have been secondarily liable).[192]

Aereo left many questions unanswered since the majority did not address volition or causation. Some questioned whether *Aereo* had even abrogated the requirement entirely.[193] But there are key differences that distinguished *Aereo* from its predecessors. Unlike in *Cablevision*, the service in *Aereo* was *replacing* rather than merely complementing the existing content access system.[194] The remote-storage DVR in *Cablevision* split the licensed stream in two, allowing consumers to continue watching programming they had paid to watch while also allowing them to record programming for later use.[195] But the remote antennae system in *Aereo* supplanted broadcast television by offering its own unlicensed access instead.[196] Even if the user selected the specific infringement, the system was set up in a manner that infringement was effectively guaranteed regardless of what the user selected. Furthermore, even if *Aereo* did abrogate volition or causation, the case was limited to the public performance right, so those

185. 573 U.S. 442–43, 451 (2014).

186. WNET, Thirteen v. Aereo, Inc., 712 F.3d 676, 692 (2d Cir. 2013).

187. *Aereo*, 573 U.S. at 441 (emphasis removed).

188. *Id.* at 441–44.

189. *Id.* at 446–47.

190. *Id.* at 452–53 (Scalia, J., dissenting).

191. *Id.* at 456.

192. *Id.* at 457.

193. *See, e.g.*, Nimmer, *supra* note 145, at 36; Kyle A. Brown, *Up in the Aereo: Did the Supreme Court Just Eliminate the Volitional Conduct Requirement for Direct Copyright Infringement?*, 46 SETON HALL L. REV. 243, 275–78 (2015).

194. *Compare* Cartoon Network LP, LLLP v. CSC Holdings, Inc., 536 F.3d 121, 131 (2d Cir. 2008), *with Aereo*, 573 U.S. at 442.

195. *Cablevision*, 536 F.3d at 124–25.

196. *Aereo*, 573 U.S. at 436.

requirements would seem to still apply to other rights, including reproduction and derivative works.[197]

Regardless, subsequent case law has reaffirmed the existence of the volition or causation requirement. The Second Circuit explicitly required volition that "causes" the infringement as recently as 2022.[198] Other courts have also ruled that the volition or causation requirement was consistent with *Aereo*.[199] For example, as recently as 2023, the Ninth Circuit reiterated that the volition or causation requirement is still valid, emphasizing that it is essential for drawing a distinction between direct and secondary infringement.[200]

III. THE CAUSATION POLESTAR

Despite courts' frequent use of volition and causation as interchangeable concepts, causation appears to be at the heart of the mechanical infringement refinement. The cases discussed above were about actions leading to infringement, not mental state. Causation has become a guide for protecting socially beneficial innovators while holding the users who undertook the infringing action liable. If merely providing the system by which infringements (and noninfringing uses) could be carried out were sufficient for direct liability, it would create expansive copyright that could stymie technological innovation. But assigning no liability even when infringement occurs would shift copyright too far the other way, permitting socially detrimental uses to deter new creativity. Trapped between Scylla and Charybdis, the elucidation of causation was far from guaranteed as the approach to mechanical infringement.

Direct liability offers courts and Congress fairly limited maneuverability compared to fair use, bespoke exceptions, licensing, and secondary liability. Fair use has four flexible, non-exclusive factors.[201] Exceptions and licensing could

197. *Cf.* Rebecca Giblin & Jane C. Ginsburg, *We (Still) Need to Talk About* Aereo*: New Controversies and Unresolved Questions After the Supreme Court's Decision*, 38 COLUM. J.L. & ARTS 109, 127, 152 (2015) ("While *Aereo* addressed only the public performance right, the majority's opinion raises questions about the continued application of the 'volition' standard to the reproduction right. . . . [However,] any impulse to disregard 'volition' in determining who commits the copyright-triggering act should be resisted."). *But see* Nimmer, *supra* note 145, at 19 ("There is no longer any room to vindicate the stance that a service provider's automatic response to user requests places it outside of liability.").

198. ABKCO Music, Inc. v. Sagan, 50 F.4th 309, 321–22 (2d Cir. 2022). *Compare* BWP Media USA Inc. v. Polyvore, Inc., 922 F.3d 42, 49 (2d Cir. 2019) (Walker, J., concurring) ("*Aereo* did nothing to disturb *Cablevision*'s volitional conduct requirement and that requirement continues to apply to cases."), *with id.* at 65 (Newman, J., concurring) ("But the status of systems that automatically caused an infringement at a customer's command remained uncertain because the Supreme Court's majority opinion said nothing about volition or causation.").

199. *See, e.g.*, Perfect 10, Inc. v. Giganews, Inc., 847 F.3d 657, 666–67 (9th Cir. 2017); BWP Media USA, Inc. v. T & S Software Assocs., Inc., 852 F.3d 436, 442 (5th Cir. 2017); Millennium Funding, Inc. v. Priv. Internet Access, Inc., No. 21-CV-01261, 2022 WL 7560395, at *13 (D. Colo. Oct. 13, 2022); Okolita v. Amazon.com, Inc., No. 22-CV-00284, 2023 WL 4847259, at *9 (D. Me. July 28, 2023).

200. Hunley v. Instagram, LLC, 73 F.4th 1060, 1074–75 (9th Cir. 2023).

201. 17 U.S.C.§ 107; *see* Matthew Sag, *God in the Machine: A New Structural Analysis of Copyright's Fair Use Doctrine*, 11 MICH. TELECOMMS. & TECH. L. REV. 381, 397 (2005) ("Congress replaced potentially limited and technologically specific rights with rights that were more broadly expressed, in order to allow copyright law

theoretically be imposed on any arrangement.[202] Secondary liability's more open-ended common law elements such as "control," "knowledge," and "material contribution" permit courts to refine their meaning in response to new technologies while considering intent.[203] But direct liability is fairly self-contained. Copying in fact and improper appropriation, in the form of a substantially similar work, results in copyright infringement.[204]

Despite these limitations, the causation requirement was a promising avenue for maintaining the incentives-access balance because it was not really an innovation as much as it was elucidating a forgotten part of copyright law. Tort law recognizes two types of causation: but-for and proximate. But-for causation means that but for the defendant's action, the injury would not have occurred.[205] Proximate causation means the defendant's conduct was sufficiently related to the plaintiff's injury to attach.[206] As the leading copyright treatise has explained, "*Netcom* simply stands for the unremarkable proposition that proximate causation historically underlines copyright infringement, same as with other torts."[207] To commit any tort, the defendant has to have taken some affirmative step. But proximate causation refers to the defendant's unlawful conduct being close enough to the plaintiff's injury for liability to attach.[208]

In addition, causation helps actualize the policy goal of the copying requirement under copyright law. This requirement is not about who caused the harm, but rather who caused the infringing copying.[209] The purpose of the copying requirement is preventing actual copying. This is why the Copyright Act, unlike the Patent Act, does not forbid the independent creation of an identical work, only the *copying* of a work.[210] Causation helps identify the actual copier by distinguishing tangential contributions to copying such as providing a service from actually pushing the button or causing the specific infringement to

to be more flexible in its treatment of new technologies."); *see also* Princeton Univ. Press v. Mich. Document Servs., Inc., 99 F.3d 1381, 1392 (6th Cir. 1996) ("Fair use [is] so flexible as virtually to defy definition." (quoting Time Inc. v. Bernard Geis Assocs., 293 F. Supp. 130, 144 (S.D.N.Y. 1968))).

202. *See* Jane C. Ginsburg, *Copyright and Control over New Technologies of Dissemination*, 101 COLUM. L. REV. 1613, 1620 (2001) ("Congress has often imposed a compromise, allowing continued exploitation of the technology, but with remuneration to the copyright owners . . . .").

203. Goodyear, *supra* note 142, at 1977.

204. Rentmeester v. Nike, Inc., 833 F.3d 1111, 1117 (9th Cir. 2018); *see* 17 U.S.C. § 106 (enumerating the exclusive rights of the copyright holder); Boisson v. Banian, Ltd., 273 F.3d 262, 267 (2d Cir. 2001) (explaining how access and "similarities that are probative of copying" can show copyright infringement).

205. Comcast Corp. v. Nat'l Ass'n of Afr. Am.-Owned Media, 140 S. Ct. 1009, 1014 (2020).

206. Lexmark Int'l, Inc. v. Static Control Components, Inc., 572 U.S. 118, 133 (2014).

207. NIMMER & NIMMER, *supra* note 123, § 13E.09.

208. *See, e.g.*, *Lexmark*, 572 U.S. at 133 ("[T]he proximate-cause requirement generally bars suits for alleged harm that is 'too remote' from the defendant's unlawful conduct.").

209. Denicola, *supra* note 145, at 1269.

210. Ty, Inc. v. GMA Accessories, Inc., 132 F.3d 1167, 1169 (7th Cir. 1997).

occur.[211] While these contributions could still qualify for secondary liability, that must be determined separately from the direct infringement analysis.[212]

Beyond copyright, causation helps actualize the goal of deterrence under tort law. According to the law and economics literature, deterrence is one prominent aim of tort law.[213] Deterrence requires a person to consider the costs and benefits of their contemplated actions before making them.[214] Liability—whether strict or otherwise—is a cost. As Guido Calabresi explains, if one could be held liable for injuries, product defects, and other harms, that should increase correlated investments in safety to avoid liability.[215] In other words, these liability structures deter harmful actions and encourage safe conduct by providing incentives for taking precautionary measures to reduce risk and the chance of injury.[216] Deterrence is a guiding principle for both fault-based and strict liability torts.[217]

Causation contributes to proper deterrence by not holding liable every party associated with an infringement. The party that presses the button to cause the infringing copying to occur is held liable, but merely providing a machine is insufficient for direct liability to attach.[218] The button-pushing copier can be deterred because they know the outcome of their action is infringement. But holding the developer of a product that could infringe in response to user action liable could risk deterring the offering of the product or force the market to effectively pay a copyright tax. This fear of finding the system to be generally infringing contributed to Judge Whyte's direct liability decision in *Netcom* and the Supreme Court's earlier *Sony* decision on secondary liability.[219]

Related to deterrence is the concept of foreseeability. At common law, courts use foreseeability to determine "the point beyond which the possibility of . . . liability . . . is unlikely to [influence the defendant's] ex ante behavior."[220] Foreseeability helps achieve the goal of deterrence by only extending liability to foreseeable acts for which the actor was responsible.[221]

---

211. Cartoon Network LP, LLLP v. CSC Holdings, Inc., 536 F.3d 121, 131 (2d Cir. 2008).

212. *See* Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1370 (N.D. Cal. 1995) (holding that while providing services "is not relevant to [Defendant's] *direct* liability for copying[,] . . . [this] may be relevant to contributory infringement").

213. STEVEN SHAVELL, ECONOMIC ANALYSIS OF ACCIDENT LAW 297 (1987).

214. Thomas C. Galligan, Jr., *Deterrence: The Legitimate Function of the Public Tort*, 58 WASH. & LEE L. REV. 1019, 1032 (2001). *See generally* GUIDO CALABRESI, THE COSTS OF ACCIDENTS: A LEGAL AND ECONOMIC ANALYSIS (1970) (proposing the idea that an individual should internalize all potential costs of an action).

