

1-9-2024

## LIMITS OF ALGORITHMIC FAIR USE

Jacob Alhadeff

Cooper Cuene

Max Del Real

Follow this and additional works at: <https://digitalcommons.law.uw.edu/wjlta>



Part of the [Computer Law Commons](#), [Entertainment, Arts, and Sports Law Commons](#), [Intellectual Property Law Commons](#), and the [Internet Law Commons](#)

---

### Recommended Citation

Jacob Alhadeff, Cooper Cuene & Max Del Real, *LIMITS OF ALGORITHMIC FAIR USE*, 19 WASH. J. L. TECH. & ARTS (2024).

Available at: <https://digitalcommons.law.uw.edu/wjlta/vol19/iss1/1>

This Article is brought to you for free and open access by the Law Reviews and Journals at UW Law Digital Commons. It has been accepted for inclusion in Washington Journal of Law, Technology & Arts by an authorized editor of UW Law Digital Commons. For more information, please contact [lawref@uw.edu](mailto:lawref@uw.edu).

LIMITS OF ALGORITHMIC FAIR USE

*Jacob Alhadeff*<sup>1</sup>, *Cooper Cuene*<sup>2</sup>, and *Max Del Real*<sup>3</sup>

ABSTRACT

In this article, we apply historical copyright principles to the evolving state of text-to-image generation and explore the implications of emerging technological constructs for copyright's fair use doctrine. Artificial intelligence ("AI") is frequently trained on copyrighted works, which usually involves extensive copying without owners' authorization. Such copying could constitute prima facie copyright infringement, but existing guidance suggests fair use should apply to most machine learning contexts. Mark Lemley and Bryan Casey argue that training machine learning ("ML") models on copyrighted material should generally be permitted under fair use when the model's outputs transcend the purpose of its inputs. Their arguments are compelling in the domain of AI, generally. However, contemporary AI's capacity to generate new works of art ("generative AI") presents a unique case because it explicitly attempts to emulate the expression copyright intends to protect. Jessica Gillotte concludes that generative AI does not illicit copyright infringement because judicial guidance requires adherence to the constitutional imperative to promote the creation of new works when technological change blurs copyright's boundaries. Even if infringement does occur, Gillotte finds that fair use would serve as a valid defense because training an AI model transforms the original work and is unlikely to damage the original artist's market for the copyrighted work. Our paper deviates from prior scholarship by exploring specific generative AI use cases in technological detail. Ultimately, we argue that fair use's first factor, the purpose of the use, and its fourth factor, the impact on the market for the copyrighted work, both weigh against a finding of fair use in generative AI use cases. However, even if text-to-image models aren't found to be transformative, we argue that the potential for market usurpation alone sufficiently negates fair use.

There is presently little specific guidance from courts as to whether using copyrighted works to build generative AI models constitutes either infringement or fair use, although several related lawsuits are currently pending. Text-to-art generative AIs present several scenarios that threaten substantial harm to the market for the copyrighted original, which tends to undercut the case for fair use. For example, a generative AI trained on copyrighted works has already enabled users to create works "in the style of" individual artists, which has allegedly caused business and reputational losses for the emulated copyright holder. Furthermore, past analyses have ignored the potential for a model to be non-transformative when its intended output has the same purpose and is of the same nature as its copyrighted inputs.

---

<sup>1</sup> Incoming Associate, Betts Patterson & Mines, 2024; J.D., University of Washington School of Law, 2024; Washington Journal of Law Technology & Art; B.A. The Evergreen State College, 2017.

<sup>2</sup> J.D. Candidate, University of Washington School of Law, Class of 2024. B.S., University of Washington, 2021, Chief Articles Editor, Washington Journal of Law, Technology, and the Arts. This author would like to extend his thanks to Professors Ryan Calo and Zahr K. Said for their assistance in developing the thesis for this article and guiding our research and analysis throughout.

<sup>3</sup> J.D. Candidate, University of Washington School of Law, 2024; *Washington Law Review*, Executive Articles Editor; B.S., Stanford University, 2011; M.S., Stanford University, 2012. We would like to thank Professors Ryan Calo and Zahr Said for their support and feedback throughout the drafting process. We're also grateful to the 2023 IP Works In Progress Conference attendees for their encouragement. Finally, thank you to Kiara Hildeman and the wonderful editors at WJLTA for helping us finalize the Article.

This article contributes to the discussion by shining a technical light on text-to-art AI use cases to explore whether some uses normatively fail to qualify as fair uses. First, we examine whether text-to-image models present a *prima facie* infringement claim. We then distinguish text-to-image generative AIs from non-image focused AIs. In doing so, we argue that when the nature of the copyrighted work and the purpose of the infringing use are the same, it is more likely that the original artist will experience market harm. This tilts the overall analysis against a finding of fair use.

**TABLE OF CONTENTS**

I. INTRODUCTION	5
II. TEXT-TO-ART GENERATIVE AI	6
A. DIFFUSION MODELS	6
B. TEXT-TO-ART AI TERMS OF SERVICE	8
C. USE CASES	8
1. “In the Style Of”	8
2. Entity-Centric Models	9
3. The Edit Function	9
III. PRIMA FACIE COPYRIGHT INFRINGEMENT	10
A. FACTUAL COPYING	11
B. UNLAWFUL APPROPRIATION	12
C. TRANSITORY COPYING	13
D. OTHER LIMITS OF COPYRIGHT INFRINGEMENT	14
E. PRIMA FACIE COPYRIGHT INFRINGEMENT IN DIGITAL SPACES	16
IV. SECONDARY LIABILITY	17
A. VICARIOUS LIABILITY	17
B. CONTRIBUTORY LIABILITY	20
1. Material Contribution	20
2. Knowledge of Infringement	22
3. Substantial Non-Infringing Uses—Imputed Knowledge/Intent Based Upon a Product’s Inherently Infringing Nature	23
V. FAIR USE	25
A. FACTOR ONE AND TRANSFORMATIVE USE	25
B. FACTOR FOUR: LICENSING, SUBSTITUTION, AND THE PRACTICAL MARKET	27
VI. APPLICATIONS TO GENERATIVE AI ART	27

A. PRIMA FACIE COPYRIGHT INFRINGEMENT	27
1. The Model Developer as the Defendant	28
2. The User as the Defendant	30
3. Database Compiler as the Defendant	30
B. SECONDARY LIABILITY	31
1. Threshold Questions of Causation	31
2. Vicarious Copyright Infringement	32
3. Contributory Liability	34
a. Material Contribution	35
b. Knowledge	36
c. Substantial Non-Infringing Uses	36
C. FAIR USE	37
1. Non-Expressive Algorithmic Fair Use	37
2. Differentiating Generative AI	40
3. Factor One	42
4. Factor Four	44
D. POLICY IMPLICATIONS	47
1. Training Set Policy Considerations	49
2. Generated Output Policy Considerations	51
VII. CONCLUSION	52

## I. INTRODUCTION

Until recently, Greg Rutkowski was a little-known illustrator whose fantasy works primarily appeared in games like *Magic: The Gathering* and *Dungeons and Dragons*.<sup>4</sup> However, Rutkowski's work has recently become known in an unexpected way: his style is one of the most sought after for AI-generated images. Beginning in Summer of 2022, a wave of text-to-image generative AIs were released in rapid succession. In July 2022, both OpenAI's DALL-E 2<sup>5</sup> and Midjourney<sup>6</sup> became publicly available as open beta projects. The next month, Stable Diffusion<sup>7</sup> was released by Stability AI. Among other features, each of these tools allowed a user to prompt the AI to create images "in the style of" an existing artist. Rutkowski's high-fantasy art style was used as a prompt over 93,000 times in the first few months, dwarfing the number of prompts mimicking other famous artists.<sup>8</sup> While Rutkowski was initially interested in the potential of the tools to market his work, he quickly became nervous that his original work would be drowned out by indistinguishable AI works.<sup>9</sup>

In order to create generative AIs, developers train models with enormous datasets such as LAION-5B, used to train Stable Diffusion.<sup>10</sup> While these datasets are typically compiled by nonprofit organizations, the datasets contain many images that are under copyright.<sup>11</sup> Naturally, artists are concerned about this trend because they are not being notified or compensated when their works are included in training datasets.<sup>12</sup> Lawsuits are already in their early stages, with one notable example being filed in January of 2023 against Midjourney, Stability AI, and Deviant Art.<sup>13</sup> In this suit, a class of artists and illustrators allege that the defendants have violated their copyrights and unfair competition laws by training AIs on copyrighted images without consent or compensation.<sup>14</sup> Although much of the complaint was dismissed at the pleadings stage, the district court found that the plaintiffs made a plausible claim that Stability AI infringed on their copyrighted works while training its model.<sup>15</sup> Should the case resume, a hearing on the merits will likely occur in the coming months.

The current collision of AI and copyright law has resulted in familiar copyright issues being raised in new contexts. The question of factual copying is complicated by a technically

---

<sup>4</sup> Melissa Heikkilä, *This Artist is Dominating AI-Generated Art. And He's Not Happy About It.*, MIT TECH. REV. (Sept. 16, 2022), <https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/>.

<sup>5</sup> DALL-E 2, OPENAI [hereinafter *DALL-E 2*], <https://openai.com/dall-e-2/> (last visited Feb. 25, 2023).

<sup>6</sup> @midjourney, TWITTER (July 12, 2022, 11:41 PM), <https://twitter.com/midjourney/status/1547108864788553729>.

<sup>7</sup> *Stable Diffusion 2.0 Release*, STABILITY.AI [hereinafter *Stable Diffusion 2.0*], <https://stability.ai/blog/stable-diffusion-v2-release> (last visited Feb. 25, 2023).

<sup>8</sup> See Heikkilä, *supra* note 1.

<sup>9</sup> *Id.*

<sup>10</sup> FAQ, LAION, <https://laion.ai/faq/> (last visited Feb. 28, 2023).

<sup>11</sup> *Id.*

<sup>12</sup> Gil Appel, Juliana Neelbauer & David A. Schweidel, *Generative AI Has an Intellectual Property Problem*, HARV. BUS. REV. (Apr. 7, 2023), <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>.

<sup>13</sup> Min Chen, *Artists and Illustrators Are Suing Three A.I. Art Generators for Scraping and 'Collaging' Their Work Without Consent*, ARTNET NEWS (Jan. 24, 2023), <https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770>.

<sup>14</sup> *Id.*

<sup>15</sup> *Andersen v. Stability AI Ltd.*, No. 23-CV-00201-WHO, 2023 WL 7132064 at \*5 (N.D. Cal. Oct. 30, 2023).

intricate and variable training method of AI models. Liability allocation is complicated by developer discretion in creating and training open-ended models that are then handed off to users for experimenting. Fair use analyses are complicated by the novel and variable impact of generative AI outputs on the markets for an original artist's work. These questions, among others, must be answered to resolve lawsuits like those recently filed against Stability AI and those certain to follow.

This paper will attempt to answer some of these key copyright-related questions raised by text-to-image generative AIs. We will evaluate the merits of the types of claims brought by the plaintiffs against Stability AI in three separate areas. First, we will discuss the extent to which use of copyrighted images in AI training datasets satisfies a *prima facie* case of copyright infringement, considering issues of factual copying, transitory copying, and unlawful appropriation. Next, we will discuss theories of liability that could apply to hold developers or end users liable, including vicarious and contributory liability. Finally, we will discuss the extent to which the use of copyrighted images in AI training data is a fair use, paying particular attention to the ways in which generative AI tools are differentiated from the AI tools currently in use with respect to the first and fourth fair use factors. Throughout each analysis, we will outline the arguments and theories that would give any potential plaintiffs the greatest likelihood of success in future cases challenging the use of their copyrighted works in AI training.

## II. TEXT-TO-ART GENERATIVE AI

In 2022, OpenAI launched DALL-E 2, which quickly gained popularity because of its significant performance improvement compared with its predecessor.<sup>16</sup> In essence, DALL-E 2 generates images from users' text-based prompts.<sup>17</sup> The prompts are effectively unrestricted—whatever users want to conjure, they can request. While novel, it is not the only tool of its kind. Among others are Stability AI's Stable Diffusion 2.0<sup>18</sup> and Google's Imagen, which is not yet publicly available.<sup>19</sup> Each AI has distinguishing nuances, but at base, they all rely on diffusion models. To best analyze how *prima facie* copyright infringement applies to diffusion models, the rest of this section will explore how the models function and the user policies behind them.

### A. DIFFUSION MODELS

Modern text-to-image AIs generate images through use of “diffusion” models. In simple terms, diffusion models take “pure noise” and turn it into a recognizable image that matches a text prompt.<sup>20</sup> Thus, there are two aspects of generative AI models that require copyright analysis: model creation and image generation.

---

<sup>16</sup> DALL-E 2's website claims that its outputs are four times the resolution compared with its predecessor and that 71.7% of users saw improvements to prompt matching. *DALL-E 2*, *supra* note 2.

<sup>17</sup> *Id.*

<sup>18</sup> *Stable Diffusion 2.0*, *supra* note 4.

<sup>19</sup> *Imagen: Unprecedented Photorealism x Deep Level of Language Understanding*, GOOGLE, <https://imagen.research.google.com/> (last visited Feb. 25, 2023).

<sup>20</sup> Here, a “noisy” image consists of randomly arranged pixels. To the human eye, it looks like static. See Louis Bouchard, *Google Brain's Answer to Dall-E 2: Imagen*, LOUIS-FRANÇOIS BOUCHARD (May 23, 2022), <https://www.louisbouchard.ai/google-brain-imagen/>.

Models are created through a process known as “training.”<sup>21</sup> Every text-to-image model begins with a “text encoder,” which is trained to recognize the relationship between images and words.<sup>22</sup> Typically, text encoders are trained with image-text data from sources like LAION-5B, a database assembled by an independent party that indexes links to images along with image-text pairs.<sup>23</sup> Based on the text encoder’s understanding of image-text relationships, diffusion model training occurs by iteratively introducing “noise” into images in the training data until they become pure noise.<sup>24</sup> In repeating this process, the model learns the inverse—that is, by creating noise from a recognized image, it simultaneously learns how to “denoise” an unrecognizable image.<sup>25</sup> Once the diffusion model is trained, it operates by first creating a low-resolution image composed of noisy pixels.<sup>26</sup> It then iteratively de-noises the image until it reaches a final product that matches the desired text-image pair.<sup>27</sup>

However, after training, the model is merely a composition of computer code if it has no further instruction. For example, DALL-E 2 requires a text prompt, which instructs the diffusion model to create an image that it estimates the user will recognize as resembling the prompt.<sup>28</sup> In the case of DALL-E 2, the model is available to anyone with an Open AI account, so text prompts are often provided by commercial users and not the developers themselves.

The difference between the model developer and the database creator is another key distinction. While the developer and database creator can be one-in-the-same, this is not always the case.<sup>29</sup> The practices of database creators vary, but they often do not store the image files for a sustained period. For example, LAION indexes links to the underlying images and associates them with image-text pairs.<sup>30</sup> Although LAION does download each image to extract the embedded image-text pair information, it quickly deletes the original images and does not store them on its servers—only the links to the originals are saved and available in the database.<sup>31</sup> Regardless, any model built from that database will need to download or otherwise read the images for training.<sup>32</sup> The nature of exactly how each model interacts with the underlying data is unclear given the opacity of each developer’s proprietary operations. Thus, it is possible that some developers download the images for a significant amount of time while others merely extract the data without ever replicating it on their servers.

The nature of text-to-art AI introduces substantial complexity into copyright analyses. The technological methods themselves create potentially novel copying scenarios. The models’

---

<sup>21</sup> *Id.*

<sup>22</sup> CHITWAN SAHARIA ET AL., PHOTOREALISTIC TEXT-TO-IMAGE DIFFUSION MODELS WITH DEEP LANGUAGE UNDERSTANDING, 2205.11487 ARXIV (May 23, 2022), <https://arxiv.org/pdf/2205.11487.pdf>.

<sup>23</sup> *Stable Diffusion 2.0*, *supra* note 4.

<sup>24</sup> Bouchard, *supra* note 17.

<sup>25</sup> *Id.*

<sup>26</sup> *Id.*

<sup>27</sup> The initial output is typically low resolution, so modern models apply additional diffusion techniques that increase the resolution. The subsequent model takes patches of the original output and corrupts them again by introducing noise, only to denoise them into a higher-fidelity version of the original patch. This method of cascading diffusion models and patching helps to keep the process computationally viable. *Id.*

<sup>28</sup> *DALL-E 2*, *supra* note 2.

<sup>29</sup> For example, Stable Diffusion is trained on approximately 5.8 billion images that were indexed and tied to image-text pairs by a separate entity. *Stable Diffusion 2.0*, *supra* note 4.

<sup>30</sup> LAION, *supra* note 7.

<sup>31</sup> *Id.*

<sup>32</sup> *Id.*



training and use involve multiple parties at every step, which could cause confusion about who is responsible for any alleged infringement. The crux of this Article is to begin cutting through this complexity so that litigants can better understand the landscape.

## B. TEXT-TO-ART AI TERMS OF SERVICE

As with all modern internet tools, developers of generative AI image models set forth terms of service for their models' use. In addition to myriad standard clauses, the developers place liability for violating "any applicable law" onto the end user.<sup>33</sup> In some cases, the terms of service clearly account for copyright claims by notifying users that they are responsible for any liability "arising out of and relating in any way to intellectual property infringement."<sup>34</sup> However, some tools' terms of service give avenues for recourse to intellectual property owners who believe their copyrights have been infringed.<sup>35</sup>

While the similarities are notable, there are also significant differences between the terms of service for the major generative AI tools. For example, DALL-E 2 explicitly grants users ownership rights of both their input and the resulting output,<sup>36</sup> whereas Stable Diffusion notes that "content created through Stable Diffusion becomes public domain."<sup>37</sup> How these contracts will hold up under legal scrutiny is unclear and outside the scope of this article. Regardless, it's clear that developers have contemplated that their generative AIs could face challenges from copyright owners.

## C. USE CASES

To focus the analysis, this article will center around three use cases. The use cases are not meant to exhaustively analyze all copyright or fair use issues for generative AI art. However, each provides insight into potential complications of applying these doctrines to generative AI. This section gives a brief description of each use case.

### 1. "In the Style Of"

The major generative AI art tools enable users to prompt the model for images "in the style of" a particular artist.<sup>38</sup> Given the tools' ability to match custom user prompts, it is natural that such a function is available. In fact, some platforms actively encourage users to ask for images modeled after specific artists.<sup>39</sup> When artists' work is included in the training data, the output images can convincingly replicate their style.

---

<sup>33</sup> *Terms of Use*, OPENAI § 3(a) [hereinafter OPENAI, *Terms of Use*], <https://openai.com/policies/terms-of-use> (last visited Feb. 28, 2023).

<sup>34</sup> *Website Terms of Service*, STABILITY.AI [hereinafter STABILITY.AI, *Website Terms of Service*], <https://stability.ai/terms-of-use> (last visited Feb. 28, 2023).

<sup>35</sup> See, e.g., OPENAI, *Terms of Use*, *supra* note 30, § 3(d), <https://openai.com/policies/terms-of-use> (last visited Feb. 28, 2023) ("If you believe that your intellectual property rights have been infringed, please send notice to the address below. We may delete or disable content alleged to be infringing and may terminate accounts of repeat infringers.").

<sup>36</sup> *Id.* § 3(a).

<sup>37</sup> STABILITY.AI, *Website Terms of Service*, *supra* note 31.

<sup>38</sup> See Heikkilä, *supra* note 1.

<sup>39</sup> For example, until recently, DALL-E 2's homepage used an example prompt of "an astronaut riding a horse in the style of Andy Warhol." DALL-E 2, OPENAI, <https://openai.com/product/dall-e-2> (last visited Dec. 9, 2022).

People are already using these prompts at scale.<sup>40</sup> As previously mentioned, Greg Rutkowski is a digital artist who primarily creates “dream fantasy landscapes” in a distinct, classical style.<sup>41</sup> His art has become well known, primarily for being featured in numerous popular video, card, and tabletop games.<sup>42</sup> Perhaps unsurprisingly, his name has been evoked thousands of times on generative AI art platforms.<sup>43</sup> In fact, one report from September 2022 showed that his name was used approximately forty-six times more than Michelangelo, Pablo Picasso, and Leonardo da Vinci.<sup>44</sup> Upon learning that his name was being used with such a high frequency, Rutkowski did an online search for his name, which returned images with his name attached that he never made.<sup>45</sup> Rutkowski has legitimate concerns about how this will impact his livelihood. Not only could misattributed AI-generated images hurt his reputation, but these tools could weaken the demand for his original art.

## 2. *Entity-Centric Models*

As previously mentioned, the largest generative AI models are trained on millions, or billions, of images from various sources. However, other models have emerged that are trained on only a single artist’s work. One such model was trained only on Rembrandt van Rijn’s works and was dubbed “The Next Rembrandt” by its developers.<sup>46</sup> The model created a work with uncanny resemblance to an original Rembrandt, causing debate about whether it should be considered a Rembrandt original.<sup>47</sup> Notably, Rembrandt’s works are in the public domain, so copyright infringement is not at issue. Still, other models trained on a specific entity’s intellectual property have emerged,<sup>48</sup> which shows the potential for similar models.

