

1-29-2024

## Robots as Pirates

Henry H. Perritt Jr.  
hperritt@kentlaw.iit.edu

Follow this and additional works at: <https://scholarship.law.edu/lawreview>



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

---

### Recommended Citation

Henry H. Perritt Jr., *Robots as Pirates*, 73 Cath. U. L. Rev. 57 (2024).  
Available at: <https://scholarship.law.edu/lawreview/vol73/iss1/7>

This Article is brought to you for free and open access by Catholic Law Scholarship Repository. It has been accepted for inclusion in Catholic University Law Review by an authorized editor of Catholic Law Scholarship Repository. For more information, please contact [edinger@law.edu](mailto:edinger@law.edu).

---

## Robots as Pirates

### Cover Page Footnote

Professor of Law (retired) and former dean, Chicago-Kent College of Law. Member of the bar: Virginia, Pennsylvania (inactive), District of Columbia, Maryland, Illinois (inactive), Supreme Court of the United States. Author of more than 100 articles and twenty-five books on dispute resolution, technology and law, and labor law. Commercial helicopter and private instrument airplane pilot. Extra class radio amateur (K9KDF). The author has followed artificial-intelligence developments since the mid-1980s and wrote an expert system to evaluate the legality of employee terminations.

# ROBOTS AS PIRATES

*Henry H. Perritt, Jr.*<sup>+</sup>

Generative AI has created much excitement over its potential to create new works of authorship in the literary and graphical realms. Its underling machine-learning technology works by analyzing the relations among elements of preexisting material in enormous databases assembled from publicly available and licensed sources. Its algorithms “learn” to predict “what comes next” in different types of expression. A complete system thus can become glib in creating new factual summaries, essays, fictional stories and images. A number of authors of the raw material used by Generative AI engines claim that the machine learning process infringes their copyrights. Careful evaluation of actual and likely claims shows that such plaintiffs claiming infringement will have a hard time proving reproduction, distribution, display, or preparation of derivative works and thus are unlikely to be able to establish copyright infringement under established doctrines. Fears of uncompensated appropriation, however, are likely to fuel the erection of more pay walls around original content and more licensing collectives.

---

<sup>+</sup> Professor of Law (retired) and former dean, Chicago-Kent College of Law. Member of the bar: Virginia, Pennsylvania (inactive), District of Columbia, Maryland, Illinois (inactive), Supreme Court of the United States. Author of more than 100 articles and twenty-five books on dispute resolution, technology and law, and labor law. Commercial helicopter and private instrument airplane pilot. Extra class radio amateur (K9KDF). The author has followed artificial-intelligence developments since the mid-1980s and wrote an expert system to evaluate the legality of employee terminations.

INTRODUCTION .....	59
A. <i>Two AI-Generated Stories</i> .....	59
1. <i>Philander's Literary Masterpiece</i> .....	59
2. <i>Debbie's Essay on AI</i> .....	60
B. <i>Critical Analysis of the Stories</i> .....	61
C. <i>Scope of Article</i> .....	65
I. TECHNOLOGY .....	66
A. <i>Machine Learning</i> .....	67
B. <i>Pattern Matching</i> .....	73
C. <i>Learning Databases</i> .....	75
D. <i>Creating</i> .....	77
E. <i>Hardware</i> .....	77
F. <i>Particular Engines</i> .....	78
II. LAW OF INFRINGEMENT .....	79
A. <i>Limited Monopoly Under § 106</i> .....	80
B. <i>Exclusive rights</i> .....	80
1. <i>Reproduction</i> .....	80
a. <i>Copying</i> .....	81
b. <i>Direct Proof</i> .....	82
c. <i>Access</i> .....	83
d. <i>Similarity</i> .....	83
e. <i>Copying Incidental to Machine Learning</i> .....	86
2. <i>Preparation of Derivative Works</i> .....	86
3. <i>Distribution</i> .....	94
4. <i>Performance</i> .....	94
5. <i>Display</i> .....	95
C. <i>Identification of Particular Works</i> .....	96
D. <i>Only in Protected Expression</i> .....	98
1. <i>Not Ideas and Facts</i> .....	98
2. <i>Scenes A Faire, Abstraction, and Filtration</i> .....	98
E. <i>Vicarious infringement</i> .....	101
F. <i>Fair use: four factors</i> .....	104
1. <i>Purpose and Character of the Use &amp; Transformativeness</i> .....	104
2. <i>Nature of the Copyrighted Work</i> .....	107
3. <i>Amount and Character of Appropriation</i> .....	108
4. <i>Impact on Market</i> .....	108
G. <i>Transformativeness</i> .....	112
H. <i>The Plaintiff's Proof Challenge</i> .....	113
III. REMOVAL OF DMCA DIGITAL RIGHTS MANAGEMENT INFORMATION .....	115
IV. CONSIDERATION OF THE PURPOSES OF COPYRIGHT .....	117
CONCLUSION .....	119

## INTRODUCTION

### A. *Two AI-Generated Stories*

#### 1. *Philander's Literary Masterpiece*

The author asked ChatGPT 3.5 to “tell me a story about how generative AI infringed the copyright of Philander, who wrote an essay that a generative AI program learned from.”

Here is the response:

Once upon a time, there was a writer named Philander who wrote an essay that was considered a masterpiece in the literary world. His essay was widely read and appreciated, and it was taught in schools and universities all around the world. Philander was proud of his work and had no idea that it would one day become the subject of a copyright dispute with a generative AI program.

The AI program in question was designed to learn from a vast database of texts, including Philander's essay. It was trained to generate new content based on the patterns and structures it had learned from the input data. The program was successful and was able to generate high-quality writing that was almost indistinguishable from that of a human.

However, there was one problem. The AI program had learned so much from Philander's essay that it had inadvertently copied some of the key phrases and ideas from the original work. This was not intentional, but rather a result of the way the program had been trained.

When Philander learned of this, he was understandably upset. He had spent years crafting his essay, and now a machine had copied his work without permission. He decided to take legal action against the developers of the AI program, arguing that they had infringed his copyright.

The case went to court, and after much deliberation, the judge ruled in Philander's favor. The developers of the AI program were ordered to pay him a substantial sum of money in damages, and they were also required to remove all content that had been generated using Philander's essay as input.

The incident raised important questions about the ownership of content generated by AI programs. While the developers argued that the AI was an independent agent and that they had no control over what it generated, the court ultimately held them responsible for any infringement that occurred as a result of their program's actions.

In the end, the case served as a cautionary tale for anyone working with generative AI. It highlighted the need for developers to be mindful of copyright law and to ensure that their programs are not inadvertently infringing on the rights of others. And for writers like

Philander, it showed that even in the age of AI, the value of original creative work cannot be overstated.<sup>1</sup>

## 2. *Debbie's Essay on AI*

The author asked ChatGPT 3.5 to “tell me a story about how generative AI infringed the copyright of Debbie, who wrote an essay that a generative AI program learned from.”

Here is the response:

Once upon a time, there was a woman named Debbie who wrote a fantastic essay on the topic of artificial intelligence. Her essay was well-researched, thoughtfully crafted, and expertly articulated. Debbie was proud of her work and shared it online for others to read and learn from.

Unbeknownst to Debbie, a company had developed a powerful generative AI program that was designed to learn from existing written content and generate new text based on that knowledge. The program had been trained on millions of documents, including Debbie's essay, which it had analyzed and dissected thoroughly.

The program was so sophisticated that it could create entirely new essays that sounded like they had been written by a human. The company that owned the program had patented its technology and was using it to create original content for various purposes, including marketing campaigns and news articles.

One day, Debbie stumbled upon an article that sounded remarkably similar to her essay. Upon closer inspection, she realized that much of the text was lifted directly from her work. She was shocked and dismayed that her hard work had been stolen and used without her permission.

Debbie decided to take legal action against the company that had used her essay to train their generative AI program. She argued that they had infringed on her copyright and that she was entitled to compensation for the unauthorized use of her work.

The company argued that their program was not designed to create exact replicas of existing works, but rather to generate new content based on a deep understanding of the language and style used in those works. They claimed that since the program had created new text, they had not violated Debbie's copyright.

The case went to court, and the judge ultimately ruled in favor of Debbie. The judge agreed that the generative AI program had indeed infringed on Debbie's copyright by using her essay as a foundation for creating new content. The company was ordered to pay Debbie a

---

1. Henry Perritt, *AI Infringes Copyright* (May 6, 2023), <https://chat.openai.com/>.

substantial sum of money as compensation for their unauthorized use of her work.

Debbie's story serves as a cautionary tale about the potential pitfalls of generative AI technology. While these programs have the potential to create amazing new content, they must be used responsibly and with respect for the rights of creators like Debbie.<sup>2</sup>

### B. Critical Analysis of the Stories

The stories are similar to ones being told every day in newspapers, magazines, blogs, radio, and television reports and in at least two complaints filed in United States district court.<sup>3</sup> Public interest in generative AI has exploded.

Reporters and commentators breathlessly report that generative AI can compose music,<sup>4</sup> write computer programming code,<sup>5</sup> write sports stories,<sup>6</sup>

---

2. Henry Perritt, *AI Infringes Copyright* (Apr. 30, 2023), <https://chat.openai.com/>.

3. Class Action Complaint & Demand for Jury Trial, *Doe 1 v. GitHub, Inc.*, No. 3:22-cv-06823 (N.D. Cal. Nov. 3, 2022) [hereinafter "GitHub Complaint"]; Class Action Complaint & Demand for Jury Trial, *Andersen et al v. Stability AI Ltd.*, No. 3:23 cv-00201 (N.D. Cal. Jan. 13, 2023) [hereinafter "Stability AI Complaint"]. The court denied standing to seek damages, but found standing for the purposes of injunctive relief. *Doe v. GitHub, Inc.*, No. 22-cv-06823-JST, 2023 U.S. Dist. LEXIS 86983, at \*10 (N.D. Cal. May 11, 2023). The court dismissed claims for violation of Section 1202(a) of the DMCA, tortious interference, fraud, false designation of origin, and violation of the CCPA for failure to state a claim upon which relief can be granted, but granted leave to amend. *Id.* at \*16. It dismissed claims for unjust enrichment as preempted by the Copyright Act. *Id.* at \*18. The Court denied defendants' motions to dismiss Plaintiffs' claim under Sections 1202(b)(1) and 1202(b)(3). Plaintiffs' claim under Section 1202(b)(2) is dismissed with leave to amend. *Id.* at \*21. It denied the motion to dismiss the plaintiffs breach-of-license claim. *Id.* at \*22. It dismissed claims for unfair competition under the Lanham Act and California common law with leave to amend. *Id.* at \*24. Discovery continued and defendant's motions to dismiss the amended complaint were pending as of November 1, 2023. See *Doe 1 v. GitHub, Inc.* (4:22-cv-06823), COURTLISTENER, <https://www.courtlistener.com/docket/65669506/doe-1-v-github-inc/> (last visited Nov. 16, 2023).

4. John Koetsier, *Generative AI Music Platform Creates Forever Songs With Artists' Unique Sounds, Melodies, And Beats*, FORBES: CONSUMER TECH (Apr. 14, 2023, 1:41 PM), <https://www.forbes.com/sites/johnkoetsier/2023/04/14/generative-ai-music-platform-creates-forever-songs-with-artists-unique-sounds-melodies-and-beats/?sh=24e202fda478>. "Generative AI in music — like generative AI in everything — is having a moment." *Id.*

5. Jeffrey Dastin, *Microsoft Attracting Users to Its Code-writing, Generative AI Software*, REUTERS (Jan. 24, 2023, 7:53 PM), <https://www.reuters.com/technology/microsoft-attracting-users-its-code-writing-generative-ai-software-2023-01-25/>.