215. CALABRESI, *supra* note 214, at 73 (describing how liability encourages safer actions).

216. RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL AND EMOTIONAL HARM § 6 cmt. d (AM. L. INST. 2010).

217. Marchisio, *supra* note 47, at 3.

218. *See supra* Subpart.II.B.

219. Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1372 (N.D. Cal. 1995); Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

220. Shyamkrishna Balganesh, *Foreseeability and Copyright Incentives*, 122 HARV. L. REV. 1569, 1591 (2009).

221. *See* Banks McDowell, *Foreseeability in Contract and Tort: The Problems of Responsibility and Remoteness*, 36 CASE W. RSRV. L. REV. 286, 299 (1985) ("Unless the factors of rational expectations, fault—

Deterrence can therefore only be achieved, and liability can only occur, where actors are aware that specific potential harms could transpire.[222] This is true even for strict liability torts, such as products liability, where liability is limited to injuries that were a foreseeable result of failing to employ an alternative design or warning.[223] Actors can attempt to internalize future harms, but if the future harms are murky, actors may not be properly deterred or could even be over-deterred, hampering socially beneficial uses. If the results of an act are unknown, it is difficult, if not impossible, to properly consider the risk and act accordingly to avoid it. Although it can be rather abstract, foreseeability accordingly limits liability to where there is a reasonable chance of the actions—and their costs and benefits—being anticipated in advance.[224]

Foreseeability also underlies some substantive requirements, including causation. In torts, causation usually "requires a close connection between the act and the injury."[225] For example, in negligence law, courts do not find causation where the injury that occurred was not foreseeable at the time of the defendant's actions.[226] Looking to expected behavioral modification through the lens of causation restricts foreseeability from extending to abstract and normatively unsupported bounds.

In copyright law, causation limits liability to the party that was most able to foresee and prevent the infringing copying. As we saw above, several courts declined to impose liability on the developer of a system with noninfringing uses, choosing instead to hold the user directly liable because they could foresee the specific infringement they caused to occur.[227] Providers of Internet services or devices may suspect that their users could use their services or products to infringe, but they cannot predict which infringements will be carried out or under what circumstances. Holding them liable for merely providing the infrastructure by which infringements are carried out would deter these services and products, harming infringing as well as noninfringing uses. The Supreme Court recognized the harm of overbroad liability for product and service providers in the contributory liability context in *Sony*, where it held that a provider does not have constructive knowledge of infringement merely because the product could be used for infringement.[228] Likewise, with direct liability, the goal is to deter specific foreseeable infringements to productively cabin liability. Causation does for direct liability what *Sony* did for secondary liability. It puts the onus on

whether legal or moral—and normative goals are to be disregarded in determining civil liability, foreseeability must remain a central inquiry.").

222. Mark A. Geistfeld, *Hidden in Plain Sight: The Normative Source of Modern Tort Law*, 91 N.Y.U. L. REV. 1517, 1553 (2016).

223. RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIAB. § 2(b)–(c) (AM. L. INST. 1998).

224. Balganesh, *supra* note 220, at 1592.

225. *In re* Boeing 737 MAX Pilots Litig., 638 F. Supp. 3d 838, 853 (N.D. Ill. 2022).

226. Balganesh, *supra* note 220, at 1595.

227. *See supra* Subpart.II.B.

228. Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

the main, and often final, actor that causes the infringement to occur to pull back from the breach and forbear from infringing.

Causation also benefits the goals of tort by providing more predictability for society. Not only can overbroad liability constrain innovation, but, as prior research has shown, uncertain liability can too.[229] This is often a risk for emerging technologies, where prior legal regimes may be inexact fits, resulting in more hypothetical legal advice until new rules finally coalesce.[230] It is therefore beneficial to identify workable liability systems for new technologies to limit this chilling effect.

There are, of course, other goals for tort law too.[231] For example, if an actor is found liable for committing a tort, there is the desire to make the victim whole by compensating him for the injury.[232] But while compensating the victim is important, compensation should not run roughshod over other limitations on liability. Fixation on only one theory of torts obscures what the law should achieve.[233]

That said, causation does not prevent the law from compensating copyright owners. The users of the devices or services who infringe can be held directly liable for proximately causing the infringement to occur. More tangentially connected parties to the infringement (such as service providers) may be held secondarily liable too.[234]

At the same time, caution about overextending liability is warranted. In addition to upsetting the incentives-access balance in copyright, overbroad liability could also constrain technological innovation and competition, potentially harming society's interests. For example, the high amount of damages available under the Copyright Act per infringed work (up to $150,000) could deter or destroy new and less-resourced market entrants, constraining the

229. *See, e.g.*, Eduardo Magrani, *New Perspectives on Ethics and the Laws of Artificial Intelligence*, INTERNET POL'Y REV., Sep. 13, 2019, at 1, 12 (warning that the risk of liability can create a chilling effect on technological innovations in autonomous AI); M. Ryan Calo, *Open Robotics*, 70 MD. L. REV. 571, 576 (2011) (describing how legal uncertainty about liability for harms caused by robots could discourage investment in robotics); Nicholas A. Ashford & Robert F. Stone, *Liability, Innovation, and Safety in the Chemical Industry*, *in* THE LIABILITY MAZE: THE IMPACT OF LIABILITY LAW ON SAFETY AND INNOVATION 367, 377 (Peter W. Huber & Robert E. Litan eds., 1991) (describing how tort liability can discourage innovation).

230. *See* Goodyear, *supra* note 142, at 1979–96 (describing how new rules for infringement liability for historical information architectures developed in response to litigation).

231. Other theories of tort, such as corrective justice and civil recourse, might be fruitful areas for future work in this space. *See* ERNEST J. WEINRIB, THE IDEA OF PRIVATE LAW 56 (1995) (describing corrective justice theory); ARTHUR RIPSTEIN, PRIVATE WRONGS 270 (2016) (describing corrective justice theory); JOHN C. P. GOLDBERG & BENJAMIN C. ZIPURSKY, RECOGNIZING WRONGS 263–64 (2020) (describing civil recourse theory).

232. John C. P. Goldberg, *Two Conceptions of Tort Damages: Fair v. Full Compensation*, 55 DEPAUL L. REV. 435, 435 (2006).

233. Guido Calabresi & Spencer Smith, *On Tort Law's Dualisms*, 135 HARV. L. REV. F. 184, 184 (2022).

234. *See* RESTATEMENT (THIRD) OF AGENCY § 7.04 cmt. b (AM. L. INST. 2006) ("Under general tort-law principles, a person who orders or induces an actor's tortious conduct is subject to liability for harm resulting to a third person when the person 'knows or should know of circumstances that would make the conduct tortious if it were his own. . . .'" (quoting RESTATEMENT (SECOND) OF TORTS § 877 (AM. L. INST. 1965)).

potential of the market to current dominant companies.[235] This is an ever-present concern, and achieving a balanced liability approach to emerging technologies is necessary to maximize inclusive innovation.[236]

IV. ARTIFICIAL INFRINGEMENT

AI-generated infringing outputs—which this Article terms "artificial infringement"—are the third stage in the evolution of copyright direct liability after corporeal and mechanical infringement. Computer and Internet systems have always relied on certain automated functions, such as keys causing corresponding actions, files being saved, and webpages responding to user clicks. But generative AI is more than merely pressing a button for a narrow, desired outcome. The identity of the infringing work is not known in advance. The AI system can generate unforeseen (and undesirable) infringing outputs despite the best efforts of the developer and user.[237] Unlike with previous technologies that created infringements, such as copy machines or VCRs, the developer and user of generative AI do not know what the output will precisely be. Breaking from the tool-like machines of the past, generative AI's unpredictable outputs present new questions for infringement.

This Article seeks to resolve this emerging crisis in copyright infringement liability by situating artificial infringement within the underexplored historical evolution of the direct infringement test in response to technological progress, using causation as a polestar. This Article proposes making the AI system itself the direct infringer to help maintain copyright's incentives-access balance. This approach is not only faithful to causation but allows courts to eschew purely strict liability for copyright law by using secondary liability (in the guise of contributory liability and inducement) as a more nuanced, intent-conscious framework for holding users and developers liable that facilitates innovation and tort goals while still enforcing copyright owners' rights.

A. DIRECT LIABILITY FOR AI

The law needs a viable artificial infringement framework to meet the new liability challenge posed by artificial infringement. Technology has continued to develop since the heyday of mechanical infringement cases like *Netcom* and *Cablevision*, leading to the emergence of today's increasingly complex generative machines. These machines exist on a spectrum from "ordinary tools" to "fully-generative machines."[238] Ordinary tools include software like

235. *See* Balkin, *supra* note 48, at 49.

236. *See, e.g.*, Clark D. Asay, *Artificial Stupidity*, 61 WM. & MARY L. REV. 1187, 1192–93 (2020) (advocating for understanding optimal innovation policies in intellectual property to allow true artificial intelligence, rather than mere artificial stupidity, to be developed).

237. I use the term developer to refer to any upstream entity involved in the AI system's creation before the end-user.

238. Jane C. Ginsburg & Luke Ali Budiardjo, *Authors and Machines*, 34 BERKELEY TECH. L.J. 343, 405 (2019).

Photoshop and AutoCAD to assist users in refining their own creative contributions.[239] Fully-generative machines are those that are capable of producing individual outputs using their own insights with only minimal user input.[240] While ordinary tools fit within the mechanical infringement paradigm, fully-generative machines such as generative AI require further analysis. Not knowing the identity of specific outputs in advance separates generative AI from the more mechanical tools that came before.

Many entities are involved in the creation of the generative AI system that makes the infringing output: the user, the parties involved in the supply chain that creates the AI system, the company that launches and maintains the product (the "developer"),[241] and the AI system itself. For the first time, developers must face the fact that the AI systems they create could, by no intent of their own, have uncontemplated and undesirable behaviors.[242] Yet courts do not consider intent in the prima facie direct infringement test—one can infringe a copyright even if they are unaware.[243] The unpredictability of AI outputs' contours is one of AI's greatest strengths, but it also creates the very real risk of liability for torts such as copyright infringement. Assigning responsibility to the wrong entity could risk unnecessarily undermining the AI system's capabilities without achieving accountability.

Like with mechanical infringement, causation—specifically proximate causation—can provide a valuable framework for assigning liability. Both the developer and the user are but-for causes of the infringing output, as the infringement would not exist without the developer training and providing the AI system and the user inputting the prompt. This Article instead focuses on proximate causation because it seeks to answer to whom or what we should attribute the infringing nature of a given output. While proximate cause has been criticized as an unclear term and a substitute for policy choices,[244] courts continue to use it to trace liability.[245] There is not a clear substitute to take its

239. *Id.* at 406.

240. *Id.* at 407.

241. For more on the eight stages in the generative-AI supply chain and the various parties involved, see Lee et al., *supra* note 5, at 285–307.

242. Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 GEO. WASH. L. REV. 1, 23 (2018); Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 365 (2016).

243. *See, e.g.*, Williams v. Gaye, 895 F.3d 1106, 1123 (9th Cir. 2018) (stating that the element of access may be established through "widespread dissemination and subconscious copying" (quoting Three Boys Music Corp. v. Bolton, 212 F.3d 477, 483 (9th Cir. 2000)); Tanksley v. Daniels, 902 F.3d 165, 173 (3d Cir. 2018) ("[I]t is no defense that a defendant copied a protected work—such as a song—subconsciously.").