## 3. *The Edit Function*

DALL-E 2’s edit function allows users to upload a photo of their choosing and prompt edits that are enabled by the platform’s robust diffusion techniques.<sup>49</sup> Using the provided image as a guide, the model can make edits to the image,<sup>50</sup> or even extrapolate additions to it.<sup>51</sup> On its website, OpenAI gives examples of “add[ing] a flamingo” of an existing image of an indoor pool and adding an entire background to a well-known portrait.<sup>52</sup> Importantly, the user always provides the baseline image.

---

<sup>40</sup> See Heikkilä, *supra* note 1.

<sup>41</sup> *Id.*

<sup>42</sup> *Id.*

<sup>43</sup> *Id.*

<sup>44</sup> *Id.*

<sup>45</sup> *Id.*

<sup>46</sup> Tim Brinkhof, *How to Paint Like Rembrandt, According to Artificial Intelligence*, DISCOVER (Aug. 23, 2021, 9:30 AM), <https://www.discovermagazine.com/technology/how-to-paint-like-rembrandt-according-to-artificial-intelligence>.

<sup>47</sup> *Id.*

<sup>48</sup> See, e.g., Justin Pinkey, *Text-to-Pokemon*, LAMBDA LABS, (Mar. 21, 2023), <https://replicate.com/lambdal/text-to-pokemon> (displaying a text-to-image model trained on Pokémon).

<sup>49</sup> See David Schnurr, *DALL-E Editor Guide*, OPENAI, (May 12, 2023), <https://help.openai.com/en/articles/6516417-dall-e-editor-guide>.

<sup>50</sup> OpenAI refers to this as “inpainting.” *DALL-E 2*, *supra* note 2.

<sup>51</sup> OpenAI refers to this as “outpainting.” *Id.*

<sup>52</sup> *Id.*

### III. *PRIMA FACIE* COPYRIGHT INFRINGEMENT

The Copyright Act affords copyright holders several rights with respect to their covered intellectual property.<sup>53</sup> Among the most important is the right to reproduction,<sup>54</sup> which is at the core of copyright infringement claims.<sup>55</sup> The reproduction right is limited to “copies and phonorecords,”<sup>56</sup> which are defined as “material objects . . . [in which a] work is fixed by any method now known or later developed, and from which the [work] can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device.”<sup>57</sup> Courts interpret material objects to be tangible representations of the original work, including data files.<sup>58</sup> Under this definition, a copy must also be “fixed,” meaning it is “sufficiently permanent to permit it to be perceived, reproduced, or otherwise communicated for a period of more than transitory duration.”<sup>59</sup> This definition has been interpreted broadly. If an unauthorized copy is created, infringement occurs even if the copy is not distributed or otherwise publicly used.<sup>60</sup>

A valid copyright also affords its holder the right “to prepare derivative works based upon the copyrighted work.”<sup>61</sup> Under the Copyright Act, a derivative work is defined as being “based upon one or more pre-existing works,” when it results in the original being “recast, transformed, or adapted.”<sup>62</sup> In practice, a derivative work infringes the copyright holder’s right to reproduction, but can be “made non-infringing through the acquisition of permission to use the underlying pre-existing expression.”<sup>63</sup> In other words, a derivative work is only non-infringing because the original work was copied with consent of the copyright owner—through licensing or other explicit permissions—or because the prior work is in the public domain.<sup>64</sup> Thus, alleged violations of the right to reproduction and the right to prepare derivative works evoke the same analytical framework, which is encompassed in the *prima facie* copyright doctrine.

A successful copyright infringement claim must show a *prima facie* case. *Prima facie* copyright infringement has two primary elements: ownership of a valid copyright with the commensurate rights and an occurrence of copying in violation of at least one of those rights.<sup>65</sup> This Article assumes the claimant has a valid copyright and wields all the corresponding rights

---

<sup>53</sup> See 17 U.S.C. § 106.

<sup>54</sup> *Id.* § 106(1).

<sup>55</sup> MELVILLE B. NIMMER & DAVID NIMMER, 2 NIMMER ON COPYRIGHT § 8.02(A) (2023).

<sup>56</sup> 17 U.S.C. § 106(1).

<sup>57</sup> *Id.* § 101 (definitions of “copies” and “phonorecords” exhibit the same language quoted).

<sup>58</sup> See *Capitol Records, LLC v. ReDigi Inc.*, 910 F.3d 649, 656 (2d Cir. 2018).

<sup>59</sup> 17 U.S.C. § 101.

<sup>60</sup> NIMMER & NIMMER, *supra* note 52, § 8.02(C).

<sup>61</sup> 17 U.S.C. § 106(2).

<sup>62</sup> *Id.* § 101.

<sup>63</sup> *Greene v. Ablon*, 794 F.3d 133, 158 (1st Cir. 2015); see also *Well-Made Toy Mfg. Corp. v. Goffa Int’l Corp.*, 354 F.3d 112, 117 (2d Cir. 2003) (claiming a work is only a derivative work if it would be considered an infringing work absent consent from the original author).

<sup>64</sup> NIMMER & NIMMER, *supra* note 52, § 3.01.

<sup>65</sup> NIMMER & NIMMER, *supra* note 52, § 13(D.02). Some commentators frame infringement as having three elements, splitting the ownership of a valid copyright prong in two: (1) claimant is the rightful holder of the copyright that (2) holds the right to reproduce at the time of the complaint. See WILLIAM F. PATRY, 3 PATRY ON COPYRIGHT § 9:4 (2023).

at the time of the claim. However, a deeper exploration of the copying component is necessary. Broadly, copying breaks down to two sub-elements: factual copying and unlawful appropriation.

### A. FACTUAL COPYING

*Arnstein v. Porter* was the first case to articulate factual copying as a distinct element.<sup>66</sup> The plaintiff, a musician, filed a copyright infringement suit alleging the defendant copied portions of his songs.<sup>67</sup> The Second Circuit bifurcated copying into two sub-elements: “(a) that defendant copied from plaintiff’s copyrighted work and (b) that the copying (assuming it to be proved) went so far as to constitute improper appropriation.”<sup>68</sup> The opinion specified that factual copying can be found either through direct evidence or circumstantial evidence.<sup>69</sup>

In *Arnstein*, the court effectively equates direct evidence to a defendant’s explicit admission of copying.<sup>70</sup> Eyewitness testimony would also suffice as direct evidence.<sup>71</sup> Unsurprisingly, direct evidence is uncommon because “[p]lagiarists rarely work in the open and direct proof of actual copying is seldom available.”<sup>72</sup> However, it’s conceivable that direct evidence of copying digital works is easier to obtain than in non-digital settings given the ubiquity of data logs in modern computer environments.

If direct evidence is unavailable, copying can still be proven through sufficient circumstantial evidence. Specifically, courts find copying when evidence proves the defendant had access to the copyrighted work and when there is probative similarity between the original and the alleged copy.<sup>73</sup> While access doesn’t have to be directly proven, there must be a reasonable possibility of access, not just a bare possibility.<sup>74</sup> Content’s existence on the internet is not dispositive of access—there must be a reasonable possibility that the defendant themselves accessed the content.<sup>75</sup>

Beyond access, the plaintiff must also prove probative similarity, which requires a minimum level of overlap between the works.<sup>76</sup> Probative similarity can be found in either copyrightable or non-copyrightable aspects of the original work, distinguishing it from substantial similarity, discussed below.<sup>77</sup> Many courts holistically compare the original with the second work and analyze probative similarity through an objective lens, and will only distinguish expressive and non-expressive elements when analyzing substantial similarity.<sup>78</sup>

---

<sup>66</sup> The case doesn’t explicitly name the element “factual copying,” but is the concept’s genesis. *See Arnstein v. Porter*, 154 F.2d 464, 468 (2d Cir. 1946).

<sup>67</sup> *Id.*

<sup>68</sup> *Id.*

<sup>69</sup> *Id.*

<sup>70</sup> *See id.*

<sup>71</sup> *NIMMER & NIMMER*, *supra* note 52, § 13(D.04).

<sup>72</sup> *Johnson v. Gordon*, 409 F.3d 12, 18 (1st Cir. 2005).

<sup>73</sup> *PATRY*, *supra* note 62, § 9:21.

<sup>74</sup> *See Building Graphics, Inc. v. Lennar Corp.*, 708 F.3d 573, 578–579 (4th Cir. 2013).

<sup>75</sup> *See Design Basics, LLC v. Lexington Homes, Inc.*, 858 F.3d 1093, 1107 (7th Cir. 2017).

<sup>76</sup> *PATRY*, *supra* note 62, § 9:21.

<sup>77</sup> *NIMMER & NIMMER*, *supra* note 52, § 13D.05.

<sup>78</sup> *See, e.g., Gates Rubber Co. v. Bando Chem. Indus.*, 9 F.3d 823, 833 (10th Cir. 1993) (“Ultimately, to prove factual copying, the plaintiff must come forward with sufficient evidence that a reasonable factfinder, taking together the evidence of access and the similarities between the programs, could find that the second work was copied from the first.”).

Ultimately, in the absence of direct evidence, sufficient access and probative similarity still amount to copying.

Sufficient circumstantial evidence of copying can also be found through “striking similarity.”<sup>79</sup> Works are strikingly similar when the works are so alike that a common source is beyond doubt.<sup>80</sup> Similarity may be striking even if not verbatim as long as the similarities are so extensive as to preclude the possibility of independent creation.<sup>81</sup> However, striking similarity is not universally endorsed as sufficient evidence to prove copying.<sup>82</sup>

## B. UNLAWFUL APPROPRIATION

In addition to factual copying, plaintiffs must prove their copying amounted to unlawful appropriation. Unlawful appropriation, also known as legal copying or substantial similarity, is a messy doctrine for which factors have never been codified.<sup>83</sup> Patry’s approach breaks unlawful appropriation down into two factors: materiality and expression.<sup>84</sup> He asserts that “the material similarities must be from similarities in expressive material, exclusively owned by the plaintiff, that appear in both the plaintiff’s and the defendant’s works.”<sup>85</sup> In other words, there must be material overlap in the expressive elements of the two works. However, he gives no guidance on the scale of materiality needed to qualify as unlawful appropriation.

While no statute expressly confines unlawful appropriation analyses, circuit courts and practitioners have provided competing frameworks.<sup>86</sup> In essence, each framework involves reconceptualizing or abstracting the works to examine whether the overlap in expressive elements is significant enough to be considered a copy. For example, some jurists have found unlawful appropriation through “fragmented nonliteral similarity,” which involves literal copying of a work’s components instead of the work in its entirety.<sup>87</sup>

Nimmer also identifies “comprehensive nonliteral similarity” as a framework for improper appropriation that judges whether there is copying of “the fundamental essence or structure of one work . . . in another.”<sup>88</sup> Comprehensive nonliteral similarity stems from the Second Circuit’s “abstractions test” that examines broad characteristics of the works’ expressions to find overlap.<sup>89</sup> As the abstractions become more general, there is a point where overlap between works does not imply substantial similarity.<sup>90</sup> To illustrate, consider two drawings of a city skyline.<sup>91</sup> If similarities can only be found at such a high level of abstraction—that they both depict city skylines—a reasonable person is unlikely to find substantial similarity. But if there are numerous similarities in more specific expressive

---

<sup>79</sup> NIMMER & NIMMER, *supra* note 52, § 13D.07.

<sup>80</sup> *Id.*

<sup>81</sup> *See* Testa v. Janssen, 492 F. Supp. 198 (W.D. Pa. 1980).

<sup>82</sup> PATRY, *supra* note 62, § 9:38 (“[Striking similarity] is a doctrine that should have been strangled at birth.”).

<sup>83</sup> NIMMER & NIMMER, *supra* note 52, § 13.03.

<sup>84</sup> PATRY, *supra* note 62, § 9:59.

<sup>85</sup> *Id.*

<sup>86</sup> NIMMER & NIMMER, *supra* note 52, § 13.03.

<sup>87</sup> *Id.* § 13.03(A)(2).

<sup>88</sup> *Id.* § 13.03(A)(1).

<sup>89</sup> *Id.* § 13.03(A)(1)(a); *see also* Nichols v. Universal Pictures Co., 45 F.2d 119, 121 (2d Cir. 1930).

<sup>90</sup> NIMMER & NIMMER, *supra* note 52, § 13.03(A)(1).

<sup>91</sup> This example is inspired by a real case. *See* Steinberg v. Columbia Pictures Indus., Inc., 663 F. Supp. 706 (S.D.N.Y. 1987).

elements like the design and layout of the buildings, the backgrounds, the framing, and the vantage point from which the drawing is made, the only reasonable conclusion is that the second work copied the original.<sup>92</sup> Although some courts still rely on variations of the abstractions test, some scholars assert that it was never meant to be a workable standard, and advise against its ongoing use.<sup>93</sup>

Perhaps the most controversial theory is the Ninth Circuit's "total concept and feel" framework, which Nimmer characterizes as a form of comprehensive nonliteral similarity.<sup>94</sup> It was first articulated in *Roth Greeting Cards v. United Card Co.*, where the court looked at the works in their entirety and found factual copying through conceptual and structural similarities.<sup>95</sup> In *Roth*, even though the greeting cards did not share exact visual characteristics, similarities were obvious between the characters depicted, the moods struck, the messages delivered, the arrangement of words, and even the lettering.<sup>96</sup> The court found that the nonliteral similarities sufficed to prove unlawful appropriation. Detractors of the total concept and feel framework claim that it puts focus onto uncopyrightable elements.<sup>97</sup> As a result, some contemporary courts repudiate this test as a standalone framework for evaluating substantial similarity.<sup>98</sup>

These frameworks provide some guidance for evaluating unlawful appropriation, but none are widely accepted. However, the frameworks' commonalities are instructive. Each framework requires analyzing individual elements that are abstracted away from the work as a whole. Plaintiffs claiming copyright infringement should still consider abstracted elements but focus more closely on expressive elements. Conceptual elements can still be acknowledged but should not be the crux of a claim.

Ultimately, unlawful appropriation remains a factually intensive and case-dependent analysis that often falls to unpredictable juries for final decision. Without further congressional guidance, the doctrine will likely remain highly fragmented and inconsistent.

### C. TRANSITORY COPYING

Beyond factual and legal copying, courts have developed a durational element to copyright infringement.<sup>99</sup> Stemming from the text of the Copyright Act, an infringing copy is only "fixed" when it meets two distinct elements: (1) being in a "sufficiently permanent," tangible medium to allow it to be perceived or reproduced and (2) that it remain embodied in that medium "for a period more than a transitory duration."<sup>100</sup> Although case law provides

---

<sup>92</sup> *Id.* at 713.

<sup>93</sup> PATRY, *supra* note 62, § 9:93.

<sup>94</sup> NIMMER & NIMMER, *supra* note 52, § 13.03(A)(1).

<sup>95</sup> See *Roth Greeting Cards v. United Card Co.*, 429 F.2d 1106, 1110 (9th Cir. 1970).

<sup>96</sup> *Id.* ("[T]he characters depicted in the art work, the mood they portrayed, the combination of art work conveying a particular mood with a particular message, and the arrangement of the words on the greeting card are substantially the same as in Roth's cards. In several instances the lettering is also very similar.").

<sup>97</sup> NIMMER & NIMMER, *supra* note 52, § 13.03(A)(1)(c).

<sup>98</sup> See, e.g., *Attia v. Society of N.Y. Hosp.*, 201 F.3d 50, 54 (2d Cir. 1999) (warning that concepts are not copyrightable and shifting analytical focus onto expressive elements). For a deeper explanation of what is copyrightable expression and what is not, see *infra* Part III.D.

<sup>99</sup> See, e.g., *Cartoon Network LP, LLLP v. CSC Holdings, Inc.*, 536 F.3d 121, 127 (2d Cir. 2008) (finding that the definition of a "copy" under 17 U.S.C. § 101 requires that it be fixed for a sufficient duration).

<sup>100</sup> *Id.* (quoting 17 U.S.C. § 101).

guideposts for defining a transitory period, the exact requirement for how long embodiment must occur remains ambiguous.

In *Cartoon Network v. CSC Holdings*, the Second Circuit determined that pieces of data stored in a buffer system for only 1.2 seconds constituted a transitory duration.<sup>101</sup> The defendant operated a television cable system that enabled users to digitally record shows at the command of its subscribers.<sup>102</sup> To effectuate the system, the defendant copied the requested show, bit by bit, in small chunks.<sup>103</sup> Each portion of the show was copied into the buffer for 1.2 seconds before being transferred to a server where it resided for the customer's viewing pleasure.<sup>104</sup> Although the court determined that the buffer was a sufficiently permanent and tangible medium for copyright purposes, the mere 1.2 seconds it remained in the buffer was deemed transitory.<sup>105</sup> Thus, no copyright infringement was found for CSC's buffer system.<sup>106</sup>

Other cases shed light on the amount of time that *doesn't* constitute a transitory period. For example, *MAI Systems v. Peak Computer* considered whether copyright infringement occurred when Peak temporarily copied MAI's proprietary operating system onto random access memory ("RAM") during its maintenance and repairs of MAI-integrated computers.<sup>107</sup> Although the opinion did not discuss the exact length of time the MAI system resided on Peak's RAM, the *Cartoon Network* court assumed that it was "for at least several minutes."<sup>108</sup> Ultimately, the Second Circuit concluded that an embodiment of several minutes clearly constitutes a non-transitory duration.<sup>109</sup> The contours of the transitory duration doctrine are fuzzy, but the caselaw at least provides outer bounds: 1.2 seconds is likely transitory, while several minutes is not.

#### D. OTHER LIMITS OF COPYRIGHT INFRINGEMENT

Although copyright holders enjoy significant benefits, the protection is not absolute.<sup>110</sup> For example, because a *prima facie* copyright infringement claim requires factual copying,<sup>111</sup> no infringement occurs when no copying occurs regardless of the level of similarity between two works.<sup>112</sup> In other words, even when a work that is substantially similar (or identical) to a previous work, there can be no infringement if the subsequent work was made with no reference to the prior work.<sup>113</sup> This is especially relevant when two similar works are based upon the same original source. If the subsequent work never referenced the first work, the first's copyright holder has no claim against the subsequent author.

---

<sup>101</sup> *See id.*

<sup>102</sup> The system is commonly known as a "Digital Video Recorder," or a DVR.

<sup>103</sup> *See id.*

<sup>104</sup> *See id.*

<sup>105</sup> *See id.*

<sup>106</sup> *Id.* at 130.

<sup>107</sup> *MAI Sys. Corp. v. Peak Computer, Inc.*, 991 F.2d 511, 517 (9th Cir. 1993).

<sup>108</sup> *Cartoon Network LP*, 536 F.3d at 127 (2d Cir. 2008) (citing *MAI Systems*, 991 F.2d at 511).

<sup>109</sup> *Id.* at 128.

<sup>110</sup> NIMMER & NIMMER, *supra* note 52, § 8.01.

<sup>111</sup> *See supra* Part III.A.

<sup>112</sup> NIMMER & NIMMER, *supra* note 52, § 8.01.

<sup>113</sup> *Id.* ("[T]he rights of a copyright owner are not infringed if a subsequent work, although substantially similar (or even identical), has been independently created without reference to the prior work.").

Copyright also does not extend to the uncopyrightable elements of a copyrighted work. Courts have long held that the facts or ideas underlying works are not protected by copyright—only the original expression of those facts or ideas.<sup>114</sup> Delineating between protectable expression and unprotectable ideas is not typically straightforward, especially considering that different types of works vary in how much expression they contain.<sup>115</sup> Because infringement requires copying of protected elements, only expressions are considered in an unlawful appropriation analysis.<sup>116</sup> Accordingly, when the original work is composed primarily of ideas or facts, copyright protection is narrow, whereas a highly expressive work enjoys broad copyright protection.<sup>117</sup> This concept is clearly articulated in *Mattel, Inc. v. MGA Entertainment, Inc.*:

If there's a wide range of expression (for example, there are gazillions of ways to make an aliens-attack movie), then copyright protection is "broad" and a work will infringe if it's "substantially similar" to the copyrighted work. If there's only a narrow range of expression (for example, there are only so many ways to paint a red bouncy ball on blank canvas), then copyright protection is "thin" and a work must be "virtually identical" to infringe.<sup>118</sup>

Thus, in the context of images, distinguishing between expressions and ideas is necessarily fact intensive. In *Mattel*, the Ninth Circuit determined that broad protection applied to Mattel-owned doll sketches that were taken by a former employee to create a competing product at MGA because there is a “wide range of expression” possible when creating this type of doll.<sup>119</sup> The court found that the sketches exhibited many expressive elements including the dolls’ hair styles, clothing, facial features, and accessories. These elements alone were the basis for the substantial similarity test, which did not require that MGA products be virtually identical to prove infringement.<sup>120</sup>

When an image is primarily composed of non-expressive elements, the analysis becomes more exacting. In *Harney v. Sony Pictures TV, Inc.*, the First Circuit found that copying did not occur when Sony produced a movie that reproduced and depicted a copyrighted photograph.<sup>121</sup> The original image and its depiction in the film showed undeniable structural similarities, like the subjects’ characteristics, the images’ arrangement, and the setting.<sup>122</sup> However, the expressive elements—the background, lighting, and religious details—were not similar enough to suggest substantial similarity.<sup>123</sup> Because the original photo was taken candidly, fewer

---

<sup>114</sup> See, e.g., *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 350, 111 S. Ct. 1282, 1290 (1991) (citing *Harper & Row, Publrs. v. Nation Enters.*, 471 U.S. 539, 548, 105 S. Ct. 2218, 2224 (1985)).