6. James Vincent, *Sports Illustrated's Publisher is Using AI to Generate Fitness Advice*, THE VERGE (Feb. 3, 2023, 9:48 AM), <https://www.theverge.com/2023/2/3/23584305/ai-language-tools-media-use-arena-group-sports-illustrated-mens-journal> (reporting on "tips" stories written by AI from archives).

report financial news,<sup>7</sup> write poetry,<sup>8</sup> write passing answers on law and business school exams,<sup>9</sup> write historical essays,<sup>10</sup> generate animated movies,<sup>11</sup> and explain esoteric engineering concepts.<sup>12</sup> But the press is also full of horror stories about how the technology infringes copyright<sup>13</sup> while it may not be entitled to copyright itself.<sup>14</sup> Moreover, critics claim that generative AI

---

7. Paul Farhi, *A News Site Used AI to Write Articles. It Was a Journalistic Disaster*, WASHINGTON POST (Jan. 17, 2023, 6:03 PM), <https://www.washingtonpost.com/media/2023/01/17/cnet-ai-articles-journalism-corrections/> (reporting on large numbers of corrections necessary in AI-written financial stories at CNET).

8. Seth Perlow, *AI is Better at Writing Poems Than You'd Expect. But That's Fine*, WASHINGTON POST (Feb. 13, 2023, 6:00 AM), <https://www.washingtonpost.com/books/2023/02/13/ai-in-poetry/> (reporting on successful poems).

9. Samantha Murphy Kelly, *ChatGPT Passes Exams from Law and Business Schools*, CNN BUSINESS (Jan. 26, 2023, 1:35 PM), <https://www.cnn.com/2023/01/26/tech/chatgpt-passes-exams/index.html> (reporting C+ law school and B to B- business school grades on robot-written essay and multiple-choice exam answers).

10. Daniel Lametti, *A.I. Could Be Great for College Essays*, SLATE (Dec. 7, 2022, 5:50 AM), <https://slate.com/technology/2022/12/chatgpt-college-essay-plagiarism.html> (reporting on use to write essays about history of London).

11. AnimationXpress Team, *How Generative AI Will Impact the Animation Industry*, ANIMATIONXPRESS (May 2, 2023), <https://www.animationxpress.com/animation/how-generative-ai-will-impact-the-animation-industry/> (explaining how generative AI speeds up the creation and improves the quality of animated movies).

12. In response to this prompt, “write a one-paragraph explanation of boundary layer in aerodynamics,” ChatGPT came up with this:

In aerodynamics, a boundary layer is a thin layer of air that forms near the surface of an object as it moves through a fluid, such as air. This layer of air experiences shear stresses due to the friction between the object’s surface and the surrounding fluid. The boundary layer affects the overall aerodynamics of the object, particularly in terms of drag and heat transfer. In laminar flow, the boundary layer is smooth and uniform, while in turbulent flow, it becomes thicker and more irregular. Understanding the behavior of boundary layers is crucial in designing efficient and effective aerodynamic systems, such as aircraft wings and turbine blades, as it allows engineers to minimize drag and optimize performance.

Henry Perritt, *AI Infringes Copyright* (May 7, 2023). The author has a bachelor’s degree in aeronautical engineering from MIT and can vouch for the accuracy of the explanation.

13. Gil Appel et al., *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> (reporting on claims that generative-AI programs infringe copyright).

14. Christopher T. Ziproli, *Generative Artificial Intelligence and Copyright Law*, CONG. RSCH. SERV. (last updated May 11, 2023), <https://crsreports.congress.gov/product/pdf/LSB/LSB10922> (exploring whether AI-generated works are entitled to copyright); Sharon Goldman, *Generative AI Output Could Be Eligible for Copyright Protection. But There’s a Catch*, VENTURE BEAT (Mar. 16, 2023, 6:37 AM), <https://venturebeat.com/ai/generative-ai-output-could-be-eligible-for-copyright-protection-but-theres-a-catch/> (reviewing position of Copyright Office).



intensifies dissemination of “harmful” content,<sup>15</sup> promotes “misinformation,”<sup>16</sup> enables manipulation,<sup>17</sup> defames innocent victims,<sup>18</sup> threatens jobs,<sup>19</sup> and generally erodes social ties.<sup>20</sup>

This article addresses one of these issues: the possibility that robot-generated content may infringe the copyrights of authors of the works in the databases from which the robot learned. It concludes that copyright owners will have a difficult time imposing liability for copyright infringement because they will have difficulty showing that protected elements of their particular works were copied, turned into derivative works, distributed to the public, displayed to the public, or performed for the public. Even if they can make such a showing, the elements of fair use militate toward excusing robot-generated material from infringement liability because it is so transformative.

That generative AI programs<sup>21</sup> and their users are unlikely to be liable for copyright infringement serves the policies of the Constitution’s Patents and Copyrights Clause and of the copyright act. Allowing breathing room for generative AI authors promotes creativity. Courts, lawmakers, and regulators should be wary of proposals to extend copyright liability to this new technology.

---

15. *NYU Law Forum—How Will the Law Handle Generative AI?*, N.Y.U. SCH. OF L., <https://its.law.nyu.edu/eventcalendar/index.cfm?fuseaction=main.detail&id=89813> (advertising April 5, 2023 conference exploring copyright infringement and generation of “harmful content”).

16. Vittoria Elliott, *It’s Way Too Easy to Get Google’s Bard Chatbot to Lie*, WIRE (Apr. 5, 2023, 12:47 AM), <https://www.wired.com/story/its-way-too-easy-to-get-googles-bard-chatbot-to-lie/> (reporting that Bard can be pushed into “denying climate change, mischaracterizing the war in Ukraine, questioning vaccine efficacy, and calling Black Lives Matter activists actors”).

17. Lance B. Eliot, *Generative AI ChatGPT As Masterful Manipulator Of Humans*, *Worrying AI Ethics And AI Law*, FORBES (Mar. 1, 2023, 8:00 AM), <https://www.forbes.com/sites/lanceeliot/2023/03/01/generative-ai-chatgpt-as-masterful-manipulator-of-humans-worrying-ai-ethics-and-ai-law/?sh=14da426a1d66> (warning that generative AI can concoct “fluently overbearing manipulations”).

18. Reuters, *Australian Mayor Prepares World’s First Defamation Lawsuit Over ChatGPT Content*, THE GUARDIAN (Apr. 6, 2023, 12:32 AM), <https://www.theguardian.com/technology/2023/apr/06/australian-mayor-prepares-worlds-first-defamation-lawsuit-over-chatgpt-content> (reporting on claim by mayor that he was falsely accused of bribery).

19. *What to Fear About Generative AI—and What to Do About It*, THE ECONOMIST (Apr. 26, 2023), <https://www.economist.com/news/2023/04/26/what-to-fear-about-generative-ai-and-what-to-do-about-it?> (reporting unease about job threats).

20. Jonathan Haidt & Eric Schmidt, *AI Is About to Make Social Media (Much) More Toxic*, THE ATLANTIC (May 5, 2023), <https://www.theatlantic.com/technology/archive/2023/05/generative-ai-social-media-integration-dangers-disinformation-addiction/673940/> (reporting on dangers of ChatGPT).

21. This article regularly refers to generative-AI systems, robots, or computers being liable for infringement. Such phrasing is intended merely to simplify the exposition; it is not meant to suggest defendant standing for inanimate systems. See Henry H. Perritt, Jr., *Copyright for Robots?*, 57 IND. UNIV. L. REV. 139 (2023) (discussing standing for computer systems and other inanimate objects).

The two robot-generated and robot-authored stories that begin the article are useful in introducing the possibility of copyright infringement liability for the output of generative AI systems. First, the difference between the stories despite nearly identical prompts shows what AI developers call *diversity* of system output and reinforces the idea that the systems do not simply retrieve pre-existing works contained in their learning databases. Second, the ChatGPT robot selected a technical essay for one story and a literary work for the other story sets up evaluation of whether the nature of the output is likely to make a difference in an infringement claim. Third, particular words and phrases used in the two different stories highlight elements that are important in the legal analysis.

Like most of the journalist stories, the robot-generated story involving Philander needs critical analysis to know whether its legal conclusion is warranted. Four phrases embedded in the story drive the legal analysis: “copied some of the key phrases and ideas;” “inadvertently” and “not intentional;” “patterns and structures it had learned from the input data;” and “the AI was an independent agent and that they had no control over what it generated.”

The first phrase describes verbatim copying, an infringement, if what was copied is protectable. The second phrase seems to frame a defense, but lack of intent to infringe is not a defense. It may, however, affect remedies, if the copyright owner establishes liability for infringement.

Under *Feist Publications, Inc. v. Rural Telephone Service Co.* “patterns and structures,” the third phrase, may constitute selection and arrangement of facts that are protectable by copyright, even if the underlying facts are outside protection.<sup>22</sup>

The fourth phrase, “independent agent,” sets up a discussion of vicarious liability.

Three phrases embedded in Debbie’s story drive the legal analysis: “much of the text was lifted directly from her work;” “They claimed that since the program had created new text, they had not violated Debbie’s copyright;” and “using her essay as a foundation for creating new content.”

The first phrase describes verbatim copying. That would, indeed, be infringement, but only if what was copied is protectable. If the robot copied only an idea or mere facts, copyright protection does not extend to it.<sup>23</sup>

The second phrase seems to frame a defense, but puts a lot of weight on the words *new* and *created*. Indeed, if the system independently created new text that was the same as Debbie’s, without copying her text, there was no infringement.

The third phrase is an incorrect statement of the law. Using protected expression as a foundation for creating new content is not infringement unless

---

22. *Feist Publ’ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 348–49 (1991).

23. 17 U.S.C. § 102(b) (2018).

the exclusive rights of an author have been violated.<sup>24</sup> Most works “stand on the shoulders” of pre-existing works.<sup>25</sup>

### C. Scope of Article

In addition to using the two robot-written stories as springboards for analysis, this article also evaluates two real lawsuits alleging that the output of generative AI programs infringed the rights of the owners of copyrighted material in their learning databases: *Doe I v. GitHub, Inc.*, and *Andersen v. Stability AI Ltd.*<sup>26</sup>

The GitHub lawsuit claims that Microsoft’s Copilot and Codex products and product distributor GitHub violated rights secured to the authors of pre-existing open-source software code by stripping it of its notices required by the open source licenses.<sup>27</sup> Most of the licenses required attribution of the author and notice of his or her copyright.<sup>28</sup>

The Stability AI lawsuit claims that Stability AI’s product, Stable Diffusion, “acquired copies of billions of copyrighted images without permission” and used them as training images.<sup>29</sup> When Stability outputs “new” images, the complaint claims, the “images are based entirely on the Training Images and are derivative works” of those images.<sup>30</sup>

This article explains the technology of generative AI, with an emphasis on how it uses pre-existing material comprising their learning databases. Then, it reviews the law of copyright infringement, including fair use, and of interference with copyright protection schemes. Those sections explore how courts draw the line between protectable expression and unprotectable ideas and facts. The next section highlights the challenging proof barriers a plaintiff faces, and argues that the constitutional purpose of copyright law—to promote creativity—is best served by limiting infringement law to a light touch. Mere threats of litigation over copyright infringement can chill innovation and creativity, however, and the best way to mitigate these threats may be development of new licensing systems for works included in deep-learning databases.

This article is the second in a quintet on legal issues raised by generative AI. The first, *Copyright for Robots?*<sup>31</sup> considers the question of whether output of a

---

24. 17 U.S.C. § 501(a) (2018) (defining copyright infringement).

25. *Quirk v. Sony Pictures Ent., Inc.*, No. C 11-3773 RS, 2013 U.S. Dist. LEXIS 47954, at \*12 (N.D. Cal. Apr. 2, 2013) (using phrase in copyright infringement case involving adaptation of preexisting work).

26. *See supra*, note 3.

27. GitHub Complaint at ¶ 9 (describing Codex and Copilot); *Id.* at ¶ 43 (describing relation of Codex, Copilot, and GitHub); *Id.* at ¶ 11 (stripping of notices violated license terms); *Id.* at ¶ 21 (describing GitHub). GitHub users make their copyrighted work available in GitHub repositories under one of several licenses defined by GitHub. *Id.* at ¶ 34 (describing licenses).