244. *See, e.g.*, DAN B. DOBBS, PAUL T. HAYDEN & ELLEN M. BUBLICK, THE LAW OF TORTS 680 (2d ed. 2011) (describing proximate causation as a "well-worn term [that] has been justly criticized for years as inaccurate, misleading, and confusing"); RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL AND EMOTIONAL HARM § 29 cmt. b (AM. L. INST. 2010) (eschewing the use of "proximate cause" because the term "is a poor one to describe limits on the scope of liability").

245. *See, e.g.*, Hunley v. Instagram, LLC, 73 F.4th 1060, 1074 (9th Cir. 2023); BWP Media USA Inc. v. Polyvore, Inc., 922 F.3d 42, 51 (2d Cir. 2019) (Walker, J., concurring). The Supreme Court has also addressed proximate cause as recently as 2023. *See* Calcutt v. Fed. Deposit Ins. Corp., 143 S. Ct. 1317, 1320 (2023).

place. The "scope of liability" term the *Restatement (Third) of Torts* adopts instead is also a difficult concept to capture.[246] Oren Bracha suggests that scope of liability—which he defines as considering "the structural position of the actors"—would be more appropriate for determining the infringer.[247] But causation also seeks to consider the structural position of the actors through the lens of who had a significant impact on the infringement occurring. Despite resistance, proximate cause remains in common usage in jurisdictions across the country—despite the *Restatements of Torts* steadfastly refusing to use the term.[248]

This Article proposes that the AI system itself should be the direct infringer because it caused the infringing expression in the output. This does not mean human actors will not face liability for artificial infringement. Secondary liability always requires a direct infringer, so where there is a direct infringer there is the possibility of secondary infringers.[249] Courts can use secondary liability—specifically contributory liability and inducement—to apply a nuanced framework for evaluating developer and user liability, as will be discussed in Subpart IV(B).

1. Limited Instances of Potential User and Developer Causation

Courts and Congress should follow the lesson of mechanical infringement by looking to causation as a polestar for maintaining the incentives-access balance in the face of artificial infringement. To determine the causal actor, we must consider each potentially responsible party. The user of the AI tool, the developer (and any other upstream participants in the tool's creation and offering to users), and the AI system itself are all potential defendants. There are circumstances when each would seem to be liable, depending on their contribution to the infringement.

First, the user. When the user provides the AI system with a prompt that clearly seeks to generate an infringing output, the user may have played the central role in the creation of an infringing output.[250] Users can have a significant influence on the output, particularly when they engage in adversarial machine learning. Recall how Sag managed to make Midjourney and Stable Diffusion generate images of Snoopy and Mickey Mouse.[251] There are even

246. RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL AND EMOTIONAL HARM § 29 cmts. a–c (AM. L. INST. 2010).

247. Oren Bracha, *Generative AI's Two Information Goods* 36 (U. Tex. L., Legal Studs. Res. Paper Mar. 11, 2025), https://ssrn.com/abstract=5115424.

248. RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL AND EMOTIONAL HARM ch. 6 special note on proximate cause (AM. L. INST. 2016).

249. *See* A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1013 n.2 (9th Cir. 2001) ("Secondary liability for copyright infringement does not exist in the absence of direct infringement by a third party.").

250. *Cf.* Rosati, *supra* note 127, at 619 (suggesting more broadly that the user may be the direct infringer).

251. Sag, *supra* note 5, at 330–31.

professional AI "whisperers" who sell specialized prompts for generating particular types of outputs.[252]

User liability has been contemplated in other areas of AI based on a causation analysis. Take, for example, so-called autonomous vehicles, which are one of the most explored areas of AI liability. These vehicles are not yet truly autonomous and always involve human input to some degree.[253] Accordingly, current autonomous vehicle liability law is driver-focused, as the vehicle remains a mere tool in the hands of the driver. For example, courts have held operators liable for accidents where the vehicle was in autopilot.[254] If the driver makes a mistake, they are liable—unless a manufacturing defect is present—because they were the cause of the accident.[255] Yet some have posited that drivers of *fully* autonomous vehicles should not be held liable for the harms inflicted by the vehicles because they would be unable to affect the car's behavior or the underlying AI system.[256] Instead, they explain, the liable party should be one who can affect the functioning of autonomous vehicles—the one who caused the accident to occur.[257]

In other cases, the developer might be the cause instead of the user. The developer may have an outsized role in causing the creation of infringing outputs if they devised a system that would generally infringe or will infringe in specific circumstances. As explained with mechanical infringement, the volition or causation dividing line separates when the developer merely provides the AI system and when it is intimately involved with the perpetration of infringement.[258] Under the mechanical infringement paradigm, the developer could be held directly liable where they devised the system in a way where it was almost guaranteed to infringe when it was used. For example, in *Aereo*, the provision of all the copied content by Aereo for profit and in competition with the broadcasters that owned the rights to the content was enough to make it directly liable even though the users chose the specific (infringing) content.[259]

252. Adi Robertson, *Professional AI Whisperers Have Launched a Marketplace for DALL-E Prompts*, VERGE (Sep. 2, 2022, at 07:30 PT), https://www.theverge.com/2022/9/2/23326868/dalle-midjourney-ai-prompt base-prompt-market-sales-artist-interview.

253. *See* Abraham & Rabin, *supra* note 9, at 130–31 (discussing the five-tiered levels of automation adopted by the National Highway Traffic Safety Administration).

254. *See, e.g.*, Brouse v. United States, 83 F. Supp. 373, 374–75 (N.D. Ohio 1949) (holding a pilot liable for an airplane collision that occurred while the plane was on autopilot mode); Came v. Micou, No. 04-CV-1207, 2005 WL 1500978, at *5 (M.D. Pa. June 23, 2005) (counting failure to disengage the cruise control as a likely act of "reckless indifference to the rights of others").

255. *See, e.g.*, Karney v. Leonard Transp. Corp., 561 F. Supp. 2d 260, 263 (D. Conn. 2008) (explaining that the plaintiff was liable for the crash, but that the manufacturer owed a duty of reasonable care in manufacture of the vehicle and would be liable for its portion of the fault).

256. *See, e.g.*, Antonio Davola, *A Model for Tort Liability in a World of Driverless Cars: Establishing a Framework for the Upcoming Technology*, 54 IDAHO L. REV. 591, 599 (2018); Abraham & Rabin, *supra* note 9, at 145.

257. Davola, *supra* note 256, at 599–600; Abraham & Rabin, *supra* note 9, at 147.

258. *See supra* Subpart.II.B.

259. Am. Broad. Cos. v. Aereo, Inc., 573 U.S. 431, 451 (2014).

In *Capitol Records, LLC v. ReDigi Inc.*, the Southern District of New York similarly found that ReDigi programmed its software to select digital files of copyrighted music, which was volitional conduct.[260] Likewise, in *Capitol Records, Inc. v. MP3tunes, LLC*, MP3tunes directly infringed where its program automatically—but at the initial direction of MP3tunes—retrieved (almost certainly copyrighted) album cover art from Amazon for music its users had uploaded.[261] These rulings are consistent with the purpose of the volition or causation requirement, which aims to hold those who select infringing works or make specific choices to copy or perform them liable, because the developers had set up their systems to copy infringing works.[262]

However, providers of automated or hands-off systems have generally not satisfied the volition or causation requirement without more. This was the case in *Netcom* and *Cablevision*, where the mere supply of Internet services and a remote DVR device did not cause the infringement.[263] Similarly, cloud storage providers and website hosts have not met the causation requirement for users' infringements when they merely provided neutral services that did not treat infringing content any differently.[264]

Based on this precedent, the mere provision of an AI system should not make the developer liable, but the developer could be the cause of infringement if it designs the AI system to infringe or its outputs are effectively guaranteed to be infringing. This could take the form of an overtly infringing AI system (e.g., a Mickey Mouse generator) or the training data could be so limited or flawed that infringement is almost guaranteed (i.e., by creating only similar outputs). The latter seems close to the situation in *GEMA v. OpenAI*, where the Munich I Regional Court held OpenAI directly liable for infringing outputs not just because they provided the models, but because the copyrighted content at issue was reproduced through simple user prompts.[265] But, more likely, the developer will obfuscate the infringing nature of the AI system if infringement is their goal. The system could generate a significant number of non-infringing outputs but, by design, also generate a sizeable number of infringing ones. This would

260. 934 F. Supp. 2d 640, 657 (S.D.N.Y. 2013).

261. *See, e.g.*, 48 F. Supp. 3d 703, 720 (S.D.N.Y. 2014) ("It was reasonable for the jury to conclude that MP3tunes' volitional conduct—creating a feature to automatically retrieve cover art from Amazon.com—caused the unauthorized reproduction of the Labels' cover art.").

262. Denicola, *supra* note 145, at 1286.

263. Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1370 (N.D. Cal. 1995); Cartoon Network LP, LLLP v. CSC Holdings, Inc., 536 F.3d 121, 131 (2d Cir. 2008).

264. *See, e.g.*, Disney Enters., Inc. v. Hotfile Corp., 798 F. Supp. 2d 1303, 1308 (S.D. Fla. 2011) (finding no volition where the defendant merely provided server storage); CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544, 556 (4th Cir. 2004) (finding no volition by a real estate listings website even where its employees cursorily screened the user-uploaded content); Fox Broad. Co. v. Dish Network L.L.C., 747 F.3d 1060, 1067 (9th Cir. 2014) ("[O]perating a system used to make copies at the user's command does not mean that the system operator, rather than the user, caused copies to be made.").

265. Landgericht München I [LG] [Munich I Regional Court] Nov. 11, 2025, Final Judgment No. 42 O 14139/24 ¶¶ 264–65 (Ger.).

complicate the analysis as the developer's role in causing the infringement is opaquer.

As machines become more autonomous, the causation line is increasingly hazy. AI systems' opacity and unpredictable outputs make it difficult at best to show causation by the developer or user of the AI system. As AI harms are often unforeseeable, it is difficult to properly deter harmful outcomes and to assign causality *ex post*.[266] This creates a potential problem where infringement occurs but there is no identifiable perpetrator to whom causation can be attributed.

With other torts, courts and scholars have started to conclude that existing law does not align with how AI functions. For example, a judge on the San Francisco County Superior Court concluded that it will be difficult to apply traditional tort doctrines to autonomous robots.[267] Others have described AI systems as defying legal categorization and traditional liability structures.[268] For example, some have concluded that strict or de facto liability for AI systems would lead to unjust outcomes.[269] These observations suggest that the novel aspects of increasingly sophisticated AI systems may necessitate a novel legal approach.

2. The Promise of AI System Causation

This Article proposes that the AI system itself could be directly liable for causing AI-generated infringements, either as an artificial person or a legal fiction. Chatterjee and Fromer previously proposed that AI systems might be said to act with volition if they can evaluate whether an output is likely to infringe upon copyright and decline to make such a copy.[270] Here, I instead show, as a historical and doctrinal matter, that AI systems could be considered

266. Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. REV. 1347, 1365–66 (2023); *see also* Babic & Cohen, *supra* note 73, at 864–65 (proposing that explainable AI cannot help us understand why an automated decision was made in a certain way); Rachum-Twaig, *supra* note 25, at 1143 ("This trend and phenomenon is coupled with the unpredictable manner in which such robots behave and the inability to foresee risks that they may inflict.").