<sup>115</sup> See *Apple Computer, Inc. v. Microsoft Corp.*, 35 F.3d 1435, 1444-47 (9th Cir. 1994) (dissecting Xerox’s graphical user interface and determining that ideas like iconic representation of familiar objects from the office environment are not protected while the exact manifestation of those icons is protected).

<sup>116</sup> See *id.* at 1445.

<sup>117</sup> See *id.* at 1444.

<sup>118</sup> *Mattel, Inc. v. MGA Entm't, Inc.*, 616 F.3d 904, 914 (9th Cir. 2010) (citations omitted).

<sup>119</sup> *Id.* at 916.

<sup>120</sup> *Id.*

<sup>121</sup> *Harney v. Sony Pictures TV, Inc.*, 704 F.3d 173, 177 (1st Cir. 2013).

<sup>122</sup> *Id.* at 186-87.

<sup>123</sup> *Id.* at 187.



elements were considered original expressions, giving the photo relatively narrow protection.<sup>124</sup>

Plaintiffs bringing copyright infringement suits should be aware of the limits of their rights. Although creative artworks are often granted broad protection, that protection will only extend to expressive elements and not to the underlying ideas.

#### E. *PRIMA FACIE* COPYRIGHT INFRINGEMENT IN DIGITAL SPACES

The Ninth Circuit's decision in *Perfect 10 v. Amazon* provides a helpful example to align this doctrine with this Article's topic. In *Perfect 10* the plaintiff asserted that one way the defendants (including Google) infringed its copyrights was by displaying thumbnails of its images when a user made a query in Google Image Search.<sup>125</sup> The court found that Google's actions amounted to *prima facie* copyright infringement because it had permanently copied the original images verbatim onto its own servers, which sufficed to show *prima facie* infringement.<sup>126</sup> Although the thumbnails were ultimately excused as fair use, this example shows that creating replicas of copyrighted digital images on a separate server amounts to copying sufficient for *prima facie* copyright infringement.

The same case also shows the limits of *prima facie* copyright infringement in digital contexts. Google Image Search's user experience had two relevant components: the thumbnail view discussed above and a full-size image view that appeared after a user clicked the thumbnail.<sup>127</sup> Google used "HTML instructions" to accomplish the full-size view, which were simply lines of code that directed the user's browser to access the third-party webpage that hosted the original image.<sup>128</sup> On top of the thumbnail claim, *Perfect 10* alleged that Google directly infringed its copyright through use of the HTML instructions.<sup>129</sup> The court disagreed, finding that this execution did not "communicat[e] a copy of the image" but merely directed a user's browser to a different site that stored the full-size images.<sup>130</sup> Contrary to the thumbnail issue, providing such instructions did not involve Google storing a copy of the original for the purposes of the Copyright Act, so no copying nor infringement occurred.<sup>131</sup> Even if Google facilitated a user's access to infringing images via a link, they would not be subject to direct infringement.<sup>132</sup> However, the court noted that such a scenario could amount to secondary liability,<sup>133</sup> which this Article will now explore.

---

<sup>124</sup> *See id.*

<sup>125</sup> *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1159 (9th Cir. 2007).

<sup>126</sup> *Id.* at 1160. Note that this issue examined the copyright holder's display right, not the right to reproduction or the right to create derivative works. However, the analysis here hinged on which entity was displaying a copy, and Google could only have displayed the copy itself if it made a copy in the first place. *Id.* Thus, the analysis of copying here will transfer to the reproduction and derivative rights. *See, e.g., Capitol Recs., LLC v. ReDigi Inc.*, 910 F.3d 649, 657 (2d Cir. 2018) (finding that a digital reproduction of a music file on a new server was "a reproduction" despite the file being deleted from the original location).

<sup>127</sup> *Perfect 10*, 508 F.3d at 1155.

<sup>128</sup> *Id.* at 1155–56.

<sup>129</sup> *Id.* at 1161.

<sup>130</sup> *Id.*

<sup>131</sup> *Id.* Similar to the thumbnail example, although this analysis involved the display right, the copying analysis will be the same for the reproduction and derivative works rights.

<sup>132</sup> *Id.*

<sup>133</sup> *Id.*

#### IV. SECONDARY LIABILITY

Secondary liability is not enumerated in the Copyright Act of 1976 but is instead a product of common law. In 1916, the 2nd circuit held that "all who unite in an infringement are . . . liable for the damages."<sup>134</sup> Since then, courts have engaged in a scoping exercise as they attempt to "identif[y] circumstances in which it is just to hold one individually accountable for the actions of another."<sup>135</sup> Expanding liability to those engaged in infringement has thus been embedded in U.S. copyright law and has granted copyright holders greater protection.

The two types of secondary liability in copyright law are vicarious and contributory infringement. "One infringes contributorily by intentionally inducing or encouraging direct infringement and infringes vicariously by profiting from direct infringement while declining to exercise the right to stop or limit it."<sup>136</sup> However, many courts have ineffectively delineated between contributory and vicarious liability. Practically, this means that courts may confuse these modes of liability "as the lines between direct infringement, contributory infringement, and vicarious liability are not clearly drawn."<sup>137</sup>

Broadly, vicarious liability requires a defendant to (1) have the right and ability to supervise the infringing conduct and (2) have a financial interest in that conduct.<sup>138</sup> On the other hand, a contributorily liable defendant must (1) have knowledge of the infringing conduct and (2) "materially contribute to the direct infringement."<sup>139</sup> Finally, both vicarious and contributory liability are necessarily predicated on the direct infringement of a third party.

##### A. VICARIOUS LIABILITY

Conventional vicarious liability principles require an agency relationship, historically based on *respondeat superior*. However, vicarious copyright liability functions as a much more "expansive form of *respondeat superior*."<sup>140</sup> In a landmark vicarious liability case from 1963, *Shapiro v. Green*, the Second Circuit explicitly fashioned a vicarious liability principle for copyright that sat outside of this conventional agentic relationship. This increased the scope and power of copyright. In *Shapiro*, Green granted a verbatim copier a license to operate in over twenty of Green's stores for over a decade.<sup>141</sup> Green had the "unreviewable discretion" to discharge any of Jalen's employees, and all of the cash from Jalen's sale of bootlegs tapes went into Green's cash registers.<sup>142</sup>

In answering the question of Green's liability, the court analyzed dance hall cases and found where a relationship like that of a landlord-tenant existed, the defendant had insufficient control to find vicarious liability.<sup>143</sup> By contrast, the owners or operators of establishments were liable when they more directly benefited from or facilitated infringement through

<sup>134</sup> *Gross v. Van Dyk Gravure Co.*, 230 F. 412, 414 (2d Cir. 1916). (Here, Rochlitz made the infringing photo, Gravure made 16,000 copies that he sold to Seligman, who then sold these copies to the public).

<sup>135</sup> *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 435 (1984).

<sup>136</sup> *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 914 (2005).

<sup>137</sup> *Id.* at 930.

<sup>138</sup> PATRY, *supra* note 62, § 21:66.

<sup>139</sup> *Id.* § 19:11.

<sup>140</sup> *Id.* § 21:44.

<sup>141</sup> *Shapiro, Bernstein & Co. v. H. L. Green Co.*, 316 F.2d 304, 307 (2d Cir. 1963).

<sup>142</sup> *Id.* at 306.

<sup>143</sup> *Id.* at 308. *See also* *Deutsch v. Arnold*, 98 F.2d 686, 688 (2d Cir. 1938) (finding that "[s]omething more than the mere relation of landlord and tenant must exist to give rise to a cause of action by plaintiffs against these defendants for infringement of their copyright on the demised premises").

knowingly receiving infringing broadcasts or hiring bands that played infringing performances.<sup>144</sup>

Those cases that looked more like a landlord-tenant relationship are less likely to be considered infringing, while those that look more like an employer-employee relationship are more likely to be considered infringing.<sup>145</sup> The court considered that the secondary infringer had “the power to police carefully the conduct of its concessionaire Jalen [direct infringer]”<sup>146</sup> and “derived an obvious and direct financial benefit from the infringement.”<sup>147</sup>

Particularly relevant to the text-to-art AI context, the *Shapiro* court accounted for the secondary infringer’s potential costs from disallowing its concessionaire from selling bootlegs.<sup>148</sup> Analyzing these relative burdens comports with the utilitarian and economic rationale of US copyright law.

While the court in *Shapiro* added to the analysis of secondary infringement, it did not provide an easily applicable definition of vicarious liability. *Gershwin v. Columbia* reaffirmed that vicarious liability does not require actual knowledge of infringing activity and held that it can occur in the absence of the employer-employee relationship if the vicariously liable had “the right and ability to supervise the infringing activity and also ha[d] a direct financial interest in such activities.”<sup>149</sup> The right and ability to supervise or control and a financial interest in infringement are now the two fundamental elements of any vicarious liability case. Moreover, the *Gershwin* court’s analysis of *Shapiro*, suggested that “the policies of the copyright law would be best effectuated if Green were held liable.”<sup>150</sup> Again we see that secondary infringement is explicitly policy focused.

In the digital context, vicarious liability has been analyzed most notably in relation to file-sharing services provided by companies such as Napster:

Through a process commonly called “peer-to-peer” file sharing, Napster allows its users to: (1) make MP3 music files stored on individual computer hard drives available for copying by other Napster users; (2) search for MP3 music files stored on other users’ computers; and (3) transfer exact copies of the contents of other users’ MP3 files from one computer to another via the Internet. These functions are made possible by Napster’s MusicShare software, available free of charge from Napster’s Internet site, and Napster’s network servers and server-side software.<sup>151</sup>

Napster applies the financial benefit prong to an Internet Service Provider (“ISP”) and illuminates what the right and ability to control or supervise looks like in a digital context.

---

<sup>144</sup> *Buck v. Jewell-La Salle Realty Co.*, 283 U.S. 191, 195 (1931) (a hotel that played an infringing radio broadcast, and provided its patrons the means of listening to the infringing broadcast, was found liable after the plaintiff had previously warned them of their infringement); *see also* *Dreamland Ball Room v. Shapiro, Bernstein & Co.*, 36 F.2d 354, 355 (7th Cir. 1929) (finding that “the owner of a dance hall at whose place copyrighted musical compositions are played in violation of the rights of the copyright holder is liable, if the playing be for the profit of the proprietor of the dance hall”).

<sup>145</sup> *Shapiro*, 316 F.2d at 307.

<sup>146</sup> *Id.* at 308.

<sup>147</sup> *Fonovisa, Inc. v. Cherry Auction, Inc.*, 76 F.3d 259, 262 (9th Cir. 1996).

<sup>148</sup> *Shapiro*, 316 F.2d at 307.

<sup>149</sup> *Gershwin Pub. Corp. v. Columbia Artists Mgmt., Inc.*, 443 F.2d 1159, 1162 (2d Cir. 1971) (according to the modern definitions enumerated in *Grokster* of contributory and vicarious liability, the *Gershwin* court conflated the two definitions finding that “a person who promoted or induced the infringing acts of the performer has been held jointly and severally liable as a ‘vicarious’ infringer.”).

<sup>150</sup> *Id.* at 1162.

<sup>151</sup> *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004, 1011 (9th Cir. 2001), *aff’d* 284 F.3d 1091 (9th Cir. 2002).

Financial benefit was satisfied because “Napster would likely charge users in the future.” Because “future revenue is directly dependent upon ‘increases in userbase,’”<sup>152</sup> Napster further demonstrates how attenuated a financial benefit may be given the balance of equities as Napster facilitated mass verbatim piracy. Napster had the right and ability to supervise infringing conduct because of their reservation of rights policy that let them terminate accounts for any reason.<sup>153</sup> The court pronounced that “[t]o escape imposition of vicarious liability, the reserved right to police must be exercised to its fullest extent. Turning a blind eye to detectable acts of infringement for the sake of profit gives rise to liability.”<sup>154</sup> The court specifically noted the boundaries of Napster’s premises, suggesting that even though the files were stored on user’s computers, users were able to find infringing material listed on Napster’s search indices and that “the file name indices, therefore are within the ‘premises’ that Napster has the ability to police.”<sup>155</sup>

The court in *In Re Aimster*, a Napster look-a-like, further explores the rationale behind expanding vicarious liability in copyright. Simply, if a plaintiff is unlikely to obtain relief against an impecunious direct infringer, then vicarious liability may be just when “effective relief is obtainable from someone who bears a relation to the direct infringers.”<sup>156</sup> Moreover, while *Grokster* is better known for the Supreme Court’s contributory infringement ruling, the Ninth Circuit’s analysis of Grokster’s vicarious liability was not overruled. As a result, it helps clarify the right and ability to control prong of vicarious liability. While Napster controlled the central song listing indices, Grokster was a truly distributed system where there was no registration or log-in process.<sup>157</sup> Furthermore, Grokster did not host a central song index, and Grokster could not terminate access without shutting down the entire system, so it was not held vicariously liable.<sup>158</sup> Creating a “dumb” system, or a system with less control, actually helped a defendant escape vicarious liability, though *Grokster* was later found contributorily liable by way of inducement.

Finally, the court in *Polygram* provides an overview of vicarious liability in the context of a Trade Show operator. The court explicitly states that vicarious liability fundamentally addresses “risk allocation” and balances incentives.<sup>159</sup> *Polygram* suggests that vicarious liability is more often found when the profiting enterprise is better able to spread the cost and distribute the loss than an innocent injured plaintiff.<sup>160</sup>

While the preceding cases all end in rulings finding vicarious liability, there are many cases that do not. In two cases involving *Perfect 10*, the courts stated that Google’s image search index linking to infringing websites and Visa’s payment card processing were

---

<sup>152</sup> Mark A. Lemley & R. Anthony Reese, *Reducing Digital Copyright Infringement Without Restricting Innovation*, 56 STAN. L. REV. 1345, 1366 (2004).

<sup>153</sup> *Napster*, 239 F.3d at 1023.

<sup>154</sup> *Id.*

<sup>155</sup> *Id.*; see also *Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc.*, 907 F. Supp. 1361, 1376 (N.D. Cal. 1995) (finding, prior to section 230, that plaintiffs raised a genuine issue of material fact as to whether a ‘computer bulletin board service’ should be vicariously liable for the potentially infringing postings of its users).

<sup>156</sup> *In re Aimster Copyright Litig.*, 334 F.3d 643, 654 (7th Cir. 2003).

<sup>157</sup> *Metro-Goldwyn-Mayer Studios, Inc. v. Grokster Ltd.*, 380 F.3d 1154, 1165 (9th Cir. 2004), *vacated and remanded*, 545 U.S. 913 (2005).

<sup>158</sup> *Id.*

<sup>159</sup> *Polygram Int'l Pub., Inc. v. Nevada/TIG, Inc.*, 855 F. Supp. 1314, 1325 (D. Mass. 1994).

<sup>160</sup> *Id.* at 1326 (“By focusing on benefit received from and control over an enterprise, a court can evaluate the defendant’s ability to spread losses and police conduct within the enterprise, as well as the underlying fairness of holding the defendant liable.”).

insufficient supervision over third-party online infringement.<sup>161</sup> Relevantly, the Ninth Circuit held that a theoretical technical solution to control infringing conduct did not satisfy the ability to control prong of vicarious liability.<sup>162</sup>

This is particularly relevant in the present context because the same questions are asked of expressive AI developers and implementers. Commentators have suggested that these companies adopt output filtration mechanisms to reduce or eliminate infringement. However, output filters are less viable in the present context because image search more often involves verbatim copies that are more easily algorithmically targeted. Here, many of the potential infringements are at least somewhat speculative as they are primarily non-verbatim copies. Therefore, substantial similarity, normally a question of fact for a jury, would have to be algorithmically analyzed, and all of Copyright's many limiting doctrines would have to be algorithmically implemented. Applying these limiting doctrines requires analog information that cannot feasibly be incorporated into an algorithm. Such an algorithmic output filtering solution is not viable and is undesirable.<sup>163</sup>

## B. CONTRIBUTORY LIABILITY

Contributory liability is rooted in a 1911 case where a defendant contributed to an infringement of the public display right for producing a derivative movie.<sup>164</sup> Contributory infringement was not initially widely used but has taken on greater significance in the digital era. *Gershwin*, though explicitly a vicarious liability case, formulated the foundational contributory liability standard. The court found that "[o]ne who, with knowledge of the infringing activity, induces, causes or materially contributes to the infringing conduct of another, may be held liable as a 'contributory' infringer."<sup>165</sup> Outside of the necessary direct infringement, from which all secondary liability attaches, there are 2 necessary components of contributory liability: (1) knowledge and (2) inducement, cause, or material contribution to an infringement.

### 1. *Material Contribution*

Contribution may be substantial or material, which means it "must bear a direct relationship to the infringing acts."<sup>166</sup> Providing the infringer site and facilities, encouraging, or inducing the infringer are frequently seen as sufficient material contributions to constitute contributory infringement, depending on the context. Like standards for vicarious liability, defining a material contribution is flexible as it ultimately depends on whether the defendant is "roughly an infringer's accomplice."<sup>167</sup>

---

<sup>161</sup> *Perfect 10, Inc. v. Visa Int'l Serv. Ass'n*, 494 F.3d 788, 805 (9th Cir. 2007) (finding that while restricting payments would make the infringing websites unprofitable, that is insufficient control over the infringing conduct); *see also* *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1173-75 (9th Cir. 2007) (finding a search engine's indexing of infringing websites was an issue of contributory liability and the ability to delist these websites did not demonstrate sufficient control as a prong of vicarious liability).

<sup>162</sup> *Perfect 10 v. Amazon*, 508 F.3d at 1174 (citations omitted) ("Perfect 10's suggestions regarding measures Google could implement to prevent its web crawler from indexing infringing websites and to block access to infringing images were not workable. Rather, the suggestions suffered from both 'imprecision and overbreadth.'").

<sup>163</sup> Dan L. Burk, *Algorithmic Fair Use*, 86 U. CHI. L. REV. 283, 306-07 (2019).

<sup>164</sup> *Kalem Co. v. Harper Bros.*, 222 U.S. 55, 63 (1911).

<sup>165</sup> *Gershwin Pub. Corp. v. Columbia Artists Mgmt., Inc.*, 443 F.2d 1159, 1162 (2d Cir. 1971).

<sup>166</sup> *Livnat v. Lavi*, No. 96 CIV. 4967 (RWS), 1998 WL 43221, at \*3 (S.D.N.Y. Feb. 2, 1998).

<sup>167</sup> *Flava Works, Inc. v. Gunter*, 689 F.3d 754, 754-55 (7th Cir. 2012).

In one instance, providing plaintiff's videos to participant schools was sufficiently material for the defendants to be "jointly and severally liable" for contributorily infringing the plaintiff's public performance right.<sup>168</sup> In a series of flea market cases, "providing the *site and facilities* for known infringing activity is sufficient to establish contributory liability."<sup>169</sup> The specific facilities included "the provision of space, utilities, parking, advertising, plumbing, and customers."<sup>170</sup> It is likely that the strong policy implications<sup>171</sup> and a high level of knowledge impacts the material contribution analysis.<sup>172</sup>

In the digital context, providing a site and facility is frequently connected to a defendant hosting infringing content on their servers. In *Perfect 10 v. Visa*, the court found that payment card processors did not provide site and facilities because defendants "do not operate the servers on which the[] [infringing content] reside[s] . . . nor does any infringing material ever reside on or pass through any network or computer Defendants operate."<sup>173</sup> Moreover, Visa's actions had "no direct connection to the infringement."<sup>174</sup> By contrast, the court in *Napster* found that providing the service that enabled users to infringe was provision of the site and facilities of infringement.<sup>175</sup> Napster did not host or provide the infringing files themselves, as the end users created the network and provided access,<sup>176</sup> but it hosted an index of songs its users made available from their home computers.<sup>177</sup>

This language is particularly pertinent to *Grokster*, which was a decentralized system that did not offer an index of songs.<sup>178</sup> Consequently, the court ruled that it did not provide the site and facilities.<sup>179</sup> Instead, *Grokster* advertised their software, which was deemed a sufficient contribution.<sup>180</sup> *Grokster's* reasoning can be traced to a seminal non-digital contributory liability case, *Screen Gems*, in which advertisers, packagers, and distributors were deemed to have materially contributed to infringement.<sup>181</sup> Specifically, *Grokster* induced infringement because when Napster was "under attack in the courts for facilitating mass infringement," it "promot[ed] its software's ability to access popular copyrighted music" through Napster's program.<sup>182</sup> In this and other contexts, advertising and providing a tool for infringement have been seen as inducement and a sufficient material contribution.<sup>183</sup>

---

<sup>168</sup> *Encyclopaedia Britannica Educ. Corp. v. Crooks*, 558 F. Supp. 1247, 1256 (W.D.N.Y. 1983).

<sup>169</sup> *Fonovisa, Inc. v. Cherry Auction, Inc.*, 76 F.3d 259, 264 (9th Cir. 1996) (citing *Columbia Pictures Indus., Inc. v. Aveco, Inc.*, 800 F.2d 59, 62 (3d Cir. 1986)).

<sup>170</sup> *Id.* .

<sup>171</sup> *Id.* at 261 (finding that the sheriff seized over 38,000 counterfeit recordings, and a year later "vendors were still selling counterfeit recordings"). Ultimately, it seems that the only practical solution to attempt to stop all of the mass infringement occurring at the flea markets was to find the owners of the markets liable.

<sup>172</sup> *Id.* at 264 ("There is no question that plaintiff adequately alleged the element of knowledge in this case.").

