



NYU | LAW

**Journal of Intellectual Property
& Entertainment Law**

VOLUME 15

NUMBER 1



Statement of Purpose

Consistent with its unique development, the New York University Journal of Intellectual Property & Entertainment Law (JIPEL) is a nonpartisan periodical specializing in the analysis of timely and cutting-edge topics in the world of intellectual property and entertainment law. As NYU's first online-only journal, JIPEL also provides an opportunity for discourse through comments from all of its readers. There are no subscriptions or subscription fees; in keeping with the open-access and free discourse goals of the students responsible for JIPEL's existence, the content is available for free to anyone interested in intellectual property and entertainment law.

The New York University Journal of Intellectual Property & Entertainment Law is published up to three times per year at the New York University School of Law, 139 MacDougal Street, New York, New York, 10012. In keeping with the Journal's open access and free discourse goals, subscriptions are free of charge and can be accessed via www.jipel.law.nyu.edu. Inquiries may be made via telephone (212-998-6101) or e-mail (submissions.jipel@gmail.com).

The Journal invites authors to submit pieces for publication consideration. Footnotes and citations should follow the rules set forth in the latest edition of *The Bluebook: A Uniform System of Citation*. All pieces submitted become the property of the Journal. We review submissions through Scholastica (scholasticahq.com) and through e-mail (submissions.jipel@gmail.com).

All works copyright © 2026 by the author, except when otherwise expressly indicated. For permission to reprint a piece or any portion thereof, please contact the Journal in writing. Except as otherwise provided, the author of each work in this issue has granted permission for copies of that article to be made for classroom use, provided that (1) copies are distributed to students free of cost, (2) the author and the Journal are identified on each copy, and (3) proper notice of copyright is affixed to each copy. A nonpartisan periodical, the Journal is committed to presenting diverse views on intellectual property and entertainment law. Accordingly, the opinions and affiliations of the authors presented herein do not necessarily reflect those of the Journal members.

The Journal is also available on WESTLAW, LEXIS-NEXIS and HeinOnline.

NEW YORK UNIVERSITY

JOURNAL OF INTELLECTUAL PROPERTY
AND ENTERTAINMENT LAW

VOLUME 15

FALL 2025

NUMBER 1

REASONABLE INTELLIGENCE: HOW SHOULD PLATFORMS BE HELD
ACCOUNTABLE FOR AI-ASSISTED COPYRIGHT INFRINGEMENT?

BEN ANDERSON*

*AI platforms present a challenge for our current doctrines of secondary copyright infringement. Currently, AI platforms' degree of exposure to secondary liability depends on the way a platform is structured and presented to users. The staple-article rule as first articulated by the Supreme Court in *Sony Corp. v. Universal City Studios* may shield the developers of models, but only if they lack an ongoing relationship with their users and do not engage in any other conduct that indicates an intent to induce infringement. This presents a unique challenge in the context of generative AI. Maintenance of an ongoing relationship between developers and end-users may be desirable as a means to implement beneficial updates to a platform, but may also expose developers to infringement liability for the acts of their users.*

To address the disconnect between the goals of copyright and a secondary liability rule that did not and could not have contemplated the advent of generative AI, this note proposes replacing the staple article rule with a dynamic system. This dynamic system involves assigning liability based on the availability of reasonable alternative design or process choices available to developers that could have mitigated the potential for a model to facilitate copyright infringement. Though it draws similarities to the framework of design defect in product liability law, a system as applied to generative AI must go beyond static design choices to be effective. The inquiry could include scrutiny

*J.D. Candidate, New York University School of Law, 2026; B.S. in Chemical Engineering and Mathematics, The University of Alabama, 2019. I would like to thank Jessica Mintz, Emily Ko, and Elyse Cox of the NYU Journal of Intellectual Property and Entertainment Law Notes Committee for their thoughtful comments and suggestions which contributed to the development of this Note. I am also grateful to Professor Christopher Sprigman for his comments and feedback, which were indispensable to the early development of this Note.

of the processes that models employ to mitigate infringement, including preventative red-teaming exercises, implementation of input and content filtering, and efforts to terminate repeat infringers' access to the model. Thinking even more broadly, an inquiry could consider as an unreasonable process choice the reluctance of model developers and platforms to allow third party research on their models, which could plausibly identify shortcomings in mitigative efforts that could be addressed.

The proposed framework would not come without challenges. The rapid development and advancement of AI technology could make it difficult to determine what kind of model design or process was available and reasonable at any particular time. These findings would be made all the more difficult in the ecosystem of secrecy that surrounds much of the current AI development. Yet still, this framework may be preferable to a rule which currently slants in favor of heavily capitalized AI developers, and may in fact affirm courts' practical approach to the issue of copyright secondary liability.

INTRODUCTION	437
I. BACKGROUND	439
A. Overview of Generative AI	439
B. Risk of Direct Copyright Infringement by Model Outputs	441
C. Doctrines of Copyright Secondary Infringement	442
D. Available Defenses, Including the Substantial Non-Infringing Use Defense	446
II. ASSESSING SECONDARY LIABILITY	449
A. Essential Elements of Substantial Non-Infringing Use	449
B. Grokster's Emphasis on Intent and Gloss on Sony	450
C. The Nature of the User-Developer Relationship	451
III. RAMIFICATIONS—DOES THE EXISTING FRAMEWORK BALANCE COMPETING INTERESTS?	457
A. Identifying and Evaluating the Public Interest Implications	457
B. Critiquing Sony and Exploring Alternative Approaches	459
IV. WHAT COULD AN EFFECTIVE REGULATORY REGIME LOOK LIKE?	462
A. Availability of the DMCA and its Limitations in Policing Generative AI	462
B. RAD-Plus as an Approach to Effective Generative AI Regulation	467
CONCLUSION	470

INTRODUCTION

The rapid advancement of generative artificial intelligence (AI) has transformed content creation, enabling users of AI to generate outputs including

text, code, images, video, and music with unprecedented ease. The capabilities of AI are difficult to keep up with, and what may be a challenge today could be commonplace in a few months' time.

The increasing capabilities of AI systems and reliance on them pose risks across society.¹ This note will be limited to addressing potential harms in the context of copyright. AI's ever-increasing ability to produce outputs that closely resemble expressions found in copyrighted works has sparked significant legal debate. One underdeveloped element of the ongoing discussion centers around the potential liability of AI model developers and platforms for secondary copyright infringement on account of any direct infringements perpetrated by the models' end-users.

In Section I, this note will first give a brief overview of the generative AI landscape, including forays into the legal doctrines pertinent to the issue of secondary liability for AI model developers and platforms. In Section II, the note will then examine which AI developers and platforms may be able to invoke the "substantial non-infringing use" defense to avoid secondary liability, as established in *Sony Corp. of America v. Universal City Studios, Inc.*, 464 U.S. 417 (1984), and why others may be left vulnerable to secondary copyright infringement liability. In Section III, a reflection on the result of applying contemporary doctrine, including the *Sony* rule, follows. The reflection includes an assessment as to whether the

¹ See, e.g., Jackson Cote, *Deepfakes and fake news pose a growing threat to democracy, experts warn*, NORTHEASTERN GLOBAL News (Apr. 1, 2022), <https://news.northeastern.edu/2022/04/01/deepfakes-fake-news-threat-democracy/> [https://perma.cc/U2N9-5CWD] (documenting the risk of social manipulation through use of deepfakes created using AI); Matthew Tokson, *The Authoritarian Risks of AI Surveillance*, LAWFARE (May 1, 2025), <https://www.lawfaremedia.org/article/the-authoritarian-risks-of-ai-surveillance> [https://perma.cc/B383-BFWF] (discussing risks of AI-powered surveillance and its use as a tool in authoritarian regimes globally); Sam Manning, *AI's impact on income inequality in the US*, BROOKINGS (July 3, 2024), <https://www.brookings.edu/articles/ais-impact-on-income-inequality-in-the-us/> [https://perma.cc/9FSX-CJDG] (noting the risk that embracement of AI perpetuates and widens income inequality by, e.g., disproportionately benefitting high-income workers); Amanda Hess, *They're Stuffed Animals. They're Also A.I. Chatbots.*, THE NEW YORK TIMES (Aug. 15, 2025), <https://www.nytimes.com/2025/08/15/arts/ai-toys-curio-grem.html> [https://perma.cc/DZK5-GGGX] (documenting risks of AI to child safety, such as by use of AI chatbots in stuffed animals and use of children's data by third-party companies); Ted A. James, *Confronting the Mirror: Reflecting on Our Biases Through AI in Health Care*, HARVARD MEDICAL SCHOOL (Sep. 24, 2024), <https://learn.hms.harvard.edu/insights/all-insights/confronting-mirror-reflecting-our-biases-through-ai-health-care> [https://perma.cc/9SDA-HVJR] (discussing risks of AI in healthcare, including through its potential to perpetuate human biases and systemic flaws in the provision of care).

result aligns with the primary goals of copyright and the public interest. A brief discussion of *Sony*'s critiques follows, and an alternative “reasonable alternative design-plus” approach to analyzing secondary infringement liability, grounded in principles of tort law, is proposed and considered. Finally, in Section IV, the note explores and evaluates existing and potential legislative and regulatory approaches to address secondary infringement liability for AI model developers and platforms.

At the outset, it's also worth noting what this note does *not* cover, at least with any rigor. This note does not attempt to answer the open question of whether the entities behind the creation and distribution of AI models are liable for direct copyright infringement through the process of training, including by curating a training dataset for a model. Likewise, this note does not fully examine secondary liability theories outside the scope of copyright, such as those stemming from potential trade dress infringement, or those outside the purview of what's conventionally considered to be intellectual property altogether. To reiterate, this note aims to assess the potential liability of AI model developers and hosts for secondary copyright infringement, stemming from the potential direct infringement liability of the model's users in creating infringing material via a generative AI model, as well as remedies through legislative action or judicial re-evaluation.

I BACKGROUND

A. *Overview of Generative AI*

Generative Artificial Intelligence (AI) refers to machine learning models designed to generate text, images, audio, and other content based on user input. These models are typically trained on vast datasets and can produce outputs, responsive to a model user's inputs, that can resemble human-created works.²

There are countless generative AI models available for use today, having varying degrees of ability, and being produced by a gamut of developers. Key players behind some of the most widely adopted generative AI models include OpenAI (ChatGPT models, DALL-E, Sora), Google (Gemini), Anthropic (Claude), Meta (Llama), Microsoft (Copilot), and Stability AI (Stable Diffusion).