267. Karnow, *supra* note 40.

268. Ignacio N. Cofone, *Servers and Waiters: What Matters in the Law of A.I.*, 21 STAN. TECH. L. REV. 167, 171 (2018); Calo, *supra* note 44, at 532–49; Balkin, *supra* note 48, at 52.

269. *See, e.g.*, Rachum-Twaig, *supra* note 25, at 1162–64 (finding strict liability too restrictive given the unpredictability of AI outputs); Lior, *supra* note 80, at 1059 ("The ambiguity surrounding AI entities' capabilities and its unpredictable actions and damages, all suggest that even though in its essence an AI entity is a product that is being sold by a company to a consumer, it is not similar enough to be treated as such."); Matthew U. Scherer, *Of Wild Beasts and Digital Analogues: The Legal Status of Autonomous Systems*, 19 NEV. L.J. 259, 280 (2018) ("[I]mposition of strict liability could lead to outcomes that seem less than just . . . liability being assigned to parties that had only a highly attenuated influence on the finished system and could not have reasonably foreseen or prevented the harm."); Curtis E.A. Karnow, *Liability for Distributed Artificial Intelligences*, 11 BERKELEY TECH L.J. 147, 188–191 (1996) (describing issues with establishing human liability for AI-generated outcomes since those outcomes are not necessarily foreseeable); Diamantis, *supra* note 9, at 590 ("Presumptive fault frameworks fail the Innovation and Safety Criteria, however. They make liability for AI as costly as it could possibly be (short of imposing exemplary damages), which in turn makes AI costly to develop.").

270. Chatterjee & Fromer, *supra* note 10.

directly liable under causation because they cause the infringement to occur. While the developer provides the initial setup and parameters, and the user provides the specific prompt, the AI system is determining the actual expression, including any potential infringement in that expression. A causation analysis suggests that making the AI system the direct infringer (but not the developer's agent, as discussed below) is the correct outcome. In addition, a direct infringer is also necessary to reach secondary liability, which offers a non-de facto framework for analyzing the developer's and user's liability.[271] This understanding of the AI system as the direct infringer allows courts to engage in a more nuanced secondary liability framework for determining developer and user liability, which will be discussed below in Subpart IV(B).

Generative AI systems are not mere mechanical tools, like a copy machine, that merely reproduce what its users provide it. We can think of two scenarios with AI: where the AI system is a tool of the user and where it can be thought of as an independent actor or agent.[272] While there is a spectrum of autonomy for AI systems, this setup is helpful because it looks to proximate causation rather than simply seeking a party to compensate the rights owner. In the autonomous vehicle context, no one contemplates holding the car liable, but this is because it is usually not the sole decision-maker: The driver (or in some cases the developer) is often able to step in and act to prevent a harm from occurring. But generative AI operates semi-autonomously to achieve a human-provided objective in a way that is unpredictable due to the AI using its own training to make decisions. This is a case where the AI system is committing torts in irreducible ways without there necessarily being a consciously responsible person.[273]

In many cases, the AI system seems to be the one that proximately causes the infringement because it is making determinations about what and how expressive content should be included in the output. Generative AI has started to shift from a mere tool to a semiautonomous system. Even though the user provides the prompt, they and the developer often cannot predict the exact output. The prompt is merely the idea that the AI system will develop into specific expression. The AI system operates as a stochastic parrot by generating the most statistically likely content in response to the user's prompt, therefore determining the exact expression that an output will contain. The AI system determining the actual (infringing) expression seems to place causation in the hands of the generative AI system itself.[274] The developer and user could foresee

271. A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1013 n.2 (9th Cir. 2001).

272. Samir Chopra & Laurence F. White, *Tort Liability for Artificial Agents*, *in* A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS 119, 121 (2011).

273. Ryan Abbott & Alex Sarch, *Punishing Artificial Intelligence: Legal Fiction or Science Fiction*, 53 U.C. DAVIS L. REV. 323, 337 (2019).

274. Christina Mulligan, *Revenge Against Robots*, 69 S.C. L. REV. 579, 590 (2018) (suggesting that autonomous robots are the proximate cause of injuries resulting from their actions); Chatterjee & Fromer, *supra* note 10, at 1910–11 (proposing that AI could "choose" whether to make a specific copy).

that an AI system could be used to infringe, much like the manufacturer of a copy machine could foresee that it might be used to make an infringing copy. But copyright law is more concerned with whether the specific infringements are foreseeable than general foreseeability of a device being used to harm, which was rejected in cases like *Sony*.[275]

The law contemplates imposing liability on non-human actors. The Copyright Act seems to consider infringement by non-humans when it says "anyone"—including state governments—can be liable for infringement.[276] This could potentially take the form of artificial legal personhood for AI systems. The law has granted other non-humans quasi-personhood as a sort of legal fiction.[277] Corporations are the most prominent example. When a corporate charter is granted, a new legal entity is recognized that can carry out certain business activities and enjoy certain constitutional rights as a "person."[278] As a legal person, corporations bear both legal entitlements and burdens and can be held liable for legally consequential actions, including torts such as copyright infringement.[279] Jurisdictions have been willing to recognize new types of non-human entities that could be liable too, such as Wyoming accepting decentralized autonomous organizations ("DAOs") in 2021.[280]

Artificial personhood for AI systems has already been suggested in other contexts, and it is an amorphous category that could encompass AI. As Judge Katherine Forrest has noted, "[w]e humans have made . . . personhood . . . flexible and malleable to the historical moment."[281] Legal theory from over thirty years ago likewise noted that the "concept of [legal] person is fuzzy at the edges" and AI, under the right circumstances, could be considered a legal person.[282] More recently, a European Parliament resolution and recommendation to the Commission on Civil Law Rules on Robotics suggested granting robots status as "electronic persons" because extant liability rules may be insufficient for increasingly autonomous robots.[283] We could think of AI

275. Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

276. 17 U.S.C. § 501(a).

277. *See* Jason Chung & Amanda Zink, *Hey Watson—Can I Sue You for Malpractice? Examining the Liability of Artificial Intelligence in Medicine*, 11 ASIA PAC. J. HEALTH L. & ETHICS 51, 77 (2018) (describing legal personhood as "a pre-existing legal fiction used to hold entities, not just natural persons, accountable"). Quasi-personhood also avoids the risk of considering an AI system property that could potentially be subject to an *in rem* action. *See, e.g.*, Petroleos Mexicanos Refinacion v. M/T KING A, 554 F.3d 99, 102–03, 108–09 (3d Cir. 2009) (describing *in rem* seizures under maritime law).

278. *See* Elizabeth Pollman, *Corporate Personhood and Limited Sovereignty*, 74 VAND. L. REV. 1727, 1732 (2021) (describing the personhood of corporations).

279. Paul B. Miller, *Corporate Personality, Purpose, and Liability*, *in* RESEARCH HANDBOOK ON CORPORATE PURPOSE AND PERSONHOOD 222, 231 (Elizabeth Pollman & Robert B. Thompson eds., 2021).

280. WYO. STAT. ANN. § 17-31-101 (West 2021).

281. Katherine B. Forrest, *The Ethics and Challenges of Legal Personhood for AI*, 133 YALE L.J.F. 1175, 1177 (2024).

282. Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231, 1285 (1992).

283. Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)), EUR. PARL. DOC. P8_TA(2017)0051 (2017).

systems as along a spectrum of personhood, where more autonomous AI systems in certain contexts are more deserving of artificial personhood.[284] Proposed personhood for AI systems could help address liability for torts perpetrated by the AI system.

Legal personhood is not, however, necessarily required to hold the AI system directly liable; we could consider AI the direct infringer under a more general legal fiction. Jack Balkin recommended approaching legal doctrine from a pragmatic approach like this to accommodate new technologies, noting that "[w]hen we employ legal fictions[,] . . . we substitute robot for human to allow the law to function effectively in the face of the legal enigmas posed by emergent behavior."[285] Congress could mandate that AI systems are the direct infringers in a statute, but courts also have some flexibility to accept this conception by themselves. After all, Judge Whyte introduced the "volition or causation" requirement for mechanical infringement *sua sponte*.[286]

3. Normative Detriments of De Facto Liability for Developers and Users

Alternatively, we could choose to ignore causation and impose de facto liability on the developer through direct liability. However, a general rule that would hold AI developers or users directly liable for all infringing outputs is normatively undesirable. Such a rule would be overreaching and unmoored from causation considerations for several reasons, including: (1) it conflicts with copyright law causation precedent, (2) the AI system, not a human, typically selects the (infringing) expression, (3) it could chill AI innovation and competition, and (4) it would likely not properly advance the traditional tort norm of deterrence and could unduly constrain AI capabilities.

First, a per se liability rule for the developer or user would likely conflict with copyright precedent. Cases such as *Parker* and *Field* held that service providers were not directly liable for automatic functions without more.[287] This is also why Judge Whyte rejected direct liability as an appropriate claim for the Internet service provider in *Netcom*, because it would effectively hold the provider liable for the mere provision of the product.[288] This is similar to the *Sony* rule in secondary liability that shields the provider of a product from constructive knowledge of resulting infringement so long as the product has

284. *See* Carla L. Reyes, *Autonomous Corporate Personhood*, 96 WASH. L. REV. 1453, 1480, 1499–1504 (2021) (proposing this spectrum in the corporate context).

285. Balkin, *supra* note 48, at 58; *see also* Daniel J. Gervais & John J. Nay, *Artificial Intelligence and Interspecific Law*, 382 SCI. 376, 376 (2023) (suggesting that the law could treat AI systems as legal persons, perhaps taking a corporate form as an AI-operated LLC).

286. Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1370 (N.D. Cal. 1995).

287. *See, e.g.*, Parker v. Google, Inc., 422 F. Supp. 2d 492, 497 (E.D. Pa. 2006), *aff'd*, 242 F. App'x 833 (3d Cir. 2007); Field v. Google Inc., 412 F. Supp. 2d 1106, 1115 (D. Nev. 2006).

288. *Netcom*, 907 F. Supp. at 1372–73.

substantial noninfringing uses.[289] The Supreme Court's recent decision in *Twitter, Inc. v. Taamneh* further supports this. There, the Court held that recommendation algorithms are merely part of the social media platform's infrastructure through which all the content on their platforms is filtered.[290] At least in the context of the terrorism claim at issue, the Court distinguished between this passive algorithmic assistance and culpable aiding and abetting.[291]

If mere automatic functions in response to users' actions did not rise to the level of causation, the even less predictable outputs of generative AI systems would seem to weigh further against such a finding of liability for the developer. Although the developer chose the data on which the AI system was trained, this is merely educating the AI system.[292] In the corporeal infringement context, merely exposing an infringer to the work they infringed does not result in liability. For example, a teacher assigning a reading in a syllabus and the student subsequently copying that work does not result in infringement by the teacher. Generative AI is the sort of scenario that Judge Forrest described as "when there is no human who made the design choice determined to have proximately caused the harm [here, the infringement]."[293]

Second, the AI system generally dictates the expressive aspects of AI-generated content, not the developer or user. In a related context, the U.S. Copyright Office has determined that AI-generated content should not be copyrightable where the AI system, not a human, is generating the expressive content.[294] In a 2025 report, the U.S. Copyright Office concluded that "[p]rompts essentially function as instructions that convey unprotectible ideas . . . [and the human user cannot] control how the AI system processes them in generating the output."[295] A district and appellate court upheld this rule, noting the unpredictability of generative AI systems' outputs versus those of cameras, which generate mechanical reproductions "only after the photographer develops a 'mental conception' of the photograph."[296] Unlike cameras, the user (and programmer) do not develop a specific "mental conception" before the work is created, just an idea expressed in a prompt. Therefore, the expressive content of the output—infringing or not—is being calculated by the AI system, not the human user or developer. The AI system dictating the creative content

289. Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

290. Twitter, Inc. v. Taamneh, 143 S. Ct. 1206, 1226–27 (2023).

291. *Id.* at 1227.