<sup>173</sup> *Perfect 10, Inc. v. Visa Int'l Serv. Ass'n*, 494 F.3d 788, 800 (9th Cir. 2007).

<sup>174</sup> *Id.* at 796.

<sup>175</sup> *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004, 1022 (9th Cir. 2001), *aff'd*, 284 F.3d 1091 (9th Cir. 2002).

<sup>176</sup> PATRY, *supra* note 62, § 21:48.

<sup>177</sup> *A & M Recs., Inc. v. Napster, Inc.*, 114 F. Supp. 2d 896, 906 (N.D. Cal. 2000), *aff'd* in part, *rev'd* in part *sub nom.* *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004 (9th Cir. 2001), as amended (Apr. 3, 2001), *aff'd* *A&M Recs., Inc. v. Napster, Inc.*, 284 F.3d 1091 (9th Cir. 2002) ("The server-side application software maintains a search index that is updated in real time as users log-on and -off of the system.").

<sup>178</sup> *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 922 (2005).

<sup>179</sup> *Id.*

<sup>180</sup> *Id.* at 937.

<sup>181</sup> *Screen Gems-Columbia Music, Inc. v. Mark-Fi Recs. Inc.*, 256 F. Supp. 399, 403 (S.D.N.Y. 1966).

<sup>182</sup> *Grokster*, 545 U.S. at 937–38.

<sup>183</sup> While the Court did not assess contributory infringement through the conventional 2-pronged analysis, they suggested that the "rule on inducement... [requires] evidence of 'active steps taken to encourage direct

Knowledge can influence a defendant's material contribution. Once the online "Bulletin Board Service" ("BBS") in *Netcom* was aware of infringement on their servers. The court determined that failing to delete this infringing material was considered a material contribution.<sup>184</sup> Therefore, with knowledge of an infringement, contributory liability imposes an affirmative duty to disallow infringing content so long as "simple measures" could reduce infringement.<sup>185</sup> In *Netcom*, the court cabins this affirmative duty within an economic policy lens to not create too onerous a requirement of technology service providers.

## 2. Knowledge of Infringement

Contributory infringement applies an actual or constructive knowledge standard because "it is sufficient that the defendant has reason to know that infringement is taking place."<sup>186</sup> "Actual knowledge" was demonstrated in *Napster*.<sup>187</sup> Courts applied a "should-have-known" standard in *Netcom*, where the court found that "evidence reveals a question of fact as to whether Netcom knew or *should have known* that Erlich had infringed plaintiffs."<sup>188</sup> Ultimately, since copyright is expressly policy-driven and secondary liability is judge-made law, courts have accepted different knowledge standards for contributory liability based on the facts of the case in light of the perceived goals of copyright.

In *Napster*, actual knowledge was found when a co-founder stated that "Napster users 'are exchanging pirated music. . . . We are not just making pirated music available but also pushing demand.'"<sup>189</sup> Similarly, in *Fonovisa*, knowledge was not an issue as thousands of counterfeit recordings had been previously seized by the Sheriff and the Sheriff later sent another letter warning of continued infringement.<sup>190</sup>

*Fleaworld* provides another illustration. The court found that the defendant, a large flea market operator, "knew or should have known" of infringement because they regularly walked through the market, personally visited vendors, checked for prohibited items, the record industry association investigators informed them of infringement, and there had been multiple raids seizing thousands of pirated records.<sup>191</sup>

In *Netcom*, the court also applied a should-have-known standard. When the plaintiff had notified the defendant that the direct infringer had posted a verbatim copy on their online BBS, the Court found it was uncertain whether "Netcom knew or should have known" the

---

infringement,' [which means defendant] 'actively and knowingly aids and abets another's direct infringement.'" *Id.* at 936. So, while *Grokster* speaks of inducing infringement, instead of knowingly and materially contributing to infringement, their rationale and legal basis rests in terms that can be easily categorized within the conventional contributory liability analysis. This is furthered by subsequent courts analyzing *Grokster*'s inducement analysis similarly to conventional contributory liability principles. *See, e.g.,* *Perfect 10 v. Amazon.com*, 508 F.3d 1146, 1171 (9th Cir. 2007) ("[U]nder *Grokster*, an actor may be contributorily liable for intentionally encouraging direct infringement if the actor knowingly takes steps that are substantially certain to result in such direct infringement.").

<sup>184</sup> *Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc.*, 907 F. Supp. 1361, 1375 (N.D. Cal. 1995).

<sup>185</sup> *Id.*

<sup>186</sup> *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 487 (1984).

<sup>187</sup> *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004, 1022 (9th Cir. 2001), *aff'd*, 284 F.3d 1091 (9th Cir. 2002).

<sup>188</sup> *Netcom*, 907 F. Supp. at 1374 (N.D. Cal. 1995) (emphasis added); *see also*, *Fonovisa v. Napster, Inc.*, No. 3:01-CV-02669, 2002 WL 398676, at \*5 (N.D. Cal. Jan. 28, 2002) (applying "the traditional formulation that either constructive or actual knowledge is sufficient to impose [contributory] liability.").

<sup>189</sup> *A & M Recs., Inc. v. Napster, Inc.*, 114 F. Supp. 2d 896, 903 (N.D. Cal. 2000), *aff'd* in part, *rev'd* in part *sub nom.* *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004 (9th Cir. 2001), as amended (Apr. 3, 2001), *aff'd* *A&M Recs., Inc. v. Napster, Inc.*, 284 F.3d 1091 (9th Cir. 2002).

<sup>190</sup> *Fonovisa, Inc. v. Cherry Auction, Inc.*, 76 F.3d 259, 261 (9th Cir. 1996).

<sup>191</sup> *Arista Records, Inc. v. Flea World, Inc.*, 2006 WL 842883 at \*15 (D.N.J. Mar. 31, 2006).

defendant infringed.<sup>192</sup> While mere allegations of infringement are insufficient notice, they are sufficient when the allegedly infringing works “contain copyright notices within them.”<sup>193</sup>

*Aimster* demonstrates the willful blindness standard where a file-sharing company encrypted all communications, which prevented the defendant “from knowing what songs were being copied.”<sup>194</sup> However, “a contributory infringer does not obtain immunity by using encryption to shield itself from actual knowledge of the unlawful purposes for which the service is being used.”<sup>195</sup> By contrast, in *Luvdarts*, a company that provided virtual greeting cards that were intended to be used only once, sued AT&T for providing the service for transferring the reproductions.<sup>196</sup> However, since AT&T was simply indifferent to infringement they were not willfully blind because they did not satisfy the two elements: (1) a subjective belief of infringement and (2) “deliberate actions to avoid learning” of infringement.<sup>197</sup>

This flexible knowledge standard mirrors the common law. *Perfect 10 v. Amazon* explicitly incorporated common law principles as “[t]ort law ordinarily imputes to an actor the intention to cause the natural and probable consequences of his conduct.”<sup>198</sup>

### 3. Substantial Non-Infringing Uses—Imputed Knowledge/Intent Based Upon a Product’s Inherently Infringing Nature

A uniquely relevant doctrine of knowledge, and contributory liability generally, is the substantial noninfringing uses standard set forth in *Sony*, the Betamax case. In *Sony*, the secondary infringer was presumed to have known about the direct infringement because of the inherently infringing design of the defendant’s product.<sup>199</sup> The *Sony* rule is then a theory of defense against an “intent to infringe... [that] is imputed from a defendant’s material contribution” based on the potentially infringing nature of the product.<sup>200</sup>

The Court in *Sony* first incorporated a new mode of liability, and then second provided a defense to this mode of liability.<sup>201</sup> First, the Court incorporated into Copyright a rule of patent infringement, 35 U.S.C. § 271(c), where sellers of *components* of a patented machine that were especially adapted to infringement were contributory infringers.<sup>202</sup> As applied to copyright, those technologies that were especially adapted to facilitate infringement could demonstrate an imputed intent to infringe.<sup>203</sup> Second, the Court included 35 U.S.C. § 271(c)’s exception for staple articles “suitable for substantial noninfringing use[s].”<sup>204</sup> Therefore, “the sale of copying equipment . . . does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes. Indeed, it need merely be *capable* of

<sup>192</sup> *Netcom*, 907 F. Supp. at 1374.

<sup>193</sup> *Id.* This may be particularly relevant for those outputs that include things such as Getty’s watermark.

<sup>194</sup> *In re Aimster Copyright Litig.*, 334 F.3d 643, 650 (7th Cir. 2003).

<sup>195</sup> *Id.* at 651.

<sup>196</sup> *Luvdarts, LLC v. AT & T Mobility, LLC*, 710 F.3d 1068, 1070–71 (9th Cir. 2013).

<sup>197</sup> *Id.* at 1073.

<sup>198</sup> *Perfect 10 v. Amazon*, 508 F.3d at 1171 (quoting *DeVoto v. Pac. Fid. Life Ins. Co.*, 618 F.2d 1340, 1347 (9th Cir. 1980)).

<sup>199</sup> *Grokster*, 545 U.S. at 931.

<sup>200</sup> *Disney Enterprises, Inc. v. Hotfile Corp.*, No. 11-20427-CIV, 2013 WL 6336286, at \*33 (S.D. Fla. Sept. 20, 2013).

<sup>201</sup> *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 435–36 (1984).

<sup>202</sup> 35 U.S.C. § 271.

<sup>203</sup> *Grokster*, 545 U.S. at 934.

<sup>204</sup> 35 U.S.C. § 271.



substantial noninfringing uses."<sup>205</sup> Since VHS recorders were capable of substantial noninfringing uses, knowledge of infringement was not imputed to Sony, so the Court did not disallow the sale of this copying equipment. The word capable includes a temporal component meaning that a technology, as is, could be used in the future for substantial non-infringing uses.<sup>206</sup>

While Sony was not found liable, the court's rationale broadened the scope of secondary liability.<sup>207</sup> The court rejected "respondents' unprecedented attempt to impose copyright liability upon the distributors of copying equipment"<sup>208</sup> by use of 35 U.S.C. § 271, but by deciding the case on 35 U.S.C. § 271(c)'s exception, rather than denying the argument outright, the court judicially incorporated patent law's 35 U.S.C. § 271 into copyright law. Now, knowledge can theoretically be imputed to distributors of inherently infringing technology without substantial noninfringing uses.

Importantly, *Sony* applies to the "manufacturer of the infringing machine, not its operator."<sup>209</sup> Sony relinquished possession of the VHS recorder upon sale and did not maintain control over the technology, so they could not actively reduce the rate of potential infringement or influence the direct infringers. This is because secondary liability is most frequently applied in instances when the infringer can influence "the use of copyrighted works by others. . . . Here, the only contact between petitioners and the users of the VTR's occurred at the moment of sale."<sup>210</sup> This contrasts with the provision of text-to-art generation, where the contact and influence are ongoing. However, the underlying rationale provided by the Sony & RCA courts may still apply to most generative AIs because of their emergent and unpredictable nature, which makes it exceptionally difficult to limit non-verbatim infringements.

While the Court decided the case on other grounds, *Grokster*'s dicta clarified the substantial non-infringing uses doctrine.<sup>211</sup> Importantly, "mere knowledge of infringing potential or of actual infringing uses would not be enough [] to subject a distributor to liability."<sup>212</sup> It is obvious that a defendant could use a copy machine to directly infringe a copyright. However, the secondary defendant knowing that a copy machine could be used to infringe is not enough to impute knowledge on the secondary defendant. Similarly, a secondary defendant knowing of someone that is using a copy machine to directly infringe would not impute knowledge or intent. Instead, "one who distributes a device with the object of promoting its use to infringe copyright, as shown by clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties."<sup>213</sup> Stated otherwise, a demonstrated intent to contribute to infringement effectively counterposes *Sony*'s

---

<sup>205</sup> *Sony*, 464 U.S. at 442.

<sup>206</sup> *Grokster*, 545 U.S. at 953–54.

<sup>207</sup> It is frequently misstated that Sony's rule immunizes products from liability that have substantial non-infringing uses. However, as shown in *Grokster* and elsewhere, an entity can be contributorily liable for the distribution of a product with the intent to facilitate infringement. See *BMG Rts. Mgmt. (US) LLC v. Cox Commc'ns, Inc.*, 881 F.3d 293, 307 (4th Cir. 2018) ("[C]ontrary to Cox's argument, the fact that its technology can be substantially employed for a noninfringing use does not immunize it from liability for contributory copyright infringement.").

<sup>208</sup> *Sony*, 464 U.S. at 421.

<sup>209</sup> *RCA Records v. All-Fast Sys.*, 594 F. Supp. 335, 339 (S.D.N.Y. 1984).

<sup>210</sup> *Sony*, 464 U.S. at 418.

<sup>211</sup> *Grokster*, 545 U.S. at 934–35 ("[N]othing in *Sony* requires courts to ignore evidence of intent if there is such evidence, and the case was never meant to foreclose rules of fault-based liability derived from the common law.").

<sup>212</sup> *Id.* at 937.

<sup>213</sup> *Id.* at 936–37.

rule. Therefore, Sony's defense does not apply when a defendant develops a technology where the intent to facilitate infringement has been substantiated.

Finally, potentially useful to practitioners and plaintiffs is Breyer's concurrence in *Grokster*. While courts have avoided strict statistical definition of substantial non-infringing uses, Breyer found that approximately 10% of VCR recordings "apparently are noninfringing . . . ." <sup>214</sup> Such a volume of potential noninfringing uses then appears sufficient to use as a defense against an imputation of knowledge. However, due to secondary liability's equitable implementation in copyright, courts will likely balance all aspects of a case as bright line rules are eschewed.

## V. FAIR USE

The fair use doctrine has long been recognized as confounding, unclear, and frustrating. <sup>215</sup> Defendants offer fair use as an affirmative defense once a plaintiff proves a prima facie case of infringement. <sup>216</sup> Although a relatively small area of case law, fair use has had an outsized impact on copyright law generally, and specifically within the realm of academia. <sup>217</sup> Given the flexible and convoluted nature of fair use, it is no surprise that much has already been written about the application of the doctrine to generative AIs, both academically <sup>218</sup> as well as in more mainstream outlets. <sup>219</sup>

### A. FACTOR ONE AND TRANSFORMATIVE USE

The first statutory factor considered in the fair use analysis is "the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes." <sup>220</sup> The first factor, along with the fourth, have taken on an outsized role in the fair use analysis, both in theory <sup>221</sup> and in practice. <sup>222</sup> The question of whether the purpose and character of the use is "fair" is rooted in Justice Story's assertion that copyright infringement

---

<sup>214</sup> *Id.* at 952.

<sup>215</sup> See generally Pierre N. Leval, *Toward a Fair Use Standard*, 103 HARV. L. REV. 1105 (1990) (discussing the inconsistency of fair use applications).

<sup>216</sup> *Id.*

<sup>217</sup> Barton Beebe, *An Empirical Study of U.S. Copyright Fair Use Opinions, 1978-2005*, 156 U. PA. L. REV. 549, 565 (2008).

<sup>218</sup> See, e.g., Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743 (2021) (arguing for broad fair use protection when training AI models).

<sup>219</sup> James Vincent, *The Scary Truth About AI Copyright Is Nobody Knows What Will Happen Next*, THE VERGE (Nov. 15, 2022, 7:00 AM), <https://www.theverge.com/23444685/generative-ai-copyright-infringement-legal-fair-use-training-data>.

<sup>220</sup> 17 U.S.C. § 107(1).

<sup>221</sup> *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 579 (1994) ("[T]he more transformative the new work, the less will be the significance of the other factors."); see also *Harper & Row Publishers, Inc. v. Nation Enterprises*, 471 U.S. 539, 566 (1985). *Harper & Row* makes the oft-cited assertion that the fourth factor is "undoubtedly the most important" of the four fair use factors. This statement, however, is conspicuously absent from the Supreme Court's own fourth-factor analysis in *Campbell* as well as in *Google LLC v. Oracle America, Inc.*, 141 S. Ct. 1183 (2021) (a recent fair use decision).

<sup>222</sup> See Beebe, *supra* note 214, at 584. Beebe's examination of fair use decisions found that the outcomes of the first and fourth factors coincided with the overall finding of the fair use analysis 84% and 81% of the time, respectively.

occurs when the use acts to “supersede the use of the original work” rather than add something new to the original.<sup>223</sup>

Over time the factor one analysis has come to revolve around the question of whether a use is transformative. As originally articulated by Judge Pierre Leval in his seminal work on fair use, “the use must be productive and must employ the quoted matter in a different manner or for a different purpose than the original. A quotation of copyrighted material that merely repackages or republishes the original is unlikely to pass the test.”<sup>224</sup> The Supreme Court adopted this standard in the *Campbell* decision when holding that parody is a transformative use.<sup>225</sup> This formulation of a transformative use has come to be the classic definition.

It is unlikely, however, that either Judge Leval or the *Campbell* court anticipated the extent to which the transformative use doctrine would expand in both scope and importance. The transformative use doctrine’s steady expansion has continued to the point where both academics<sup>226</sup> as well as the circuit courts<sup>227</sup> have begun to call for a reconsideration of the role transformative use plays in the fair use analysis. Courts have trended toward ruling even minute changes as transformative, and such a finding of transformative use is increasingly likely to be dispositive for the first factor’s overall analysis. As a byproduct, the exact contours of what types of use may or may not be transformative, while crucial to the overall analysis, are elusive to determine concretely.<sup>228</sup>

An infringing use can be transformative either in content or purpose, and finding either has a significant influence on the first factor in favor of fair use. For example, a use can have a transformative purpose even when copying the original work verbatim if the context is changed. Defendants can make particular use of this prong of the doctrine because whether or not a purpose is transformative is interpreted widely. For example, in *Fox News Network, LLC v. TVEyes, Inc.*, a service that allowed a user to watch 10-minute segments of a news broadcast had the purpose of enhancing efficiency, which was transformative in relation to the original broadcast’s purpose of disseminating news programming.<sup>229</sup> Even though segments of up to 10 minutes were viewable in their original format without alteration, the purpose of isolating these segments from an “ocean of programming” was still transformative enough.<sup>230</sup>

The leniency given to the defendant in *TVEyes* is indicative of a low bar for transformative uses generally, particularly with respect to transformative purposes. Holdings like in *TVEyes* infer that the transformative use requirement has been reduced to labeling a use transformative as long as it does something more than merely repackaging or republishing.<sup>231</sup> While the *TVEyes* court eventually found that the defendant’s use was not fair, the decision relied on a fourth factor analysis where the defendant’s generation of millions of dollars in

---

<sup>223</sup> *Folsom v. Marsh*, 9 F.Cas. 342, 345 (C.C.D. Mass. 1841)(No. 4,901).

<sup>224</sup> Leval, *supra* note 212, at 1111.

<sup>225</sup> *Campbell*, 510 U.S. at 579.

<sup>226</sup> NIMMER & NIMMER, *supra* note 52, § 13.05(B)(6) (“It would seem that the pendulum has swung too far in the direction of recognizing any alteration as transformative”).

<sup>227</sup> *E.g.*, *Kienitz v. Sconnie Nation LLC*, 766 F.3d 756, 758 (7th Cir. 2014). The 7th Circuit expresses concern that relying exclusively on the transformative use doctrine may override the statutory fair use factors.

<sup>228</sup> Clark D. Asay, Arielle Sloan & Dean Sobczak, *Is Transformative Use Eating the World?*, 61 B.C. L. REV. 905, 955 (2020).

<sup>229</sup> *Fox News Network, LLC v. TVEyes, Inc.*, 883 F.3d 169, 178 (2d Cir. 2018) (finding TVEyes’ watch function “somewhat transformative”).

<sup>230</sup> *Id.* at 177.

<sup>231</sup> Rebecca Tushnet, *Content, Purpose, or Both?*, 90 WASH. L. REV. 869, 877 (2015).

revenue after disregarding the plaintiff's refusal to grant a license tipped the overall analysis against fair use.

#### B. FACTOR FOUR: LICENSING, SUBSTITUTION, AND THE PRACTICAL MARKET

The fourth fair use factor requires an evaluation of the effect of the use upon the potential value of or market for the original work.<sup>232</sup> Within this broad analysis, a number of subfactors have developed to ensure that factor four is aligned with the actual market harm suffered by a copyright owner rather than attenuated or hypothetical harms. Among these are the extent to which the infringing use serves as a market substitute for the original work,<sup>233</sup> whether the infringing use harms an existing licensing market for the original,<sup>234</sup> and a balancing of the public benefits the infringing use generates with the harm it causes to the original author.<sup>235</sup> In the following section, we will apply both factors one and four to a fair use analysis of the inclusion of copyrighted works in generative AI datasets.

### VI. APPLICATIONS TO GENERATIVE AI ART

Given the novelty of text-to-art AIs, there is little by way of case law to guide courts and practitioners on its implications for copyright law. This section aims to fill that gap by analyzing three elements that are especially challenging when applying copyright to this new context: direct *prima facie* copyright infringement, secondary liability, and fair use. Regardless of how courts ultimately apply copyright doctrines to generative AI, the intersection is fraught with policy impacts, which this section discusses after the descriptive analysis.

#### A. *PRIMA FACIE* COPYRIGHT INFRINGEMENT

Many generative AI use cases are innocuous. The websites of the most popular developers show uncontroversial prompt examples, like “an astronaut riding a horse in photorealistic style.”<sup>236</sup> But some uses may be considered harmful, as evidenced by recent lawsuits filed by artists whose works were used in training popular generative AI models.<sup>237</sup> This Article focuses on the previously enumerated use cases to explore the potential limits of copyright infringement and fair use, which requires a threshold evaluation of whether these tools even engage in *prima facie* copyright infringement. To accomplish this, we will first

---

<sup>232</sup> 17 U.S.C. § 107(4).