28. GitHub Complaint at ¶ 64.

29. Stability AI Complaint at ¶¶ 1–2.

30. Stability AI Complaint at ¶ 4.

31. Henry H. Perritt, Jr., *Copyright for Robots?*, 57 U. IND. L. REV. 139 (2023).

generative AI program qualifies for copyright protection. The third article in the trilogy<sup>32</sup> considers the question of liability for harmful content such as defamation, invasion of privacy, intentional infliction of emotional distress, or misappropriation of trade secrets, communicated by a generative AI program. It analyzes the role of Section 230 of the Communications Decency Act<sup>33</sup> in shielding AI systems from such liability. The fourth article questions the rush to regulate AI and offers a taxonomy of regulatory alternatives.<sup>34</sup> The fifth concludes that the impact of generative AI on the total number of jobs in the economy will be modest or positive.<sup>35</sup>

## I. TECHNOLOGY

Generative AI systems<sup>36</sup> use statistical analysis, massive computing power, and huge databases of pre-existing material to learn how to create new content that resembles but supplements and extends the expression in the database. These systems can write computer programming code and explain it,<sup>37</sup> tell stories, write essays about history or technology, write music, and create still photographs and feature-length full-motion videos.<sup>38</sup> To do this, they synthesize features of language, video depiction, music, or other subjects from millions, sometimes billions, samples of pre-existing works contained in their databases. This process of synthesis is called *machine learning*.<sup>39</sup> Then the systems use sophisticated *pattern matching* to associate semantic concepts they have learned with user-provided prompts and ultimately generate creative material.<sup>40</sup>

The underlying mathematical models and statistical algorithms are fifty to one hundred years old, but their power has been increased enormously by dramatic advances in computing power, low-cost storage, and broadband communication links; software advances, including improved techniques, new models, and opensource toolboxes; complemented by “Big Data”—the Internet

---

32. Henry H. Perritt, Jr., *Robot Slanderer*, 46 UNIV. OF ARK. LITTLE ROCK L. REV. 169 (2023).

33. 47 U.S.C. § 230 (2018) (immunizing communications intermediaries from liability for content on their systems).

34. Henry H. Perritt, Jr. *Robot Regulations*, 75 S.C. L. REV. (forthcoming 2023).

35. Henry H. Perritt, Jr., *Robot Job Destroyer*, 84 LA. L. REV. (forthcoming 2023).

36. This article also uses the terms “generative AI,” “creative robot,” “robot-created,” and “computer-generated” to mean the same thing.

37. See Alexander Amini, *MIT Introduction to Deep Learning 6.S191*, YOUTUBE (Mar. 10, 2023),

[https://www.youtube.com/watch?v=QDX-1M5Nj7s&list=PLtBw6njQRU-rwp5\\_\\_7C0oIVt26ZgjG9NI&index=2](https://www.youtube.com/watch?v=QDX-1M5Nj7s&list=PLtBw6njQRU-rwp5__7C0oIVt26ZgjG9NI&index=2).

38. See discussion *supra* Section INTRODUCTION.B (discussing news stories on achievements of generative AI).

39. See discussion *infra* Section I.A.

40. *Id.*

Revolution—which has made most of the world’s information available in digital form at the click of a computer mouse.<sup>41</sup>

#### A. Machine Learning

“Machine learning [is the] ability [of a computer] to learn without being explicitly programmed.”<sup>42</sup> “Deep learning extracts patterns from data using neural networks.”<sup>43</sup> Generative AI uses deep learning to develop a tool kit of very finely grained elements of expression which it then can reassemble into coherent output shaped by a user’s instructions. In all cases, the machine learning works from a huge database of training samples. It uses statistical techniques and mathematical equations associated with them to deconstruct the data into quantitative features and to predict how those features relate to each other.

All deep learning involves multiple layers of feature extraction. In face recognition, for example, the lowest layer recognizes and extracts clusters of a half-dozen or so pixels representing edges and lines.<sup>44</sup> The next higher layer recognizes combinations of those as corners and curves. A higher intermediate layer recognizes colors, shadows, and texture. Still higher layers discern eyes, noses, ears, corners, colors, shadows and texture. The highest layer recognizes facial structure.<sup>45</sup> Each progressively higher layer associates lower-level features with semantic aspects of pictures, with increasing levels of generality as one proceeds upward through the layers.<sup>46</sup>

Supplied with enough images in its learning database with the requisite subject matter and diversity,<sup>47</sup> the system can learn how particular aspects of all kinds of pictures can be represented digitally. It also learns how those aspects and features can be assembled so as to create an entirely new image of something meaningful: a face, a helicopter, or a wildfire. The state-of-the-art has been built up from techniques developed for more specialized types of image processing such as those associated with computer-aided tomography,<sup>48</sup> automated

---

41. See *MIT Introduction to Deep Learning 6.S191*, *supra* note 37.

42. *Id.* at 6:25 (modified original quote).

43. *Id.* at 7:11.

44. See *MIT Introduction to Deep Learning 6.S191*, *supra* note 37.

45. See *MIT Introduction to Deep Learning 6.S191*, *supra* note 37 (illustrating what happens at different layers).

46. See *Method and System for End-to-end Image Processing*, U.S. Patent No. 10,997,690 (explaining a method for processing input image, extracting features, and generating image in final layer of deep learning network).

47. See *generally* *Image Classification & Info Retrieval Over Wireless Digit. Networks & the Internet*, U.S. Patent No. 2014/0105467 A1, at paras. 0017-0026 (filed Dec. 2, 2013) (identifying preceding face recognition patents).

48. See Jie-Zhi Cheng et al., *Computer-Aided Diagnosis with Deep Learning Architecture: Applications to Breast Lesions in US Images and Pulmonary Nodules in CT Scans*, 6 *SCI. REP.* (2016), <https://www.nature.com/articles/srep24454>; see also *Computer Visualization of Anatomical Items*, Japan Patent No. 2019-166326.

animation for video games,<sup>49</sup> face recognition,<sup>50</sup> and object recognition. Image-processing, language-processing, and music-processing deep learning systems resemble each other in that they all work their way up from the finest grained elements of their subject matter and use statistical techniques to build a semantic tree from which new image, language, or music artifacts can be generated.<sup>51</sup>

The content of the learning databases is completely different, however, as are the semantic representations. The content of the databases and the particular techniques used vary depending on subject matter. A face-recognition program uses a database of faces. A cattle-herding program uses images of cattle. A music program uses a database of musical works. A computer coding program uses sections of computer code. A graphical image-creation program uses millions of images, and a text-program generation program uses millions of stories, articles, and tables of information.

With some types of data, it is sufficient if the machine learns to predict what should come next after a particular element of information. Given the name of a procedure, a computer coding program can predict the lines of code that should follow to implement the procedure. Given instructions about key signature and chord progression, a music composition program can write the baseline for a song by predicting what sequence of chords follow the tonic chord. Told the time signature and instructed to syncopate, it knows not to start a note on the first beat of the measure.

The systems learn by minimizing their *loss functions*—their errors in computing what comes next.<sup>52</sup> In image recognition systems,<sup>53</sup> for example, some of the samples contain the target image, and others contain something else. Thus, a robocowboy might be trained to recognize cattle by being presented with hundreds of thousands of images of different kinds of animal, tagging only those that represent cows, bulls, steers, and calves.<sup>54</sup> A computer system uses statistical analysis implemented through neural networks to evaluate which images match the tagged exemplars. The loss function represents quantitatively how many times it tags a snake as a cow and a rabbit as a bull. By recalculating the values in its layers, a model minimizes its loss function and converges on a set of quantitative templates that represent cattle. Machine-learning techniques accommodate challenges associated with recognizing the target image despite

---

49. See *supra* Section I.E (explaining the genesis of dominant deep-learning chip in video game processing.)

50. Alexander Amini, *MIT Introduction to Deep Learning 6.S191*, YOUTUBE (Mar. 10, 2023), at 12:30 (offering illustrations from face recognition.)

51. See *infra* text accompanying notes 69–77.

52. *Id.* at 32:56 (explaining loss functions in neural network learning).

53. The article explains machine learned by reference to image-recognition rather than natural-language processing, because image recognition is easier to understand.

54. See Henry H. Perritt, Jr., *The 21st Century Cowboy: Robots on the Range*, 43 U. ARK. LITTLE ROCK L. REV. 149 (2020) (exploring feasibility of robot cowboy who herds cattle; describing machine learning aimed at recognition of cattle).

different orientations, different lighting conditions, and different backgrounds.<sup>55</sup> The multiple layers in a robocowboy system would allow the robot to learn what a bovine is; learn how to reorient an image so that it more easily can compare it with others; and learn what features uniquely define a particular animal. It refines algorithms and templates so that it can apply these steps to an arbitrary set of images in production systems.

In image recognition generally, the techniques work by scanning the lines of an image, much as a laser printer or office scanner does, and looking for discontinuities in brightness and color. A model of an image can be constructed, starting with the locations of those discontinuities.<sup>56</sup> Then, a complex hierarchy of layers in neural networks,<sup>57</sup> implementing statistical algorithms, can compare the location of different types of discontinuities between images, and thus identify images that are most similar, working its way up through the layers. The indicia of similarity are the particular features at the highest levels that discriminate a cow from a wolf—or one bull from another.<sup>58</sup>

The “machine-learning” label applies to the process of identifying the distinguishing features that have statistical significance.<sup>59</sup> There is nothing magical about the analysis: it is factor analysis,<sup>60</sup> which has been used as a social

---

55. This involves the second step in most typologies: *alignment*.

56. A line connecting the discontinuities represents an “edge” in the image, such as the edge of a cheek in a human face.

57. See Divyansh Dwivedi, *Face Recognition for Beginners*, TOWARDS DATA SCI. (Apr. 28, 2018), <https://towardsdatascience.com/face-recognition-for-beginners-a7a9bd5eb5c2> (explaining how neural networks can use statistical techniques such as Principal Component Analysis, Linear Discriminant Analysis, Independent Component Analysis, Discrete Cosine Transforms, Gabor Filters, and Markov Models for face recognition); see generally Arun Alvappillai & Peter Neal Barrina, *Face Recognition Using Machine Learning*, available at <http://noiselab.ucsd.edu/ECE285/FinalProjects/Group7.pdf> (brief but formal paper on face recognition algorithms).

58. This is the *feature extraction* step. Some commentators call this identifying “landmarks.” See Adam Geitgey, *Machine Learning is Fun! Part 4: Modern Face Recognition with Deep Learning*, Medium (Jul. 24, 2016), <https://medium.com/@ageitgey/machine-learning-is-fun-part-4-modern-face-recognition-with-deep-learning-c3cffc121d78> (explaining machine learning for face recognition, step by step, including use of landmarks). The literature refers to features of human faces. Logically, it can be applied to distinguish among different species.

59. One popular method is the Viola/Jones approach. See Paul Viola & Michael Jones, *Rapid Object Detection using a Boosted Cascade of Simple Features*, CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (2001), <https://www.cs.cmu.edu/~efros/courses/LBMV07/Papers/viola-cvpr-01.pdf> (describing method in a more formal paper); see generally *The Viola/Jones Face Detector*, U. BRITISH COLUMBIA (2001), <https://www.cs.ubc.ca/~lowe/425/slides/13-ViolaJones.pdf> (slides explaining method);

60. See generally *Factor Analysis*, STATISTICS SOLUTIONS: COMPLETE DISSERTATION <https://www.statisticssolutions.com/factor-analysis-sem-factor-analysis/> (last visited Aug. 24, 2023).

science methodology for more than one hundred years.<sup>61</sup> Statistically predicting what comes next in an information artifact is not new, either. It is a fundamental principle of compression algorithms such as MP3 for audio, JPEG for photographs, and MPEG for full-motion video.<sup>62</sup> What has changed is computing power, the availability of digital storage, cheap digital cameras, and an enormous inventory of digital representations of faces.

In human face-matching, the process is the same, except that the foundational layer examines a large inventory of images of human faces (the *training database*) to learn what a face is.<sup>63</sup>

The most important measurements for face recognition programs are the distance between the eyes, the width of the nostrils, the length of the nose, the height and shape of the cheekbones, the width of the chin, the height of the forehead and other parameters. After that, the obtained data are compared with those available in the database, and, if the parameters coincide, the person is identified.<sup>64</sup>

The feature measurements can be expressed as a “feature vector,” or “faceprint,” which represents a particular face.<sup>65</sup>

Machine learning works similarly when the domain of interest is natural language rather than images.<sup>66</sup> Text programs parse the text in their databases and associate words, phrases, and sentences with each other and with a hierarchy of semantic concepts.<sup>67</sup> Then, given prompts that trigger particular semantic concepts, the programs assemble appropriate words and phrases according to grammar rules that it also has learned from its samples to output coherent stories and essays.<sup>68</sup>

---

61. See generally Charles Spearman, *Demonstration of Formulæ for True Measurement of Correlation*, 18 AM. J. PSYCH. 161 (Apr. 1, 1907) (documenting the use of factor analysis in the social sciences).