292. *See* Lemley & Casey, *supra* note 5, at 776–77 (arguing this in the fair use context).

293. Forrest, *supra* note 281, at 1203.

294. U.S. COPYRIGHT OFF., COPYRIGHT AND ARTIFICIAL INTELLIGENCE PART 2: COPYRIGHTABILITY iii (2025) [hereinafter COPYRIGHTABILITY REPORT]; *see also* Thaler v. Perlmutter, 130 F.4th 1039, 1044–45 (D.C. Cir. 2025). An individual may only obtain a narrow copyright in their modification of AI-generated work after the fact or by using AI as a tool over which they exercise expressive control. COPYRIGHTABILITY REPORT, *supra* note 294, at 24–27.

295. COPYRIGHTABILITY REPORT, *supra* note 294, at 18.

296. Thaler v. Perlmutter, 687 F. Supp. 3d 140, 146 (D.D.C. 2023), *aff'd*, 130 F.4th 1039 (D.C. Cir. 2025).

means that its actions, rather than a human user's prompt, are likely causally creating the substantially similar, infringing expression in the output.

Admittedly, the legal standards for authorship and infringement are different, but the respective tests are premised on the same expressive content in the work. At their root, authorship and infringement both require the causal creation of the expressive material at issue. In the AI context, authorship and infringement may also be entangled from a normative standpoint. Rights in AI-generated outputs should entail accepting responsibility for outputs too.[297] Despite their different tests, AI authorship and models' capacity to infringe should be linked due to the single creation of expression.

Third, holding developers or users liable for all infringements generated by an AI system would violate longstanding copyright incentives-access balance norms and could have a chilling effect on AI competition and innovation. As explained above, traditional liability rules are not always a perfect fit for new technologies.[298] When applying or refining copyright law to accommodate new technologies, courts and Congress must be careful to avoid upsetting copyright's carefully crafted incentives-access balance.[299] A careless analysis could misbalance copyright too far in favor of incentives and make developers or users bear the cost of any infringements despite their best efforts.

In this case, holding the developer or user de facto liable would disincentivize new entrants in the generative AI industry, harming competition.[300] Such a rule would create significant liability exposure ranging from $200 up to $150,000 per work infringed.[301] While insurance could potentially help mitigate these costs, insurance premiums could still impose a significant expense on less well-resourced market entrants—if they could even cover the full scope of damages a copyright judgment may bring.[302] A broad liability rule and resulting costs could therefore significantly stifle the latent potential of the generative AI industry. In particular, it would likely limit the viability of new, smaller players entering the field. Incumbents can absorb the financial damages of infringement liability.[303] Some of the major players have even promised to indemnify users for any copyright infringement claims.[304] New startups cannot compare in financial terms, making copyright infringement

297. Daniel J. Gervais, *The Machine as Author*, 105 IOWA L. REV. 2053, 2087 (2020).

298. *See supra* Parts.I–II.

299. *See supra* Subpart.I.C.

300. *See* Abbott, *supra* note 242, at 21 (expressing concern that treating AI as a product for purposes of tort liability would discourage automation).

301. 17 U.S.C. § 504(c) (providing statutory damages of $750 to $30,000, up to $150,000 for willful infringement, and as low as $200 for innocent infringement).

302. *See* Anat Lior, *Insuring AI: The Role of Insurance in Artificial Intelligence Regulation*, 35 HARV. J.L. & TECH. 467, 501–02 (2022) (discussing the benefits and vast costs associated with insurance for AI).

303. *See* Eric Goldman, *Generative AI Is Doomed*, 29 MARQ. INTELL. PROP. & INNOVATION L. REV. 149, 165 (2025) ("[Generative AI] incumbents can use industry regulation to raise new entrants' costs and deepen their own competitive moat.").

304. *See, e.g.*, *Commercial Terms of Service*, ANTHROPIC (June 17, 2025), https://www.anthropic.com/legal/commercial-terms.

liability a proportionally greater risk to those companies whether they are held liable or expected to bear the litigation costs for their users. Other barriers to entry are lowering, such as the price of developing a bespoke AI system decreasing as companies use existing base models.[305] Permitting an absolute direct liability rule against developers would insulate the current positions of major technology companies at the vanguard of generative AI at the expense of new, innovative would-be entrants, harming overall competition and innovation.

Fourth, strict attribution of causation to developers or users would not fully serve tort law's goal of deterrence. Calabresi's cheapest cost avoider theory could suggest that the developer is in the best position to prevent the infringements by modifying the underlying model or the system.[306] In some cases, the user may be best situated to avoid infringement if they are engaging in adverse machine learning. Yet even if the cheapest cost avoider theory were the only consideration (and it should not be), an overarching goal is to incentivize use only when the benefits outweigh the costs.[307] In the case of copyright infringement, where successful claims could result in damages of $150,000 per infringed work, costs could quickly equal or even outpace the benefits.[308] A lack of foreseeability into when and how many infringements will occur substantially muddies the cost-benefit analysis for potential market entrants.

In practice, de facto liability for all artificial infringement would subject developers to liability despite the fact that they cannot effectively deter unforeseeable infringements due to the AI system's nondeterministic nature.[309] Deterrence requires foresight that one's acts would result in misconduct, but the nature of AI systems makes this unlikely because a user or developer would not know if or when infringements would occur.[310] Alternatively, deterrence could incentivize precautions, which requires the defendant to have been able to foresee the risk of injury and its magnitude in advance.[311] Foreseeability can be an abstract concept, even when grounded in causation. But the sheer unpredictability of the generative AI outputs presents a clearer case of insufficient foreseeability.

Even if the developer internalized the harm, the exact harms and their costs would be opaque. Researchers may be able to detect, reverse engineer, and remedy some problematic results in advance, but that is an elusive luxury.[312] It would be foolhardy to expect that *ex ante* testing is a panacea. No amount of

305. Sage Lazzaro, *AI's Big Shift from 'Model-Forward' Innovation to 'Product-Forward*,' FORTUNE (Jan. 25, 2024, at 14:26 ET), https://fortune.com/2024/01/25/ai-model-forward-innovation-product-forward.

306. CALABRESI, *supra* note 214, at 156.

307. CHOPRA & WHITE, *supra* note 272, at 128–29.

308. 17 U.S.C. § 504(c).

309. *See supra* Part.III.

310. Lemley & Casey, *supra* note 14, at 1361.

311. Calo, *supra* note 44, at 542.

312. Lemley & Casey, *supra* note 14, at 1337.

testing can make a developer certain about what an AI system would do in response to a particular prompt, let alone an untested prompt.[313] Even experts often miss eventual problems in algorithmic outputs.[314] This issue is compounded by machine learning constantly improving and changing AI models.[315] The sheer number of possibilities and the ever-changing nature differentiate AI models from current technologies with more circumscribed and static operations. Even if a software program such as Apple's iOS occasionally has updates, those are controlled by humans and have been human-tested. The iterative updating of machine learning, however, is largely independent from humans, further rendering any potential harms flowing from AI unforeseeable. An absolute liability rule for developers or users would therefore be untethered from deterrence because they would be held liable regardless of their best mitigation efforts or the foreseeability of the harms.

If the developer is unable to appreciate all the potential harms, it would be difficult for the developer to ascertain how much to restrict the AI system. Developers may limit their AI systems to avoid liability. AI can go beyond human capabilities due to its machine learning and nondeterministic nature, which necessarily results in some unpredictability in specific outputs. To avoid liability, AI developers may unduly restrict AI's capabilities, limiting what benefits it could produce for society. This is especially true for smaller market entrants, who may be more hesitant to risk unknown liability. The ability of AI to go beyond human possibilities is due to its unpredictable outputs. Restricting AI too far could limit positive externalities for society.

These problems highlight why de facto liability for copyright infringement in the generative AI context is normatively undesirable. Others have gone further, suggesting that strict liability is always inappropriate for copyright infringement, as it can lead to undesirable and unfair results for accidental infringements.[316] While this may go a step too far, this instinct is especially poignant with AI. Applying causation to hold the AI system directly liable serves a similar purpose by extracting copyright law from purely strict liability. In this context, it achieves a more nuanced, fault-based liability framework by utilizing secondary liability for developers and users, as will be discussed below in Subpart IV(B).

4. Parallel Problems with a Principal-Agent Relationship

Making the AI system directly liable does not mean that the developer should be de facto liable for the AI system's infringements under a principal-agent relationship, as this would effectively impose the same de facto liability

313. Kroll et al., *supra* note 50, at 650.

314. *Id.* at 647.

315. *Id.* at 660.

316. Patrick R. Goold, *Moral Reflections on Strict Liability in Copyright*, 44 COLUM. J.L. & ARTS 123, 145 (2021); Oren Bracha & Patrick R. Goold, *Copyright Accidents*, 96 B.U. L. REV. 1025, 1029 (2016).

on developers without the limit of causation. Agency law requires that a principal assents to the agent acting "on the principal's behalf and subject to the principal's control [and that] the agent manifests assent or otherwise consents so to act."[317] The Supreme Court has described control as one of the "basic features of an agency relationship."[318] Generally, agency takes the form of actual authority, apparent authority, or *respondeat superior*. Actual authority requires that "the agent reasonably believes, in accordance with the principal's manifestations to the agent, that the principal wishes the agent so to act."[319] Apparent authority is when "a third party reasonably believes the actor has authority to act on behalf of the principal and that belief is traceable to the principal's manifestations."[320] Finally, *respondeat superior* is the principle that "[a]n employer is subject to liability for torts committed by employees while acting within the scope of their employment."[321] Underlying all three of these relationships is a consensual relationship between principal and agent and control by the principal.[322]

Some have previously proposed agency or *respondeat superior* as the optimal liability framework for AI systems, albeit largely outside of the copyright context.[323] The primary justification for these proposals is that AI systems themselves would be judgment-proof because they do not have financial resources of their own.[324] Similarly, in the copyright context, Chatterjee and Fromer suggest the AI system could potentially be treated as an agent of the developer or user who employs them because even if a machine could be held directly liable, it would have no assets, rights, or legal status—effectively nullifying any practical benefits of such a ruling.[325]

These scholars correctly identify the problem of holding the AI system liable when it does not have assets to compensate injured parties, but, nonetheless, agency law is a poor doctrinal and normative fit for generative AI.[326] The developer and generative AI system lack the consensual relationship

317. RESTATEMENT (THIRD) OF AGENCY § 1.01 (AM. L. INST. 2006).

318. Hollingsworth v. Perry, 570 U.S. 693, 713 (2013).

319. RESTATEMENT (THIRD) OF AGENCY § 2.01 (AM. L. INST. 2006).

320. *Id.* § 2.03.