² *How Does Generative AI Work?*, MICROSOFT, <https://www.microsoft.com/en-us/ai/ai-101/how-does-generative-ai-work> [https://perma.cc/T8PD-NMAA] (last visited Nov. 9, 2025).

This is far from an exhaustive list—there are countless other models available for a number of use cases, often tailored to more specialized or niche markets (AI song or audio stem [drums, vocals, etc.] generation, for example).³

A key element for generative AI to properly function is software. Software enables end-users to interact with the model, regardless of where the necessary computation occurs. Take OpenAI’s LLM offerings as an example. The majority of these models are available to end-users through a web browser or a native application, whether accessed on a mobile device or computer. For most of OpenAI’s LLMs, however, the user’s hardware—the mobile device or the computer—does not carry out the actual computation associated with the user’s input and subsequent model output. That computation is typically completed at a datacenter, by a computer equipped with robust hardware that enables computation that the user’s device is simply not capable of performing locally.⁴ In these cases, model developers and platforms must provide software that links users’ own devices to the computing power and other resources required to carry out their requests, a model comporting with what’s generally known as software as a service (SaaS).⁵ Even for smaller-scale models where compute occurs locally on an end-user’s device (and thus not provided as SaaS), a software interface is generally necessary.

Generative AI platforms vary in product offerings, licensing models, terms of use, and the degree of openness in prompting their AI systems permit, among other differences. The models which capture the largest AI market share among consumers are by-and-large closed, proprietary models.⁶ These models are typically closely monitored by the providers, with user inputs and model outputs subject to greater scrutiny. Some other models, particularly open-source models (though what it means to be “open-source” in the realm of generative AI is itself up for debate), allow users greater flexibility in model prompting, and provide

³ Mark Wilson, *Suno explained: How to use the viral AI song generator for free*, TECHRADAR (Feb. 14, 2025), <https://www.techradar.com/computing/artificial-intelligence/what-is-suno-ai> [https://perma.cc/2W3E-RRGN].

⁴ Mary Zhang, *ChatGPT and OpenAI’s use of Azure’s Cloud Infrastructure*, DGTL INFRA (Jan. 26, 2023), <https://dgtdinfra.com/chatgpt-openai-azure-cloud/> [https://perma.cc/Y95V-WUGV].

⁵ TOD GOLDING, BUILDING MULTI-TENANT SaaS ARCHITECTURES 14 (O’Reilly Media 2024).

⁶ Assad Abbas, *The AI Monopoly: How Big Tech Controls Data and Innovation*, UNITE.AI (Dec. 27, 2024), <https://www.unite.ai/the-ai-monopoly-how-big-tech-controls-data-and-innovation/> [https://perma.cc/9L6A-2N6J].

subsequent developers greater freedom to implement modifications to the model.⁷ Popular “open-source” models include Meta’s Llama series of models, Stability AI’s Stable Diffusion, and xAI’s Grok.

B. Risk of Direct Copyright Infringement by Model Outputs

On account of their training and responsive to user inputs, generative AI models can produce or direct users to content that may infringe the exclusive rights associated with existing copyrighted works, whether it be the right of reproduction, or the derivative work right.⁸ This may occur either through direct replication of training data—a problem known as overfitting—or through their normal operation, often at the behest of creative prompting by users.⁹

These risks span the creative arts implicating protected literary works (e.g., AI generating an article that mimics a copyrighted newspaper editorial or linking to unauthorized copies of protected books), visual works (see OpenAI’s new GPT-4o image generation capabilities for a current example, and the explosion of visual content raising numerous questions of infringement),¹⁰ musical works (e.g., AI composing melodies that mirror existing copyrighted songs), or even the narrower copyright protection afforded to computer code (e.g., AI reproducing proprietary software code snippets). Though it’s unlikely model developers or platforms possess the required volition to bring claims of direct infringement against them, they remain attractive targets to content owners in part due to their large public presence and financial capacity to pay out potential claims, being positioned similarly to internet service providers from this perspective.¹¹

⁷ *The Open Source AI Definition – 1.0-RC2*, OPEN SOURCE INITIATIVE, <https://opensource.org/ai/drafts/the-open-source-ai-definition-1-0-rc2> [https://perma.cc/7788-6DKP] (last visited Nov. 9, 2025).

⁸ Brian Moriarty et al., *Digital Image Creation Using AI Risks Copyright Infringement*, BLOOMBERG LAW (Sept. 16, 2024), <https://news.bloomberg.com/us-law-week/digital-image-creation-using-ai-risks-copyright-infringement> [https://perma.cc/9SG3-5YJ7].

⁹ See Matthew Sag, *Copyright Safety for Generative AI*, 61 Hous. L. Rev. 295, 312 (2023) (“If the model memorizes the training data, it might communicate original expression from the training data via its output.”); Joseph D’alfonso, *Generative artificial intelligence outputs, copyright infringement, and the assignment of liability*, 5 AI and Ethics 5295, 5300 n.16 (2025) (noting end-users’ ability to circumvent content guardrails).

¹⁰ See Maxwell Zeff, *OpenAI’s viral Studio Ghibli moment highlights AI copyright concerns*, TECHCRUNCH (Mar. 26, 2025), <https://techcrunch.com/2025/03/26/openais-viral-studio-ghibli-moment-highlights-ai-copyright-concerns/> [https://perma.cc/GK9X-TD2A].

¹¹ *Commerce and Cyberspace – Understanding Why ISPs Are Frequently Sued for Copyright Infringements*, THELAW.INSTITUTE (Dec. 27, 2023), <https://thelaw.institute/commerce-and-cyberspace/isps-sued-for-copyright-infringements/> [https://perma.cc/WBY8-9BAD].

The risk of direct infringement may be mitigated in some respects by model developers. As to problems associated with overfitting, one identified solution has been to engage in de-duplication of specific works in a model's training data, directly remediating the risk that duplicates in a model's training data poses for overfitting or model memorization of specific works.¹²

Model platforms may also regulate their models in action in various ways. One example is by blocking certain user inputs that would lead to a heightened risk of generating outputs that could infringe existing works, or augmenting them to reduce the likelihood of generating infringing output.¹³ Platforms may also review outputs prior to transmitting them to an end-user, preventing access to a potentially infringing output.¹⁴ Developers may employ a “red-teaming” strategy, attacking a model with novel and creative prompting strategies in an attempt to manipulate AI models and identify weaknesses, which may be addressed by the developer.¹⁵ These protections may broadly be considered generative AI “guardrails”.¹⁶ Courts, regulators, and AI stakeholders continue to debate how to assess infringement-related risks and best respond to them.

C. Doctrines of Copyright Secondary Infringement

Secondary copyright infringement comes in a number of different flavors. Liability may be imposed under theories of contributory copyright infringement, vicarious copyright infringement, and active inducement of infringement. Though none of these doctrines are found expressly in the Copyright Act, they all have gained recognition in copyright jurisprudence.¹⁷

¹² See Sag, *supra* note 9, at 338–39 (“Deduplication will not only reduce the likelihood of downstream copyright infringement, it will also mitigate privacy and security risks and reduce the cost of training.”).

¹³ Shunchang Liu et al., *CopyJudge: Automated Copyright Infringement Identification and Mitigation in Text-to-Image Diffusion Models*, in ARXIV, at 2 (2025), <https://arxiv.org/pdf/2502.15278.pdf> [<https://perma.cc/MSF2-QFXS>] (discussing the “CopyJudge” approach and methods to prevent backdoor access to disallowed generations).

¹⁴ See *Content Filtering*, MICROSOFT (Sept. 16, 2025), <https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/content-filter?tabs=warning%2Cuser-prompt%2Cpython-new> [<https://perma.cc/J8XD-XJQN>].

¹⁵ Leon Derczynski et al., *Defining LLM Red Teaming*, NVIDIA DEVELOPER (Feb. 25, 2025), <https://developer.nvidia.com/blog/defining-llm-red-teaming/> [<https://perma.cc/D4H8-GA75>].

¹⁶ Drishti Shah, *What are AI guardrails?*, PORTKEY BLOG (Jan. 6, 2025), <https://portkey.ai/blog/what-are-ai-guardrails/> [<https://perma.cc/XR3Q-JTFK>].

¹⁷ See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 434 (1984) (“The Copyright Act does not expressly render anyone liable for infringement committed by another.”). *But see* Peter S. Menell

In the area of copyright, contributory infringement may be defined as assigning liability when a party (1) has knowledge of the infringing activity, and (2) induces, causes, or materially contributes to the infringing conduct of another.¹⁸ What constitutes “material” contribution is an open question in many cases, but must be more than “mere quantitative contribution.”¹⁹ Resolution of the issue depends upon a determination of the function that the alleged contributory infringer plays in the total reproduction process.²⁰ Specific theories or subcategories of contributory copyright infringement will be expanded on throughout this note.

Vicarious infringement has a long history in common law jurisprudence, and is derived from the doctrine of agency, holding a superior responsible for the acts of their subordinate. In modern jurisprudence, vicarious liability may be imposed on a party despite a lack of knowledge of specific infringing acts, so long as the party has the right and ability to supervise and derives a direct financial benefit from the infringement.²¹

When assessing potential vicarious infringement liability, analogies can be made to either the “landlord-tenant” model, where a lack of supervision and no direct benefit from infringement to the secondary party cautions against imposing liability,²² or an expanded “employer-employee” model, where “a right and ability to supervise coalesce with an obvious and direct financial interest in the exploitation of copyrighted materials” may suggest finding liability.²³ In today’s digital environment, the framework has expanded to cover models where a business

and David Nimmer, *Unwinding Sony*, 95 CALIF. L. REV. 941, 977 (2007) (“[T]he language of the Copyright Act and its legislative history establish that the Copyright Act does expressly render some actors liable for infringement committed by another.”).

¹⁸ See *Gershwin Publ’g Corp. v. Columbia Artists Mgmt.*, 443 F.2d 1159, 1162 (2d Cir. 1971).

¹⁹ See *Fortnightly Corp. v. United Artists Television, Inc.*, 392 U.S. 390, 396–97 (1968). The standard for “material contribution” is currently before the Supreme Court, in the context of an internet service provider being held liable for the infringing acts of its users. *Cox Commn’s, Inc. v. Sony Music Ent.*, No. 24-171 (U.S. 2025).

²⁰ See *Fortnightly Corp.*, 392 U.S. at 396–97.

²¹ See *Shapiro, Bernstein and Co. v. H.L. Green Co.*, 316 F.2d 304, 307 (2d Cir. 1963).