62. See Jeffrey S. Vitter, A Google TechTalk, Prediction as Data Compression, YOUTUBE (Aug. 20, 2019), <https://www.youtube.com/watch?v=wSQo2xUiSf0>.

63. See generally U.S. Patent Application No. 14/094,739, Publ'n No. 2014/0105467 A1, at paras. 0009-0015 (published Apr. 17, 2014) (Background of the invention, describing detection, alignment, normalization, representation, and matching steps in pattern matching).

64. See Divyansh Dwivedi, *Face Recognition for Beginners*, TOWARDS DATA SCIENCE (Apr. 28, 2018), <https://towardsdatascience.com/face-recognition-for-beginners-a7a9bd5eb5c2> (presenting computer program in the Python language that performs the basic steps).

65. See, e.g. U.S. Patent No. Application No. 14/094,739, Publ'n No. US20140105467A1, at 0016, 0034 (published Apr. 17, 2014) (explaining feature vector).

66. See, e.g. Julia Hirschberg & Christopher D. Manning, *Advances in Natural Language Processing*, 349 SCIENCE 261 (2015) (explaining evolution of computerized natural language processing).

67. See Alexander Mamaev, *Generating Music with AI (or Transformers Go Brrrr)*, MEDIUM (Jan. 28, 2021) <https://alxmamaev.medium.com/generating-music-with-ai-or-transformers-go-brrrr-3a3ac5a04126> (offering Python code illustrating parsing and tokening for low levels of music machine learning).

68. The assertion in the text is a logical extrapolation of the process explained in the patent discussed in the text and accompanying footnotes 69–77.



It is a bottom-up learning process: first, distinguishing individual words, then evaluating the frequency with which words appear together in phrases, then associating the phrases with concepts, and then building a hierarchy of concepts—a semantic tree—<sup>69</sup>not unlike a conventional thesaurus.<sup>70</sup> A loss function quantifies the accuracy of the system’s predictions. Sophisticated networks feed back their results to compare them with actual data, permitting loss functions to be computed and then improved.<sup>71</sup>

A recent patent for generative AI explains how generative AI text systems work at the technical level.<sup>72</sup> This 51-page patent explains in some detail how machine learning is used to train a system for natural language output in a style that “talks like me.” At a fundamental level, natural language analyzers and generators use an ontological network: a sophisticated, computerized thesaurus, which classifies linguistic concepts and organizes them according to their relationship with each other. The result is a semantic framework for a particular language. Particular words are slotted into their appropriate ontological classes, thus organizing the entire vocabulary of the language.

The system parses training text into sentences,<sup>73</sup> uses pattern matching to classify concepts expressed in each sentence, and then assigns semantic tokens accordingly.<sup>74</sup> Deictic context is developed by the use of anchor words, which signify that the surrounding syntactical units should be processed as teaching examples. Examples of anchor words signifying comparison include “increase, reduction, decrease, decline, rise, fall, raise, and lower.”<sup>75</sup> Anchor words are specified for each semantic concept. A complete system may use scores of separately patented methods.<sup>76</sup>

---

69. Jesús Giménez & Lluís Márquez, *Linguistic Measures for Automatic Machine Translation Evaluation*, 24 MACH. TRANSLATION 209, 213–15, 232–36 (2010) (noting use of semantic trees in machine translation of languages); see Hyejin Youn, et al., *On the Universal Structure of Human Lexical Semantics*, 113 PROC. OF THE NAT’L ACAD. OF SCI. OF THE U.S. OF AM. 1766 (2016); see generally Jean-Pierre Koenig & Anthony R. Davis, *The Key to Lexical Semantic Representations*, 42 J. LINGUISTICS 71, 73, 87, 106 (2006).

70. Warren Sack, *Conversation Map: An Interface for Very Large-Scale Conversations*, 17 J. MGMT. INFO. SYS. 73, 79–80 (2000-2001) (explaining analysis of email messages and Usenet posts with reference to semantic trees and thesauri).

71. See Jason Brownlee, *Loss and Loss Functions for Training Deep Learning Neural Networks*, MACHINE LEARNING MASTERY (last updated Oct. 23, 2019), <https://machinelearningmastery.com/loss-and-loss-functions-for-training-deep-learning-neural-networks/> (explaining loss functions).

72. Applied Artificial Intelligence Technology for Using Natural Language Processing to Train A Natural Language Generation System, U.S. Patent No. 11,042,713 (issued June 22, 2021) [hereinafter “‘713 Patent”].

73. *Id.* at col. 16 l. 37–col. 17 l. 38.

74. *Id.* at col. 17 l. 38–col. 18 l. 3, col. 18 l. 5–col. 19 l. 18.

75. *Id.* at col 11 l. 55–col. 16 l. 35.

76. The ‘713 Patent lists 20 patents and 26 patent applications. *Id.* at [56].

The system described in the patent uses metadata from natural language processing to identify concept expressions, ultimately expressing them in the form of an annotated template structurally compatible with the transformer architecture. Concepts are identified by their inclusion of anchor words: single or compound words globally unique to a particular concept. The patent gives twenty-three examples of anchor words for the *change* concept.<sup>77</sup> It gives ten anchor words for the *compare* concept, beginning with “more” and “less.”<sup>78</sup>

It summarizes the process flow as:

1. Tokenize a document into sentences
2. For each sentence:
  - A. Preprocess dependency and constituency criteria for named entity recognition
  - B. Leverage user data in the systems ontology to identify and flag known resources (entities and attributes)
3. For each pre-processed sentence:
  - A. Pass the sentence through a separate pattern matcher for each concept
  - B. For each template extractor, apply a set of heuristics to extract the relevant subtree from the parsed sentence and parameterize the sentence into a form compatible with the transformer.<sup>79</sup>

Subtree extraction works by moving up the tree from the anchor word to the complete clause containing the anchor word.<sup>80</sup>

The most sophisticated machine learning systems employ transformers—the “T” in ChatGPT. Transformers are a type of neural network architecture that “remember” what came before, maybe long before, in a sequence of words, ideas, sounds, or images. The enable better and more efficient prediction of “what comes next” that the common alternative, recurrent neural networks. Transformers use convolutional neural networks combined with attention models.<sup>81</sup>

---

77. *Id.* at col. 12 l. 30–52.

78. *Id.* at col. 12 l. 61–col. 13 l. 3.

79. *Id.* at col. 16 l. 7–35.

80. *Id.* at col. 18 l. 5–34.

81. See Giuliano Giacaglia, *How Transformers Work*, TOWARDS DATA SCI. (Mar. 10, 2019), <https://towardsdatascience.com/transformers-141e32e69591> (explaining transformers and comparing them with recursive neural networks and giving examples of language translation systems). “A transformer model is a neural network that learns context and thus meaning by tracking relationships in sequential data like the words in this sentence.” Rick Merritt, *What Is a Transformer Model?*, NVIDIA: BLOGS (Mar. 22, 2022), <https://blogs.nvidia.com/blog/2022/03/25/what-is-a-transformer-model/> (explaining how transformers “pay attention”); see Alexander Amini, *MIT 6.S191: Recurrent Neural Networks, Transformers, and Attention*, YOUTUBE (Mar. 17, 2023), [https://www.youtube.com/watch?v=ySEx\\_Bqxvvo](https://www.youtube.com/watch?v=ySEx_Bqxvvo).

When the domain is music, the learning process involves parsing a recording of music or a musical score into its most basic components of pitch, duration, loudness, sequencing, and simultaneity of notes. Then, at progressively higher levels, it learns to associate patterns of these components into rhythm patterns, time signatures, and tonal keys. At still higher levels, it learns patterns that involve repetition of musical phrases and sentences, such as the exposition, development, and recapitulation sections of a sonata<sup>82</sup> or the verse-chorus-verse-chorus-bridge structure of a rock song.<sup>83</sup> In addition to predicting the probability of what might come next in a particular piece of music in its database, generative AI systems like MuseNet, develop statistically based relationships among various works and multiple composers such as between Green Day and Pachelbel or between the Beatles and Beethoven.<sup>84</sup>

Image and text generating programs require more sophistication than mere prediction. Image programs use a combination of diffusion and interpolation to represent images in their database and then to reassemble small pieces of them in response to user prompts. The process works only when some degree of human tagging of the data exists.

Deep learning models use neural-network transformer architecture to learn the characteristics of their subject-matter from enormous amounts of data. MuseNet, for example, uses the recompute and optimized kernels of Sparse Transformer<sup>85</sup> to train a 72-layer network with 24 attention heads—with full attention over a context of 4096 tokens. “This long context may be one reason why it is able to remember long-term structure in a piece . . . .”<sup>86</sup>

### B. Pattern Matching

AI systems use pattern matching at multiple levels. They use the technique at the bottom layer of their learning stages, to detect primitive features of relevant information, such as edges in an image, or delimiters like spaces or punctuation marks in text. They also use pattern matching at the highest level of their processing to make use of information that now is slotted into position in a

---

82. MasterClass, *Sonata Form in Music: A Basic Guide to Sonata Form*, MASTERCLASS (June 7, 2021),

<https://www.masterclass.com/articles/sonata-form-in-music#11Q8IBOeI7QJBInNbumtcv> (explaining sonata in Baroque, Classical, and Romantic periods).

83. The Tunedly Team, *All About Rock Music*, TUNEDLY (Oct. 9, 2019), <https://www.tunedly.com/blog/allaboutrockmusic.html#:~:text=Like%20other%20popular%20genres%2C%20the,so%20feel%20free%20to%20experiment> (describing common structure of rock song.)

84. Christine Payne, MuseNet, OPENAI (Apr. 25, 2019), [openai.com/blog/musenet](https://openai.com/blog/musenet) (displaying interactive chart of composers).

85. Rewon Child & Scott Gray, *Generative Modeling with Sparse Transformers*, OPEN AI (Apr. 23, 2019), <https://openai.com/research/sparse-transformer>.

86. Christine Payne, *MuseNet*, OPENAI (Apr. 25, 2019), [openai.com/blog/musenet](https://openai.com/blog/musenet) (displaying interactive chart of composers) (giving example of a sample imitating Chopin).

semantic tree. Semantic pattern matching involves determining whether two artifacts have the same meaning. The artifacts may be sentences; they may be passages of music; they may be images. The pattern matching system encodes the artifacts as vectors whose variables correspond to what the system has learned from a large database through machine-learning techniques. The variables reflect entries in a high-level semantic tree. It then uses statistical measures of similarity to determine if they have the same meaning.<sup>87</sup>

Face matching programs perform pattern matching at their highest level to select the closest possible matches with prompt faces. The semantic possibilities are the actual faces from the systems' enrolled databases.<sup>88</sup> None of the other generative AI applications, however, selects entire works from a database—even the face-matching programs do not select images from the learning database. Rather, they aggregate properties that they infer from the individual learning-database entries. Then, they assemble new creations based on matching those properties with user prompts. It is possible, of course, that someone might build a generative AI program that appropriates particular inputs and presents them in its output, but that is not the norm for how the technology works.

One author has given an example of a semantic search engine coded in the Python language.<sup>89</sup> The same basic techniques embodied in the program are used to organize information derived from large learning databases,<sup>90</sup> to generate questions pertinent to database information,<sup>91</sup> and to enable conversational computerized tutors for students.<sup>92</sup> The basic idea of search based on pattern

---

87. Ken Gu & Rohit Saha, *An Introduction to Semantic Matching Techniques in NLP and Computer Vision*, MEDIUM (Sept. 8, 2021), <https://medium.com/georgian-impact-blog/an-introduction-to-semantic-matching-techniques-in-nlp-and-computer-vision-c22bf3cee8e9> (explaining semantic pattern matching); Kaveti Naveenkumar & Shrutendra Harsola, *Deep Learning for Semantic Text Matching*, MEDIUM (2006), <https://towardsdatascience.com/deep-learning-for-semantic-text-matching-d4df6c2cf4c5> (giving examples); Brian Gallagher, *Matching Structure and Semantics: A Survey on Graph-Based Pattern Matching* (2006), <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=ca593017e255771a0600a054a0c2ce75abf921ca>.