321. *Id.* § 2.04.

322. *Id.* § 1.01 cmt. c.

323. *See, e.g.*, Lior, *supra* note 80, at 1071–73 (proposing *respondeat superior* for AI models and their developers); CHOPRA & WHITE, *supra* note 272, at 128 (proffering agency liability); Duffourc, *supra* note 82, at 28 (suggesting agency liability for AI physicians); Aslihan Asil & Thomas G. Wollmann, *Can Machines Commit Crimes Under U.S. Antitrust Laws?*, 3 U. CHI. BUS. L. REV. 1, 6 (2023) (proposing *respondeat superior* for assigning liability for AI acts in the criminal antitrust context); Gabriel Hallevy, *The Criminal Liability of Artificial Intelligence Entities—from Science Fiction to Legal Social Control*, 4 AKRON INTELL. PROP. J. 171, 191 (2010) (suggesting that the AI system could be held directly liable for a crime, but that the developer could also face liability because the AI is its agent).

324. Lior, *supra* note 80, at 1046; CHOPRA & WHITE, *supra* note 272, at 128; Asil & Wollmann, *supra* note 323.

325. Chatterjee & Fromer, *supra* note 10.

326. In addition, joint and several liability is not a viable alternative. It would present the same issue of de facto liability, albeit at a lower damages threshold for the developer or user than if they were solely liable.

and degree of control that are usually the *sine qua non* of agency law. There is no consensual relationship between an AI system developer and the AI system itself. While generative AI is becoming increasingly autonomous, it does not yet have the ability to manifest assent. From a technical perspective, software cannot form a consensual relationship.[327] Relatedly, an AI system seemingly could not perform the fiduciary duty of loyalty because it is a machine, not a human.[328]

In addition, the level of control by a generative AI developer over the system is limited. The developer could terminate the system, but its control over specific outputs of the AI system is circumscribed by the AI's nondeterministic nature. The Supreme Court rejected an agency theory where the purported agents "decide for themselves, with no review, what arguments to make and how to make them."[329] Similarly, generative AI creates unpredictable outputs without developer review at the time of creation. In the copyright context, the Ninth Circuit found that a principal-agent relationship could exist where the principal exercised "substantive supervision[,] . . . select[ed] and remove[d] [alleged agents] on the basis of their performance, [and] control[led] . . . the[ir] . . . work schedule."[330] This level of control is lacking for generative AI, where there are no human employees and the creation of outputs occurs independently of individualized supervision by the developer.

Holding the developer liable as the principal of the AI agent is also normatively undesirable. The same reasons strict liability should be disfavored more broadly for artificial infringement apply here too. As described in Subpart IV(A)(3) above, strict liability in this context conflicts with copyright precedent on causation; the AI system is deciding the infringing expression; it could harm AI innovation and competition; and it would likely not properly advance deterrence but would unduly constrain AI's capabilities. In addition, de facto liability under an agency theory would remove the causation limitation, which, while limited, still helps balance copyright's competing goals of incentivizing creative works and providing public access to those works.

Regardless, joint and several liability does not fit the facts of AI-generated infringements. The developer and user are typically not acting in concert with the AI system to achieve the infringement or independently causing an indivisible injury. *See* RESTATEMENT (THIRD) OF TORT: APPORTIONMENT OF LIAB. §§ 12, 15, A18 (AM. L. INST. 2000). Negligence-based joint and several liability may seem more attractive, but assigning liability between defendants still considers "the strength of the causal connection between the person's risk-creating conduct and the harm." RESTATEMENT (THIRD) OF TORT: APPORTIONMENT OF LIAB. § 8 (AM. L. INST. 2000). As explained previously, the developer and user typically did not proximately cause the artificial infringement. *See supra* Subpart.IV.A.1.

327. Deven R. Desai & Mark Riedl, *Responsible AI Agents* 20 (Ga. Tech Scheller Coll. of Bus., Paper No. 5147666), https://ssrn.com/abstract=5147666.

328. *See* RESTATEMENT (THIRD) OF AGENCY § 1.01 cmt. e (AM. L. INST. 2006).

329. Hollingsworth v. Perry, 570 U.S. 693, 713 (2013).

330. Mavrix Photographs, LLC v. LiveJournal, Inc., 873 F.3d 1045, 1055 (9th Cir. 2017).

B. SECONDARY LIABILITY FOR HUMANS

Holding an AI system directly liable and eschewing an AI principal-agent relationship does not mean the developer or user is never liable. Some have proposed that infringing AI-generated outputs create a lacuna in copyright law that requires a holistic shift to negligence to address systemic social risk without the de facto liability of copyright infringement.[331] But such a broad change to copyright law is unnecessary because secondary liability already operates as an intent-sensitive framework that considers and punishes actors for their contributions to infringement. There must be a direct infringer for there to be a secondary one.[332] Imposing direct liability on the AI system (or, failing that, the user) allows courts to apply a more nuanced liability determination for developers and users under secondary liability. Secondary liability would help assuage concerns about compensating injured parties if judgment-proof AI systems were liable.[333] It would also not require remaking copyright doctrine or introducing new causes of action, such as through negligence.

Even if lawmakers and courts declined this Article's proposal to hold AI systems directly liable, courts could instead impose direct liability on the user to reach the more nuanced secondary liability framework for developers.[334] Holding the user de facto liable for artificial infringement is less normatively attractive than holding the AI system liable due to the chilling effects on AI usage (or the pressure on developers to indemnify users). However, the practical risk of user liability for artificial infringement may be circumscribed by existing liability for their subsequent use and rightsholder reticence to pursue such claims. If users subsequently share or display the infringing output, they could be directly liable for that use anyway.[335] Rights owners may also be reticent to sue individual AI users, given the backlash that occurred when the Recording Industry Association of America ("RIAA") sued individuals (including children and the deceased) who used Grokster and other peer-to-peer file sharing systems in the early 2000s.[336] Rights owners have also previously brought systemic

331. Bracha, *supra* note 247, at 49.

332. *See* A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1013 n.2 (9th Cir. 2001) (noting that there must be a direct infringer before there can be a secondary one); *see also* Limelight Networks, Inc. v. Akamai Techs., Inc., 572 U.S. 915, 922 (2014) ("[W]here there has been no direct infringement [of a patent], there can be no inducement of infringement under § 271(b).").

333. *See supra* notes 323–325 and accompanying text.

334. *See* Rosati, *supra* note 127, at 621 (suggesting that this liability structure could be proper in certain situations); Lee et al., *supra* note 5, at 357 (same).

335. Of course, personal uses of AI-generated outputs may be lawful under copyright's fair use doctrine. *See, e.g.*, Am. Geophysical Union v. Texaco Inc., 60 F.3d 913, 916 (2d Cir. 1994) ("We do not deal with the question of copying by an individual, for personal use in research or otherwise (not for resale), recognizing that under the fair use doctrine or the *de minimis* doctrine, such a practice by an individual might well not constitute an infringement."). Rights owners are also less likely to discover and more likely to tolerate such uses. *See* Tim Wu, *Tolerated Use*, 31 COLUM. J.L. & ARTS 617, 619–20 (2008) (describing a category of likely infringing but tolerated uses).

336. *See, e.g.*, John P. Mello, Jr., *RIAA Settles First Lawsuit Against 12-Year-Old Brianna LaHara*, TECHNEWSWORLD (Sep. 11, 2003, at 13:59 PT), https://www.technewsworld.com/story/riaa-settles-first-lawsuit-

architectural infringement claims against technology developers and providers rather than users because of the ease of suing them compared to millions of individual users and their much larger coffers.[337] Regardless, holding the user directly liable for artificial infringement would still allow courts to engage in a secondary liability analysis for developer liability.

Secondary liability, premised on AI system or user liability, provides a more nuanced framework that can consider both developers' (and users') actions relating to a specific infringement rather than imposing de facto liability on them even when they did not intend to infringe and the AI system is really causing the infringement. The distinction between direct and secondary liability—and the flexibility of secondary liability—"has long provided welcome breathing room for new technologies and uses."[338] Direct liability was not contemplated for simpler tool-like machines, where the mechanical infringement paradigm was effective and causation could be more steadily deduced. Using the causation requirement to hold AI systems (or, failing that, users) directly liable can help solve AI-related liability questions by maintaining the role of causation and permitting courts to utilize a secondary liability framework for analyzing claims against developers and users.

Copyright law has developed three forms of secondary liability: vicarious liability, contributory liability, and inducement. Vicarious liability requires that the defendant have the right and ability to control the infringement and receives a financial benefit from it.[339] Contributory liability attaches when the defendant knows of the infringement and materially contributes to it.[340] Inducement requires one to distribute a device and promote its infringing capabilities.[341] Secondary liability comes from common law, and judges have previously crafted and refined the doctrine in response to technologies without the need for new laws.[342]

At present, courts have not addressed how secondary liability should apply with generative AI. In my prior work, I elucidated intent as a historical polestar courts and Congress used to refine secondary liability doctrine to maintain copyright's incentives-access balance.[343] In this context, intent means an actor's action or inaction, once aware of a particular alleged infringement, that is substantially certain to facilitate or further the infringement.[344] By using intent

---

against

-12-year-old-brianna-lahara-31561.html; Eric Bangeman, *"I Sue Dead People . . . ,"* ARS TECHNICA (Feb. 4, 2005, at 12:43 PT), https://arstechnica.com/uncategorized/2005/02/4587-2.

337. *See* Goodyear, *supra* note 142, at 1983–84.

338. Giblin & Ginsburg, *supra* note 197, at 143.

339. Perfect 10, Inc. v. Giganews, Inc., 847 F.3d 657, 673 (9th Cir. 2017).

340. *Id.* at 671. This was the state of the law prior to *Cox Communications, Inc. v. Sony Music Entertainment*, 146 S. Ct. 959 (2026). This Article was written and edited prior to the publication of the Court's decision.

341. Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd., 545 U.S. 913, 936–37 (2005).

342. Goodyear, *supra* note 142, at 1977–78.

343. *Id.* at 1975–79.

344. *Id.* at 1976.

as the balancing polestar once again, courts and Congress can refine secondary liability to accommodate generative AI under copyright law. While vicarious liability is a strict liability system plagued by similar problems to those discussed above in Subpart IV(A), refinement to contributory liability, plus inducement, could maintain the incentives-access balance by imposing technologically feasible obligations on developers.

1. Vicarious Liability

Courts impose vicarious copyright liability on an entity when it has the right and ability to control infringing content and receives a financial benefit directly attributable to said infringement.[345] Vicarious liability in copyright law descends from agency law and *respondeat superior* and, like them, imposes de facto liability based on control.[346]

Vicarious liability for artificial infringement suffers from doctrinal and normative flaws. First, it is not clear that the developer could be vicariously liable under the doctrinal test. One of the two requirements for vicarious liability is having the right and ability to control the infringement.[347] The definition of control is somewhat nebulous, but courts have started to coalesce around the ability to delete or remove online content and are reluctant to find control where there is less.[348] A developer cannot predict, let alone remove, AI-generated outputs, suggesting that a court may not find the right and ability to control in the AI context. The AI system uses its own developed insights to create.[349] AI system operators could disable the entire system or filter the outputs,[350] but requiring the disabling of the entire product goes much further than copyright precedent has held.[351] Likewise, here, a finding of requisite control would seem

345. *Giganews*, 847 F.3d at 673 (quoting Perfect 10, Inc. v. Visa Int'l Serv. Ass'n, 494 F.3d 788, 802 (9th Cir. 2007)).