²² See *Deutsch v. Arnold*, 98 F.2d 686, 688 (2d Cir. 1938).

²³ See *Shapiro*, 316 F.2d at 307–08.

may be held liable for the acts of its customers, so long as the business derives a direct financial benefit.²⁴

Active inducement of infringement, as articulated by the court in *MGM Studios, Inc. v. Grokster, Ltd.*, may constitute a third discrete theory of secondary liability.²⁵ However, the court itself in *Grokster* noted that “[o]ne infringes contributorily by intentionally inducing or encouraging direct infringement.”²⁶ It may be better understood as a class of contributory infringement, re-asserting the principle laid out in *Gershwin* that a party’s inducement (or encouragement) of infringing conduct by another may be the basis for finding liability.²⁷ “Mere knowledge of infringing potential or of actual infringing uses would not be enough[] to subject a distributor to liability. Nor would ordinary acts incident to product distribution The inducement rule, instead, premises liability on purposeful, culpable expression and conduct, and thus does nothing to compromise legitimate commerce or discourage innovation having a lawful promise.”²⁸

In the context of generative AI, model developers and platform hosts may be liable for secondary infringement stemming from the direct infringement of model outputs prompted by users. Though it is context-dependent, such direct infringement by users could form the basis for any of the preceding doctrines of secondary copyright infringement.

Several high-profile lawsuits currently working their way through the courts include claims against various generative AI platforms for secondary infringement: some have survived defendants’ efforts to dismiss, and some have not. In *Kadrey v. Meta*, plaintiffs initially alleged vicarious copyright infringement on a theory that “every output of [Meta’s] LLaMA language models is an infringing derivative work” while Meta has the right and ability to control the output of its LLMs while benefitting financially from the infringing outputs.²⁹ All claims in the original complaint—aside from direct infringement—were dismissed, with the

²⁴ See A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1022–23 (9th Cir. 2001) (stating “vicarious liability extends beyond an employer/employee relationship” and finding Napster was likely vicariously liable for the infringing acts of its users); *id.* at 1027.

²⁵ 545 U.S. 913, 936 (2005).

²⁶ *Id.* at 930.

²⁷ See *Gershwin Publ’g Corp. v. Columbia Artists Mgmt.*, 443 F.2d 1159, 1162 (2d Cir. 1971).

²⁸ *Grokster*, 545 U.S. at 937.

²⁹ See Complaint at ¶ 44, *Kadrey v. Meta Platforms, Inc.*, No. 3:23-cv-03417 (N.D. Cal. Jul. 7, 2023).

court discrediting the novel theory underpinning the vicarious infringement claim.³⁰

In *Andersen v. Stability AI*, plaintiff artists similarly alleged that every output image of the Stable Diffusion AI system is necessarily a derivative work of the works in its training data.³¹ Similar to the outcome in *Kadrey*, skepticism has been leveled at this novel theory of liability.³² In an amended complaint, the *Andersen* class action plaintiffs subsequently added claims of inducement of copyright infringement,³³ which have survived Defendants' motions to dismiss.³⁴

Lastly, in *New York Times v. Microsoft*,³⁵ the plaintiffs alleged contributory copyright infringement attributed to the direct infringing acts perpetrated by model end users, premised on the capability of GPT-based products to infringe, as well as subsequent measures taken by the defendants pertaining to design choice.³⁶ The plaintiffs' claims of contributory infringement survived the district judge's April 2025 court order.³⁷

³⁰ See *Kadrey v. Meta Platforms, Inc.*, No. 23-CV-03417-VC, 2023 WL 8039640, at *1 (N.D. Cal. Nov. 20, 2023) ("[T]he complaint offers no allegation of the contents of any output, let alone of one that could be understood as recasting, transforming, or adapting the plaintiffs' books. Without any plausible allegation of an infringing output, there can be no vicarious infringement."). No amended claims for secondary copyright infringement have been brought forth in the time since.

³¹ See Complaint at ¶ 95, *Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201 (N.D. Cal. Jan. 13, 2023).

³² See *Andersen v. Stability AI Ltd.*, 700 F. Supp. 3d 853, 868 (N.D. Cal. 2023) ("Defendants make a strong case that I should dismiss the derivative work theory without leave to amend because plaintiffs cannot plausibly allege the Output Images are substantially similar or re-present protected aspects of copyrighted Training Images, especially in light of plaintiffs' admission that Output Images are unlikely to look like the Training Images.").

³³ See First Amended Complaint at ¶¶ 232–37, 354–59, *Andersen v. Stability AI Ltd.*, No. 3:23-cv-00201 (N.D. Cal. Nov. 29, 2023).

³⁴ See *Andersen v. Stability AI Ltd.*, 744 F. Supp. 3d 956, 969 (N.D. Cal. 2024); *see also id.* at 975 (denying co-defendant Runway AI's motion to dismiss claims of induced copyright infringement).

³⁵ In April 2025, this case was consolidated into MDL No. 3143 — *In re: OpenAI, Inc. Copyright Infringement Litigation*. *In re OpenAI, Inc., Copyright Infringement Litig.*, 776 F. Supp. 3d 1352 (U.S. J.P.M.L. 2025)

³⁶ See Complaint at ¶ 179, *New York Times Co. v. Microsoft Corp.*, No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023).

³⁷ See *New York Times Co. v. Microsoft Corp.*, 777 F. Supp. 3d 283, 309 (S.D.N.Y. 2025). The Times' subsequent amended complaint in the MDL maintained its claims for both vicarious and contributory infringement. *See Second Amended Complaint at ¶¶ 176–80, 185–87, In re OpenAI, Inc., Copyright Infringement Litig.*, No. 1:25-md-03143 (S.D.N.Y. May 28, 2025).

D. Available Defenses, Including the Substantial Non-Infringing Use Defense

Since its release, the holding in *Sony* has remained an influential and attractive defense to a finding of secondary liability based on the theory that a defendant's sale of a product that facilitates direct infringement of a copyright renders them secondarily liable.

The court grounded its holding in *Sony* on ideas borrowed from patent law, where it was recognized that the typical remedy for a finding of contributory infringement is an injunction on the sale of the product facilitating infringement, despite there being a lack of intellectual property right covering the facilitating article in question.³⁸ An injunction would impair the public interest in access to the facilitating product, giving the patentee whose rights were infringed control over sale of the product.³⁹ Given this practical expansion of the scope of the underlying monopoly granted by the patent beyond its specific grant, courts have denied patentees any right to control the distribution of unpatented articles unless they are "unsuited for any commercial noninfringing use."⁴⁰ This approach to secondary patent infringement is codified at 35 U.S.C. § 271(c).

The *Sony* court reasoned that despite "substantial differences between the patent and copyright laws", a liability rule must be struck in both areas that balances effective protection of the statutory monopoly and the rights of others to engage in substantially unrelated areas of commerce.⁴¹ The court ultimately held that "the sale of ... articles of commerce[] does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes. Indeed, it need merely be capable of substantial noninfringing uses."⁴² This rule has come to be known as the "staple-article rule."

In the context of generative AI, questions remain as to the staple-article rule's effectiveness in insulating model developers and platforms from secondary infringement liability. Some of these questions will be explored later in this note.⁴³

³⁸ See *Dawson Chem. Co. v. Rohm & Hass Co.*, 448 U.S. 176, 198 (1980).

³⁹ *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 440 (1984).

⁴⁰ *Dawson Chem.*, 448 U.S. at 198.

⁴¹ *Sony*, 464 U.S. at 442.

⁴² *Id.*

⁴³ See *infra* Section II.

The staple-article rule was invoked by OpenAI and Microsoft in their ongoing litigation against the New York Times. In support of their initial motion to dismiss, the defendants argued that the characteristics of an accused product, standing alone, cannot impute the culpable intent required to find defendants contributorily liable.⁴⁴ Denying defendants' motions to dismiss the Times' claims of contributory infringement, Judge Stein reasoned that both *Sony* and *Grokster* did not foreclose the Times' contributory claims, indicating that a "material contribution" theory is available while also noting the presence and importance of the "ongoing relationship between direct infringer and contributory infringer at the time the infringing occurred".⁴⁵

Beyond *Sony*, other defenses against secondary liability are available to parties accused of secondary copyright infringement. A defendant may argue, for example, that they had no knowledge of the infringing acts forming the basis for a contributory liability claim.⁴⁶ The specific metes and bounds of the requisite knowledge required to fulfill this element, however, is a matter of disagreement and ongoing debate.

This lack of knowledge defense was also invoked by OpenAI and Microsoft against the New York Times. In their initial motion to dismiss,⁴⁷ the defendants argued that they lacked actual knowledge of specific acts of infringement, which they argue are baseline requirements, with "generalized knowledge" of "the possibility of infringement" not being enough.⁴⁸

⁴⁴ See Memorandum of Law in Support of Defendants' Motion to Dismiss at 16, New York Times Co. v. Microsoft Corp., No. 1:23-cv-11195 (S.D.N.Y. Feb. 26, 2024).

⁴⁵ See New York Times Co. v. Microsoft Corp., 777 F. Supp. 3d 283, 309 (S.D.N.Y. 2025) ("[I]n *Sony* there was no ongoing relationship between the direct infringer and the contributory infringer at the time the infringing conduct occurred. Here, however, an ongoing relationship exists between defendants and end users.") (internal quotations and citations omitted). In the MDL, Microsoft and OpenAI have maintained the "substantial noninfringing uses" defense in their most recent answer to the Times' amended complaint. See Microsoft Answer to N.Y. Times Co. Amended Compl. at 30, In re OpenAI Copyright Infringement Litig., No. 1:25-md-03143 (S.D.N.Y. Jun. 11, 2025); see also OpenAI Answer to N.Y. Times Co. Second Amended Compl. at 36, In re OpenAI Copyright Infringement Litig., No. 1:25-md-03143 (S.D.N.Y. Jun. 11, 2025).

⁴⁶ See *Gershwin Publ'g Corp. v. Columbia Artists Mgmt.*, 443 F.2d 1159, 1162 (2d Cir. 1971) (indicating knowledge of infringing activity is required to be held liable as a contributory infringer).

⁴⁷ See Memorandum of Law in Support of Defendants' Motion to Dismiss at 16–17, New York Times Co. v. Microsoft Corp., No. 1:23-cv-11195 (S.D.N.Y. Feb. 26, 2024).

⁴⁸ See *Luvdarts LLC v. AT & T Mobility, LLC*, 710 F.3d 1068, 1072 (9th Cir. 2013).