<sup>233</sup> *Campbell*, 510 U.S. at 591 (discussing the relationship between the degree to which a use is transformative and its likelihood of acting as a market substitute for the original work).

<sup>234</sup> *Id.* at 591–92.

<sup>235</sup> *Google LLC v. Oracle America, Inc.*, 141 S. Ct. 1183, 1208 (2021). If Oracle's copyright was enforced, it would hamper the ability of programmers to develop new technologies, which the Court observed ran counter to the creativity-furthering objectives of copyright. Thus, the public interest at issue in this case was the ability of the programmers that had come to rely on the APIs Google had used.

<sup>236</sup> *DALLE-E 2*, OPENAI, <https://openai.com/dall-e-2/> (last visited Feb. 25, 2023).

<sup>237</sup> *See, e.g.*, Complaint at 1, *Andersen et al. v. Stability AI et al.*, No. 3:23CV201 (N.D. Cal. filed Jan. 13, 2023); Complaint at 1, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23CV135 (D. Del. filed Feb. 03, 2023).

evaluate whether each potential defendant—the developer, the database creator, and the user—is engaging in direct *prima facie* copyright infringement.

### 1. *The Model Developer as the Defendant*

Both the training and output stages of generative AI art tools warrant copying analysis.<sup>238</sup> When considering model training, a copying analysis depends on exactly how the training is conducted. As noted previously, it is feasible that the developer never downloads images during training, so there would be no copy made in a fixed medium to constitute *prima facie* copyright infringement for lack of factual copying. But assuming downloading does occur (which it did during Stable Diffusion’s training), a verbatim copy of the original images is recreated on the servers of the model developer, which would constitute legal copying and factual copying, per *Perfect 10*. Still, it is possible that the verbatim copy would reside on the developer’s servers for only a transitory duration—a detail that’s unknowable without inside, intimate knowledge of the training process. Depending on whether the copying is transitory, it’s possible that the act of training models constitutes copying in an infringement analysis.<sup>239</sup> Note that because this stage comes before any output has been generated from the model, it does not differ between use cases.

The Northern District of California recently found that direct infringement allegations against model developers for conduct during the training phase are plausible. In *Andersen v. Stability AI*, Judge Orrick largely granted defendant Stability AI’s motion to dismiss with leave to amend.<sup>240</sup> However, one claim in the initial complaint was not dismissed—that Stability AI directly infringed the plaintiff’s copyrighted work when training Stable Diffusion.<sup>241</sup> The court cautioned that a final determination on this claim could not be reached at the motion-to-dismiss stage, but the facts provided were sufficient to plausibly allege direct infringement by the model developer, Stability AI.<sup>242</sup>

When analyzing infringement upon model output, the analysis is murkier and will likely depend on the use case. The “in the style of” and the entity-centric use cases will have similar analyses. In both cases, it is undisputed that the model developer has access to copyrighted inputs. Consequently, the analysis will hinge on the exact output in dispute, whether it has probative similarity to the input, and whether it amounts to unlawful appropriation. When the output shows no resemblance to any input, there can be no probative similarity, so there would be no factual copying. Even if the plaintiff’s work was undoubtedly included in the model’s training data, the output in question would lack any similarity with the original, so there would be no unlawful appropriation upon which to build an infringement claim. This is likely to be the case even though both use cases seek to capture the essence or style of a single artist or entity. Because style and similar attributes are akin to ideas, not expression, they don’t fall under the protection of copyright. Unless an output image contains protected elements of the original, there is probably no *prima facie* copyright infringement.

The analysis changes if an output *does* resemble the input images. Consider an output that is a verbatim (or near verbatim) copy of an input to the “in the style of” or entity-centric

---

<sup>238</sup> This article assumes the input images are under protectable copyright.

<sup>239</sup> Regardless, it is unlikely that any *prima facie* copyright infringement that occurs during the model’s training would beget liability because it would be excused by fair use. *Infra* Part VI.C.

<sup>240</sup> *Andersen v. Stability AI Ltd.*, No. 23-CV-00201-WHO, 2023 WL 7132064 at \*17 (N.D. Cal. Oct. 30, 2023).

<sup>241</sup> *Id.* at \*5.

<sup>242</sup> *Id.*

use case, which is generated in response to a prompt that elicits a specific image.<sup>243</sup> If the resulting image is indeed part of the training set, it is likely that this output would constitute factual copying because there is direct evidence that the model copied the input image. Even if there is no direct evidence of copying, there is at least probative similarity and sufficient evidence of access, so factual copying is clear.<sup>244</sup> Given the undeniable similarities between the input and output, unlawful appropriation is likely also present under the comprehensive nonliteral similarity doctrine. Furthermore, transitory copying is unlikely to be a bar to the *prima facie* case against the developer, assuming the output remains on its server for longer than a few minutes.<sup>245</sup>

Despite the colorable argument that a developer engages in *prima facie* copyright infringement during the output phase, it may not be the proper liability holder. When the prompt elicits a specific input image, the output can be considered entirely at the behest of the user. That is, without the user specifically prompting the AI with text associated with an input image, the output probably would not resemble such an image. As a result, a *prima facie* copyright case may not stand against the developer in this scenario. Instead, the user may bear the liability.<sup>246</sup>

This analysis changes with a slight factual tweak. Assume that instead of prompting the model with specific text tied to an input, the user enters generic language that isn't closely linked with an input image. For example, suppose the user prompted "a picture of a professional woman" and a specific copyrighted input appeared. In this scenario, the user is no longer directing the model to generate the infringing content. Instead, it seems more accurate to say the model is creating the copy on its own accord. This hypothetical keeps the finding of factual copying and unlawful appropriation intact while raising legitimate questions about which entity is engaging in direct infringement. If a plaintiff can successfully characterize the model as outputting a copy, it's possible that they can succeed on a *prima facie* copyright infringement claim against the developer.<sup>247</sup>

As this example highlights, there is a critical difference between text-to-art AIs and other technologies that enable copying at scale. In previous sections, this Article discussed the photocopier as a potential means for copyright infringement. When someone uses a copy machine to replicate a copyrighted work, they open themselves to liability for infringement, but not the machine's designer or manufacturer. As previously discussed, just because a machine *can* be used for infringing purposes doesn't mean it should be held responsible when a separate party effectuates such infringement. However, a photocopier's output is deterministic—that is, it will only produce a replica of what its user puts in. On the other hand, text-to-art AI outputs are probabilistic. While the models account for the inputs, the results are not often something the user envisioned (if they envisioned anything). This characteristic is unique to generative AIs and makes it possible for the machine to produce infringing content

---

<sup>243</sup> A recent study showed that, given the right prompt, some popular text-to-art models can generate outputs that are nearly identical to input images. NICHOLAS CARLINI ET AL., EXTRACTING TRAINING DATA FROM DIFFUSION MODELS, 2301.13188 ARXIV 6 (Jan. 30, 2023), <https://arxiv.org/pdf/2301.13188.pdf>.

<sup>244</sup> See *supra* Part III.A.

<sup>245</sup> However, it is conceivable that the image would be stored on the developer's server for a longer period. Without specific knowledge of how the developer's system works, we cannot know whether transitory copying will factor into the analysis.

<sup>246</sup> The doctrines of secondary liability and contributory liability could feasibly implicate the developer in this scenario. *Infra* Part VI.B.

<sup>247</sup> Whether this claim would survive the affirmative defense of fair use is a separate question. *Infra* Part VI.C.

without the user ever intending to do so. Thus, basing a theory of liability for text-to-art AI infringement on theories for old technologies may be misguided.

The way diffusion models generate images also creates ambiguity. As outlined, diffusion models respond to prompts by first generating a set of noisy pixels before de-noising them until a recognizable image results. This conflicts with the plain meaning of “copying,” which conjures thoughts of photocopiers that replicate original content in whole cloth. With diffusion models, no “copying” exists as such. Instead, the model essentially rearranges pixels until it determines its output matches the user’s desire with sufficiently high probability. However, the copying process does not factor into the *prima facie* infringement analysis—all that matters is that a factual and legal copy is created. As a result, regardless of their novel functions, diffusion models’ denoising processes are unlikely to negate an otherwise sound *prima facie* copyright claim.

## 2. *The User as the Defendant*

If a lawsuit names a model’s user as the defendant, different considerations arise. For example, if the infringement claim focuses on the model’s training, the user will not factor into the analysis. But, if the claim is made upon a particular output, the user may well be held liable.

For “in the style of” and entity-centric scenarios, the exact output and the originating prompt will determine the analysis. As noted in the previous section, an output that bears no resemblance to any copyrighted input will not satisfy probative or substantial similarity standards, so the plaintiff cannot build a *prima facie* infringement case regardless of who the defendant is. However, as an output approaches striking similarity, the prompt given by the user could influence whether they can be targeted for liability. If a user’s prompt clearly seeks infringing output and the output meets the elements of *prima facie* copyright, they can undoubtedly be held liable even if other parties may be vicariously liable.<sup>248</sup> However, as discussed above, if a generic prompt yields an output that constitutes *prima facie* copyright infringement, the user’s being liable is not certain.

On the other hand, under the right circumstances, the user is likely to be the sole holder of liability for the edit function. The nature of the edit function is that users can upload pictures of their choosing and use diffusion techniques to alter the images as they see fit. The user can upload a copyright protected photo without any action taken by the model developer, so any copying that occurs can only be attributed to the user.<sup>249</sup> Even if the edits demanded by the user significantly change the photo, it could be considered a derivative work that was enabled by direct, verbatim copying by the user upon the photo’s upload. Thus, liability for infringing activity through the copy function is likely to be placed on the user.

## 3. *Database Compiler as the Defendant*

Assuming the database compiler is a separate entity from the developer, they are the least likely holder of liability in any scenario. The database compiler does not directly participate in either training or output, so they are unlikely to be deemed liable for any actions at those stages. However, they might be engaging in *prima facie* copyright infringement even

---

<sup>248</sup> This is analogous to the DVR systems in *Cartoon Network v. CSC Holdings*. There, the DVR users were responsible for copying that happened onto the “Arroyo server” because that copying only happened upon the request of the user. *Cartoon Network LP, LLLP v. CSC Holdings, Inc.*, 536 F.3d 121, 127 (2d Cir. 2008).

<sup>249</sup> See *supra* note 128, and accompanying text.

before the training begins. For example, LAION's 5B database is essentially an index of links to images and image-text pairs. Importantly, LAION does not host the images themselves—they only provide locations to images at other locations on the internet. However, to extract the image-text embeddings, LAION temporarily downloads the images before deleting them shortly thereafter. So, *prima facie* copyright infringement may occur when the database compiler creates the index, but this depends on how long LAION keeps the downloaded images and if a court would consider the storage period transient. Overall, however, database compilers have limited exposure to liability in the landscape of text-to-art AI.

## B. SECONDARY LIABILITY

Beyond direct infringement claims, text-to-art AI will test secondary liability theories in several ways. There are hurdles that each lawsuit will likely have to overcome, and key aspects of the law that courts will consider.

### 1. *Threshold Questions of Causation*

Secondary copyright liability requires an associated direct infringement. Therefore, secondary liability theories in text-to-art AI outputs require the user to have directly infringed on a copyright when prompting an output that qualifies as *prima facie* infringement.<sup>250</sup> However, this is not a decided outcome. It is uncertain which actions are the direct versus indirect cause of an infringing output when expressive AI is involved. To explore this issue, consider the following user prompts (1) "Mickey Mouse," (2) "a cartoon mouse in the style of Walt Disney," or (3) "a cartoon mouse." Each scenario implicates unique issues of reasonable foreseeability<sup>251</sup> and basic notions of fairness,<sup>252</sup> and affect which entity is best positioned to avoid harm.

A unique feature of generative AI is that it complicates the previously implied causation analysis of authorship<sup>253</sup> and infringement for outputs. This uncertainty is demonstrated by suggestions that authorship should alternately be attributed to the end-user-prompter, the AI developer, finding outputs joint works, finding outputs works for hire, or even to the AI itself.<sup>254</sup> Clearly, to what the output "owes its origin"<sup>255</sup> is up for debate. In this vein, the Copyright Office recently decided the case of Kashtanova's comic book, in which the visual

---

<sup>250</sup> It is cognizable that in open-sourcing these LLMs and Diffusion Models, a court could find direct infringement by the subsequent implementer, at which point, the original trainer of the model could be vicariously liable.

<sup>251</sup> Boyden suggests that authorial causation should be partially based on whether an author could reasonably foresee the protectable elements of an output. Bruce E. Boyden, *Emergent Works*, 39 COLUM. J.L. & ARTS 377, 383 (2016).

<sup>252</sup> ANTONY M. HONORÉ, RESPONSIBILITY AND FAULT 14–15 (1st ed. 2002) ("By allocating credit for the good outcomes of actions and discredit for bad ones, society imposes outcome responsibility; though often the rewards it attaches and, outside the law, the sanctions it imposes are informal and vague. Under a system of outcome responsibility, we are forced, if we want to keep our social account in balance, to make what amounts to a series of bets on our choices and their outcomes. Provided we have a minimum capacity for choosing and acting, we win the bets and get credit for good outcomes more than we lose them and incur discredit for bad ones. We have to take the risk of harmful outcomes that may be sheer bad luck and not our fault; but that does not make the system unfair to people who are likely to be winners overall.").

<sup>253</sup> Boyden, *supra* note 248, at 383.

<sup>254</sup> See generally, AVIV GAON, THE FUTURE OF COPYRIGHT IN THE AGE OF AI (2021).

<sup>255</sup> Feist Publications, Inc. v. Rural Tel. Serv. Co., 499 U.S. 340, 346 (1991) (quoting Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53, 58 (1884)).



art that was largely AI-generated was ultimately uncopyrightable.<sup>256</sup> Here, the Office stated that "it was Midjourney—not Kashtanova—that originated the ‘traditional elements of authorship.’”<sup>257</sup> The implication is that the output owed its origin to Midjourney, which was behind the creative aspects of the work. According to the Copyright Office, allegations of direct infringement against an AI art company might be a more appropriate pleading. However, this oversimplifies intertwined causation for AI-outputs. Instead, a spectrum of causation is more accurate because "[t]he circumstances under which computer-generated works are prepared seem too varied to permit a single solution."<sup>258</sup> Therefore, courts determining whether to sue an AI art company for direct or secondary infringement may require an initial threshold causation inquiry.

As such, Andersen et al.’s recent lawsuit against Stability AI claims that Stability is both directly<sup>259</sup> and vicariously liable for outputs.<sup>260</sup> The complaint's vicarious liability claim must exist under the premise that the end-user is directly liable for an infringing output. Therefore, the complaint demonstrates that there are colorable arguments for end-user direct infringements upon which claims of secondary liability may be based.

## 2. Vicarious Copyright Infringement

Vicarious liability is a theoretically colorable avenue for litigation against the model developer or the model implementer, depending on the facts of the case. Vicarious liability requires an ability to supervise infringing conduct and potential financial benefits.

The model implementer and the model developer will have varied commercial interests and various degrees of control over the potentially infringing output based on their specific actions. Given the effectively limitless number of scenarios for development and implementation, a court will determine a given entity’s ability to supervise or control the outputs on a case-by-case basis.

Model implementers conventionally charge a fee to end-users, and the financial benefit prong of vicarious liability is likely satisfied in most instances. If an implementer decided to provide their model for free, then the financial benefit prong may be materially attenuated depending on the scope of infringement and the degree to which the implementer was able to control or supervise outputs. For example, *Napster* demonstrates a more attenuated financial benefit analysis, where potential future income was sufficient. For a free-to-use generative AI, this may only be relevant in circumstances where a model is found to generate massive numbers of infringing outputs, like how *Napster* facilitated mass verbatim piracy. Entity-centric models embody this idea. For example, if the text-to-Pokémon AI generated massive numbers of verbatim copies of Pokémon, then the financial benefit analysis may be substantially relaxed.

---

<sup>256</sup> U.S. COPYRIGHT OFFICE, ZARYA OF THE DAWN, LIBR. OF CONG. (2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf> (last visited Mar 15, 2023).

<sup>257</sup> *Id.*

<sup>258</sup> Boyden, *supra* note 238, at 383.

<sup>259</sup> Complaint at 31, Andersen et al. v. Stability AI et al., No. 3:23-CV-00201-WHO, 2023, WL7132064 (N.D. Cal. Oct. 30, 2023) ("Defendants directly infringed Plaintiffs' and the Class's rights because they have: prepared Derivative Works based upon one or more of the Works in violation of 17 USC 106(2). . . . Defendants are using copies of the training images interconnected with their AI Image products to generate digital images and other output that are derived exclusively from the Training Images, and that add nothing new.").

<sup>260</sup> *Id.* at 32. ("Individuals have used AI Image Products to create works using the names of Plaintiffs and the Class in prompts and passed those works off as original works by the artist whose name was used in the prompt. Such individuals are referred to herein as 'Imposters[.]' By using a particular artist's name, Imposters can cause the AI Image Product to rely more heavily on that artist's prior works to create images that can pass as original works by that artist. These output images are referred to herein as 'Fakes.' . . . The Defendant-owner of the AI Image Product used to create each Fake is vicariously liable for any infringements committed by Imposters.").

Several expressive AIs have been open-sourced, meaning that implementers can freely finetune, alter, and implement models. In such circumstances, the financial benefit to the original foundation model developer will be more attenuated, but still possible. Given the courts' willingness to accept a fairly limited financial benefit if other facts of the case present a compelling equitable reason to find infringement, this should not be seen as an insurmountable barrier to a plaintiff suing the developer rather than the implementer. Moreover, since open-sourcing provides open access to a tool that could facilitate more widespread infringement, the creator of an open-sourced foundation model may have a more relaxed financial benefit prong depending on the circumstances.

The core debate over vicarious liability for model implementers or developers will likely be whether they have the right and ability to supervise infringing conduct. This follows pre-*Shapiro* case law where liability was largely distinguished by the degree of control the defendant had over the premises of infringement. The primary factors that would inform a court of the defendant's level of control are: (1) the scale of facilitated infringement; (2) that AI implementers have the ability to terminate end-user accounts; (3) that model implementers have demonstrated an ability to control general types of outputs; (4) that model developers have not demonstrated an ability to control specific outputs; and (5) that potential infringements are most often not verbatim, which means that plaintiffs ought to be wary of policy arguments suggesting that it is onerous and anti-expressive to require developers to police infringement for non-verbatim outputs.

In many ways the implementer has the right and ability to supervise outputs. The Copyright Office has implied that outputs 'owe their origin' to the AI, which appears to demonstrate an exceptionally unusual and great degree of control over the outputs. Moreover, the potential infringement occurs on the implementer's servers, within their premises, and the AI company likely has the ability to terminate user accounts. The court in *Napster* stated that the company had sufficient supervision because they could terminate user accounts and they hosted a centralized index of pirated media. This was found even though the direct infringement occurred on user servers. Text-to-art AI is unlike *Napster* primarily because of *Napster*'s massive scale of verbatim piracy. In *Napster*, the court found that the index of hosted songs could be traced to millions of verbatim pirated copies. With generative AI art, there is uncertainty over whether a given output is substantially similar to a copyrighted original, or even whether a limiting doctrine should apply, such as *scènes à faire*, *de minimis* copying, fair use, merger, and idea versus expression. While expressive AI generation occurs directly on the implementer's premises (i.e., its server), and defendants can terminate user accounts, uncertainty over direct infringements fundamentally alters vicarious liability analysis.

There have been substantial successes regarding text-to-art AIs controlling certain types of outputs. For example, users of DreamStudio are unable to prompt the AI to create pornographic images. When an end-user prompts terms such as 'porn' or 'sex,' then Stability AI's DreamStudio will not generate imagery. When an end-user attempts to prompt potentially violent or harmful content, such as 'Joe Biden committing a war crime' or 'killing,' DreamStudio will generate imagery, but the output does not respond directly to the prompt and is effectively harmless. These are ways model implementers and developers have shown their ability to control general categories of outputs.

However, this level of developer or implementer control over outputs is not specific enough to ensure generations avoid infringement. Developers do not know what will be generated prior to generation. This is demonstrated by the countless examples of users jailbreaking LLMs, prompting around guidelines, and AI 'hallucinations.'<sup>261</sup> These

---

<sup>261</sup> Karen Weise & Cade Metz, *When A.I. Chatbots Hallucinate*, N.Y. TIMES (May 9, 2023), <https://www.nytimes.com/2023/05/01/business/ai-chatbots-hallucination.html>.

hallucinations result in the AI outputting content that may run counter to the rules developers have attempted to implement. Developers may attempt to ensure that their outputs do not misinform, manipulate users, output-sensitive information, or output graphic imagery. These attempts to control the outputs are inconsistently effective largely because the developers themselves do not know how or why expressive AIs generate the content that they do. This current degree of uncertainty does not embody the ability to supervise or control outputs conventionally required under vicarious infringement. The uncertainty in controlling outputs is multiplied by the uncertainty of substantial similarity analysis. Therefore, most expressive AI developers or implementers do not currently have the ability to forbid specific outputs that would be necessary to eliminate non-verbatim copies and incorporate copyright's limiting doctrines.<sup>262</sup> However, depending on the specific details surrounding the operation of an expressive AI or ongoing technological development, there may be a greater likelihood that expressive AI companies would have the ability to control outputs.