346. Fonovisa, Inc. v. Cherry Auction, Inc., 76 F.3d 259, 261–62 (9th Cir. 1996).

347. *Giganews*, 847 F.3d at 673 (quoting *Visa Int'l Serv. Ass'n*, 494 F.3d at 802).

348. *See, e.g.*, Bus. Casual Holdings, LLC v. YouTube, LLC, No. 22-CV-3007, 2023 WL 6842449, at *3 (2d Cir. Oct. 17, 2023) (per curiam) ("YouTube did not decline to exercise its right to stop TV-Novosti's alleged infringement, and instead removed the three videos shortly after learning about their alleged infringement."); Luvdarts, LLC v. AT & T Mobility, LLC, 710 F.3d 1068, 1072 (9th Cir. 2013) (holding that AT&T was not vicariously liable because it could not supervise its networks for infringing activity at that time).

349. *Cf.* Gervais, *supra* note 297, at 2059 (offering the same perspective in the copyrightability context).

350. *See* Lee et al., *supra* note 5, at 357 (suggesting that this is sufficient for satisfying the right and ability to control).

351. *See, e.g.*, A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1020–22 (9th Cir. 2001) (imposing liability when, upon receiving knowledge and having the ability to remove infringements, a platform fails to purge the infringing content); Oppenheimer v. Allvoices, Inc., No. 14-cv-00499, 2014 WL 2604033, at *7 (N.D. Cal. June 10, 2014) (discussing the removal of infringing content and disabling of access to sites with infringing content); Dunham v. Lei, No. 20-CV-03716, 2021 WL 4595808, at *7 (C.D. Cal. June 7, 2021) (discussing the liability of a storefront selling infringing products); Sony Music Ent. v. Cox Commc'ns, Inc., 464 F. Supp. 3d 795, 813–14 (E.D. Va. 2020) (finding liability where the defendant had the "actual ability to stop or limit ongoing infringement by modifying or terminating an account"); Agence France Presse v. Morel, 934 F. Supp. 2d 547, 575 (S.D.N.Y. 2013) (stating that the ability to remove infringing photos and block accounts can indicate the right and ability to control required for finding vicarious liability).

uncertain at best. The developer's lack of control over the contours of the expression in the AI-generated output and the inability to retract the infringement undermine the vicarious liability argument.

Even if a court did find control, that would trigger a normative problem: vicarious infringement, like direct infringement and principal-agent relationships, would effectively impose de facto liability.[352] Vicarious liability holds a party liable for the tortious actions of another based solely on their relationship without regard to fault.[353] Therefore, it would hold the developer liable for any resulting infringements just like direct liability and agency law. This de facto—and potentially widespread—liability would lead to the same normative problems described in Subpart IV(A)(3) above, including conflicting with copyright precedent, ignoring causation, chilling innovation and competition, and not achieving proper deterrence, all while restricting the capabilities of AI.

Normatively, vicarious liability may be an even worse paradigm than direct liability for artificial infringement. To be directly liable, a defendant would still need to have caused the infringement.[354] A vicarious party, however, is held liable for the direct infringer's infringements *regardless* of their causation.[355] The lack of foreseeability still means that none of the stakeholders are able to fully assess the risks of the AI system and act appropriately in response, undermining deterrence.[356] This has led some to conclude that automated AI requires a paradigm shift away from strict liability and similar de facto rules such as vicarious liability altogether.[357]

2. Contributory Liability and Inducement

A novel but promising alternative use of holding the AI system directly liable is applying contributory liability and inducement for artificial infringement. Contributory copyright infringement liability requires knowing about the specific infringement and materially contributing to it.[358] Since generative AI has substantial noninfringing uses, knowledge of infringement cannot be imputed to the developer but must be determined based on specific

352. Alfred C. Yen, *Third-Party Copyright Liability After* Grokster, 91 MINN. L. REV. 184, 214 (2006).

353. PETER CANE, RESPONSIBILITY IN LAW AND MORALITY 175 (2003).

354. Chatterjee & Fromer, *supra* note 10, at 1910.

355. EMI Christian Music Grp., Inc. v. MP3tunes, LLC, 844 F.3d 79, 99 (2d Cir. 2016).

356. Rachum-Twaig, *supra* note 25, at 1163.

357. *See, e.g.*, *id.* at 1153 (examining how the unpredictability of AI raises material normative problems for applying direct or vicarious liability); Marchisio, *supra* note 47, at 11 ("[T]his proposal . . . shows the need to shift '*obligations*' away from producers and [developers] when robots [can act] autonomously from their original design."); Beckers & Teubner, *supra* note 9, at 80–81, 83–85, 94–95 (rejecting one-size-fits-all approaches to AI liability and proposing looking at the specific use of the AI system and the involvement of the human to determine the liability regime that should be applied).

358. Perfect 10, Inc. v. Giganews, Inc., 847 F.3d 657, 670 (9th Cir. 2017) (quoting Perfect 10, Inc. v. Visa Int'l Serv. Ass'n, 494 F.3d 788, 795 (9th Cir. 2007)).

infringements.[359] Contributory liability is a fault-based standard that allows for more flexibility and more consideration of intent and tort law goals since it focuses on fault with regard to specific infringements.[360] Courts and Congress have crafted a contributory liability framework that usually considers platforms' actions in response to specific instances of infringements rather than infringement more generally because it focuses on platforms' specific intent.[361] Unlike negligence, which looks to *ex ante* impositions on generative AI systems that could stifle innovation and do not necessarily consider intent,[362] contributory liability scrutinizes parties' actions to determine if they sought to further the infringement once aware of it. Indeed, the Supreme Court barred holding a developer contributorily liable merely because it was generally foreseeable that its device could be used to infringe.[363] Contributory liability therefore appears best suited to maintain balanced copyright in the face of the novel unpredictability of AI-generated outputs by considering developers' and users' own actions.

Once an AI system developer is aware that there is a specific infringement—whether through the rights owner or anyone else reporting it, or otherwise learning of it—it would be required to take a certain action in response to negate its material contribution to that infringement. Different courts apply different material contribution tests, including removing the infringement (i.e., notice-and-takedown, specifically the takedown component), providing the "site and facilities" for the infringement, or failing to take "simple measures" to prevent the infringement.[364] Notice-and-takedown is perhaps the most prominent due to its parallel with the requirements under the Digital Millennium Copyright Act ("DMCA"). The DMCA provides a safe harbor for online service providers for the copyright infringements of their users in exchange for them meeting a series of requirements.[365] At the heart of these requirements is an obligation for the service provider to expeditiously remove infringing content on its platform once it is reported by the rights owner or the service provider otherwise learns it is infringing.[366]

Notice-and-takedown would not work for AI systems because the developer could not simply remove the infringing output once it has been reported. The generated work already exists in the world and was likely downloaded at that point. It may also be difficult to prevent a specific work from being infringed again. So-called instance unlearning is a nascent approach, and

359. *See* Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

360. Yen, *supra* note 352, at 214–15.

361. *See* Goodyear, *supra* note 142, at 1976.

362. *See* Bracha, *supra* note 247, at 49, 52 (proposing a negligence standard for AI-generated outputs).

363. *Sony*, 464 U.S. at 442.

364. A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1022 (9th Cir. 2001); Perfect 10, Inc. v. Giganews, Inc., 847 F.3d 657, 671 (9th Cir. 2017). For a discussion about these different standards, see Goodyear, *supra* note 142, at 2021–22.

365. 17 U.S.C. § 512.

366. *Id.* § 512(c)(1)(C).

although it may have long-term promise, it currently incurs high computational cost and inaccuracies.[367] It could even compromise the AI system's performance.[368] A developer could instead retrain the entire model without the allegedly infringed work—the so-called "gold standard" for machine unlearning—but that would be prohibitively expensive and may not prevent the emergence of similar latent information in outputs.[369]

A more promising approach may be output suppression through output or input filters, or system prompts, which would act as system guardrails to reduce the likelihood of specific content.[370] However, mechanical determinations of copyright infringement are far from perfect and can impede free speech.[371] Indeed, AI-generated outputs could implicate AI systems' users' First Amendment rights to listen and speak.[372] While a filter could be required, it is quite difficult to devise filters that are robust enough to capture more than mere surface-level matches but not so broad as to be proscriptively overinclusive of fair uses and other noninfringing outputs.[373] Indeed, in devising the DMCA system, Congress expressly did not require service providers to engage in proactive activities such as filtering for this reason.[374]

An alternative that builds upon notice-and-takedown while considering developers' technical capabilities and secondary liability's intent polestar for refining the law would be a more promising vehicle for maintaining copyright's incentives-access balance. Some computer scientists and law professors noted the potential viability of a notice-and-takedown-like regime for AI-generated libel.[375] The fact that there are several different material contribution tests allows

---

367. A. Feder Cooper et al., *Machine Unlearning Doesn't Do What You Think: Lessons for Generative AI Policy and Research*, ARXIV 2 (Oct. 31, 2025), https://arxiv.org/pdf/2412.06966. Likewise, while instance attribution could show when an output relied on a particular material, reducing copyright infringement claims where the copyrighted work was not actually relied upon, scaling costs and inaccuracies are high at present. Henderson et al., *supra* note 5, at 30.

368. Lee et al., *supra* note 5, at 358.

369. Cooper et al., *supra* note 367, at 5. One could also engage in structural removal, which can achieve a rough approximation of retraining from scratch, but this technique would still be very expensive and has largely been explored with simpler traditional AI rather than large machine learning-based generative AI systems. *Id.*

370. *Id.* at 6.

371. *See* Dan L. Burk, *Algorithmic Fair Use*, 86 U. CHI. L. REV. 283, 285 (2019) (concluding that algorithmic determination of fair use is practically and technically fraught).

372. Eugene Volokh, Mark A. Lemley & Peter Henderson, *Freedom of Speech and AI Output*, 3 J. FREE SPEECH L. 651, 654, 658–59 (2023); Cass R. Sunstein, *Artificial Intelligence and the First Amendment*, 92 GEO. WASH. L. REV. 1207, 1223 (2024). *But see* Dan L. Burk, *Asemic Defamation, or, the Death of the AI Speaker*, 22 FIRST AMEND. L. REV. 189, 228–29 (2024) (concluding that LLM outputs are not protected speech because there is no (human) speaker or meaning other than that attributed by the listener alone); Peter N. Salib, *AI Outputs Are Not Protected Speech*, 102 WASH. U.L. REV. 83, 87 (2024) (arguing that AI-generated outputs should be placed in less protected First Amendment categories).

373. Henderson et al., *supra* note 5, at 28–30.

374. 17 U.S.C. § 512(m).