The Times predictably disputed this interpretation of the knowledge requirement, and in its opposition⁴⁹ argued that knowledge of specific infringements is not required to support a finding of contributory infringement, with a finding that the defendant knew or should have known that its service would encourage infringement.⁵⁰ Ultimately, the Times's contributory infringement claims survived the defendants' motions to dismiss, with Judge Stein finding the allegations sufficient at the pleading stage to establish a plausible inference that defendants possessed actual or constructive knowledge of third-party infringement.⁵¹

Given that knowledge of infringing acts is not a requirement for a finding of vicarious liability,⁵² a lack of knowledge is cabined to defending against allegations of contributory liability. However, that does not render knowledge irrelevant in an assessment of vicarious liability. For example, a lack of knowledge could plausibly serve to corroborate a claim that a defendant lacked the requisite supervision or degree of control necessary to support a vicarious liability claim.

Speaking of supervision and control, *Sony* also holds in part that contributory and vicarious infringement liability is conditioned on the secondary infringer being in a position to control the use of copyrighted works by others, and therein authorizing such use without permission from the copyright owners.⁵³ The court in *Sony* unfortunately does not draw clear distinctions between what constitutes "contributory" and "vicarious" infringement, going so far as to admit the fuzziness in demarcation of each avenue of infringement.⁵⁴ In modern parlance, this discussion can be deemed to apply most directly to what we know as vicarious copyright infringement.

In the generative AI context, this means that some developers may claim they lack an ongoing relationship with end-users who generate infringing content, and do not have the ability to effectively supervise the acts of end-users, in order to

⁴⁹ See Opposition to Motion to Dismiss at 9, *New York Times Co. v. Microsoft Corp.*, No. 1:23-cv-11195 (S.D.N.Y. Mar. 11, 2024).

⁵⁰ See *Arista Recs. LLC v. Usenet.com, Inc.*, 633 F. Supp. 2d 124, 154 (S.D.N.Y. 2009); *see also Capitol Recs., LLC v. ReDigi Inc.*, 934 F. Supp. 2d 640, 658 (S.D.N.Y. 2013).

⁵¹ See *New York Times Co. v. Microsoft Corp.*, 777 F. Supp. 3d 283, 307 (S.D.N.Y. 2025).

⁵² See *Shapiro, Bernstein and Co. v. H.L. Green Co.*, 316 F.2d 304, 307 (2d Cir. 1963).

⁵³ See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 437 (1984).

⁵⁴ See *id.* at 435 n.17.

escape vicarious liability. In effect, developers would be arguing their models are more akin to a discrete product (such as a videocassette recorder), rather than provision of a service. Although generative AI models take many different shapes and applications, open-source models with a comparative lack of oversight are the most likely to make such arguments.

The foregoing list is non-exhaustive. Other potential defenses abound, the availability of which will of course depend on facts pertinent to a particular dispute.

II

ASSESSING SECONDARY LIABILITY

A. *Essential Elements of Substantial Non-Infringing Use*

The cornerstone of the *Sony* rule's availability to generative AI model developers and platforms centers on the question of whether AI systems meet the standard of "substantial non-infringing use."

First, are AI models and services "staple articles introduced into the stream of commerce"? We may look to *Grokster* for clarity here, as that case was centered around software, like any generative AI case would be. *Grokster* was concerned with software which facilitated peer-to-peer (P2P) file sharing. The court found no issue in deeming the P2P software was a product, and was distributed by *Grokster*.⁵⁵ In concurrence, Justice Breyer noted that *Grokster*'s "product"—peer-to-peer software—passes *Sony*'s test.⁵⁶ Implicit in this conclusion is a finding that the P2P software was an article introduced into the stream of commerce. Nothing in this reasoning suggests that AI software would fare any differently as to classification as an article introduced into the stream of commerce.

Next, the actual use cases for generative AI models must be accounted for. *Sony* only requires that the product in question is capable of commercially significant noninfringing uses. Unless the untenable position initially advanced by the plaintiffs in *Kadrey* or *Andersen* is taken—that *all* generative AI outputs are necessarily infringing derivative works (a position which was discredited by the

⁵⁵ See *MGM Studios, Inc. v. Grokster, Ltd.*, 545 U.S. 913, 919–20 (2005).

⁵⁶ See *id.* at 952 (Breyer, J., concurring).

court and subsequently withdrawn by the plaintiffs)⁵⁷—generative AI models will be capable of noninfringing uses.

To determine whether those noninfringing uses are commercially significant, the proportion of infringing uses to noninfringing uses may be considered. Concurring in *Grokster*, Justice Breyer found it sufficient that around 10% of exchanged files were non-infringing, a figure close to the “9% or so of authorized time-shifting uses of the VCR that the Court faced in *Sony*.⁵⁸ In *Sony*, this amount of programming was, by itself, deemed to be significant.⁵⁹

Though there are no analogous figures to reference in the context of generative AI, it’s likely a safe bet that *at least* an equal proportion of uses of generative AI models (and probably far greater) result in non-infringing outputs, given the wide range of applications for generative AI, from assisting in academic research, to creating original artwork, to generating and troubleshooting computer programming, and beyond. Thus, it’s likely that most generative AI models meet the threshold laid out in “substantial non-infringing use” jurisprudence and would be eligible for defense under *Sony*’s staple-article rule. However, any platforms that are purposed to explicitly facilitate content replication—that is, services that are “good for nothing else” but infringement—would be liable for contributory copyright infringement under *Sony*.⁶⁰

B. *Grokster’s Emphasis on Intent and Gloss on Sony*

The Supreme Court’s 2005 ruling in *MGM Studios v. Grokster* clarified the *Sony* analysis by holding that the *Sony* safe harbor does not preclude liability for contributory copyright infringement on any theory; rather, it shields against liability based on presuming or imputing intent to cause infringement solely from the design or distribution of a product capable of substantial lawful use, which the distributor knows is in fact used for infringement.⁶¹

⁵⁷ See *Kadrey v. Meta Platforms, Inc.*, No. 23-CV-03417-VC, 2023 WL 8039640, at *1 (N.D. Cal. Nov. 20, 2023); *see also* *Andersen v. Stability AI Ltd.*, 700 F. Supp. 3d 853, 868 (N.D. Cal. 2023).

⁵⁸ *Grokster*, 545 U.S. at 952 (Breyer, J., concurring).

⁵⁹ *See Sony*, 464 U.S. at 444.

⁶⁰ *See Grokster*, 545 U.S. at 932 (citing *Canada v. Michigan Malleable Iron Co.*, 124 F. 486, 489 (6th Cir. 1903)).

⁶¹ *Id.* at 933.

As a result, a product manufacturer/distributor or service provider cannot argue that *Sony* precludes *any* finding of secondary liability simply because the product or service offered is capable of substantial noninfringing use. If evidence shows statements or actions directed to promoting infringement, *Sony*'s staple-article rule will not preclude liability.⁶²

In the context of generative AI, the key is thus determining whether an AI developer's actions could be seen as "inducing" infringement. For example, if an AI company markets its platform as a tool for generating content similar to copyrighted works, *Grokster*'s holding may suggest imposing contributory liability under a theory of active inducement. Conversely, if an AI system is positioned primarily for research, education, or general creative use, a *Sony*-based defense could be stronger. Courts have interpreted acts pertinent to inducement as a fairly broad category, which has included choices to implement or abstain from implementing design features.⁶³

Ultimately, applying *Sony* to generative AI will likely vary in its effectiveness on a case-by-case basis, heavily influenced by developer intent, product design, use cases, and preventative measures taken to curb or mitigate infringement. What is clear, however, is that *Sony* does not serve as a panacea to AI model developers and platforms hoping to make use of it to shield themselves from any instance of secondary liability.

C. *The Nature of the User-Developer Relationship*

Some proponents of AI argue that the secondary liability question should begin and end with *Sony*. In comments to the Copyright Office, for example, some groups argued that liability for AI-generated material should lie on the end-user,

⁶² *Id.* at 935.

⁶³ *Id.* at 916 (finding the fact that *Grokster* did not attempt to develop filtering tools or other mechanisms to diminish the infringing activity using their software to be evidence of intent to induce infringement); *see also* Arista Recs. LLC v. Lime Grp., 784 F.Supp.2d 398, 430 (S.D.N.Y. 2011) (finding Lime's decision to turn off hash-based filtering by default to be a conscious design choice, resulting in a failure to mitigate infringement and serving as evidence of Lime's intent to induce infringement); Andersen v. Stability AI Ltd., 744 F. Supp. 3d 956, 969 (N.D. Cal. 2024) ("The plausible inferences at this juncture are that Stable Diffusion by operation by end users creates copyright infringement and *was created to facilitate that infringement by design.*"') (emphasis added).

while analogizing AI systems to consumer products such as VCRs.⁶⁴ Other groups, often representing the interests of artists and other rights holders, stressed the need to look beyond the AI “product” itself, given the ongoing relationship between developers and users.⁶⁵

By any measure, those groups touting *Sony* as the be-all, end-all haven’t laid out a complete picture. A crucial factor in *Sony*’s applicability—or perhaps more specifically in its effectiveness as a defense—is the relationship between users and AI developers or platforms at the moment direct infringement occurs. *Sony* dealt with a technology (the videocassette recorder) where users had near-total autonomy in how they utilized the product. After the point of sale, *Sony* did not remain involved with its customers in any material way.⁶⁶ Factual findings at the district court also indicated that *Sony* did not market its product in a manner so as to induce or encourage its customers to use it to create unauthorized copies of protected works.⁶⁷ The court contrasted these facts with prior holdings affirming findings of secondary liability, where a “contributory” infringer (we would likely deem them a vicarious infringer today) was in a position to control the use of copyrighted works by others and had authorized the use without permission from the copyright owner.⁶⁸

In short, due to the lack of an ongoing relationship, the only way the *Sony* plaintiffs could plausibly allege a salient theory of secondary infringement was

⁶⁴ See U.S. Copyright Off., Comment Letter of CCIA Pursuant to Request for Comments on Artificial Intelligence and Copyright at 21 (2023), <https://ccianet.org/wp-content/uploads/2023/10/CCIA-Comments-to-Copyright-Office-on-AI.pdf> [https://perma.cc/2ATN-QBND]; U.S. Copyright Off., Comment Letter of EFF Pursuant to Request for Comments on Artificial Intelligence and Copyright at 5–6 (2023), https://www.eff.org/files/2023/11/08/comments_of_eff_to_copyright_office_re_generative_ai.pdf [https://perma.cc/XMD9-ZEAW].