The degree to which an implementer or developer is able to control outputs is based on their technical skill.<sup>263</sup> This leads to a paradox of incentives that secondary liability should not promote. The more developers are able to understand and control the expressive AI they develop, the more they demonstrate the ability to control potential infringements, and the higher their likelihood of being held vicariously liable for outputs. Therefore, if a court were to rule against a finding of vicarious liability exclusively on the basis that a company lacks the ability to control outputs, then expressive AI companies would be incentivized to ensure they are unable to sufficiently understand and control outputs. This is particularly undesirable because, while copyright distributes an entitlement for original expression that ought not be breached, there are substantially greater near-term harms that expressive AI poses. Copyright should not disincentivize research about control over generative AI by finding no vicarious liability exclusively because the developer does not have the ability to control what the AI generates. Instead, a court should be wary of finding conventional expressive AI not vicariously liable for non-verbatim generations, both because substantial similarity analysis is non-deterministic and for public policy reasons.<sup>264</sup>

So far, we have discussed the hurdles and advantages encountered by plaintiffs suing expressive AI-companies. Ultimately, unique edge cases will push the bounds of conventional vicarious liability standards. The most likely edge case is a model fine-tuned to expressly emulate the work of a given artist or company. These edge cases still face the hurdles already mentioned, most notably in the form of secondary liability for non-verbatim infringements. Given the sheer volume of expressive AI models presently implemented, it is not unlikely that such an AI could provide facts sufficient to merit a vicarious liability claim.

### 3. *Contributory Liability*

Contributory liability requires an underlying direct infringement, knowledge of the infringement on the part of the defendant, and material contribution to that infringement by the defendant. Since the material contribution prong of contributory liability is less contained than the ability-to-control prong of vicarious liability, there is a higher likelihood of unpredictable legal outcomes for claims of contributory liability. However, any contributory liability claims will have to overcome the same structural hurdles that a vicarious liability claim might face,

---

<sup>262</sup> There is a likelihood that they have the ability to disallow generation of verbatim copies.

<sup>263</sup> Secondary liability has never explicitly incorporated a reasonable copier standard, though these circumstances may merit such a discussion.

<sup>264</sup> It is also relevant that courts can both find vicarious liability and subsequently take a functionalist approach finding an entity has designed their system to effectuate the harmful result, but through otherwise descriptively legal means. *See* Am. Broad. Companies, Inc. v. Aereo, Inc., 573 U.S. 431, 451 (2014).

namely that most outputs are non-verbatim copies. Like vicarious liability, since the roles and actions of potential defendants are highly variable, it is more appropriate to analyze the factors a court should consider in a contributory liability analysis.

### **a. Material Contribution**

Expressive AI implementers generate the content based on an end-user prompt. Expressive AI is a necessary and fundamental component in the generation of potentially infringing outputs, so it may seem obvious that the material contribution prong would be satisfied. This is far from settled law, as no court has decided what entity should be found as the cause for the generation. A camera or photocopier are often necessary and fundamental tools that facilitate infringement, but camera manufacturers are not held contributorily liable. However, a copier is a ‘dumb’ tool. Expressive AI is distinct from a simple tool because our mental model for volition and causation is challenged as the technology bridges the uncanny valley towards something recognizably agentic and human-like. This mental model for causation will be further challenged as developers integrate multiple expressive AI systems, create recursive multi-modal expressive AIs that build on their own diffusion generations through LLM prompting, and create systems that can collectively generate their own synthetic data. At some point, human input or prompting may become so attenuated from the ultimate generation that courts will be forced to cede direct causation to the AI or find an entity to be the constructive cause of the generation. The rapid pace of development makes it impossible to definitively determine what mental model a court should apply based upon highly variable fact-patterns. In other words, the technology may seem most like a camera today, whereas in five years the facts of generation may look different. For this reason, there is a strong argument for a prospective functionalist approach that incorporates policy considerations, focuses on substance over form, and clarifies the law for technologists and plaintiffs.

Currently, commentators have analogized expressive AI to tools such as a camera, to works for hire, to joint works, and to conventional ISPs like bulletin boards, even suggesting that the AI deserves authorship rights itself. Such a view attributes causation to the AI rather than the prompter.<sup>265</sup> Each of these metaphors are imperfect and have potentially undesirable outcomes. Unpacking the results of each of these attributions is presently unnecessary, but they do suggest that the law and its commentators do not know how best to categorize, or even describe, what expressive AI materially contributes to generated outputs. Suffice it to say, expressive AI may not meet the bar for material contribution if it is analogized to a camera, but it may meet the bar if it is analogized to an ISP bulletin board. In flea markets, providing the site and facilities was sufficient; in file sharing, advertising or hosting a central index of songs was sufficient; and in online bulletin boards, hosting the content, being made aware of infringement, and permitting the infringement to remain on its system was sufficient material contribution if there were *simple* measures to prevent further harm. Unique to expressive AI is that not only is the potentially infringing work hosted on the defendant’s servers, but their software does a substantial portion of the generation. Therefore, under most circumstances, it is likely that a developer or implementer’s actions would satisfy the material contribution prong of contributory liability.

However, it is not unlikely that a fact-pattern would arise where an implementer does little more than take another’s open-source foundation model, create a user-interface, and provide access to another’s model. In such circumstances, then the best argument for

---

<sup>265</sup> Under such theory, the AI could more likely be held directly liable for infringements, though this is also unproven in the courts. However, implicit in copyright is a causal assessment for authorship that is distinct from causation in copying. *See Generally*, Shyamkrishna Balganesh, *Causing Copyright*, 117 COLUM. L. REV. 1, 72 (2017).

contribution may be providing or advertising the tool, in which case, contribution may be more speculative.<sup>266</sup> In *Screen Gems*, the court found advertisement and distribution to be a sufficient material contribution.<sup>267</sup> Therefore, depending on the overall context, such an AI art implementer may have sufficiently contributed to infringement.

The primary legal argument against finding material contribution is that multiple cases have cabined material contribution in a digital context to require simple<sup>268</sup> or reasonable<sup>269</sup> solutions to infringement.<sup>270</sup> Therefore, both the overarching policy context and the material contribution standard may be determined by whether there is a simple, reasonable, or feasible means of identifying and remedying potential infringements. We are then brought back to verbatim versus non-verbatim underlying direct infringement and the hypothetical technical solution of output filters, which are highly problematic for copyright and free speech.

### **b. Knowledge**

A court's finding of knowledge will depend heavily on the facts of a given case. Key factors to consider include the conversations among defendant representatives, the structure of a given technology and the statistical likelihood of an infringing generation, communications with authorities or plaintiffs, and a defendant's advertisements. Particularly relevant is that since outputs are not frequently verbatim generations and are unpredictably emergent, a defendant may not have actual or constructive knowledge of specific underlying direct infringements.

Courts have found sufficient actual and constructive knowledge and have demonstrated a willingness to impute knowledge based on an inherently infringing technology. In most circumstances, knowledge in the context of expressive AI will follow conventional common law tort liability standards and factual analysis. It is important to note that the required level of knowledge can depend on other factors in the analysis, including the degree of contribution, the scale of infringement, and how the decision implicates the goals of copyright.

### **c. Substantial Non-Infringing Uses**

Given the rate at which new model developments and implementations are being developed, it is possible that a novel use case will challenge Sony's substantial non-infringing uses doctrine. It is arguable that certain implementations, such as the text-to-Pokémon generator or those models fine-tuned to generate work emulating a single artist, already implicate this doctrine. However, current implementations we have encountered do not appear to massively copy copyrightable elements. Implicating secondary liability in these current implementations would incorporate aspects of trademark law into copyright, and would likely protect styles and ideas that explicitly exist outside of copyright. Therefore, these implementations of expressive AI models do not appear to challenge Sony's substantial non-infringing uses doctrine.

However, a text-to-Mickey generator that nearly always outputs generations substantially similar to Mickey Mouse would likely challenge Sony's doctrine. A uniquely

---

<sup>266</sup> There is a strong policy argument that suggests that implementers should be similarly liable for harms from providing access to a system. This policy analysis is like supply chain liability in products liability.

<sup>267</sup> *Screen Gems-Columbia Music, Inc. v. Mark-Fi Recs., Inc.*, 256 F. Supp. 399, 403–04 (S.D.N.Y. 1966).

<sup>268</sup> *Religious Tech. Ctr. v. Netcom On-Line Commc'n Servs., Inc.*, 907 F. Supp. 1361, 1375 (N.D. Cal. 1995).

<sup>269</sup> *Perfect 10, Inc. v. Google, Inc.*, No. CV 04-9484 AHM SHX, 2010 WL 9479060, at \*4 (C.D. Cal. July 30, 2010), *aff'd*, 653 F.3d 976 (9th Cir. 2011).

<sup>270</sup> *VHT, Inc. v. Zillow Grp., Inc.*, 918 F.3d 723, 745 (9th Cir. 2019) (“Thus, Zillow did not have appropriately ‘specific’ information necessary to take ‘simple measures’ to remedy the violation.”).

relevant distinguishing factor that would favor plaintiffs against expressive AI is their continued connection to the potentially infringing product. In *Sony*, the defendant sold end users a physical product, at which point the defendant had no continuing engagement with the product or the direct infringer. Here, defendant SAAS expressive AI companies have continuous contact with the direct infringer and implement the infringing service. While finding aspects of copyright in characters, such as style, challenges copyright's idea-expression distinction, specifically expressed elements of Disney's manifestations of Mickey Mouse are protected. Therefore, there is a distinct possibility that Sony's substantial non-infringing uses doctrine could be challenged by an expressive AI. In other words, an expressive AI could be found to be so inherently infringing that knowledge of contributory infringement could be imputed to the AI implementer.

In such a use-case, the social utility of the expressive AI would be decreased, so a court's ruling would not as strongly implicate the balancing of copyright's private versus public interests. A primary rationale for secondary liability is that suing many direct infringers may be impractical or too costly. When a product is exclusively infringing, this economic rationale for applying secondary liability is strong. However, when there are substantial noninfringing uses, then this economic rationale is weakened. For example, computer manufacturers know, or should know, that computers are frequently used to infringe copyright, and that suing computer manufacturers would save in enforcement costs. However, computers provide substantial noninfringing social benefits, and the economic rationale for secondary liability is muted. So, a court is unlikely to find computers inherently infringing. Conventional expressive AI has substantial noninfringing uses, is not normally inherently infringing, and has potential social benefits a court will need to balance. In contrast, an expressive AI that largely generates infringing versions of Mickey Mouse provides limited social benefit, so a court would not have to consider that factor. Therefore, such an edge use-case could reasonably lead a court to hold specific expressive AI technology to be inherently infringing and impute knowledge of infringement to the developer or implementer.

### C. FAIR USE

In this section, we will survey the existing scholarship discussing fair use in the AI context, focusing on the extent to which that scholarship analyzes AIs with non-expressive purposes. Next, we will examine the differences between the generative AIs considered by our use cases and those at issue in past scholarly works. Finally, we will discuss the divergence of the fair use analysis between expressive and non-expressive AIs with respect to the first and fourth factors, with a particular focus on whether generative AIs represent a transformative use and the extent to which they serve as substitutes for the works and authors that make up their training data sets.

#### 1. *Non-Expressive Algorithmic Fair Use*

The ascendance of AI in the past year has been nearly inescapable. Earlier iterations of AI technology were relatively limited tools that could perform simple tasks, such as crawling a written submission and checking it for plagiarism.<sup>271</sup> Since early 2021, however, the steady development and release of AI tools has turned into a veritable flood of new applications, ranging from human-like chatbots to the text-to-image generators that are the focus of this Article. As a result, the disruptive potential of AI tools has become a trending topic in subject areas ranging from philosophy to venture capital to national security.

---

<sup>271</sup> See *A.V. ex rel. Vanderhye v. iParadigms, LLC*, 562 F.3d 630 (4th Cir. 2009).

As these AI tools become more powerful and ubiquitous, the fact that their training data contains huge amounts of copyrighted work has become a pressing issue. Within a fair use analysis, the ability of generative AIs to create new content based on the works contained in their training data sets serves as the key factor that differentiates them from older tools that had merely functional or analytical capabilities. As we will explain, this distinction jeopardizes existing arguments which posit that the inclusion of copyrighted works in AI training data sets is (or should be) a fair use.

Much of the existing literature analyzing AI tools has trended toward a consensus that the inclusion of copyrighted material in AI datasets should be considered a fair use.<sup>272</sup> These arguments are multifaceted, but at the core of the argument in favor of fair use is the observation that many traditional AI tools use their copyrighted training data for a purpose unrelated to the expressive purpose of the authors of the original works.<sup>273</sup> According to this reasoning, AIs do not copy the works in their training data sets for the expressive aspects of those works, but instead copy merely to learn something about the way that the work is assembled. A facial recognition AI, for example, analyzes photographs of faces in order to map out the facial geometry contained in the photos – not for the purpose of appreciating the creative arrangement and aesthetics of the portrait.

It is important that AIs have historically used copyrighted works to learn something non-expressive about those works because the non-expressive aspects of a work are not part of the rights conferred to the original author by copyright law.<sup>274</sup> In the portrait example described above, the artistic expression of the photographer or the styling of the subject may be copyrighted, but the facial structure of the portrait’s model falls outside of the traditional scope of copyright. If an AI only has the purpose of training on the latter, then its use is more likely to be fair.

In practice, this means that traditional AI tools are using artistic and literary works as data to be learned from rather than copying their expressive components. Traditional AIs have no interest in a work’s expression—the style of a visual artist, the prose of an author, and the performative quirks of a musician—they simply are interested in identifying the objects in an artwork, the informational content of a novel, or the identity of a song’s performer.

*Sega Enterprises, Ltd. v. Accolade, Inc.*<sup>275</sup> is an oft-cited case example of the application of the fair use doctrine to the use of a copyrighted work for its non-expressive components. Accolade, a videogame developer, wanted to develop games for the Genesis, a gaming console owned by Sega, but sought to avoid becoming a Sega licensee. To circumvent the need to obtain an official developer’s license from Sega, Accolade purchased a Genesis console and used it to reverse engineer segments of source code required to produce a game compatible with the console.<sup>276</sup> Then, using the functional components of the source code, Accolade created wholly original games that worked on the Genesis.

The Ninth Circuit found that Accolade only copied the parts of Sega’s code that provided the functional requirements required to operate a game on Sega’s consoles. These functional requirements are the parts of Sega’s programs that were not protected by

---

<sup>272</sup> See, e.g., Lemley & Casey, *supra* note 215; See also Amanda Levendowski, *How Copyright Law Can Fix Artificial Intelligence’s Implicit Bias Problem*, 93 WASH. L. REV. 579 (2018).

<sup>273</sup> Lemley & Casey, *supra* note 215, at 64.

<sup>274</sup> 17 U.S.C. § 102(b); see also *supra* Part III.D.

<sup>275</sup> 977 F.2d 1510 (9th Cir. 1992).

<sup>276</sup> *Id.* at 1514–15.

copyright.<sup>277</sup> In particular, the court noted the absence of “evidence in the record that Accolade sought to avoid performing its own creative work.”<sup>278</sup> In other words, Sega’s copyright protected the creative expression of the games it had developed, but not the functional code that was required to allow those games to run in the first place. This observation was core to the finding that Accolade’s use of Sega’s code was transformative. This process – referred to as intermediate copying – is analogous to the way that AIs typically learn from data, and to the way that many machines process data in general.<sup>279</sup> This observation was core to the Ninth Circuit’s finding that Accolade’s use of Sega’s code was transformative.

Beyond a misalignment in the protected aspects of an original work and those that are used by an AI, a number of policy considerations have also been cited in support of considering use of copyrighted works in AI training a fair use. Among the most compelling is the need for high-quality training data in order to avoid bias against underrepresented or underprivileged populations, which might arise from having incomplete or otherwise insufficient data.<sup>280</sup> While this idea can be summed up succinctly with the well-known programming adage “garbage in, garbage out,”<sup>281</sup> in practice, the effects of incomplete training data for AI can have significant consequences. Levendowski cites the proliferation of AI into use cases involving the courts, cars, and banks – all cases in which AIs that discriminate on the basis of race or gender can cause significant harm to members of disadvantaged groups.<sup>282</sup> As facial recognition AIs are used by law enforcement with increasing frequency,<sup>283</sup> the need to remedy existing AIs’ troubling lack of accuracy when identifying faces belonging to those of minority communities is of particular concern.<sup>284</sup>

Another argument in support of allowing AI models to train on copyrighted data is the need to facilitate technological innovation in a competitive discipline. AIs are not the first technology to pose a novel copyright law question, and many past issues have been resolved in favor of fair use. The most notable of these is *Sony Corp. of America v. Universal City Studios Inc.*<sup>285</sup> Sony dealt with VCR recording technology which, at the time, was as novel as AIs are today. The Supreme Court resolved the case firmly in favor of fair use, noting that “when technological change has rendered its literal terms ambiguous, the Copyright Act must be construed in light” of its basic purpose to stimulate the creation of creative works for the public good.<sup>286</sup> The finding that time-shifting TV programs was a fair use allowed a vibrant market for new VCR recordings and other home video to develop in the following years. This technological advance would not have been possible without a finding of fair use facilitating the VCR’s development.<sup>287</sup>

---

<sup>277</sup> *Id.* at 1522 (citing 17 U.S.C. § 102(b)).

<sup>278</sup> *Id.*

<sup>279</sup> See James Grimmelmann, *Copyright for Literate Robots*, 101 IOWA L. REV. 657 (2016).

<sup>280</sup> See Levendowski, *supra* note 269.

<sup>281</sup> Margaret Rouse, *Garbage In, Garbage Out*, TECHOPEDIA (Jan. 4, 2017), <https://www.techopedia.com/definition/3801/garbage-in-garbage-out-gigo>.

<sup>282</sup> Levendowski, *supra* note 26, at 583–85.

<sup>283</sup> Clare Garvie, Alvaro M. Bedoya & Jonathan Frankle, *The Perpetual Line-Up, Unregulated Police Face Recognition in America*, GEO. L. CTR. ON PRIV. & TECH. (Oct. 18, 2016), <https://www.perpetuallineup.org/>.

<sup>284</sup> Natasha Singer & Cade Metz, *Many Facial Recognition Systems are Biased, Says U.S. Study*, N.Y. TIMES (Dec. 19, 2019), <https://www.nytimes.com/2019/12/19/technology/facial-recognition-bias.html>.

<sup>285</sup> This case is colloquially known as the Betamax case. 464 U.S. 417 (1984).

<sup>286</sup> *Id.* at 433.

<sup>287</sup> Edward Lee, *Technological Fair Use*, 83 S. CAL. L. REV. 797, 799 (2010).



The development of AI is also not the first time in which the practical implementation of a novel technology relies on its use of copyrighted works being a fair use. Not only is the fair use question a legal issue for the developers of these new technologies, a holding of fair use, or lack thereof, can determine whether a new technological product is economically viable at all.<sup>288</sup> One has to look no farther than Napster, whose business model proved completely impractical following the failure of its fair use defense.<sup>289</sup> On the flipside, Google's successful reliance on fair use when developing its Android operating system was a massively profitable decision that facilitated the development of a core part of its business.<sup>290</sup> Unlike Google, however, many smaller AI developers cannot afford to "bet the company"<sup>291</sup> on fair use, meaning that clear and definitive fair use jurisprudence is essential to facilitating the ability of AI developers to innovate. Of course, any such clear and predictable articulation of fair use has proven elusive.

All of this to say that there exists a variety of convincing reasons that the use of copyrighted data for AI training should definitively be considered fair use in functional use cases. Indeed, this paper does not seek to argue that the use of copyrighted data in AI training datasets should not be a fair use in the general case. On the contrary, we merely seek to define the outer limits of the arguments outlined in papers such as *Fair Learning* and others. In that paper, Lemley & Casey note that the use of expressive data to create an expressive output can complicate the fair use analysis.<sup>292</sup> We now seek to explore that complication more thoroughly, beginning with outlining how generative AIs differ from those discussed in prior works.

## 2. Differentiating Generative AI

Generative AIs complicate the analysis applicable to traditional AIs because they *are* designed with the purpose of learning and mimicking the copyrightable aspects of the works in their training datasets. That is, generative AIs copy the creative aspects of the works of artists. Until recently, the majority of AIs have accepted an input, such as a face or piece of writing, analyzed the factual content of that input, and produced as an output some piece of information about that input, such as the name of a person or the likelihood of that piece of writing being plagiarized. A generative AI, however, accepts a textual prompt as input and generates a new creative output based on what it's learned from the creative content of its training data.

Generative AIs operate by gleaning information about the expression present in the works in their training data in order to generate and recognize those expressions in their output. Stated otherwise, generative AIs have a different purpose in their use of training data than that occurring in more functional AIs. This creates a difficulty: while generative AIs extract non-copyrightable information about the creative works in their training data, that same information is about copyrightable artistic expression.

---

<sup>288</sup> *Id.* at 808.

<sup>289</sup> *A&M Records, Inc. v. Napster, Inc.*, 239 F.3d 1004 (9th Cir. 2001). Following this decision, Napster ceased their operations and later filed for bankruptcy in June of 2002 as the result of a string of lost lawsuits.

<sup>290</sup> *Google, LLC v. Oracle America, Inc.*, 141 S. Ct. 1183(2021). When developing Android, Google made extensive use of APIs owned by Sun Microsystems, which was later purchased by Oracle.