375. *See* Eugene Volokh, *Large Libel Models? Liability for AI Output*, 3 J. FREE SPEECH L. 489, 514–15 (2023); Peter Henderson, Tatsunori Hashimoto & Mark Lemley, *Where's the Liability in Harmful AI Speech?*, 3 J. FREE SPEECH L. 589, 642, 649 (2023).

some flexibility in creating a new refinement to the contributory liability test rather than relying on one of the existing options.

A promising refinement option could be what I term "notice-and-revision." Instead of requiring takedowns, this structure would require expeditious *revision* of the model or tool to address the learned-of infringement. Revision could take a variety of forms, including filtering code, adjusting the model, adding training content to lower the chance of a similar output, or providing instructive feedback to the tool. Even if developers endeavor to take some of these actions, these choices are not part of the liability structure at present but merely voluntary *ex post* decisions to prevent similar infringements (and liability for those infringements) in the future. By interpreting notice-and-revision as the material contribution requirement for contributory liability, developers avoid liability for unforeseeable infringements in the first instance and must then revise the system to prevent similar infringements in the future. While the revision will not necessarily prevent all future infringements, even of the same copyrighted work, it puts the onus on the AI system developer to improve the system once they are aware of its specific infringement capabilities. Once a developer is on notice of the specific infringement problem, it is then capable of taking action to adjust the model or system to prevent reoccurrence.[376] The burden would be on the developer to determine which technique would best strike a balance between reducing the likelihood of specific infringements while not unduly restricting related functions. Indeed, this is seemingly what happened with ChatGPT, which mitigated its "Snoopy" problem once it became aware of it.[377] This solution may not be perfect, but it offers a reasonably copyright-protective solution. Thus, notice-and-revision offers a balanced approach that places certain reasonable obligations on the developer only once they are aware of the problem—holding the developer liable where they have intentionally acted or not acted in response to knowledge of specific infringements.

Notice-and-revision provides at least five benefits that help maintain the incentives-access balance: (1) it applies *ex post* restrictions, (2) it accounts for the unpredictability of AI-generated outputs, (3) it targets the most harmful uses of infringing outputs, (4) it provides a more nuanced path to liability that incorporates deterrence goals, and (5) it does not preclude courts or Congress imposing additional *ex ante* industry best practices on developers.

First, direct, principal-agent, or vicarious liability would impose *ex ante* restrictions on AI, and the developer would be held liable for any infringements that occur. But, as explained above, this categorical liability rule would create a significant chilling effect on AI innovation for all but the wealthiest and most

376. *See* Volokh, *supra* note 375, at 514–15 (suggesting that a developer, once on notice, could add code to prevent defamation from reoccurring).

377. Matthew Sag (@matthewsag), X (Feb. 21, 2024, at 14:24 ET), https://twitter.com/matthewsag/status/1760429736344383671.

risk-tolerant companies.[378] Notice-and-revision, however, is more akin to an *ex post* restriction where liability would only be triggered once the developer has *specific* knowledge of the infringement problem and fails to act. General knowledge that an AI system could be used for infringement is insufficient for contributory liability so long as there are substantial noninfringing uses (which should typically be the case with AI systems).[379] Although the Munich I Regional Court found OpenAI directly liable, the court appeared to suggest a problematic rule contrary to specific knowledge, finding that OpenAI already had knowledge of the infringements because it generally knew its models could memorize data since 2021.[380]

Second, to realize the benefits of AI, we must accept that AI-generated outputs will sometimes be unpredictable and account for that in crafting liability rules. But direct, principal-agent, and vicarious liability fail to do this because they would always hold the developer or user liable despite the unpredictability of the outputs. Contributory liability strikes a balance by incorporating ongoing support duties in the sense that the developer must respond to reported infringement or risk being held liable.[381] It does not incorporate monitoring requirements, and more broadly, copyright law has typically placed the duty to monitor on the rights owner. For example, under the DMCA, service providers do not have a duty to monitor their platforms for infringing content.[382] The previously mentioned filter for copyrighted content could be a possible solution, but it would still have its inherent issue of being too blunt an instrument to capture infringement without also ensnaring lawful free expression.[383] Contributory liability accounts for the unpredictable nature of AI-generated outputs and intent by looking to fault rather than imposing a static, stringent test.

Third, the difficulty in identifying infringing outputs may be advantageous for society because only the most harmful uses would be implicated. While AI systems generate millions of outputs, rights owners and others may only be aware of specific infringements in a fairly small subset of circumstances. Those cases are likely to be instances when the output is being used in a prominent and public way, such as for commercial purposes. These are the sort of uses that pose the greatest threat to the copyright owners' own rights, as demonstrated, for example, by the importance of commerciality and market effect in copyright

378. *See supra* notes 300–305 and accompanying text.

379. Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 442 (1984).

380. Landgericht München I [LG] [Munich I Regional Court] Nov. 11, 2025, 42 O 14139/24 ¶¶ 232, 265 (Ger.).

381. Perfect 10, Inc. v. Amazon.com, Inc., 508 F.3d 1146, 1172 (9th Cir. 2007) (quoting Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc., 907 F. Supp. 1361, 1375 (N.D. Cal. 1995)) (holding that contributory liability would lie where a service provider failed to "'take simple measures to prevent further damage' to copyrighted works, . . . yet continue[d] to provide access to [known] infringing works").

382. 17 U.S.C. § 512(m)(1).

383. *See* Ard, *supra* note 113, at 547 (noting the public costs of filtering).

law's fair use analysis.[384] This structure would allow many personal, noncommercial uses to go undetected. These pose less danger to rights owners, and indeed, similar uses outside of the AI context often go undiscovered and unlitigated. Rights owners and courts have even been sympathetic to these types of uses, as shown in *Cablevision* and the Supreme Court's earlier *Sony* decision, where individuals who copied televised content for their own personal enjoyment were not parties to the litigation.[385] Imposing de facto liability rules compared to the more intent-focused contributory liability standard could chill even these types of uses.

Fourth, contributory liability provides a more nuanced path for liability that can help deterrence. The knowledge requirement is premised on the belief that it would not be feasible to require a platform to affirmatively seek knowledge of whether billions of outputs could infringe millions of copyrighted works, but that once the platform knows of a specific infringement, it is reasonable to expect the platform to act. The notice-and-revision structure would require that the developer have knowledge of the specific infringement that is occurring and then revise the algorithms to mitigate this danger. The developer's action or inaction provides indicia of their intent. There are a variety of tools developers can use to limit infringements, but revising the model or system is always an option.[386] If the developer fails to revise the model or system or take some other action, they would be liable for infringement. This would require specific intent to not act knowing that infringement could likely recur, rather than de facto liability for the unforeseeable infringements of an AI system. This encourages deterrence because it focuses on what is known and foreseeable to happen again. It is also a better structural fit for AI technology, which generates unpredictable outputs and is constantly improving, both through machine learning and, critically, human intervention.[387]

Finally, notice-and-revision for contributory liability does not preclude courts or Congress from crafting additional requirements for AI companies to comply with industry best practices to lower the amount of infringements. Prior studies have shown that AI systems have different rates of infringements and regurgitation.[388] Others have discussed potential technical *ex ante* measures

384. *See* Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith, 143 S. Ct. 1258, 1266 (2023) (emphasizing the importance of the use being of a commercial nature under the first fair use factor); Dr. Seuss Enters., L.P. v. ComicMix LLC, 983 F.3d 443, 458–59 (9th Cir. 2020) (discussing the impact of the market effect under the fourth factor).

385. *See* Cartoon Network LP, LLLP v. CSC Holdings, Inc., 536 F.3d 121, 133 (2d Cir. 2008) (holding that RS-DVR customers "made" copies, not Cablevision); Sony Corp. of Am. v. Universal City Studios, Inc., 464 U.S. 417, 434 (1984) ("The two respondents in this case do not seek relief against the Betamax users who have allegedly infringed their copyrights.").

386. Cooper & Grimmelmann, *supra* note 98, at 147.

387. *See* Abraham & Rabin, *supra* note 9, at 141–42 (noting that AI improves all the time, making design defect claims a poor fit).

388. *Introducing CopyrightCatcher*, *supra* note 106; Cooper et al., *supra* note 100, at 12; Ahmed et al., *supra* note 100.

developers could take.[389] While it is beyond the scope of this Article, developers should consider implementing these best practices to the extent they help reduce infringements and improve the novelty of AI systems' outputs. Even if developers do not implement them, they would still have to undertake best efforts to improve their system once they are aware of the specific infringement problem. Notice-and-revision only provides a one-time safe harbor of sorts per infringement issue rather than an excuse to not comply with industry practices.

Beyond the notice-and-revision structure for contributory liability, the inducement doctrine could also be used to capture more nefarious actors that could also qualify as causes of infringement subject to direct liability if the AI system were not held liable. While the AI system could be the direct infringer in all cases, as contemplated above, intentionally infringing developers and users would still be liable under the inducement doctrine. As established by the Supreme Court in *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, one is liable for infringement if they were inducing it through "purposeful, culpable expression and conduct."[390] Therefore, a developer that intends their AI system to be an infringement generation device or a user who actively seeks to have the AI system infringe could be held liable for inducement. While these were identified above as limited cases in which the developer or user could be directly liable because they caused the infringement, inducement would also capture this nefarious conduct.

Overall, contributory liability and inducement provide a promising liability framework for artificial infringement. While it is a departure from the pure direct liability of corporeal and mechanical infringement, the imposition of direct liability on the AI system is grounded in the causation polestar courts and Congress have previously used in response to new technologies. It also operates in tandem with secondary liability refinements to, together, maintain copyright's incentives-access balance despite the acute challenges of artificial infringement.

CONCLUSION

Artificial infringement is the next frontier of copyright law. When artificial infringement litigation advances past fair use considerations—and it will—this interpretation of who should be held liable and under what circumstances will hopefully serve as a useful framework for courts and policymakers. How the law crafts the liability framework for infringing AI-generated outputs could have significant ramifications for the industry and its burgeoning future, as well as copyright law.

As courts and Congress address ever more novel and complex technologies, they must use all the tools at their disposal to maintain copyright's incentives-access balance. Refinements to the direct liability test are an

389. *See, e.g.*, Henderson et al., *supra* note 5, at 3.
390. *Id.* at 937.

underconsidered possibility that previously led to the emergence of the "volition or causation" requirement in mechanical infringement cases involving complex machines such as Internet access and remote DVRs. However, as a more static doctrine, direct liability can only go so far. A proper reading of causation suggests that the AI system (or, failing that, the user) should be held liable as the direct infringer. In addition, making the AI system the direct infringer facilitates the use of more nuanced secondary liability tests and a notice-and-revision refinement to determine user and developer liability for infringing AI-generated outputs. Together with fair use, courts can use direct and secondary liability refinements to maintain copyright's incentives-access balance and not unduly stymie new technologies.

More broadly, this analysis suggests that solutions to new technologies' challenges to copyright need not be cabined to one particular doctrine. Artificial infringement will likely not be the final stage in the evolution of copyright liability. As history shows, new technologies are likely to emerge and once again strain the bounds of copyright law. While it is impossible to predict the future, new advances in AI or sentient technologies could necessitate further refinements to maintain copyright's incentives-access balance. If and when that time comes, courts and Congress should consider how to innovatively combine copyright doctrines like direct liability, secondary liability, fair use, and others to best achieve balanced copyright law.