⁶⁵ See U.S. Copyright Off., Comment Letter of NMA Pursuant to Request for Comments on Artificial Intelligence and Copyright at 17–18 (2023), <https://www.newsmediaalliance.org/wp-content/uploads/2023/12/NMA-Reply-to-USCO-AI-Notice-December-2023.pdf> [https://perma.cc/H5UZ-QC63]; U.S. Copyright Off., Comment Letter of RIAA Pursuant to Request for Comments on Artificial Intelligence and Copyright at 24–25 (2023), <https://copyrightalliance.org/wp-content/uploads/2023/11/A2IM-and-RIAA-INITIAL-COMMENTS-ON-AI-NOI-Filed-version-1.pdf> [https://perma.cc/N4YC-MRZA].

⁶⁶ *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 438 (1984) (“The only contact between *Sony* and the users of the Betamax that is disclosed by this record occurred at the moment of sale.”).

⁶⁷ *Id.*

⁶⁸ See *id.* at 437 n.18.

through the design of the distributed videocassette recorder itself.⁶⁹ This principle was highlighted by the court in *Grokster*, which found the staple article rule as discussed in *Sony* to serve as a defense only to secondary infringement theories which presume or impute intent to cause infringement *solely from* the design or distribution of a product capable of substantial lawful use.⁷⁰

To recap, *Sony* does not provide a catch-all defense to secondary liability. *Grokster* elucidated the example of specific acts taken by a product distributor to induce infringement as one end-around the *Sony* defense. Likewise, the presence of an ongoing relationship between an end-user and an AI model platform may form the basis for a claim of vicarious infringement, irrespective of the *Sony* staple-article rule.⁷¹

Generative AI presents a more complex landscape than the *Sony* court faced. For one, the specific relationship between the party behind a generative AI model and the users of the model will vary considerably, and as a result the analysis will be quite different between models. As discussed previously, generative AI models are often classified as closed or open systems, which may roughly correlate with the degree of contact/interaction with a model and its users.⁷²

In reality, the degree of “openness” a model possesses will fall on a sliding scale. A truly “open-source” model, released with outside access to source code and with no restrictions on use or reproduction, falls at one end of the spectrum. No widely adopted model fits this definition, with the closest example being the academically-focused BigScience Large Open-science Open-access Multilingual Language Model, or BLOOM for short.⁷³ Although users are permitted to fine-tune, train, evaluate, or re-parametrize BLOOM, some restrictions on use remain, and BigScience reserves the right to restrict usage of the model or modify its

⁶⁹ See *id.* at 439.

⁷⁰ MGM Studios, Inc. v. Grokster, Ltd., 545 U.S. 913, 933 (2005).

⁷¹ See, e.g., A&M Recs., Inc. v. Napster, Inc., 239 F.3d 1004, 1022–23 (9th Cir. 2001) (“*Sony*’s ‘staple article of commerce’ analysis has no application to Napster’s potential liability for vicarious copyright infringement.”).

⁷² See *supra* Section I.A.

⁷³ Elizabeth Gibney, *Open-Source Language AI Challenges Big Tech’s Models*, 606 NATURE 850 (2022), <https://www.nature.com/articles/d41586-022-01705-z> [<https://perma.cc/NX3E-827B>].

outputs based on updates.⁷⁴ Many models which claim to be open or open-source, including Meta’s Llama models, for example, do not allow outside users to inspect the model’s training data or code base.⁷⁵

The importance of a model’s “openness” in the context of potential liability for secondary copyright infringement lies in its impact on a party’s right and ability to supervise infringing conduct—an essential element any plaintiff must make out to successfully present a vicarious infringement claim. “The ability to block infringers’ access to a particular environment for any reason whatsoever is evidence of the right and ability to supervise[;]⁷⁶ if a defendant has the right to block access, such right must be “exercised to its fullest extent” to “escape imposition of vicarious liability.”⁷⁷

In theory, some generative AI models could argue they lack the right and ability to supervise infringing conduct, insulating them from claims of vicarious infringement. The developers and providers of decentralized, open-source peer-to-peer software have successfully argued against vicarious liability in the past by doing so.⁷⁸

Conversely, if a developer retains control over how users interact with the model (e.g., through API restrictions, moderation tools, or outright bans from using the service), they may be more susceptible to such claims. Analogizing again to peer-to-peer software, Napster, which was built on a proprietary and centralized indexing software architecture, was found liable for vicarious copyright infringement.⁷⁹ In any case, however, a defendant’s failure to police the system’s “premises” which it “controls and patrols” may lead to imposition of vicarious

⁷⁴ BigScience RAIL License v1.0 (May 19, 2022), <https://huggingface.co/spaces/bigscience/license> [<https://perma.cc/VG6U-UPUX>].

⁷⁵ Elizabeth Gibney, *Not all ‘open source’ AI models are actually open: here’s a ranking*, NATURE Ref. 1 (June 19, 2024), <https://www.nature.com/articles/d41586-024-02012-5> [<https://perma.cc/M5XY-7Q66>].

⁷⁶ Arista Recs. LLC v. Usenet.com, Inc., 633 F. Supp. 2d 124, 157 (S.D.N.Y. 2009) (quoting *Napster*, 239 F.3d at 1023).

⁷⁷ *Napster*, 239 F.3d at 1023.

⁷⁸ See *MGM Studios, Inc. v. Grokster, Ltd.*, 380 F.3d 1154, 1164–65 (9th Cir. 2004), *rev’d on other grounds*, 545 U.S. 913 (2005).

⁷⁹ See *Napster*, 239 F.3d at 1024; *see also Grokster*, 380 F.3d at 1158–59 (“[W]e found Napster had the right and ability to supervise Napster users because it controlled the central indices of files, users were required to register with Napster, and access to the system depended on the validity of a user’s registration.”).

liability for copyright infringement.⁸⁰ As the provider of a digital service/product, the “right and ability” reserved by a generative AI defendant “is cabined by the system’s current architecture.”⁸¹ As such, a respective AI model developer or platform’s right and ability to supervise infringing conduct will be a fact-intensive inquiry dealing with the specifics of that model’s architecture.

OpenAI’s terms of use may be an illuminating example.⁸² The terms stipulate “[w]e reserve the right to suspend or terminate your access to our Services or delete your account if we determine: [y]ou breached these Terms or our Usage Policies[; w]e must do so to comply with the law[; or y]our use of our Services could cause risk or harm to OpenAI, our users, or anyone else.”⁸³ Such a policy clearly indicates OpenAI’s reserved right to block infringers’ access to a particular environment. Thus, if OpenAI failed to exercise its right to block infringer’s access to its services to its fullest extent as permitted by its model’s capabilities, vicarious liability remains on the table.

A model’s degree of openness is just one element of note when assessing the nature of a developer or platform’s relationship with its end-users. The presence or lack of an ongoing pecuniary relationship between end-users and the platform undoubtedly factors into any vicarious liability assessment. Access to many popular generative AI models is offered on a paid subscription basis.⁸⁴

The pecuniary relationship between model and user speaks directly to the second element of vicarious infringement, in which a direct financial benefit from infringing activity must be established. “[F]inancial benefit exists where the availability of infringing material ‘acts as a draw for customers.’”⁸⁵ Further, the size of said draw in proportion to a defendant’s overall business is immaterial.⁸⁶

⁸⁰ See *Napster*, 239 F.3d at 1023–24 (citing *Fonovisa, Inc. v. Cherry Auction, Inc.*, 76 F.3d 259, 262–63 (9th Cir. 1996)).

⁸¹ *Id.* at 1024.

⁸² *Terms of Use*, OPENAI (Dec. 11, 2024), <https://openai.com/policies/row-terms-of-use/> [https://perma.cc/E85Y-TV9A].

⁸³ *Id.*

⁸⁴ See Charles Rollet, *The hottest AI models, what they do, and how to use them*, TECHCRUNCH (Mar. 30, 2025), <https://techcrunch.com/2025/03/30/the-hottest-ai-models-what-they-do-and-how-to-use-them/> [https://perma.cc/9KD7-6FX4] (documenting several recently released AI models requiring a monthly paid subscription for access).

⁸⁵ *Napster*, 239 F.3d at 1023 (quoting *Fonovisa*, 76 F.3d at 263–64).

⁸⁶ *Perfect 10, Inc. v. Giganews, Inc.*, 847 F.3d 657, 673 (9th Cir. 2017).

Not all services that consumers value act as a draw, however. To establish a direct financial benefit under the consumer draw theory, evidence must demonstrate that a service provider attracted or retained subscriptions because of the infringement it facilitates, or lost subscriptions because of eventual obstruction of infringement.⁸⁷

In the context of generative AI, it's not difficult to imagine that at least some subscriptions are motivated by the potential for use of models to create infringing content. The recent explosion of interest in GPT-4o's image generation capability⁸⁸—and contemporaneous removal of content moderation safeguards by OpenAI⁸⁹—is one timely example, with the tool being notoriously used to create images in the iconic artistic style of Hayao Miyazaki, the renowned animator behind Studio Ghibli. Though artistic style itself is an idea and not protectable expression under copyright, it's entirely plausible, if not probable, that GPT-4o users will create images that would infringe the derivative work right attached to, say, the iconic characters of Ghibli animations, or perhaps of those attached to other IP franchises.⁹⁰

It's worth considering the fact that many popular LLMs have free and paid tiers, allowing users to access a selection of models at no cost, or on limited bases. Could a model developer argue that they lack any financial benefit from infringing acts taken by users in the free tier? Perhaps, but the stronger argument is likely that some users who are drawn to the service given its potential for creating potentially infringing content will end up purchasing a paid subscription. The free tier is in effect a marketing tool for the platform, and likely would not serve as an escape valve for vicarious infringement liability; rather, it would be seen as incentivizing further subscriptions and revenue.

⁸⁷ See *Ellison v. Robertson*, 357 F.3d 1072, 1079 (9th Cir. 2004) (finding no vicarious copyright infringement in light of a lack of evidence that the defendant online service provider attracted, retained, or lost subscriptions because of facilitating or obstructing infringement).

⁸⁸ See Zeff, *supra* note 10.

⁸⁹ See Maxwell Zeff, *OpenAI peels back ChatGPT's safeguards around image creation*, TECHCRUNCH (Mar. 28, 2025), <https://techcrunch.com/2025/03/28/openai-peels-back-chatgpts-safeguards-around-image-creation/> [<https://perma.cc/9KD7-6FX4>].