88. Face-recognition terminology calls the learning databases the “training” databases and the databases of actual persons whose faces might match a candidate image the “enrolled” database. See Henry H. Perritt, Jr., *Defending Face-Recognition Technology (And Defending Against It)*, 25 J. TECH. L. & POL'Y 41, 50–51 (2020) (describing two types of databases).

89. See Omar Boufeloussen, *How to Build a Semantic Search Engine Using Python*, MEDIUM (Jan. 6, 2023), <https://medium.com/mllearning-ai/how-to-build-a-semantic-search-engine-using-python-5c68e8442df1>.

90. See *Semantic Graphs and Conversational Agents*, U.S. Patent No. 9,196,245 B2 (issued Nov. 24, 2015) (describing semantic clustering techniques to group utterances according to topic and importance).

91. See *Automated Tool for Question Generation*, U.S. Patent No. 10,614,106 B2 (issued Apr. 7, 2020).

92. Arthur C. Graesser et al, *Learning by Communicating in Natural Language with Conversational Agents*, 23 CURRENT DIRECTIONS IN PSYCH. SCI. 374, 375-76 (2014) (describing computerized natural language tutor).

matching has been used for a decade or more to improve free-text search.<sup>93</sup> Both Westlaw and Lexis have natural language search features.<sup>94</sup> The advances in text-based systems permit search for concepts (semantic categories) rather than just literal strings of text.

### C. Learning Databases

Large language models learn from enormous databases of text such as CC-100<sup>95</sup> or Pile.<sup>96</sup> The databases are constructed from crawling the Web, for specialized sources such as FreeLaw and PubMed, from online discussion groups like Reddit, from the Gutenberg books collection, and from Wikipedia.<sup>97</sup> NVIDIA uses:

- Web pages, include articles, blogs, and news reports.
- Books, including works of fiction, non-fiction, and academic texts.
- Scientific papers from various fields, including physics, biology, chemistry, and more.
- Social media text data from platforms such as Twitter, Facebook, and Instagram.
- News articles various sources covering current events and trends.
- Chat logs and customer service interactions to understand natural language better and to improve conversational abilities.<sup>98</sup>

Models to write computer programming code can learn from publicly available databases. GitHub allowed open-source developers to make their open-source code available to other developers around the world, by putting it into public repositories hosted on GitHub's internet-accessible servers for free or at low cost.<sup>99</sup> The GitHub Complaint is vague on the sources of computer

---

93. See *Method and Apparatus for Semantic Pattern Matching for Text Retrieval*, European Patent No. 0601550 B1 (granted Aug. 22, 2001) (using semantic pattern recognition to locate relevant text rapidly).

94. Paul Callister, *Introduction to Online Legal Research (Westlaw and Lexis)*, U. OF ILL. COLL. L.: ALBERT E. JENNER JR. MEM'L LIBR. (2002), <https://www1.law.umkc.edu/faculty/callister/bootcamp/ot/nl01.html>.

95. Alexis Conneau, et al., *Unsupervised Cross-lingual Representation Learning at Scale*, FACEBOOK AI (Apr. 2020), <https://paperswithcode.com/dataset/cc100>.

96. *The Pile: An 800GB Dataset of Diverse Text for Language Modeling* (last visited Nov. 2, 2023), <https://pile.eleuther.ai/>.

97. See Matt Rickard, *LAION, The Pile, and More Datasets*, MATT RICKARD (Dec. 4, 2022), <https://matt-rickard.com/laion-the-pile-and-more-datasets> (listing and evaluating different training datasets); Kevin Schaul et al., *Inside the Secret List of Websites that Make AI Like ChatGPT Sound Smart*, WASHINGTON POST (Apr. 19, 6:00 AM), <https://www.washingtonpost.com/technology/interactive/2023/ai-chatbot-learning/> (identifying major sources as patents.google.com, wikipedia.org, scribd.com, fool.com, and kickstarter.com).

98. Henry H. Perritt, Jr., ChatGPT response to prompt, "What sources do large language models use?" CHATGPT (Apr. 16, 2023).

99. GitHub Complaint at ¶¶ 113–115 (describing founding a purpose of GitHub).

program code used to train the accused generative AI system. It says, “GitHub and OpenAI have offered shifting accounts of the source and amount of the code or other data used to train and operate Copilot.”<sup>100</sup> Elsewhere it says, “Though Defendants have been cagey about what data was used to train the AI, they have conceded that the training data includes data in vast numbers of publicly accessible repositories”<sup>101</sup> The complaint briefly explains how the generative AI software works: “Codex and Copilot are assistive AI-based systems offered to software programmers. These AI systems are each trained on a large corpus of publicly accessible software code and other materials, including all the Licensed Materials.”<sup>102</sup>

Image models construct learning databases from a number of different sources with hundreds of thousands of images, some of them tagged.<sup>103</sup> The Stability AI complaint alleges that over five billion images scraped from websites constituted the training database for Stable Diffusion.<sup>104</sup>

LAION (“Large-Scale Artificial Intelligence Open Network”), based in Hamburg, Germany, has a dataset, LAION-400M, containing 400 million training images that include text captions.<sup>105</sup> The training images in LAION-400M and their text captions were copied or scraped from web pages or other sources without the consent of the image owners or website operators. LAION-400M was extended in LAION-5B, a new dataset of 5.85 billion training images.<sup>106</sup> Vendors, such as Telus International,<sup>107</sup> offer data collection, bounding boxes, semantic segmentation, landmark and key point annotation, and image classification and categorization.

OpenAI’s MuseNet<sup>108</sup> learns from a database of thousands of musical works contained in sources such as:

- ClassicalArchives
- BitMidi
- Online collections of jazz, pop, African, Indian, and Arabic styles, and the

---

100. *Id.* at ¶ 14.

101. *Id.* at ¶ 10.

102. *Id.* at ¶ 46.

103. Pranav Dar, *25 Open Datasets for Deep Learning Every Data Scientist Must Work With*, ANALYTICS VIDHYA (last modified Aug. 27, 2021), <https://www.analyticsvidhya.com/blog/2018/03/comprehensive-collection-deep-learning-datasets/>.

104. Stability AI Complaint at ¶ 57.

105. *Id.* at ¶¶ 101–104 (describing LAION and its databases).

106. *Id.*

107. TELUS International, *Image Data*, <https://www.telusinternational.com/solutions/ai-data-solutions/ai-training-data/image-data> (last visited Aug. 29, 2023).

108. Christine Payne, *Muse Net*, OPENAI (Apr. 25, 2019), <https://openai.com/research/musenet>.

- MAESTRO dataset.<sup>109</sup>

Open AI says, “MuseNet was not explicitly programmed with our understanding of music, but instead discovered patterns of harmony, rhythm, and style by learning to predict the next token in hundreds of thousands of MIDI files.”<sup>110</sup> Other, more specialized, machine learning systems, use databases of road videos taken from a driver’s perspective, to build autonomous cars and autopilots for cars with drivers.<sup>111</sup>

#### D. Creating

When a generative AI system creates, it starts at its top layers and uses pattern matching to associate user prompts with concepts contained in the top layer of its semantic tree. It then works its way down through the layers, from top to bottom, pulling increasingly fine fragments of language, music, or imagery, depending on its purpose. It is important to understand that when the system works its way downward through the transformer layers, it does not simply follow the same paths that it used working its way up when it was learning. Rather, it makes paragraph, sentence, and word choices (or their equivalents in music or imagery) indicated by the values determined statistically in its learning process as to what comes next. It is extremely unlikely that it would follow a path back to verbatim content in the learning database.

#### E. Hardware

Ordinary computer memory and CPU chips<sup>112</sup> are incapable of handling all the computations necessary for machine learning at anything close to an acceptable speed. So, deep learning systems use hundreds of specialized chips developed initially for video gaming. The market leader is NVIDIA, which enjoyed an 80% market share in 2020. Intel and Advanced Micro Devices are aggressively moving to catch up, and new entrants like Cerebras Systems claim to be leapfrogging NVIDIA in the capability of new chips.<sup>113</sup>

Cerebras says that its WSE-2 (Wafer Scale Engine) is the largest chip ever built.<sup>114</sup> It has “2.6 trillion transistors, 40GB of high performance on-wafer

---

109. *Id.*

110. *Id.*

111. Alexander Amini, *MIT 6.S191: Recurrent Neural Networks, Transformers, and Attention*, YOUTUBE (Mar. 17, 2023).

112. A *chip* is an integrated semiconductor circuit. CPU is a central processing unit.

113. Marisa Garanhel, *Top 20 Artificial Intelligence Chips of Choice in 2022*, AI ACCELERATOR INST. (Oct. 14, 2022), <https://www.aiacceleratorinstitute.com/top-20-chips-choice-2022/>; see Asa Fitch, *Nvidia, Other Chip Companies Race to Cash In on ChatGPT Frenzy*, THE WALL STREET JOURNAL, (Feb. 23, 2023 9:38 AM), <https://www.wsj.com/articles/chip-makers-see-chatgpt-stirring-strong-demand-for-advanced-processors-76f152d1>.

114. Marisa Garanhel, *Top 20 Artificial Intelligence Chips of Choice in 2022*, AI ACCELERATOR INST. (Oct. 14, 2022); see generally CEREBRAS, <https://www.cerebras.net/> (last visited Sept. 13, 2023); see generally Nolan Day et al., *Cerebras-GPT: Open Compute-Optimal*

memory, and 850,000 AI-optimized cores to accelerate AI work.”<sup>115</sup> The WSE-2 is “independently programmable and optimized for the tensor-based, sparse linear operations underpinning inference and neural network training for deep learning.”<sup>116</sup> The scope of the chip permits programming teams to build models that are large-scale without wrestling with the distributed programming necessary for GPU clusters.<sup>117</sup>

The NVIDIA chip has “76.3 billion transistors and 18,432 CUDA Cores capable of running at clocks over 2.5 GHz, while maintaining 450W TGP” electric power draw. Its architecture was developed to calculate things like lighting effects in a typical videogame,<sup>118</sup> which requires multiple retracing operations per pixel, ranging as high as 600 in NVIDIA’s latest chips’ RT overdrive mode.<sup>119</sup> The most sophisticated deep-learning engines use hundreds or thousands of these chips.<sup>120</sup> That has implications for both cost and power consumption.

#### F. Particular Engines

In 2023, the largest large-language models included: WuDao 2.0 Beijing Academy of Artificial Intelligence, with 1.75 trillion parameters; MT-NLG NVIDIA and Microsoft, with 530 billion parameters; GPT-3, Open AI, with 175 billion parameters; LaMDA Google, with 137 billion parameters; and ESMFold Meta AI, with 15 billion parameters.<sup>121</sup> These models develop statistical predictions of what text comes next. Their numbers of parameters quantify the number of factors they consider in making predictions and generating output. Each parameter uses a variable, the value of which the model can vary as it

---

*Language Models Trained on the Cerebras Wafer-Scale Cluster*, ARVIX (Apr. 6, 2023), <https://arxiv.org/abs/2304.03208>.

115. *Id.*

116. Marisa Garanhel, *Top 20 Artificial Intelligence Chips of Choice in 2022*, AI ACCELERATOR INST. (Oct. 14, 2022).

117. *Id.*

118. *NVIDIA Ada GPU Architecture*, NVIDIA, <https://images.nvidia.com/aem-dam/Solutions/geforce/ada/nvidia-ada-gpu-architecture.pdf> (last visited Sept. 1, 2023); see Sharon Goldman, *How Nvidia Dominated AI — and Plans to Keep It That Way as Generative AI Explodes*, VENTUREBEAT (Feb. 23, 2023, 5:00 AM), <https://venturebeat.com/ai/how-nvidia-dominated-ai-and-plans-to-keep-it-that-way-as-generative-ai-explodes/> (reporting on how chips developed for video games drove machine learning to new levels).