<sup>291</sup> Lemley & Casey, *supra* note 215, at 769.

<sup>292</sup> Lemley & Casey, *supra* note 215, at 777.

This distinction between generative and functional AI tools can be described using the framework of expressive vs. non-expressive use of copyrighted materials.<sup>293</sup> Matthew Sag defines an expressive use as one in which a work is copied with the intention of allowing a human to appreciate those expressive qualities, which can be as simple as downloading a movie with the purpose of watching it for entertainment.<sup>294</sup> A non-expressive use, on the other hand, occurs when a work is reproduced in a way that is not intended to allow humans to enjoy or appreciate the expressive aspects of the original.<sup>295</sup> If the purpose of downloading the movie was instead to analyze the dialogue to ascertain information about the characters' speech patterns, that use would be non-expressive.

Within this framework, the way generative AIs learn from copyrighted material falls squarely into the expressive category. Generative AIs use their training data to learn to mimic the expression of the copyrighted works they train on, meaning that their outputs allow human users to enjoy or appreciate the same creative content as an original artist's unique work. AIs that perform functional tasks, on the other hand, fall into the non-expressive category. The output of a self-driving car's AI, for example, only yields information that helps the car navigate safely and avoid collisions. In this case the AI's output is the performance of a functional task, which does not contain any of the expressive content unique to the original photographs that would be used for the AI to train on – a user cannot enjoy the creative framing of a photograph if the way the AI makes use of that photograph is the way a self-driving car functions when faced with the task of avoiding pedestrians while navigating an intersection.

Differences in the way expressive vs. non-expressive AIs use training data are observable in their outputs. On the level of individual copyrighted works, rather than retaining the creative expression of the original works, the AI uses their data as “raw material”<sup>296</sup> or “grist for the mill.”<sup>297</sup> As noted in the example of the self-driving car AI, the expressive content of the training photos is not observable in the car's function. This is in stark contrast with the operation of generative AI. If a user prompts Stable Diffusion or DALL-E to generate an image in the style of Pablo Picasso, for example, the traits that are a hallmark of Picasso's work are clearly observable in the generated image. While Picasso's work is particularly distinct, less-recognizable pieces of work can also be reverse-engineered in their entirety through use of clever prompts.<sup>298</sup> In other words, a core way in which expressive AIs are differentiated from non-expressive AIs is the extent to which the inputs are visible in, and can be consumed using, the outputs.

An AI's classification as expressive or non-expressive influences the fair use analysis in two important ways. We discuss the transformative use doctrine at length in a later subsection, but suffice to say here that the fact that an AI's creative purpose makes its use of copyrighted works less likely to be transformative under the first fair use factor. Additionally, the copying of expressive works to create an expressive AI is more likely to act as a substitute for the original artist's works – or for the artist's creative output as a whole – when analyzing

---

<sup>293</sup> Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC'Y U.S.A. 291 (2019). Sag's work focuses on the expressive vs. non-expressive distinction in the context of text data mining.

<sup>294</sup> *Id.* at 301.

<sup>295</sup> *Id.*

<sup>296</sup> Leval, *supra* note 212, at 1111.

<sup>297</sup> Matthew Sag, *Copyright and Copy-Reliant Technology*, 103 NW. U. L. REV. 1607 (2009).

<sup>298</sup> CARLINI ET AL., *supra* note 240.

the fourth fair use factor. More broadly, as discussed in section one above, the use of the expressive elements of a work rather than the non-expressive elements pushes the AI's use toward the copyrightable side of the ideas/expression dichotomy.<sup>299</sup> These three key effects of the expressive/non-expressive distinction combine to support a significant divergence in the fair use analysis when applied to an expressive AI compared to a non-expressive AI, providing plaintiffs with a key argument when bringing suit against expressive AI tools.

### 3. *Factor One*

As previously discussed, a finding of a transformative use is often dispositive for the factor one analysis. In the context of data collection and AI, transformative purposes have been found where the service at issue canvasses a dataset to provide new information about that data. A leading example is found in the Google Books cases, which show that even the verbatim use of vast quantities of copyrighted works can be transformative where the purpose is to provide information about those works.<sup>300</sup> Google compiled a massive database containing over 20 million books in order to design a tool that could categorize, search, and mine data about the works.<sup>301</sup> The Second Circuit found that the purpose of Google's database – namely, the ability to search the text of a database of books and glean certain information from those works – was “highly transformative” compared to the original creative purpose of those books.<sup>302</sup> What Google effectively designed was a tool that provided users information *about* the books rather than the information and creative expression *contained* in the books. In other words, it's highly unlikely that an author would create a work with a purpose that aligns with Google's – that is, the purpose of creating a large database that can be used to generate information.<sup>303</sup> The use of machine learning to read and categorize the books is a quintessential example of a use having transformative purpose despite extensive verbatim copying.

Whether a work is transformative in content is a decidedly different question. A new work is content-transformative when it imbues the original work with a new expression, meaning, or message.<sup>304</sup> For visual works, this new expression can be found in the different aesthetic, artistic style, or format of the new work in comparison to the original. Each of these factors are considered in *Cariou v. Prince*, where the Second Circuit analyzed whether the appropriation works of a defendant artist were content-transformative. Each of the defendant's works were composed of altered portions of works originally taken from a little-known photography book. The court concluded that the majority of the works at issue met the content-transformative bar by altering their size and artistic meaning, among other factors listed above.<sup>305</sup>

The content-transformative standard suffers from the fact that the interpretation of creative works can be subjective. Criticism of the content-transformative doctrine, for example, has expressed concern that whether a use is content-transformative is “entirely within the eye

---

<sup>299</sup> Pamela Samuelson, *Why Copyright Law Excludes Systems and Processes from the Scope of its Protection*, 85 TEX. L. REV. 1921, 1968 (2007).

<sup>300</sup> Authors Guild v. Google, Inc., 804 F.3d 202, 207 (2d Cir. 2015).

<sup>301</sup> *Id.* at 208-09.

<sup>302</sup> *Id.* at 218.

<sup>303</sup> Tushnet, *supra* note 228, at 889.

<sup>304</sup> Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569, 579 (1994).

<sup>305</sup> Cariou v. Prince, 714 F.3d 694, 706 (2d Cir. 2013).

of the judicial beholder.”<sup>306</sup> Even in *Cariou*, Judge Wallace’s concurring opinion expressed serious skepticism that the judges of the Second Circuit could competently distinguish between whether or not the artistic works at issue in the case had been sufficiently transformed.<sup>307</sup>

Ambiguity in the concept of transformative use is of particular consequence, given the tendency for the finding of transformative use to sway the overall fair use analysis. While the transformative use doctrine is absent from the statutory foundation of the fair use test, it has become the defining consideration of factor one. Between the first mention of the transformative use doctrine in the case law and 2017, nearly ninety-one percent of all cases that find a use transformative have found that use to be fair overall.<sup>308</sup> Thus, it is crucial that a plaintiff challenging the use of their work be prepared to argue that the defendant’s use of their work is not transformative in either content or purpose.

Despite the growth of the transformative use doctrine, a plaintiff challenging the use of their works in a generative AI dataset still has an avenue to argue that the use of their works is not transformative. Especially in the context of the author-centric use case, the argument that the defendant’s use is not transformative in purpose is straightforward. The name-to-Pokémon generator, for example, has only one feature: to generate images of Pokémon based on the input provided by a user. These images, in turn, have the purpose of providing entertainment to the user prompting the AI. This purpose aligns with the purpose of the original works, which is also to provide entertainment for fans of the Pokémon franchise and the games, TV shows, and other media it encompasses. Because of this alignment, Author-centric generative AIs may not have a transformative purpose. This conclusion is significantly influenced by the fact that these author-centric use cases have a narrow purpose in that they can generate only one type of image.

The extent to which author-centric AIs are transformative in content is less clear. Even an author-centric AI is unlikely to produce output that is a verbatim copy of a work by the original artist, although it is a possibility.<sup>309</sup> Typically, however, an author-centric AI will only copy the style of an author or artist and the output will not be substantially similar to the original work. Style is only “one ingredient”<sup>310</sup> of expression, and simply mimicking the style of an artist is akin to copying an idea, where what is copied is not protected by copyright.<sup>311</sup>

While mimicking an artist’s style is not itself infringement, the extent to which an already infringing use mimics the style of an original could be considered under the first fair use factor. Especially when an infringing use has a similar target audience as the original work, a more significant change in the content may be necessary to find that the use is content-

---

<sup>306</sup> Jane C. Ginsburg, *Letter from the U.S.: Exclusive Rights, Exceptions, and Uncertain Compliance with International Norms – Part II (Fair Use)*, 242 REVUE INTERNATIONALE DU DROIT D’AUTEUR [RIDA] 175 (2014) (Fr.).

<sup>307</sup> *Cariou*, 714 F.3d at 713 (Wallace, J., concurring). Wallace criticized specifically the fact that the distinction between the five non-transformative works and the twenty-five transformative works was made without any additional evidence that would be presented on remand.

<sup>308</sup> Asay et al., *supra* note 225, at 941.

<sup>309</sup> See Giorgio Franceschelli & Micro Musolesi, *Copyright in Generative Deep Learning*, 4 DATA & POL’Y e17 (2022). Franceschelli and Musolesi observe that despite the probabilistic model of generative AIs, a chance still exists that an AI’s output may “reconstruct idiosyncrasies of input data instead of reflecting underlying trends” to an extent that would mimic a specific input work to an infringing degree.

<sup>310</sup> *Steinberg v. Columbia Pictures Industries, Inc.*, 663 F.Supp. 706, 712 (S.D.N.Y. 1987).

<sup>311</sup> 17 U.S.C. § 102(b).

transformative.<sup>312</sup> Thus, in the case of the name-to-Pokémon generator, the artistic style used makes many of the generated images nearly indistinguishable from an authentic Pokémon to those unfamiliar with the franchise. While the analysis will vary based on the facts of an individual case, there is a stronger argument to be made that the output of an author-centric generative AI is not content-transformative.

The functions of Expressive AIs are also distinguishable from those in *Sega v. Accolade* and the Google Books cases. For starters, the intermediate copying in *Sega*<sup>313</sup> resulted in a different end use than in the expressive AI case. As mentioned, the Second Circuit in *Accolade* approvingly noted the fact that Accolade did not copy the Genesis source code simply to avoid doing their own creative work in developing a game. Rather, Accolade copied the Genesis source code only to determine how to make their game work on the console from a functional perspective. This may not hold for expressive AIs. While the copied works in the training dataset are used similarly as an intermediate step in the creation of an end product, the purpose of expressive AIs align far more closely with the purposes of the original works than the end products at issue in either *Accolade* or the Google Books cases. As a result, the argument that intermediate copying provides a strong case in favor of fair use is questionable, as the end product enabled by the intermediate copying is an expressive use that aligns with the same purpose of the copied works.

When taken as a whole, the fact that generative AIs are an expressive use of the original copyrighted works substantially alters the first factor analysis from that concerning a non-expressive AI. Expressive AIs are less likely to have a transformative purpose because their expressive nature is more likely to align with the original creative works that comprise their datasets. While the extent to which expressive AIs are transformative in context is variable, there are cases, such as the name-to-Pokémon generator, in which they would likely not be. In either case, the first factor argument in favor of fair use is weaker than in the case of non-expressive AIs.

#### 4. Factor Four

One does not need to have an overly expansive view of the future of AI to foresee the threat that generative AI tools pose to creative professions. As these tools improve, they have the potential to act as substitutes not just for individual works of artists, but for the artists themselves.<sup>314</sup> This potential for substitution as market harm contributes directly to the fourth factor in its analysis of whether or not the infringing use would adversely impact the market for the original work should that use become widespread.<sup>315</sup>

Already, generative AI has shown to have the capability to rival human-created works in multiple different cases. In September 2022, an art competition at the Colorado State Fair caused controversy when it was won by an AI-generated work.<sup>316</sup> AI-generated art has sold for hundreds of thousands of dollars at auctions, shattering the expectations that had been set for

---

<sup>312</sup> Tushnet, *supra* note 228, at 881 (Tushnet makes this observation while discussing the lack of a content-transformative use in the appropriation artwork at issue in *Cariou v. Prince*).

<sup>313</sup> *Sega*, 977 F.2d 1510 (9th Cir. 1992).

<sup>314</sup> Benjamin L. W. Sobel, *Artificial Intelligence's Fair Use Crisis*, 41 COL. J. L. & ARTS 45 (2017).

<sup>315</sup> See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417 (1984).

<sup>316</sup> Kevin Roose, *An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy*, N.Y. TIMES (Sept. 2, 2022), <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>.

its sale price.<sup>317</sup> In the less extravagant world of graphic design, tools are already available to create mockups for websites and apps.<sup>318</sup> As these applications continue to proliferate, they threaten to leave few areas of visual art untouched.

The varied nature of generative AI applications can have a major impact on the analysis of the fourth fair use factor. The type of generative AI use case at issue in a lawsuit can swing the fourth factor from strongly supporting a finding of no fair use to being far more ambiguous. In the author-centric use case, for example, there is a clear connection between proliferation of AI generated Pokémon and the disruption of the market for the original, copyrighted Pokémon. It's plausible that a user would opt not to pay for original merchandise where free AI generated alternatives are indistinguishable. Even in the general case, when a generative AI has the capability to mimic one artist's style, it can become impossible to tell whether a specific work is original to that artist or if it is a similar-looking AI-generated work. We have discussed the story of Greg Rutkowski at length, but this is exactly the situation he's found himself in with respect to Stable Diffusion AI: while Rutkowski originally thought that his AI-art popularity may be good for his brand, he quickly found that his original works were almost impossible to find amidst the deluge of AI-generated mimicry.<sup>319</sup>

While some artists have seen their markets begin to be displaced, the broader impact of generative AI on different types of visual art markets remains highly conjectural at this point. While we have pointed to the headline-grabbing sale of AI-generated art at auction as evidence that AI art has the potential to supplant existing fine artists, it remains to be seen whether art collectors and dealers will remain interested in AI art as it becomes a common occurrence rather than a novel phenomenon. Accordingly, any artists wishing to challenge the use of their works in an AI's training dataset must assess the factual circumstances of the market surrounding their works and the impact of the AI on that market when bringing a lawsuit. Frustratingly, such potential plaintiffs may simply have to wait for time to tell what the true scope and magnitude of these disruptions will be.

One fourth factor consideration that is coming into focus is the harm done to the licensing markets of copyright owners. The imperfections of the licensing market analysis included in the fourth factor has been written about at length. Especially when a fully-fledged licensing market is yet to develop, arguing that a licensing market *could* develop poses the issue of circular reasoning, in which a plaintiff argues that a potential market exists based on the use a defendant claims is fair.<sup>320</sup> In response, when considering lost licensing revenue, only harm to licensing markets that are "traditional, reasonable, or likely to develop"<sup>321</sup> count against fair use. When a licensing market does exist, however, the analysis is likely to weigh against fair use if a defendant has declined to obtain that license where other similar actors have.<sup>322</sup> When viewed together, the more vibrant and established a licensing market for a

---

<sup>317</sup> *Is Artificial Intelligence Set to Become Art's Next Medium?*, CHRISTIE'S (Dec. 11, 2018),

<https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>

<sup>318</sup> See, e.g., *About Us*, DESIGNS.AI, <https://designs.ai/about-us> (last visited May 25, 2023).

<sup>319</sup> Heikkilä, *supra* note 1.

<sup>320</sup> NIMMER & NIMMER, *supra* note 52, § 13.05(A)(4).

<sup>321</sup> *American Geophysical Union v. Texaco*, 60 F.3d 913, 930 (2d Cir. 1994). *Texaco* is far from a perfect example of the licensing market analysis and has faced criticism that the court reached too far in finding that Texaco harmed the licensing market for the journals containing the individual articles at issue in the case. See Mark A. Lemley, *Should a Licensing Market Require Licensing?*, 70 L. & CONTEMP. PROBS. 185 (2007).

<sup>322</sup> *Associated Press v. Meltwater U.S. Holdings, Inc.*, 931 F.Supp.2d 537 (S.D.N.Y. 2013).

copyrighted work is, the more likely any harm to that market suffered by an author will sway the fourth fair use factor.

The practicality of a licensing market for AI training images is mired by the colossal size of the data sets used by popular generative AI tools. The data set underpinning DALL-E, for example, contains hundreds of millions of image and text pairs.<sup>323</sup> Proponents of treating training data as fair use point out the impracticality of tracking down the copyright owners for each and every image and obtaining a license for each one.<sup>324</sup> Doing so would impose a large administrative burden and require significant resources. Indeed, Wendy Gordon's influential economic theory of fair use would characterize the massive transaction costs involved as a motivation for finding the use at issue fair.<sup>325</sup>

There are indications, however, that such a licensing market is beginning to develop. Massive depositories of images, such as Getty, have begun to license the millions of text and image pairs at their disposal to AI developers that need training data. For example, a recently filed complaint by Getty against Stability AI argues that Stability AI's use of Getty's images injured Getty because Getty engaged in the business of licensing its images to AI and machine learning developers.<sup>326</sup> The images in question were either created by Getty staff photographers, acquired by Getty via assignment of copyrights, or, crucially, licensed to Getty by the original photographers.<sup>327</sup> The existence of an emergent licensing market rebuts the argument that the licensing harm suffered by a plaintiff is merely hypothetical and hinges on the use being determined to be unfair.

Continued future development of licensing markets is also likely because of the potential for unique datasets to provide immense value to competing generative AI tools. Industry insiders have posited that obtaining proprietary datasets from companies like Google, Quora, and—in the case of generative AIs—Getty images and similar services, will serve as keys to differentiate one company's AI service from another.<sup>328</sup> IT giants have begun to offer services helping clients develop AI training datasets.<sup>329</sup> Not only is the market for AI training data—and the licenses associated with the copyrighted material thereof—likely to develop, it's already developing. The rapid pace of progress in this industry means that the potential for such a market has come into focus quickly, and the concrete evidence of this development refutes the assertion that licensing markets for AI datasets are too impractical or hypothetical to weigh against fair use for the fourth factor analysis.

While the fourth factor also demands an evaluation of the public benefits derived from an infringing use,<sup>330</sup> those benefits are difficult to balance here. Naturally, while many artists will be displaced, the overall volume of creative works possible with the aid of generative AI

---

<sup>323</sup> Ryan O'Connor, *How DALL-E 2 Actually Works*, ASSEMBLYAI (Apr. 19, 2022), <https://www.assemblyai.com/blog/how-dall-e-2-actually-works/>.

<sup>324</sup> Lemley & Casey, *supra* note 215.

<sup>325</sup> Wendy J. Gordon, *Fair Use as Market Failure: A Structural and Economic Analysis of the "Betamax" Case and Its Predecessors*, 82 COL. L. REV. 1600 (1982) (note that she backed off the consideration of transaction costs being dispositive).

<sup>326</sup> Complaint at 2, Getty Images (US), Inc. v. Stability AI, Inc., No. 1:23-CV-135 (D. Del. Feb. 3, 2023).

<sup>327</sup> *Id.*

<sup>328</sup> David Sacks, *The Give-to-Get Model for AI Startups*, SUBSTACK (Mar. 29, 2023), <https://sacks.substack.com/p/the-give-to-get-model-for-ai-startups>.

<sup>329</sup> See, e.g., *Our AI Training Data Platform*, TELUS INT'L, <https://www.telusinternational.com/solutions/ai-data-solutions/platform> (last visited May 5, 2023).

<sup>330</sup> Google, LLC v. Oracle America, Inc., 141 S. Ct. 1183, 1206 (2021).

will skyrocket. How does one balance the artistic merit of AI-generated works with the creative value of artistic works created by human artists? The comparative value of the two is a philosophical question that is outside of the scope of this paper.

Concededly, even in a world where AI generated art is ubiquitous, it's plausible that current artists will be able to adapt to the new artistic landscape. Rather than being entirely replaced by generative AIs, a world exists where AI tools are harnessed by existing artists to improve their efficiency, lower the costs of their projects, and generate more revenue overall. There are already examples of this change taking place: even outside of the generative AI context, many artists are experimenting with combining AI with their own creative talents to create new and exciting art.<sup>331</sup> However, the adaptation of some artists to an AI-guided future doesn't preclude the fourth factor arguments made by others.

These considerations aside, there are legitimate fourth factor arguments to be made against fair use as it relates to expressive generative AIs. Not only do these AIs have the potential to serve as a substitute for the works of the original artists, but they also have the potential to act as a substitute for the artists themselves. Regardless of the eventual scale of this substitution in the long term, the licensing harm done to copyright owners by the permissionless inclusion of their works in datasets is rapidly becoming concrete. These two divergences from the fourth factor analysis compared to non-expressive AI mean that plaintiffs have a legitimate argument that expressive, generative AIs should not be a fair use under factor four.

#### D. POLICY IMPLICATIONS

Expressive AI portends dramatic social change. Thus, the degree to which copyright curtails expressive AI will affect this social change and courts will at least implicitly address policy considerations. Therefore, the centuries old balancing act between the rights of monopoly entitlement holders and the interests of scientific<sup>332</sup> and artistic<sup>333</sup> progress will continue.<sup>334</sup> This is particularly relevant as the purpose of copyright's monopoly entitlement "is the conviction that encouragement of individual effort by personal gain is the best way to advance public welfare."<sup>335</sup> While the appropriate balance of copyright's monopoly entitlement is debated,<sup>336</sup> the general purpose of this constitutional entitlement is to "promot[e] broad public availability of literature, music, and the other arts... [and] to stimulate artistic

---

<sup>331</sup> Christopher McFadden, *7 of the Most Important AI Artists That Are Defining the Genre*, INTERESTING ENG'G (Nov. 10, 2019, 6:07 AM), <https://interestingengineering.com/innovation/7-of-the-most-important-ai-artists-that-are-defining-the-genre>.