⁹⁰ See also Victor Tangermann, *Lawyer Says Studio Ghibli Could Take Legal Action Against OpenAI*, FUTURISM (Mar. 28, 2025, at 17:34 ET), <https://futurism.com/lawyer-studio-ghibli-legal-action-openai> [<https://perma.cc/H5QX-GEZZ>] (suggesting Studio Ghibli may have actionable claims against OpenAI related to false advertising, trademark infringement, and unfair competition).

The requirement of a direct financial benefit from infringing activity may work to the advantage of academic or research-focused models, so long as they do not charge for use of their models. When paired with the knowledge that such models tend to fall closer to the open end of the open-source/proprietary spectrum, the developers of such models may find themselves shielded against any claims of vicarious infringement, meaning the *Sony* staple-article rule may act more effectively for these models.

III

RAMIFICATIONS—DOES THE EXISTING FRAMEWORK BALANCE COMPETING INTERESTS?

A. Identifying and Evaluating the Public Interest Implications

Though important to remember that the specific assessment of secondary liability for any particular AI model will depend entirely on the facts specific to that particular model and its users, a few general trends emerge when applying doctrines of secondary copyright infringement to the generative AI ecosystem.

First, the potential for secondary liability solely based on the design or distribution of an AI model will in most cases be negligible, given generative AI's capability for substantial lawful use, despite developers'/platforms' knowledge that their models may be used for infringement.⁹¹

Second, other theories of contributory copyright infringement apart from those based on design of the model also remain available, including active inducement, though they are not the focus of this note.⁹²

Third, the presence of a continuing relationship between model developers/platforms and end-users at least raises the possibility of holding generative AI developers/platforms liable for vicarious copyright infringement, contingent on a showing that all required elements of such a claim are met. Any potential effect of the current statutory framework under the DMCA on vicarious infringement theories will be examined in greater depth later in this note.⁹³

⁹¹ MGM Studios, Inc. v. Grokster, Ltd., 545 U.S. 913, 933 (2005).

⁹² *Id.* at 936.

⁹³ See *infra* Section IV.A.

Within the purview of the relationship between model and end-user, a few generalizations may be drawn. Model developers with a more attenuated relationship to their end-users, often those that exist closer to the open-source end of the open/proprietary spectrum, will fare better against claims of vicarious infringement, often due to their relative lack of control or ability to supervise and restrict the infringing acts of end-users. Developers of proprietary models may be left more vulnerable to claims of vicarious infringement, exemplified by OpenAI's reservation of the right to restrict access to users whose use of OpenAI's services could cause risk or harm to OpenAI, their users, or anyone else—a seemingly low bar.

On the one hand, open-source and open-access is central to scientific thinking and progress, and it seems fitting that the law serves to protect such models against secondary liability. The essence of scientific reasoning is reproducibility,⁹⁴ and without models being open to scrutiny, it's not clear how they may be objectively evaluated. Further, open-sourcing of models may bring additional benefits, such as by allowing easier identification and mitigation of bias in training data, as well as for facilitating market entry by potential competitors.⁹⁵ Given the current legal battles over the use of copyrighted works in AI training data,⁹⁶ it may be difficult to entice model developers to fully disclose their training datasets publicly until there is a clear resolution to the question of whether such use is fair.

Yet, given its lack of supervision, open-sourcing may allow a greater propensity for misuse of models down-stream. Some academics have advocated against open-sourcing for any LLMs that are likely to be used to generate outputs that infringe on protected rights in a material fashion.⁹⁷ It may be a grave mistake to allow generative AI developers to fall back on *Sony* to wash their hands of responsibility for injuries stemming from the use of repurposed, powerful AI tools based off of their own and used for illicit purposes.

⁹⁴ Gibney, *supra* note 75, at Ref. 1.

⁹⁵ See EFF Comments, *supra* note 64, at 6.

⁹⁶ See, e.g., *Case Tracker: Artificial Intelligence, Copyrights and Class Actions*, BAKERHOSTETLER, <https://www.bakerlaw.com/services/artificial-intelligence-ai/case-tracker-artificial-intelligence-copyrights-and-class-actions/> [https://perma.cc/9Y9Q-VSBW] (last visited Oct. 29, 2025) (surveying key litigation raising copyright issues related to AI platforms).

⁹⁷ See, e.g., Matthew Sag, *Copyright Safety for Generative AI*, 61 Hous. L. Rev. 295, 340 (2023).

Taking this into account, it may also feel unfair to punish models which take greater efforts to regulate or safeguard their models and thus necessarily assume some degree of control over their services, in the hopes of bringing them into legal compliance. This confluence may serve to reward a lack of oversight, and disincentivize responsible AI governance.

B. Critiquing Sony and Exploring Alternative Approaches

The *Sony* doctrine has had its share of criticism leveled at it, though it has persisted to this day. One criticism of the doctrine is its permissive nature: in her concurring opinion in *Grokster*, Justice Ginsburg noted that she would have found the defendants had not met the evidentiary standard needed to take advantage of the *Sony* safe harbor.⁹⁸ Justice Breyer, disagreeing, noted that the lax evidentiary burden attached to the *Sony* rule is meant to protect technological innovation, and that an increased burden would undercut the protection *Sony* affords.⁹⁹

Other critiques have noted the peculiarity of a judicial decision which elected not to examine Congressional intent in reforming the Copyright Act of 1976, just eight years earlier.¹⁰⁰ In a preceding article, Menell and Nimmer also note this flawed historical basis of a “historic kinship” between patent and copyright—a basis by which a provision of section 271(c) of the Patent Act was imported into the Copyright Act.¹⁰¹ They conclude that any “historic kinship” between patent and copyright does not justify use of patent law’s departure from the common law as a blueprint for enactment of a copyright statute.¹⁰² The authors continue by evaluating Congress’ intent in drafting and enacting specific provisions of the Copyright Act of 1976, finding that only features of the act that survived to enactment where Congress explicitly drew on patent law should be interpreted by analogy to patent law.¹⁰³ The question of secondary liability for copyright

⁹⁸ MGM Studios, Inc. v. Grokster, Ltd., 545 U.S. 913, 945–46 (2005) (Ginsburg, J., concurring).

⁹⁹ *Id.* at 957–60 (Breyer, J., concurring).

¹⁰⁰ See Peter S. Menell and David Nimmer, *Legal Realism in Action: Indirect Copyright Liability’s Continuing Tort Framework and Sony’s De Facto Demise*, 55 UCLA L. Rev. 1, 15 (2007).

¹⁰¹ See Menell and Nimmer, *supra* note 17, at 943–44.

¹⁰² *Id.* at 987.

¹⁰³ *Id.* at 993.

infringement had been developed in the common law prior to enactment, and Congress specifically declined to alter from established case law.¹⁰⁴

Speaking of secondary liability's history of development in the common law, Menell and Nimmer note courts' recognition of copyright's roots in tort as early as 1869, continuing through the time period leading up to the enactment of the 1976 Copyright Act.¹⁰⁵

Menell and Nimmer continue by suggesting that had the court followed congressional intent and imposed a different rule under *Sony*, one based in the tort principles from which copyright sprung, the outcome would have been the same.¹⁰⁶ Such an approach would also have created a more flexible framework going forward for addressing the challenges of new technology.¹⁰⁷ They propose evaluating indirect copyright infringement liability under a reasonable alternative design (RAD) framework, derived from products liability.¹⁰⁸ Under such an approach, the key questions to be answered are: was a proposed alternative design available—that is, was the design feasible at the time of manufacture? And, does the reduction in risk of harm outweigh the loss in utility?¹⁰⁹ Ultimately, they conclude that the design alternatives that were feasible at the time of manufacture of the Betamax would have exhibited significant adverse effects on the legitimate interests of users, with little to gain in reduction of cognizable harm to copyright owners.¹¹⁰

Though Menell and Nimmer's first article focuses on the facts of *Sony* and its associated Betamax technology, they also reason that a tort-based framework is more accurate to the practical realities in the time since.¹¹¹ Principles of RAD have made their way into other theories of secondary liability, including inducement

¹⁰⁴ *Id.*; see also H.R. REP. No. 94-1476, 61, 159–60 (1976) (acknowledging potential liability for contributory infringement while refraining from elaborating or further defining what constitutes contributory infringement, and later expressly declining to alter existing vicarious copyright infringement doctrine).

¹⁰⁵ See Menell and Nimmer, *supra* note 100, at 9 (citing *Lawrence v. Dana*, 15 F. Cas. 26, 61 (C.C. Mass. 1869) (No. 8,136), and *Ted Browne Music Co. v. Fowler*, 290 F. 751, 754 (2d Cir. 1923)).

¹⁰⁶ Menell and Nimmer, *supra* note 17, at 1022.

¹⁰⁷ *Id.*

¹⁰⁸ *Id.* at 1018.

¹⁰⁹ *Id.* at 1018.

¹¹⁰ *Id.* at 1020.

¹¹¹ See Menell and Nimmer, *supra* note 100, at 3–4.

and vicarious liability analyses.¹¹² Courts have even in some cases reintroduced RAD principles when assessing theories of contributory copyright infringement under the *Sony* staple article rule.¹¹³ The principles articulated—particularly with respect to a products liability-styled RAD approach—fit soundly in the context of generative AI, speaking to the truth of the authors’ point that such an approach creates a more flexible framework to address the challenges of new technology.

In applying such an approach, parties would need to first assess whether the AI system in question could be designed to minimize infringement risks—that is, determine whether alternative designs were available to the developer at the time of release which would have been feasible to adopt, given the state of technology and costs associated with adoption. Those design choices found to be feasible would then be weighed on their relative decrease in social harm as compared to the original design, as well as their relative impediments to utility of the model as compared to the original design. For developers whose models could have been developed and distributed in an alternative manner that would have prevented infringement effectively while incurring only minor setbacks in utility, indirect infringement liability may be imposed.

Though a liability framework based on the availability of design features to mitigate infringement could provide a more holistic approach to assessing indirect copyright infringement claims, it may be necessary to take the approach a step further in the context of SaaS and generative AI platforms, where the distinction between product and service is blurred. A framework considering only design choices relating to software architecture may not capture the other levers that platforms may have available to them in mitigating infringement. A comprehensive “RAD-plus” assessment should include consideration of a firm’s policies relating

¹¹² See Arista Recs. LLC v. Lime Grp., 784 F.Supp.2d 398, 429–31, 435–36 (S.D.N.Y. 2011) (“[A] failure to utilize existing technology to create meaningful barriers against infringement is a strong indicator of intent to foster infringement.”); *In re Aimster Copyright Litigation*, 334 F.3d 643, 654 (7th Cir. 2003) (finding that vicarious liability could conceivably have been found in *Sony*, “on the theory that while it was infeasible for the producers of copyrighted television fare to sue the viewers who used the fast-forward button on Sony’s video recorder to delete the commercials and thus reduce the copyright holders’ income, *Sony could have reduced the likelihood of infringement ... by a design change.*”) (emphasis added).