119. *NVIDIA Ada GPU Architecture*, at 4; see *Graphics Cards*, NVIDIA, <https://www.nvidia.com/en-us/geforce/graphics-cards/> (last visited Aug. 31, 2023).

120. Sharon Goldman, *How Nvidia Dominated AI — and Plans to Keep It That Way as Generative AI Explodes*, VENTUREBEAT (Feb. 23, 2023, 5:00 AM) (reporting that OpenAI used 10,000 NVIDIA GPUs to train ChatGPT).

121. Cam Dilmegani, *Large Language Models: Complete Guide in 2023*, AIMULTIPLE (last updated June 21, 2023), <https://research.aimultiple.com/large-language-models/>.



learns.<sup>122</sup> The top image generation engines are Open AI's DALL-E 2,<sup>123</sup> Midjourney,<sup>124</sup> and Stability AI's<sup>125</sup> DreamStudio and Stable Diffusion.<sup>126</sup>

Top music creators Magenta and MuseNet,<sup>127</sup> can generate new music from prompts in the form of short MIDI files, and MusicLM<sup>128</sup> can generate music from text prompts. The results are moderately pleasing, but lack structure, surprise, and resolution.<sup>129</sup> The models typically undergo expensive database and computationally intensive basic training by their developers and then a simpler "fine-tuning" by customers on their own, much smaller databases.

## II. LAW OF INFRINGEMENT

Derivative works are not new.

Around 1730, Johann Sebastian Bach wrote an arrangement of Antonio Vivaldi's *Concerto for Four Violins*. Bach changed the key of the concerto from B minor to A minor, and the four featured instruments from violins to harpsichords. More importantly, he enriched Vivaldi's score with melodic elaborations and harmonic colorations. Bach first performed the concerto with his sons at Zimmerman's Coffee House in Leipzig. Neither Bach nor Zimmerman charged for the performance, although Zimmerman benefited from increased coffee sales. Both the Vivaldi concerto and Bach's arrangement of it are frequently performed today. While the popularity of Bach's concerto has arguably undermined the market for Vivaldi's original work, it is just as likely that the market for Vivaldi's

---

122. *Id.*

123. See *Dalle-E2*, OPEN AI, (Aug. 31, 2022), <https://openai.com/product/dall-e-2>.

124. See MIDJOURNEY, <https://www.midjourney.com/home/?callbackUrl=%2Fapp%2F> (last visited Aug. 29, 2023).

125. See STABILITY AI LTD., <https://stability.ai/> (last visited Aug. 28, 2023).

126. Harry Guinness, *The Best AI Image Generators in 2023*, ZAPIER (July 17, 2023), <https://zapier.com/blog/best-ai-image-generator/> (reviewing top image generator picks).

127. Ezra Sandzer-Bell, *Google Magenta Studio: A Free AI Music Generator*, AUDIOCIPHER (Jan. 26, 2023), <https://www.audiocipher.com/post/google-magenta> (comparing Magenta with MuseNet).

128. Ezra Sandzer-Bell, *Google Magenta Studio: A Free AI Music Generator*, AUDIOCIPHER (Jan 26, 2023); Ezra Sander-Bell, *5 Reasons Google's MusicLM AI Text-to-Music App is Different*, AUDIOCIPHER (Feb. 2, 2023), <https://www.audiocipher.com/post/musiclm> (giving examples of text prompts and musical results); see Andrea Agostinelli et al, *MusicLM: Generating Music From Text*, GOOGLE RSCH., <https://google-research.github.io/seanet/musiclm/examples/> (last visited May 6, 2023) (giving examples of music generated from different text prompts and different music generated from the same text prompts).

129. The author has written some 100 songs and is familiar with music theory and music composition.

concerto has been enhanced through association with the work of a musician of much greater renown.<sup>130</sup>

There was no copyright law in Germany when Bach committed what would today, in the United States, almost certainly be considered piracy and, legally, a violation of the Copyright Act.<sup>131</sup>

#### A. Limited Monopoly Under §106

The Copyright Act gives the holder of a copyright the exclusive right to:

- (1) Reproduce the copyrighted work;
- (2) Prepare derivative works;
- (3) Distribute *copies* to the public;
- (4) Perform the copyrighted work *publicly*;
- (5) *Display* the copyrighted work *publicly*; and
- (6) *Perform* sound recordings *publicly* by means of a digital audio transmission.<sup>132</sup>

The statute does not confer an exclusive right to use, read, watch, listen to, or otherwise to enjoy a copyrighted work.

The Stable Diffusion complaint alleges violation of all of these exclusive rights, except the last.<sup>133</sup> The robot-generated story about Debbie says that the judge issued judgment in her favor based on his finding that the AI system “us[ed] her essay as a foundation for creating new content.”<sup>134</sup> The third phrase is an incorrect statement of the law. Using protected expression as a foundation for creating new content is no infringement unless the exclusive rights of an author have been violated. Most works “stand on the shoulders” of pre-existing works.

#### B. Exclusive rights

##### 1. Reproduction

The copyright owner has the exclusive right to “to reproduce the copyrighted work in copies or phonorecords.”<sup>135</sup>

To prevail on a copyright infringement claim, a plaintiff must show that (1) he or she owns the copyright in the infringed work, and (2) the

---

130. Charles Cronin, *I Hear America Suing: Music Copyright Infringement in the Era of Electronic Sound*, 66 HASTINGS L. J. 1187, 1245 (2015) (internal footnotes omitted).

131. The United States did not exist when Bach wrote the work described in the text. If the 2023 U.S. Copyright had been in effect in 1730 and applicable to Bach’s composing activities, his adaptation of Vivaldi’s work would have been creation of a derivative work, infringing Vivaldi’s exclusive right under 17 U.S.C. § 106(2).

132. 17 U.S.C. § 106 (2018).

133. Stability AI Complaint at ¶ 160.

134. See *supra* text accompanying note 2.

135. 17 U.S.C. § 106(1).

defendant copied protected elements of the copyrighted work.<sup>136</sup> A copyright plaintiff may prove copying with circumstantial, rather than direct, evidence. ‘Absent direct evidence of copying, proof of infringement involves fact-based showings that the defendant had ‘access’ to the plaintiff’s work and that the two works are ‘substantially similar.’<sup>137</sup>

Intent to infringe is not an element. The robot-generated story about Philander says that the AI system he accuses of infringement “inadvertently” copied his expression.<sup>138</sup> The second phrase seems to frame a defense, but lack of intent to infringe is not a defense.<sup>139</sup> Lack of intent, however, would shield the defendant from liability for vicarious infringement.<sup>140</sup> It also may affect remedies if the copyright owner establishes liability for infringement.

Facts establishing infringement of the reproduction right likely exist with most robot-generated expression, given the nature of the way that any computer system moves digital information back and forth, copying it from storage into RAM and processor chips and back to RAM and storage. But the plaintiff will have to show that these acts of reproduction occurred with respect to her particular copyrighted works, as § III.G shows.

The Stability AI complaint alleges direct infringement of the reproduction, derivative-work, performance, and display rights.<sup>141</sup> It alleges that the defendants use copyrighted images in the database to generate images “that are derived exclusively from the Training Images, and that add nothing new.”<sup>142</sup>

#### *a. Copying*

The way in which machine learning works<sup>143</sup> means that a copyright infringement plaintiff is not going to be able to show verbatim copying as between copyrighted input owned by the plaintiff and particular output from the program. There might, however, be intermediate ephemeral copying in the early stages of the machine learning process as contents of that learning database are manipulated.

---

136. *Williams v. Gaye*, 895 F.3d 1106, 1119 (9th Cir. 2018). The court’s use of the word “copied” makes it obvious that the claim was based on infringement of the reproduction right, as opposed to one of the other rights enumerated in § 106.

137. *Id.* at 1119 (internal citation and quotations omitted).

138. *See supra*, note 1.

139. *Perfect 10, Inc. v. Giganews, Inc.*, 847 F.3d 657, 666–67 (9th Cir. 2017) (volitional conduct, but not intent, is requirement for direct infringement liability).

140. *See infra* Section II.E (discussing vicarious infringement); *MGM Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 930–34 (2005) (analyzing intent as an element of vicarious infringement).

141. Stability AI Complaint at ¶ 160.

142. *Id.* at ¶ 164.

143. *See supra* Section I.A (explaining machine learning and giving examples).

In *Rimini Street, Inc. v. Oracle International Corp.*,<sup>144</sup> the district court, citing *MAI Systems Corporation v. Peak Computer, Inc.*,<sup>145</sup> held that copying computer codes into representation in random access memory (“RAM”) constitutes copying within the scope of the reproduction right, notwithstanding that the copies are ephemeral; each time the program ran, a new copy was created.<sup>146</sup> The holding in *MAI*, however, has been called into doubt in some more recent cases.<sup>147</sup> The consensus is that a representation in RAM must be more than transitory to qualify as embodiment for the reproduction right.<sup>148</sup>

The robot-generated story about Debbie says that “[t]hey [the accused infringer] claimed that since the program had created new text, they had not violated Debbie’s copyright.”<sup>149</sup> The second phrase seems to frame a defense, but puts a lot of weight on the words *new* and *created*. Indeed, if the system independently created new text that was the same as Debbie’s, without copying her text, there was no infringement.

*b. Direct Proof*

*Rogers v. Koons*,<sup>150</sup> involved a sculpture copied from a photograph of a group of puppies. The court of appeals affirmed a permanent injunction against the further display or possession of the infringing sculpture. The court of appeals found direct evidence of copying:

Here, the trial court found original elements of creative expression in the copyrighted work were copied and that the copying was so blatantly apparent as not to require a trial. We agree that no reasonable juror could find that copying did not occur in this case. First, this case presents the rare scenario where there is direct evidence of copying. Koons admittedly gave a copy of the photograph to the Italian artisans with the explicit instruction that the work be copied. Moreover, the importance of copying the very details of the photograph that embodied plaintiff’s original contribution—the poses, the shading, the expressions—was stressed by Koons throughout the creation of the

---

144. *Rimini St. v. Oracle Int’l Corp.*, 473 F. Supp. 3d 1158, 1202 (D. Nev. 2020).

145. *MAI Sys. Corp. v. Peak Computs., Inc.*, 991 F.2d 511, 518–19 (9th Cir. 1993).

146. *Rimini St. v. Oracle Int’l Corp.*, 473 F. Supp. 3d 1158, 1203–04 (D. Nev. 2020).

147. *Compare* *Cartoon Network LP v. CSC Holdings, Inc.*, 536 F.3d 121, 127–30 (2d Cir. 2008) (suggesting that brief transitory copies may not qualify for “embodiment” aspect of reproduction right) *and* *CDK Global LLC v. Brnovich*, 16 F.4th 1266, 1276 (9th Cir. 2021) (holding that more than transitory duration is required for embodiment) *with* *Stenograph LLC v. Bossard Assocs.*, 144 F.3d 96, 101–02 (D.C. Cir. 1998) (discussing criticisms of *MAI* but finding it to be the general rule).

148. The phrase “in copies” in the statutory statement of the reproduction right in 17 U.S.C. § 106(a) is known as the “embodiment requirement.” *See* *Matthew Bender & Co. v. West Publ’g Co.*, 158 F.3d 693, 703 (2d Cir. 1998) (explaining embodiment in copies as part of reproduction right).

149. *See supra*, note 1.

150. *Rogers v. Koons*, 960 F.2d 301 (2d Cir. 1992).

sculpture. His instructions invariably implored that the creation must be designed “as per photo.” This undisputed direct evidence of copying is sufficient to support the district court’s granting of summary judgment.<sup>151</sup>

The robot-generated story about Debbie says that “much of the text was lifted directly from her work” by the AI system she accuses of infringement.<sup>152</sup> The first phrase describes verbatim copying. That would, indeed, be infringement, but only if what was copied is protectable. If the robot copied only an idea or mere facts, copyright protection does not extend to it.