<sup>332</sup> See *Grant v. Raymond*, 31 U.S. 218, 219 (1832) (finding that an accurate description of the invention "is necessary in order to give the public, after the privilege shall expire, the advantage for which the privilege is allowed, and is the foundation of the power to issue a patent").

<sup>333</sup> See *Fox Film Corp. v. Doyal*, 286 U.S. 123, 127 (1932) ("The sole interest of the United States and the primary object in conferring the monopoly lie in the general benefits derived by the public from the labors of authors.").

<sup>334</sup> See *Google v. Oracle*, 141 S. Ct. at 1203 (finding that Google's verbatim copying "of the Sun Java API . . . was consistent with the creative 'progress' that is the basic constitutional objective of copyright itself").

<sup>335</sup> *Mazer v. Stein*, 347 U.S. 201, 219 (1954); see also *Feist Publications, Inc. v. Rural Tel. Serv. Co.*, 111 S. Ct. 1282, 1290 (1991) (finding "the primary objective of copyright is not to reward the labor of authors, but 'to promote the Progress of Science and useful Arts.'").

<sup>336</sup> See *Eldred v. Ashcroft*, 537 U.S. 186, 213 (2003) (Breyer, J., dissenting) (discussing the purpose of copyright's entitlement and the social value of extending the length of copyright's monopoly term).



creativity for the public good."<sup>337</sup> In this way, Copyright is explicitly utilitarian<sup>338</sup> and practitioners on both sides will need to be able to articulate why their claims advance the public good.

Unique to Text-to-Art AI is the degree to which it intertwines function and expression. No other technology has heretofore so greatly depended upon original creative works while simultaneously promoting both functional and creative endeavors. Relevantly, the IP clause conjunctively requires that intellectual property law "promote the Progress of Science and useful Arts."<sup>339</sup> When past courts have adjudicated copyright issues impacting the progress of both science and art, they often enumerate these policy concerns and distinguish between functional and creative domains. Intermediate software copying cases,<sup>340</sup> search engine cases,<sup>341</sup> and the Supreme Court's recent ruling on copying APIs are particularly pertinent to this analysis.<sup>342</sup> When copying implicates a socially valuable use or function other than conventional expression, copyright is challenged<sup>343</sup> and protection granted to rights holders is more limited.<sup>344</sup>

Since secondary liability is an equitable doctrine that exists at the edges of copyright's entitlement, the policy considerations of courts are more apparent.<sup>345</sup> Within secondary liability, many otherwise absent aspects of copyright present themselves, such as good faith<sup>346</sup> or the volume of infringement.<sup>347</sup> Courts have shown that these equitable factors can be uniquely balanced against the defendant's ability to limit that harm. The court in *Netcom* found that the defendant may be contributorily liable "assuming Netcom is able to take simple

---

<sup>337</sup> *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 156 (1975).

<sup>338</sup> See Peter S. Menell, *Intellectual Property: General Theories*, in 2 *ENCYCLOPEDIA OF L. & ECON.* 129, 130 (Boudewijn Bouckaert & Gerrit De Geest eds., 2001) ("the utilitarian framework has been particularly central to the development of copyright law in the United States").

<sup>339</sup> U.S. Const. art. I, § 8, cl. 8.

<sup>340</sup> E.g., *Sega Enterprises Ltd. v. Accolade, Inc.*, 977 F.2d 1510, 1527 (9th Cir. 1992); *Sony Computer Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 603 (9th Cir. 2000).

<sup>341</sup> E.g., *Perfect 10 v. Amazon, Inc.*, 508 F.3d at 1168.

<sup>342</sup> See *Google LLC v. Oracle Am., Inc.*, 209 L. Ed. 2d 311, 141 S. Ct. 1183, 1186 (2021).

<sup>343</sup> *Id.* at 1208 ("The fact that computer programs are primarily functional makes it difficult to apply traditional copyright concepts in that technological world").

<sup>344</sup> *Id.* at 1202 (finding that verbatim copying of implementing code is fair because the "value in significant part derives from the value that... programmers, invest of their own time and effort to learn," the copying was done "only insofar as needed to include tasks that would be useful in smartphone programs," it "seek[ed] to expand the use and usefulness of Android-based smartphones... and its use was consistent with that creative 'progress' that is the basic constitutional objective of copyright itself.").

<sup>345</sup> See *Elf-Man, LLC v. Brown*, 996 F. Supp. 2d 1056, 1060 (E.D. Wash. 2014) (quoting *Sony*, 464 U.S. at 432) ("[C]ourts faced with novel theories of vicarious liability must be 'circumspect' in construing the scope of a copyright owner's monopoly. The guiding principle is that the benefit conferred by the copyright owner's monopoly must 'ultimately serve the cause of promoting broad public availability of literature, music and other arts.'").

<sup>346</sup> See *Screen Gems-Columbia Music, Inc. v. Mark-Fi Records*, 256 F. Supp. 399, 401 (S.D.N.Y. 1966) (finding that infringers were "fly-by-night companies organized by unscrupulous individuals who . . . sell the bootlegged product."); see also *Fonovisa v. Cherry Auction*, 76 F.3d at 261 ("The Fresno County Sheriff's Department raided the Cherry Auction swap meet and seized more than 38,000 counterfeit recordings. The following year, after finding that vendors at the Cherry Auction swap meet were still selling counterfeit recordings.").

<sup>347</sup> *Grokster*, 545 U.S. at 913. (*Grokster's* liability for inducement to copy was colored by the mass verbatim infringement where "billions of files [were] shared across peer-to-peer networks each month.").

measures to prevent further damage.”<sup>348</sup> In *Perfect 10 v Google*, the court found a similar possibility of contributory liability so long as Google “could take simple measures to prevent further damage.”<sup>349</sup> Here, the courts are tending towards a kind of cost-benefit analysis not dissimilar to negligence calculus.

Outside of secondary liability, courts also consistently consider social benefits when considering emerging technologies. A core reason why transformative uses are fair, even in the face of mass verbatim copying, is because they serve copyright’s social “objective of contributing to public knowledge.”<sup>350</sup> For example, search engines were fair use in no small part because of their social benefit.<sup>351</sup> As such, practitioners ought to be equipped to discuss those affordances that expressive AI grants society, the scope of copyright harm caused, and reasonable methods to ameliorate this harm.

In the context of expressive AI, there are two primary potential infringements. The first is the mass verbatim copying of original works for AI training. The second is substantially similar outputs, regardless of whether the output is deemed to originate from the AI or the prompter.

Policy considerations that could be applied to both training data and outputs include the AI’s promotion of copyright’s purpose, increased access to creativity, and AI’s impact on the jobs of conventional artists. Expressive AI promotes the broad availability of art, which is a foundational purpose of copyright.<sup>352</sup> Moreover, it could be argued that expressive AI promotes access to the arts, particularly to marginalized communities without the means to create.<sup>353</sup> Expressive AI may also promote the progress of both art and science, which could weigh in a court’s decision-making. Conversely, sympathetic plaintiffs will benefit from the fact that potentially wealthy companies have freely copied the works of small artists, which may strain their job prospects.

### 1. Training Set Policy Considerations

Fair use itself is an equitable doctrine that legally incorporates policy,<sup>354</sup> so this analysis will briefly overview those policy considerations which we have not touched on.

Over the last several decades as copyrightable works have increasingly intersected with functional technologies, courts have frequently found that those uses that promote progress are more likely fair.<sup>355</sup> By contrast, when a technology does not promote progress, but simply

<sup>348</sup> *Religious Tech. Ctr. v. Netcom On-Line Comm’n Servs., Inc.*, 907 F. Supp. 1361, 1375 (N.D. Cal. 1995).

<sup>349</sup> *Perfect 10 v. Amazon, Inc.*, 508 F.3d at 1146, 1172.

<sup>350</sup> *Authors Guild v. Google, Inc.*, 804 F.3d 202, 214 (2d Cir. 2015); *see also id.* at 212 (stating that the “goal of copyright is to expand public knowledge”).

<sup>351</sup> *Perfect 10 v. Amazon*, 508 F.3d at 1166 (“‘Google Image Search provide[s] great value to the public’ . . . [which] outweighed the significance of Google’s superseding use.”).

<sup>352</sup> *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 156 (1975).

<sup>353</sup> *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 101 (2d Cir. 2014) (public access to art for marginalized communities was a factor discussed at length in this litigation, including extensive congressional evidence that access was a major concern of recent copyright legislation).

<sup>354</sup> Examples of policy oriented fair use include the enumerated categories of likely fair uses, the doctrine of transformation furthering the goals of copyright, and those rulings such as *HathiTrust* that have furthered access to art for marginalized communities.

<sup>355</sup> Justin Hughes, *The Sub Rosa Rules of Copyright Fair Use*, 64 ARIZ. L. REV. 1, 27–47 (2022) (finding that the fair use standard has developed de facto rules around intermediate copying for reverse engineering, searchable database creation, and user interface copying).

circumvents copyright, then equitable<sup>356</sup> and functional<sup>357</sup> considerations support a finding of infringement.

Other policy considerations that may at least implicitly impact lawmaking decisions include AI's apparently insatiable need for creative work, the difficulty in determining copyrighted material, market failures and transaction costs of obtaining licenses, barriers to entry for small entities, and incumbency. Such policy considerations may be outside the scope of the courts and the legislature may be better positioned to act on these concerns.

Expressive AI has become viable in large part due to the massive number of works it ingests. The LAION database for text-to-art AI provides an open dataset of 5.85 billion text-image pairs,<sup>358</sup> which Stability AI uses for its open-sourced Stable Diffusion model. LAION created this dataset by crawling publicly available images from the internet.<sup>359</sup> Given copyright's incredibly low threshold of creativity<sup>360</sup> and that registration is valid retroactively, it is highly likely that this and other datasets contain massive amounts of copyrightable material. However, given the volume of works used, and the internet's anonymity, accurately determining work attribution for licensing becomes nearly impossible.

In this way, the necessary volume of copying, the low-threshold of creativity, and the anonymity of 'artists' results in a market failure that creates transaction costs to licensing markets that appear astronomically high. Many publicly available images do not effectively identify the internet's billions of authors, or identify participation in the creative commons, in such a way that an entity could effectively transact with each of them to license their works. Moreover, given the volume of necessary works, each individual work is of limited value,<sup>361</sup> so the licensing fee for an individual photo would have to be exceptionally low. Given these high transaction costs, only companies such as Getty, Shutterstock, or Disney could meaningfully participate in the industry.<sup>362</sup> With such marginal licensing fees, there is little incentive for the many billions of authors of the internet to go through even a streamlined licensing process.<sup>363</sup> Moreover, if such licensing fees were required, then copyright's monopoly entitlement could put a major barrier to small entities participating in the creation of

---

<sup>356</sup> *A&M Recs., Inc. v. Napster, Inc.*, 239 F.3d 1004, 1014 (9th Cir. 2001), *aff'd sub nom.* *A&M Recs., Inc. v. Napster, Inc.*, 284 F.3d 1091 (9th Cir. 2002).

<sup>357</sup> *Am. Broad. Companies, Inc. v. Aereo, Inc.*, 573 U.S. 431, 450 (2014).

<sup>358</sup> Romain Beaumont, *LAION-5B: A New Era of Open Large-Scale Multi-Modal Datasets*, LAION (Mar. 31, 2022), <https://laion.ai/blog/laion-5b/>.

<sup>359</sup> *Id.*

<sup>360</sup> *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 345 (1991) ("[T]he requisite level of creativity is extremely low; even a slight amount will suffice. The vast majority of works make the grade quite easily, as they possess some creative spark, 'no matter how crude, humble or obvious' (citation omitted) it might be.") In this way, all the unthinking photos posted to the web are copyrightable regardless of their potentially limited artistic expression.

<sup>361</sup> *Uses of Copyrighted Materials for Machine Learning*, STATE OF ISRAEL MINISTRY OF JUSTICE 7 (Dec. 18, 2022), <https://www.gov.il/BlobFolder/legalinfo/machine-learning/he/18-12-2022.pdf> ("[E]ach individual work is a single component in an enormous dataset and holds an immaterial weight in the dataset. The market value of each work will undoubtedly reflect this immateriality and be infinitesimal—probably significantly lower than the costs required for obtaining licenses for each work.").

<sup>362</sup> *Id.* at 22 ("Obviously, startup companies cannot afford excessive licensing. Moreover, defensive licensing practices of incumbents might create negative externalities on startups . . .").

<sup>363</sup> *Id.* ("Determining that ML dataset composition requires a license can increase the centralization in content markets. Such a conclusion could strengthen the market status of centralized content entities at the expense of individual authors.").

foundation models while providing major incumbent content holders such as Facebook and Google a substantial advantage.

There are of course exceptions, such as Shutterstock and Getty Images. These centralized online image companies have databases of text-image pairs uniquely valuable for machine learning. Moreover, their images have clearly identifiable watermarks that identify an entity with which an AI company could effectively bargain or avoid their incorporation into an ML training set altogether. Bargaining with these institutional entities reduces the market failure associated with high transaction costs.

## 2. *Generated Output Policy Considerations*

A major policy consideration that substantially favors plaintiffs in secondary liability claims are the foundational premises of cost spreading and the unfathomably high cost of broadly suing conventional end-user-promoters.<sup>364</sup> However, given courts' recent trend of finding technological uses of works fair,<sup>365</sup> it seems likely that the balance of courts' policy arguments may favor defendants.

Most importantly, as secondary liability has been tested by socio-technical change, secondary liability has not been found in cases of non-verbatim direct infringement.<sup>366</sup> This is likely because it may not seem reasonable to hold an entity liable for another's infringement when the underlying infringement is uncertain. Moreover, finding secondary liability in such circumstances may deter future artists and unduly burden expression as was the effect of the DMCA combined with the resultant implementation of automated content filtration algorithms.

While text-to-art AI have output near verbatim copies of original works,<sup>367</sup> expressive AI does not normally generate exact copies of copyrighted originals.<sup>368</sup> Given the large volume of outputs<sup>369</sup> and the emergent unpredictability of AI-generations, the ability to reduce infringements may be limited to the implementation of algorithmic output filters.<sup>370</sup> However, empirical evidence of output filters has demonstrated that they hurt more than they help. In the context of the DMCA, empirical evidence suggests algorithmic filters resulted in one-third of take-down requests that were either "fundamentally flawed" or "questionable."<sup>371</sup> As a result, "[e]rrors in automated anti-piracy scripts have generated hundreds of DMCA notices 'not related at all' to any of their copyrighted content, and have resulted 'in censorship of

<sup>364</sup> *Polygram Int'l Pub., Inc. v. Nevada/TIG, Inc.*, 855 F. Supp. 1314, 1325 (D. Mass. 1994) (describing the jurisprudential history and basis of copyright's secondary liability).

<sup>365</sup> See *supra* Part VI.C.

<sup>366</sup> The Ninth Circuit in *Perfect 10 v. Amazon* suggests that Google could be contributorily liable for their image search function if the district court found that Google had sufficient knowledge and a reasonable means of reducing direct infringement. While the court found Google's direct copying and transformation into thumbnail-sized replicas was transformative, the secondary liability analysis addressed the verbatim copies of directly infringing websites, to which Google redirected searchers. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146 (9th Cir. 2007).

<sup>367</sup> CARLINI ET AL., *supra* note 240.

<sup>368</sup> *Generative AI: Advantages, Disadvantages, Limitations, and Challenges*, FACT PROTOCOL (Jan. 17, 2023, 10:09 AM), <https://fact.technology/learn/generative-ai-advantages-limitations-and-challenges/>.

<sup>369</sup> *DALL-E Now Available Without Waitlist*, OPENAI (Sept. 28, 2022), <https://openai.com/blog/dall-e-now-available-without-waitlist>.

<sup>370</sup> Peter Henderson et al., *Foundation Models and Fair Use*, 20 (Apr. 13, 2023), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4404340](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4404340); see also Burk, *supra* note 160, at 300 (finding that "algorithms are currently deployed . . . to manage the overwhelming job of policing digital content").

<sup>371</sup> Amanda Reid, *Considering Fair Use: DMCA's Take Down & Repeat Infringers Policies*, 24 COMM. L. & POL'Y 101, 111 (2019).

perfectly legal content.”<sup>372</sup> While proponents of output filters, or algorithmic enforcement, admit their potential flaws,<sup>373</sup> they do not fully appreciate the extent to which the DMCA has shown that “algorithmic copyright enforcement may be[come]. . . a tool for global censorship.”<sup>374</sup> Beyond purely algorithmic enforcement, the same incentive structure of the DMCA, that promoted the centralized corporate over-policing of expression, would be present if expressive AI companies were found liable for infringing outputs.<sup>375</sup> Such is the inevitable balance among expression, copyright’s monopoly entitlement, and mass enforcement mechanisms.

A primary reason that algorithmic enforcement via output filters fails is because it cannot adequately apply copyright’s complex limiting doctrines. Algorithmic fair use “threatens to degrade the exception into an unrecognizable form”<sup>376</sup> and habituate artists to a more restrictive copyright norm.<sup>377</sup> Structural and technical limitations disallow algorithmic application of not only fair use but also “idea/expression distinction, exhaustion, functionality, [scenes a faire,] and other doctrines that likely do far more than fair use to control the shape and scope of copyright.”<sup>378</sup> As such, any policy arguments regarding the feasibility of reducing harm through algorithmic means should be aware of these pitfalls and lawmakers should not prescribe or incentivize the use of algorithmic copyright enforcement.<sup>379</sup>

## VII. CONCLUSION

Significant questions remain about how copyright law applies to text-to-art AI. The technology’s nature introduces complexity throughout an infringement analysis, even in foundational questions like who is justly liable when infringement occurs. The various stakeholders involved in generative AI art suggest that secondary liability can be found, but exactly how remains unclear. Perhaps most importantly, a fair use defense looms large for any plaintiff seeking remuneration through an infringement suit. Given the judicial and academic trend of expanding fair use to cover new technological contexts, it is possible that development, implementation, and use of text-to-art AI will be mostly free from infringement liability.

Still, text-to-art AI feels different compared with previous technologies. The fact that the model can produce an image based on minimal input from the user makes it seem like it

---

<sup>372</sup> *Id.*

<sup>373</sup> Henderson et al., *supra* note 367, at 20 (finding that fair use judgment cannot be reduced to an algorithm).

<sup>374</sup> Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 STAN. TECH. L. REV. 473, 490 (2016).

<sup>375</sup> See generally Charles W. Hazelwood, Jr., *Fair Use and the Takedown/Put Back Provisions of the Digital Millennium Copyright Act*, 50 IDEA 307, 307 (2010) (finding that the DMCA has misaligned incentives to the point that copyright holders have sent successful notices “for cases of questionable infringement, de minimis infringement, or in clear cases of fair use” including even high school video projects), [https://ipmall.law.unh.edu/sites/default/files/hosted\\_resources/IDEA/idea-vol50-no2-hazelwood.pdf](https://ipmall.law.unh.edu/sites/default/files/hosted_resources/IDEA/idea-vol50-no2-hazelwood.pdf)

<sup>376</sup> Burk, *supra* note 160, at 306.

<sup>377</sup> *Id.*

<sup>378</sup> *Id.* at 307.

<sup>379</sup> *Perfect 10 v. Amazon*, 508 F.3d 1146, 1174 (9th Cir. 2007) (regarding algorithmic filters for search engines, the Ninth Circuit intelligently found Google not vicariously liable in part because they “lack[] the ability to analyze every image on the internet, compare each image to all the other copyrighted images that exist in the world ... and determine whether a certain image on the web infringes someone’s copyright.”); see also *Religious Tech. Ctr. v. Netcom On-Line Comm.*, 907 F. Supp. 1361, 1374 (N.D. Cal. 1995) (finding “it is beyond the ability of a BBS operator to quickly and fairly determine when a use is not infringement where there is at least a colorable claim of fair use.”).

has a mind of its own. The Copyright Office nearly stated as much when it rescinded a copyright for an author that used AI generated art, stating “it was Midjourney—not Kashtanova—that originated the ‘traditional elements of authorship.’”<sup>380</sup> Even if this is an overstatement, the model can produce infringing images without any specific direction from a user to do so. As a result, it seems appropriate to explore liability for those that created the model, namely the developers, either for the copying that occurred upon the model’s training or for individual infringing outputs. Given that a text-to-art AI model’s purpose is to generate images with similar characteristics to its input data, fair use protection may be less warranted than in cases like *Author’s Guild*<sup>381</sup> where the model served a fundamentally different function than its inputs.

Of course, halting the expansion of fair use or extending secondary liability doctrines to this use case could have meaningful consequences for the generative AI industry. The mere threat of secondary liability reduces the incentive for developers to better understand their models. And denying a fair-use defense to copying that occurs during a models’ training could pose an existential threat to current development methods. Quick adoption, rapid innovation, and more lawsuits involving text-to-art AI will paint a better picture before long.

---

<sup>380</sup> U.S. COPYRIGHT OFF., ZARYA OF THE DAWN: LIBR. OF CONG. (2023), <https://copyright.gov/docs/zarya-of-the-dawn.pdf> (last visited Mar 15, 2023).

<sup>381</sup> *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015).