¹¹³ See *In re Aimster*, 334 F.3d at 653 (“Even when there are noninfringing uses ... if the infringing uses are substantial then to avoid liability as a contributory infringer *the provider of the service must show that it would have been disproportionately costly for him to eliminate or at least reduce substantially the infringing uses.*”) (emphasis added).

to infringement mitigation, including efforts to track use of offered generative AI tools by repeat offenders, and to curtail access at appropriate thresholds. Under a process- and policy-inclusive approach, internal red-teaming efforts (or a lack thereof) may also be considered in determining whether a generative AI platform defendant is liable for contributory infringement.

To be fair, a RAD-plus approach would assuredly hold model developers and platforms to a higher standard than the traditional staple-article rule of *Sony*. This approach could plausibly chill development in AI, though it may encourage more responsible AI governance across the entire landscape. Given the immense sources of capital behind AI technologies, the risk to chilling AI development may be overblown.¹¹⁴ The following section will in part contemplate a regulatory approach which incorporates adoption of a RAD-plus framework in exchange for a degree of protection from claims of vicarious infringement liability for model developers and platforms. Such a framework could balance innovation with accountability more effectively than the current doctrine, or a legislative solution styled similarly to the Digital Millennium Copyright Act (DMCA) safe harbor.

IV

WHAT COULD AN EFFECTIVE REGULATORY REGIME LOOK LIKE?

A. *Availability of the DMCA and its Limitations in Policing Generative AI*

In the same breath as the *Sony* court adopted the staple-article rule, it also acknowledged Congress's authority to reexamine and tailor the Copyright Act to better accommodate the adoption of new, disruptive technology.¹¹⁵ The Supreme Court has reiterated its reluctance to expand the protections of copyright absent an explicit directive from Congress. "Sound policy, as well as history, supports our consistent deference to Congress when major technological innovations alter the market for copyrighted materials. Congress has the constitutional authority and

¹¹⁴ See Samantha Subin, *Tech megacaps plan to spend more than \$300 billion in 2025 as AI race intensifies*, CNBC (Feb. 8, 2025 at 11:02 ET), <https://www.cnbc.com/2025/02/08/tech-megacaps-to-spend-more-than-300-billion-in-2025-to-win-in-ai.html> [https://perma.cc/MC3F-SEBS].

¹¹⁵ See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 456 (1984) ("It may well be that Congress will take a fresh look at this new technology, just as it so often has examined other innovations in the past.").

the institutional ability to accommodate fully the varied permutations of competing interests that are inevitably implicated by such new technology.”¹¹⁶

In the time since *Sony*, the Court has reiterated that “the legislative option remains available.”¹¹⁷ Supporters of varied groups of technology have advocated for and received modification of the nation’s copyright laws, as they pertain to their specific industries. The DMCA, for one, is a prominent example in the digital economy.

The need for the DMCA arose from the express carve-out of protection from intellectual property infringement liability which was placed into Section 230 of the Communications Decency Act of 1996, which generally provides immunity for online computer services with respect to third-party content generated by its users.¹¹⁸

The most pertinent element of the DMCA as it pertains to secondary copyright infringement is found in the Online Copyright Infringement Liability Limitation Act (OCILLA), widely known as the “safe harbor” provision or Section 512. The safe harbor institutes a notice-and-takedown procedure which most online service providers (OSPs) may adhere to in exchange for protection from direct and indirect copyright infringement liability for the actions of an OSP’s users.¹¹⁹

Some protections of the DMCA safe harbor may apply to generative AI platforms today.¹²⁰ Save for models which run entirely on local hardware, most generative AI platforms fit the safe harbor’s definition of a “service provider” which in relevant part states “the term ‘service provider’ means a provider of online services or network access, or the operator of facilities therefor.”¹²¹

¹¹⁶ *Id.* at 431.

¹¹⁷ *MGM Studios, Inc. v. Grokster, Ltd.*, 545 U.S. 913, 965 (2005) (Breyer, J., concurring).

¹¹⁸ 47 U.S.C. § 230(e)(2) (“Nothing in this section shall be construed to limit or expand any law pertaining to intellectual property.”); *see also Gucci Am., Inc. v. Hall & Assocs.*, 135 F.Supp.2d 409, 412–17 (S.D.N.Y. 2001) (finding an ISP has no immunity for contributory liability for trademark infringement under Section 230).

¹¹⁹ 17 U.S.C. § 512.

¹²⁰ *See Uzma Chaudhry, AI and digital governance: Platform liability laws in the US*, IAPP (Sep. 18, 2024), <https://iapp.org/news/a/ai-and-digital-governance-platform-liability-laws-in-the-u-s> [https://perma.cc/L3WM-39D6].

¹²¹ 17 U.S.C. § 512(k)(1)(B).

Beyond meeting the permissive definition of a service provider, entities wishing to take advantage of the safe harbor must also implement the conditions stipulated in 512(i). These conditions require service providers to adopt and reasonably implement a policy that provides for the termination of subscribers and account holders of the service provider's system or network who are repeat infringers, while providing notice to subscribers and account holders of said policy.¹²² "Making the entrance into the safe harbor too wide would allow service providers acting in complicity with infringers to approach copyright infringement on an image by image basis without ever targeting the source of these images [S]ervice providers are meant to have strong incentives to work with copyright holders. The possible loss of the safe harbor provides that incentive and furthers a regulatory scheme in which courts are meant to play a secondary role to self-regulation."¹²³

This policy requirement may render many generative AI platforms ineligible for the safe harbor at the outset, at least under current business practices. A platform that allows use of its service without registration, a feature of several leading models such as OpenAI's ChatGPT and Google's Gemini,¹²⁴ would have no way to monitor and terminate use by repeat infringers. However, the 512(i) conditions could serve to suggest what type of policy threshold a court may consider necessary to clear under a hypothetical RAD-plus system for assigning indirect liability.

Further, protection under the DMCA is not absolute. If a service provider possesses actual and "red flag" knowledge of infringements through use of its service, it must act expeditiously to remove or disable access to infringing material to remain in the graces of the safe harbor.¹²⁵

Per the § 512(c)(1)(B) exception, as another example, the safe harbor does not provide a shield from liability for hosting infringing works when a party "receive[s]

¹²² 17 U.S.C. § 512(i)(1)(A).

¹²³ Perfect 10, Inc. v. Cybernet Ventures, Inc., 213 F.Supp.2d 1146, 1177–78 (C.D. Cal. 2002).

¹²⁴ See Christoph Schwaiger, *You can use ChatGPT without an account – here's how*, Tom's GUIDE (July 11, 2024), <https://www.tomsguide.com/ai/chatgpt-you-can-connect-to-chatgpt-without-an-account-heres-how-it-works> [https://perma.cc/SS4H-6WCK].

¹²⁵ 17 U.S.C. § 512(c)(1)(A); *see also* Mavrix Photographs, LLC v. LiveJournal, Inc., 873 F.3d 1045, 1057 (9th Cir. 2017) (defining "actual knowledge" as subjective knowledge of the service provider, and "red flag knowledge" as awareness of facts that would have made the specific infringement objectively obvious to a reasonable, non-expert person).

a financial benefit directly attributable to the infringing activity, in a case in which the service provider has the right and ability to control such activity.”¹²⁶ A similar exception exists with respect to linking to infringing works.¹²⁷ These exceptions draw similarities to common law vicarious infringement liability.¹²⁸ The § 512(c)(1)(B) financial benefit prong has generally been interpreted by courts in a manner equivalent with the common law vicarious liability standard.¹²⁹

There is some disagreement among the circuit courts as to the correct interpretation of the § 512(c)(1)(B) exception’s “right and ability to control” requirement. In the 2nd Circuit, the current understanding of the § 512(c)(1)(B) exception to the safe harbor requires a “right and ability to control” that is “something more than the ability to remove or block access to materials posted on a service provider’s website.”¹³⁰ Though it reaches beyond the requirements of common law vicarious liability, with identified positive examples meeting the § 512(c)(1)(B) exception “involv[ing] a service provider exerting substantial influence on the activities of users,” the Second Circuit does not impart a requirement of specific knowledge of infringing acts.¹³¹ In the 9th Circuit, a showing of actual knowledge of specific infringing acts *is* required to fall within the § 512(c)(1)(B) exception.¹³²

Perfect 10 v. Cybernet is the rare case in which a court expressly found that a service provider had the right and ability to control infringing activity under § 512(c)(1)(B) and was cited affirmatively by the Second Circuit in *Viacom*. Cybernet was a corporation running a service which permitted its subscribers access to a collection of websites for a fee.¹³³ By prescreening and refusing to allow sites that do not comply with its dictates, giving the sites extensive advice on issues of layout, appearance, and content, monitoring images to make sure celebrity images do not oversaturate the content Cybernet’s service provided access to, and prohibiting the

¹²⁶ 17 U.S.C. § 512(c)(1)(B).

¹²⁷ 17 U.S.C. § 512(d)(2).

¹²⁸ *See supra* Section I.C.

¹²⁹ *See* Perfect 10, Inc. v. CCBill LLC, 488 F.3d 1102, 1117 (9th Cir. 2007) (“[D]irect financial benefit should be interpreted consistent with the similarly-worded common law standard for vicarious copyright liability.”) (internal quotations omitted).

¹³⁰ Viacom Intern., Inc. v. Youtube, Inc., 676 F.3d 19, 38 (2d Cir. 2012).

¹³¹ *Id.*

¹³² *See* UMG Recordings v. Shelter Capital Partners LLC, 667 F.3d 1022, 104 (9th Cir. 2011).

¹³³ Perfect 10, Inc. v. Cybernet Ventures, Inc., 213 F.Supp.2d 1146, 1181–82 (C.D. Cal. 2002).

proliferation of identical sites, *Cybernet* “exhibit[ed] precisely this slightly difficult to define ‘something more.’”¹³⁴

The aim of this note is not to settle the split between the 2nd and 9th Circuits. It suffices for our purposes to conclude that, at the very least, service providers, such as some generative AI models, could remain vulnerable to a vicarious-like type of secondary copyright infringement liability even under any section 512 safe harbor, cabined by the relevant model’s architecture.