The GitHub Complaint alleges facts that would support a claim for direct infringement: “Copilot often simply reproduces code that can be traced back to open-source repositories or open-source licensees.”<sup>153</sup> GitHub conceded that Copilot reproduces passages of code verbatim.<sup>154</sup> The complaint alleges, that the output may contain code snippets of more than 150 characters that match code from training data.<sup>155</sup> Nevertheless, the GitHub claimants did not plead copyright infringement, probably because their code, though copyrighted, was open source.<sup>156</sup>

#### *c. Access*

Access to the copyrighted work is a logical element of indirect proof of copying, because one cannot copy something if he does not have access to it. So even if an infringement plaintiff proves similarity between the copyrighted work and the accused work, she also must show access. In the generative AI context, that would involve showing that the copyrighted work was included in the database from which the robot learned. If the defendant had no access, it was not possible for him to copy. If the accused and the original work are not similar, one was not copied from the other.

#### *d. Similarity*

Similarity of the accused work to the underlying work can be probative of copying and infringement of the reproduction right. In *Rogers v. Koons*,<sup>157</sup> the court of appeals affirmed a permanent injunction against the further display or possession of a sculpture copied from a photograph of a group of puppies.

---

151. *Id.* at 307.

152. *See supra*, note 2.

153. GitHub Complaint at ¶ 13.

154. GitHub Complaint at ¶ 90; *see also id.* at ¶¶ 74, 87 (alleging verbatim or nearly verbatim copies).

155. GitHub Complaint at ¶ 90.

156. *See supra* Section INTRODUCTION.C, describing GitHub litigation; *see also* GitHub Complaint at ¶ 64 (characterizing license for plaintiff’s material as “open-source”).

157. *Rogers v. Koons*, 960 F.2d 301, 301 (2d Cir. 1992).

[E]ven were such direct evidence of copying unavailable, the district court's decision could be upheld in this case on the basis that defendant Koons' access to the copyrighted work is conceded, and the accused work is so substantially similar to the copyrighted work that reasonable jurors could not differ on this issue.<sup>158</sup>

Substantial similarity does not require literally identical copying of every detail.<sup>159</sup> Such similarity is determined by the ordinary observer test: the inquiry is "whether an average lay observer would recognize the alleged copy as having been appropriated from the copyrighted work."<sup>160</sup>

Stated another way, the question is whether an "ordinary observer, unless he set out to detect the disparities, would be disposed to overlook them, and regard their aesthetic appeal as the same."<sup>161</sup>

The Koontz court went on to say:

We recognize that ideas, concepts, and the like found in the common domain are the inheritance of everyone. What is protected is the original or unique way that an author expresses those ideas, concepts, principles or processes. Hence, in looking at these two works of art to determine whether they are substantially similar, focus must be on the similarity of the *expression* of an idea or fact, not on the similarity of the facts, ideas or concepts themselves.<sup>162</sup>

The court applied the basic principles to the fact on trial:

It is not therefore the idea of a couple with eight small puppies seated on a bench that is protected, but rather Roger's *expression* of this idea—as caught in the placement, in the particular light, and in the expressions of the subjects—that gives the photograph its charming and unique character, that is to say, makes it original and copyrightable.

Thus, had appellant simply used the *idea* presented by the photo, there would not have been infringing copying. But here Koons used the identical expression of the idea that Rogers created; the composition, the poses, and the expressions were all incorporated into the sculpture to the extent that, under the ordinary observer test, we conclude that no reasonable jury could have differed on the issue of substantial

---

158. See *id.*, at 307 (citing *Warner Bros., Inc. v. Am. Broad. Cos., Inc.*, 654 F.2d 204, 207 (2d Cir. 1981)).

159. See *id.* (citing 3 David Nimmer, *Nimmer on Copyright*, § 13.03[A] (2003)); see also *Comptone Company Co. Ltd. v. Rayex Corp.*, 251 F.2d 487, 488 (2d Cir. 1958).

160. *Id.* at 307 (citing *Ideal Toy Corp. v. Fab-Lu Ltd.*, 360 F.2d 1021, 1022 (2d Cir.1966)).

161. *Id.* at 307–08 (citing *Peter Pan Fabrics, Inc. v. Martin Weiner Corp.*, 274 F.2d 487, 489 (2d Cir. 1960)).

162. See *id.* at 308 (citing *Durham Indus., Inc. v. Tomy Corp.*, 630 F.2d 905, 912 (2d Cir. 1980)).

similarity. For this reason, the district court properly held that Koons “copied” the original.<sup>163</sup>

The existence of dissimilarities does not save an accused work:

Moreover, no copier may defend the act of plagiarism by pointing out how much of the copy he has not pirated.<sup>164</sup> Thus, where substantial similarity is found, small changes here and there made by the copier are unavailing. It is only where the points of dissimilarity exceed those that are similar and those similar are—when compared to the original work—of small import quantitatively or qualitatively that a finding of no infringement is appropriate.<sup>165</sup> This is not the case here. Koons’ additions, such as the flowers in the hair of the couple and the bulbous noses of the puppies, are insufficient to raise a genuine issue of material fact with regard to copying in light of the overwhelming similarity to the protected expression of the original work.

Because of Koons’ extensive use of the same expression of the idea that Rogers’ created, it was properly held that he “copied” the protected features of the original. No genuine issue of material fact exists with respect to this finding; “String of Puppies” was copied from the photograph “Puppies” based either on the direct evidence of copying or on proof of access and substantial similarity. In light of this summary judgment was properly granted on this issue.<sup>166</sup>

Lucille Ponte argues that the concept of substantial similarity is an uneasy fit for digital sampling music cases, where small amounts of copyrighted works are taken literally and reassembled into new works, often with aspect of the sound manipulated electronically.<sup>167</sup> The owner of copyrighted works included in a learning database will have a difficult time showing that any particular copyrighted work was used by the computer program to generate any particular output. He will have an even more difficult time proving copying. Direct proof of copying would require getting into the intricacies of the algorithms used by the AI program and the runtime logs showing access by those algorithms to one or more particular copyrighted works.

Thwarted in obtaining direct proof of infringement of the reproduction right, he will fall back on indirect proof by showing access and substantial similarity. Access will be far less of a problem. He only needs to show that his copyrighted work was included in the learning database employed by the accused program.

---

163. See *Sheldon v. Metro-Goldwyn Pictures Corp.*, 81 F.2d 49, 56 (2d Cir. 1936), *cert. denied*, 298 U.S. 669 (1936).

164. *Id.*

165. See 3 David Nimmer, *Nimmer on Copyright* § 13.03[B][1][a].

166. *Rogers*, 960 F.2d at 307–08.

167. Lucille M. Ponte, *The Emperor Has No Clothes: How Digital Sampling Infringement Cases Are Exposing Weaknesses in Traditional Copyright Law and the Need for Statutory Reform*, 43 AM. BUS. L.J. 515, 528–529 (2006) (arguing for new statutory protection for copyright owners of sampled music).

Substantial similarity is a much higher obstacle. The accused work is unlikely to be similar in its entirety to any particular work in the learning database, and even if it is similar, similarity is likely to involve only unprotected aspects of the copyrighted work.

Image generation systems are more prone to be liable for infringement of the reproduction right because they are more likely than text, music, or computer coding systems to extract entire images from the training database and present them as part of their output. For example, the author asked DALL-E2 to generate a picture of a DC-3 aircraft landing on a dirt airstrip for the cover of one of the author's novels. The program presents a photograph of a DC-3 over a dirt aircraft. It is obvious to casual observation that the depiction of the aircraft is an actual photograph rather than a rendering of the semantic characteristics of a DC-3 model. One can presume that the photographs are found somewhere in DALL-E2's learning database.

Substantial similarity also is at issue in adjudication of the derivative-work right.<sup>168</sup>

#### *e. Copying Incidental to Machine Learning*

The way in which machine learning works<sup>169</sup> means that a copyright infringement plaintiff is not going to be able to show verbatim copying as between copyrighted input owned by the plaintiff and particular output from the program. There might, however, be intermediate ephemeral copying in the early stages of the machine learning process as contents of that learning database are manipulated.

### *2. Preparation of Derivative Works*

The copyright owner has the exclusive right “to prepare derivative works based upon the copyrighted work.”<sup>170</sup> Commentators and courts disagree on whether the derivative-work right adds anything to the reproduction right, or is merely a historical anomaly.

Factually, generative AI outputs are derived from the systems learning databases. But “derivative work” is a term of art under the Copyright Act.<sup>171</sup> All works are in some sense derived from what is gone before; artists stand on the shoulders of their predecessors.<sup>172</sup> “[C]reating original content by reusing preexisting content is nothing new.”<sup>173</sup> So, a claim of infringement of the

---

168. See *infra* Section II.B.2.

169. See *supra* Section I.A, explaining machine learning and giving examples.

170. 17 U.S.C. § 106(2) (2018).

171. See 17 U.S.C. § 101 (2018) (defining derivative work as being “based upon” pre-existing work and giving examples of how a copyrighted work may be “recast, transformed, or adapted”).

172. *Rosenfeld v. W.B. Saunders*, 728 F. Supp. 236, 244 (S.D.N.Y. 1990) (using phrase in denying injunction in trademark infringement case).

173. Daniel Gervais, *The Tangled Web of UGC: Making Copyright Sense of User-Generated Content*, 11 VAND. J. ENT. & TECH. L. 841, 844 (2009).



derivative work right has to show that the new, infringing, work was derived from protectable elements of the preceding copyrighted work. To do that, the claimant must show close similarity, the standard for indirect proof of copying.<sup>174</sup>

In *Twin Peaks Productions, Inc., v. Publications International, Ltd.*,<sup>175</sup> the Second Circuit, quoting *Nimmer*, said that “the right to make derivative works is ‘completely superfluous’” when the plaintiff can prove infringement of one of the other exclusive rights, such as reproduction or performance.<sup>176</sup> It nevertheless went on to affirm the district court’s finding of derivative work.

Professor Gervais argues that the derivative work right should be disentangled from the reproduction right:<sup>177</sup> “[t]he distinction between reproduction and derivation is thus that, rather than being merely copied, parts of the protected expression of the primary work are transferred to, and transformed in, the product of the derivative use.”<sup>178</sup>

Put differently, operationally the two inquiries are often distinct; normatively, they almost always are. The qualitative part of the reproduction inquiry focuses chiefly on the form of what was taken while the derivation inquiry looks at a deeper level of appropriation, namely at whether the creative choices that made the primary work worthy of copyright protection were taken.<sup>179</sup>

He cites as an example the Harry Potter lexicon as an infringement of the derivative work right, but not the reproduction right.<sup>180</sup>

The derivative-work right also is deprived of independent operation by conflating its transformative elements with the transformative aspect of fair use analysis. Under this approach, if a work is transformative, it is derivative work, but it is privileged by fair use. If it is not privileged by fair use, it is not derivative. In *Clean Flicks v. Soderbergh*,<sup>181</sup> the district court entered summary judgment for the plaintiff movie studios on a suit for infringement against a

174. *But see* Daniel Gervais, *The Derivative Right, or Why Copyright Law Protects Foxes Better Than Hedgehogs*, 15 VAND. J. ENT. & TECH. L. 785, 800 (2013) [hereinafter *Hedgehogs*] (many U.S. courts apply the “substantial similarity” test to derivative works in the same way they do to copied works); *see* *Berkic v. Crichton*, 761 F.2d 1289, 1294 (9th Cir. 1985) (finding that movie was not substantially similar to novel and rejecting infringement of derivative work right claim). “We are confident that no reasonable reader or moviegoer would recognize either the book or movie *Coma* as a “picturization” or “dramatization” of *Reincarnation*.” *Id.*

175. *Twin Peaks Prods. v. Publ’ns Int’l, Ltd.*, 996 F.2d 1366 (2d Cir. 1993).

176. *Id.* at 1373. (quoting 2 *Nimmer* § 8.09[A], at 8-114); *see Hedgehogs, supra* note 174, at 799–800 (noting *Nimmer*’s agreement that current U.S. law subsumes derivative work right under reproduction right, thus leaving little scope for derivative work right).

177. *See generally Hedgehogs, supra* note 174.

178. *Id.* at 806.

179. *Id.* at 807 (internal footnotes omitted).

180. *Id.* at 841.