Section 512(b), (c), and (d) provide a liability shield against infringement resulting from the caching, hosting, or linking to protected works, respectively. As they are normally used, a generative AI platform’s creation of an output and storage within its network for retrieval by the user would be properly categorized under 512(c) as an example of information residing on systems or networks at the direction of a user or under 512(d) if linking to sources requested by an end-user, though perhaps in a context somewhat differently than courts have been presented with in the past. Typical media or content hosting sites, such as YouTube, Vimeo, or social media platforms, host works where they may be viewed by many, whereas any work hosted by a generative AI platform is in most cases viewable only by the user. Facing a similar issue of fitting a square peg into a round hole, the *Cybernet* court noted Cybernet’s system functionality made it a poor fit for the categories established by the DMCA.¹³⁵ Yet the court proceeded with analyzing Cybernet’s affirmative defense under 512(c) as if it qualified, “because Cybernet does run a web-page . . . and maintains computers to govern access to the [web-page] family’s websites there is good reason to believe that it is an ‘provider of online services’ under 512(k)(1)(B).”¹³⁶ Given the historical permissive categorization of service providers and their eligibility for the § 512 safe harbor, generative AI platforms may fare no differently.

As stated earlier, the § 512 safe harbor provisions require service providers to implement a notice-and-takedown procedure. Should generative AI platforms have a policy meeting the § 512(i) requirements and implement a notice-and-takedown system, it would be ineffective in protecting the interests of rights-holders and may be fully incompatible with the architecture of a typical AI platform. Potentially

¹³⁴ *Id.* at 1173–74, 1181–82.

¹³⁵ See *id.* at 1175 n.19.

¹³⁶ *Id.*

infringing material created with the help of a model will be redistributed outside the control of model developers. Rights holders who are alerted to their works being infringed will not be directed to the generative AI model that produced the work, at least not initially. They would be directed to the site, for example, which hosts the infringing material. For the procedure to be effective in curbing use of platforms to generate infringing material, rights holders must have knowledge of the actual source. The source is self-evident in most instances under the DMCA notice-and-takedown approach—think of a username attached to a posted YouTube video or a specific user posting content on a social media platform. In the generative AI context, however, though potentially infringing works are in a sense hosted by a platform (through chat histories, for example), they are generally not accessible by those other than the user prompting its generation. It would be near impossible for a right holder to obtain the URL at which the hosted material is found, which is typically required for those wishing to file a DMCA notice.¹³⁷

To analogize, generative AI does not fit cleanly within the § 512 safe harbor as it is not a bulletin board, the likes of a YouTube, Reddit, or other social media site or media aggregator. It's more so a tool of creation. In short, due to its ex-post nature and the position of generative AI platforms within the creative ecosystem, the notice-and-takedown system would have little to no effect, while providing any AI model that can qualify with an unwarranted benefit in the form of a liability shield.

B. RAD-Plus as an Approach to Effective Generative AI Regulation

Academics and outside commentators have called for legislative solutions to pressing issues of the digital age.¹³⁸ Rather than relying on the DMCA, potential alternative regulatory measures could include federal legislation incentivizing responsible AI development while limiting liability for compliant platforms. One possible blueprint could be built around the RAD-plus approach.¹³⁹

¹³⁷ *How can I file a DMCA Takedown Notice?*, DMCA.com (Nov. 22, 2023), <https://www.dmca.com/FAQ/How-can-I-file-a-DMCA-Takedown-Notice> [https://perma.cc/VL8Q-SHKN].

¹³⁸ See Menell and Nimmer, *supra* note 17, at 1023 (“Congress needs to take up questions [of balancing liability with innovation in the digital age] and consider the full range of institutional regimes available to guide copyright as technology advances.”).

¹³⁹ See *supra* Section III.B.

Consider legislation that imposes a RAD-plus framework in consideration for some degree of insulation against common law vicarious infringement liability for generative AI developers and platforms. In exchange for adopting reasonable safety precautions and policies meant to minimize the risk of generating infringing outputs, AI platforms could be insulated from liability stemming from their continued relationship with users. Recall that under their current business practices, there are serious questions as to whether leading generative AI model platforms implement the policies necessary to qualify for the DMCA 512 safe harbor, and to accordingly take advantage of the heightened standard for vicarious liability under 512(c)(1)(B).

In terms of responsible AI governance, there are reasons to encourage an ongoing relationship between AI model developers/platforms and end-users. The relationship may allow developers to easily implement a content moderation policy, which could include the monitoring and restriction of user inputs and/or model outputs that are likely to lead to acts of infringement. Retaining the ability to restrict model access to repeat infringers may also necessitate a system architecture which requires ongoing contacts between the model developer and user. Likewise, maintaining the conduit between model developers and end-users will allow for on-the-fly modifications to AI tools, to seamlessly implement improvements to a model's design. Unique to the digital context, the ongoing connection may also serve to remove from circulation outdated versions of the model which do not follow contemporary safety protocols, a valuable ability which would be near impossible to implement with respect to tangible consumer goods, like a Betamax. A limit to vicarious infringement liability could also buttress such a legislative proposal against claims of harming or chilling innovation, merited or not.

Of course, no solution comes without potential drawbacks. One could argue that even with a system in place that incentivizes developers to incorporate the most prudent safeguards into their models, the potential for misuse would still be too great. Essentially, the rate of development in AI safeguards is being outpaced by the capabilities of AI systems. Taken to its farthest point, one could argue RAD-plus would chill the development of safety measures themselves, for fear of being saddled with them in the future. The merits of such an argument would require a deeper investigation of the development landscape: are there capable groups, whose incentives are misaligned with generative AI models, that are developing safeguards? Or are these innovations being drawn up mostly internally? I'd consider

this kind of argument a bit stronger in the context of copyright, as product “safety” means something entirely different than in the tort context. In short, product safety is itself a strong draw for consumers in tort. When equating safety to potential for infringing the right of another, however, the respective draw is mitigated, if not reversed altogether for some consumers.

The following is an attempt to sketch out what a RAD-plus framework could look like in practice. When assessing contributory copyright infringement claims against generative AI platforms, the capability of commercially significant non-infringing uses would no longer be the lodestar. Rather, in analogizing to tort principles, a product or service may be deemed defective in design and therefore render the developer/platform behind said product or service liable for contributory copyright when its foreseeable risks of harm (e.g. perpetuating direct copyright infringement) could have been reduced or avoided through the adoption of a reasonable and available design, policy, or process by any predecessor in the commercial chain of distribution, and when the omission of the alternative feature renders the product not reasonably safe (e.g. prone to facilitating direct copyright infringement). A stronger form of this rule could provide for some kind of penalty on the developers or hosts of generative AI models for a failure to comply with acceptable design standards, even in the absence of alleged direct infringements.

Digital products and services may already be living in a RAD-esque world, with adherence to the DMCA being the prerequisite for what’s considered the requisite degree of “safety” when it comes to service providers. Even so, the safe harbor provisions were surely not enacted with the specter of advanced artificial intelligence platforms and their capabilities looming over Congress. What’s reasonable in 1998 may not be in 2025. Accordingly, simple enactment of static conditions (analogous to reasonable design choices) may not go far enough to implement a dynamic standard with a threshold that rises with advances in technology.¹⁴⁰

It may be wishful thinking to imagine a Congress that takes up the issue itself, especially considering the Executive’s current position on AI.¹⁴¹ Perhaps

¹⁴⁰ See Menell and Nimmer, *supra* note 100, at 13.

¹⁴¹ See *Removing Barriers to American Leadership in Artificial Intelligence*, Exec. Order No. 14179 (2025), <https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/> (revoking existing AI policies that

the more likely avenue for relief is a simple judicial abdication of the *Sony* staple-article rule in favor of a RAD-plus framework. Though this solution would leave some generative AI developers and platforms potentially vulnerable to claims of vicarious infringement and perhaps unfairly so, it would still incentivize responsible imposition of safety measures and monitoring of platforms, and frankly, the groups behind leading models should have little trouble financing their exposure. The situation today is unlike that of the 1990s, when the DMCA was ushered in prior to a full realization of the implications of a rapidly digitizing economy, and the immense leverage big tech firms would come to possess.¹⁴² Imposition of RAD-plus would also influence the design and operation of open-source models, whose developers may otherwise be insulated from much secondary liability of either the contributory or vicarious type.

It's worth remembering the judicial reluctance to extend the grant of exclusive rights in the face of evolving technology mentioned at the outset of this section. Perhaps this inertia would prove too strong for our current judiciary to feel comfortable "expanding" the exclusive rights copyright provides. The legislative option surely remains available, but finding the consensus required to exercise it, especially against objections of harming or chilling innovation, may be a challenge too tall to overcome. The RAD-plus approach seeks to balance innovation, accountability, and the rights of creators, ensuring that generative AI models remain both legally compliant and functionally beneficial—reaching an optimal balance of acceptable risk and utility.

CONCLUSION

It is still too early to say how courts will definitively assess claims of secondary copyright infringement brought against generative AI model developers and platforms. Any claims for contributory copyright infringement, if based solely on the design of the model, would likely fail when run through the staple-article rule as articulated in *Sony*.

Even under the current framework, plaintiffs will often have other options in bringing forth their claims. The trouble lies in a disconnect between these available

"act as barriers to American AI innovation," and marking the promotion of "human flourishing, economic competitiveness, and national security" as key AI policy goals).

¹⁴² See Matthew Sag, *Internet Safe Harbors and the Transformation of Copyright Law*, 93 NOTRE DAME L. REV. 499, 506 (2017).

pathways for plaintiffs, and the incentives they impute upon the generative AI industry.

If the courts or Congress were to reconsider the staple-article rule and replace it with a framework grounded in copyright's history, a dynamic system based on the availability of reasonable alternative design and process choices to mitigate the potential for generative AI models to facilitate infringement could assume a central role in resolving copyright disputes pertaining to the outputs of generative AI systems. Though shouldering a slightly more onerous standard upon generative AI model providers, the new system could effectively balance the potentially massive utility in using a generative AI model against an acceptable threshold of risk for misuse, given the available safety measures and precautions of the day.

Safeguarding and encouraging innovation while equitably enforcing and protecting the rights of copyright holders remains a challenge, although just one of many AI-related quandaries that society currently faces. Any legal framework must protect creative industries while fostering the continued responsible development of AI technologies. Ongoing legal debates, some in the sphere of artificial intelligence, will shape the future of copyright law in the years ahead. Perhaps the winds of change will lead copyright down a new path, one inspired by its roots.