181. *Clean Flicks of Colo., LLC v. Soderbergh*, 433 F. Supp.2d 1236, 1237, 1243 (D. Colo. 2006).

defendant who had cleaned up copyrighted movies by removing scenes containing sex, profanity, and violence and redistributed them on DVDs. It held that the accused DVDs were not derivative works because they were not transformative, and “that [the] same question arises under the fair use defense.”<sup>182</sup> “It is undisputed that the edits are a small percentage of most of the films copied and the use is clearly for commercial gain. There is nothing transformative about the edited copies.”<sup>183</sup>

Anthony Reese critically evaluates the *Clean Flicks v. Soderbergh* case and explores the facts of other cases involving transformativeness and derivative works.<sup>184</sup> Several of them involved defendants who took only excerpts from plaintiff’s works and incorporated them into the defendant’s own, larger work.

Notwithstanding Professor Gervais’s effort to breathe new life into the derivative work right, distinguishing it from the reproduction right, it may be just as well to let sleeping dogs lie. The derivative work right developed when the reproduction right was understood much more narrowly than it is in the 21st century.<sup>185</sup> If the reproduction right was limited to literal copying or something very close to it, copyright would be unduly circumscribed, and so the derivative work right was invented to cover other forms of close appropriation that might not qualify as infringement of reproduction right at the time.<sup>186</sup> Now, the reproduction right is construed quite broadly, well beyond literal copying. Courts – and Gervais himself – struggle to carve out a role for the derivative work right that is different from the reproduction right without cutting back significantly on fair use.

The best solution is to avoid the intellectual acrobatics and conclude that the derivative work right adds little to the reproduction right. *Kellman v. Coca-Cola Co.*,<sup>187</sup> however, explains why there may be room for separate derivative-work claims. In that case, the accused work was a three-dimensional hat, and the pre-existing copyrighted works were purely decorative two-dimensional images. The hats were derivative works “for which there might be a separate demand that” the original copyright holder might someday see to exploit.<sup>188</sup> The differences in form by preclude a finding of infringement of the reproduction, display, or distribution right, leaving only infringement of the derivative-work right.

---

182. *Id.* at 1241.

183. *Id.*

184. See generally R. Anthony Reese, *Transformativeness and the Derivative Work Right*, 31 COLUM. J.L. & ARTS 467 (2008).

185. *Hedgehogs*, *supra* note 174, at 792–96.

186. *Id.*

187. *Kellman v. Coca-Cola Co.*, 280 F. Supp.2d 670 (E.D. Mich. 2003).

188. *Id.* at 677 (citing *Ty Inc. v. Publications Intern. Ltd.*, 292 F.3d 512, 519 (7th. Cir. 2002)).

*Warner Bros. Entertainment Inc. v. RDR Books*,<sup>189</sup> found sufficient copying to infringe the reproduction right. It rejected, however, the claim that the lexicon infringed the derivative work right:

Given that the Lexicon’s use of plot elements is far from an “elaborate recounting” and does not follow the same plot structure as the *Harry Potter* novels, Plaintiffs’ suggestion that these portions of the Lexicon are “unauthorized abridgements” is unpersuasive. Second, and more importantly, although the Lexicon “contain [s] a substantial amount of material” from the *Harry Potter* works, the material is not merely “transformed from one medium to another,” as was the case in *Twin Peaks*.<sup>190</sup> By condensing, synthesizing, and reorganizing the preexisting material in an A-to-Z reference guide, the Lexicon does not recast the material in another medium to retell the story of *Harry Potter*, but instead gives the copyrighted material another purpose. That purpose is to give the reader a ready understanding of individual elements in the elaborate world of *Harry Potter* that appear in voluminous and diverse sources. As a result, the Lexicon no longer “represents [the] original work[s] of authorship.”<sup>191</sup> Under these circumstances, and because the Lexicon does not fall under any example of derivative works listed in the statute, Plaintiffs have failed to show that the Lexicon is a derivative work.<sup>192</sup>

It also found that evaluation of the four factors failed to support a fair use defense, a part of the holding considered in § III.E.<sup>193</sup>

*Penguin Random House LLC v. Colting*,<sup>194</sup> distinguished the *Harry Potter* case and held that:

[A]dd[ing] additional material at the end, specifically a few brief pages of “Analysis,” “Quiz Questions,” and information about the author, [did not change the conclusion that] they are primarily dedicated to retelling plaintiffs’ stories. Two pages of analysis do not convert the Guides overall—which are largely composed of “Story Summaries”—into something that no longer “represents the original work of authorship.” Like a translation, dramatization, or motion picture adaptation (three categories explicitly delineated by Congress as derivative works).<sup>195</sup>

In another case, the court cautioned that mere similarity does not amount to derivation:

---

189. *Warner Bros. Ent. Inc. v. RDR Books*, 575 F. Supp. 2d 513 (S.D.N.Y. 2008).

190. *Twin Peaks Prods., Inc. v. Publ’ns Int’l, Ltd.*, 996 F.2d 1366, 1373 (2d. Cir. 1993).

191. 17 U.S.C. § 101 (2018).

192. *Warner Bros. Ent. Inc. v. RDR Books*, 575 F. Supp. 2d 513, 539 (S.D.N.Y. 2008).

193. *Id.* at 554.

194. *Penguin Random House LLC v. Colting*, 270 F. Supp.3d 736, 548 (S.D.N.Y. 2017).

195. *Id.* (internal citation omitted).

Best Brands' conclusory assertion that the Squeezamals are derivative works reflects a fuzzy understanding of copyright law. It is not enough to argue, as Best Brands do, that the Squeezamals are derivative works because they were generally "based on the hot trends from Japan." Nor is it sufficient to assert that Beverly Hills purportedly "wanted the designers to create products like the BC Mini and Silly Squishy" products or directed designers to look at those products "for inspiration." These facts do not show that Beverly Hills based its designs on anything more than the unprotectable *scenes a faire* of the plush toy world—*i.e.*, elements "which are as a practical matter indispensable, or at least standard, in the treatment of a given topic."<sup>196</sup> These indispensable elements belong to the public domain.<sup>197</sup>

The court denied referral to the Register of Copyright in an infringement action.<sup>198</sup>

The test for "substantial similarity of ideas" compares, not the basic plot ideas for stories, but the actual concrete elements that make up the total sequence of events and the relationships between the major characters. The extrinsic test for similarity of ideas looks beyond the vague, abstracted idea of a general plot and instead "focuses on . . . the objective details of the works . . . . The extrinsic test requires a comparison of plot, theme, dialogue, mood, setting, pace, and sequence."<sup>199</sup>

Asher Mitchell<sup>200</sup> reviews the history of the derivative work right<sup>201</sup> and argues that its application stifles creativity in folk music.<sup>202</sup> He argues for a less-inclusive ordinary observer test for substantial similarity.<sup>203</sup>

In *Micro Star v. Formgen Inc.*,<sup>204</sup> the Ninth Circuit reversed the district court and held that an accused work was an infringing derivative work, not entitled to fair use privilege. The court found that "the audiovisual displays generated

---

196. *Id.* at 747 (quoting *Hoehling v. Universal City Studios, Inc.*, 618 F.2d 972, 979 (2d Cir. 1980)).

197. *Beverly Hills Teddy Bear Company v. Best Brands Consumer Products, Inc.*, No. 19-cv-3766, 2020 U.S. Dist. LEXIS 99346, at \*8 (S.D.N.Y. June 4, 2020) (internal citations to record omitted).

198. *Id.* at \*2.

199. *Berkic v. Crichton*, 761 F.2d 1289, 1293–294 (9th Cir. 1985) (finding that movie was not substantially similar to novel and rejecting infringement of derivative work right claim).

200. Asher Mitchell, Note, *Cause We Don't Give a Darn: The Fundamental Tension Between Modern Copyright Law and American Folk Music*, 67 WASH. U. J. L. & POL'Y 443 (2022).

201. *Id.* at 446–451.

202. *Id.* at 451–456 (emphasizing the role of borrowing in folk music).

203. *Id.* at 464–65.

204. *Micro Star v. Formgen Inc.*, 154 F.3d 1107 (9th Cir. 1998).

when the player chooses the N/I levels come entirely out of D/N-3D's source art library."<sup>205</sup>

Micro Star further argues that the MAP files are not derivative works because they do not, in fact, incorporate any of D/N-3D's protected expression. In particular, Micro Star makes much of the fact that the N/I MAP files reference the source art library, but do not actually contain any art files themselves. Therefore, it claims, nothing of D/N-3D's is reproduced in the MAP files. In making this argument, Micro Star misconstrues the protected work. The work that Micro Star infringes is the D/N-3D story itself—a beefy commando type named Duke who wanders around post-Apocalypse Los Angeles, shooting Pig Cops with a gun, lobbing hand grenades, searching for medkits and steroids, using a jetpack to leap over obstacles, blowing up gas, avoiding radioactive slime. A copyright owner holds the right to create sequels, and the stories told in the N/I MAP files are surely sequels, telling new (though somewhat repetitive) tales of Duke's fabulous adventures. A book about Duke Nukem would infringe for the same reason, even if it contained no pictures.<sup>206</sup>

Dan Burke evaluates *Micro Star*<sup>207</sup> in the context of avatar creation by users of video games. He explains why application of the derivative work right is not easy:

Although the game elements displayed or heard during play are drawn from those provided by the game designers, the potential number of combinations is enormous. In controlling the game play, the player makes a large number of selections including which maps to traverse, what avatar movements to enact, what objects to interact with or acquire, what nonplayer characters (NPCs) to engage, and what comments to address to other players. In the kinds of games common to e-sports tournament play, it is unlikely that the designers foresaw all or even most of the possible play combinations.<sup>208</sup>

He suggests that the best solution is to consider that the work created by the user is a derivative work, the creation of which was authorized by the video game designer.<sup>209</sup>

Burk's solution will not work for most generative AI systems because the content of the learning databases was contributed, not by the system designed but by others to did something to make their content practically available on the

---

205. *Id.* at 1112.

206. *Id.* at 1112 (internal footnotes and citations omitted).

207. Dan L. Burk, *Owning E-Sports: Proprietary Rights in Professional Computer Gaming*, 161 U. PA. L. REV. 1535, 1547 (2013).

208. *Id.* at 1548.

209. *Id.* at 1549.

Internet, or otherwise. In some cases, they consented to use of their content by the generative AI system; in most cases they did not explicitly consent.

If the owners of copyrights in all of the samples used in a learning database band together and sue for infringement of their derivative-work rights, their claims might be more interesting. While they could not show that any particular output was the derivative work of any particular input, they could show that it was the derivative work of the inputs collectively.

Suppose someone took two short stories from the same author and created a new work by interposing sentences from the two-pre-existing works alternatively. The resulting selection and arrangement would be bizarre and probably not very pleasing, but it likely would satisfy the test of originality under *Feist*. It also would seem to infringe the derivative work right of the author of the two-pre-existing works. A derivative work is something *based on* a pre-existing work, and the hypothetical mishmash surely is based on the two works whose language comprises it.

Now, suppose the pre-existing works are owned, not by a single author but by multiple authors. The latecomer, as before, systematically interposes sentences taken alternatively from each of the pre-existing works. It is not clear why the introduction of multiple pre-existing works by multiple authors should defeat the derivative work characterization of the latecomer's creation.

In *Video Pipeline, Inc. v. Buena Vista Home Entertainment, Inc.*,<sup>210</sup> the district court held that clips taken from copyrighted movie trailer, digitized, and shown on Internet Web sites infringed the derivative-work rights of the copyright holders.<sup>211</sup> Because each of the clip previews and trailers was a derivative work of a copyrighted movie, the collection of clips was a derivative work and thus infringing. However, “[u]se of merely the still photographic images of the copyrighted movies [would cause] its argument [to] hold more force.”<sup>212</sup>

*Rogers v. Koons*, involved a sculpture copied from a photograph of a group of puppies.<sup>213</sup> The court of appeals affirmed a permanent injunction against the further display or possession of the infringing sculpture.<sup>214</sup> The court of appeals found both direct and indirect evidence of copying.<sup>215</sup>

Appellant gave his artisans one of Rogers' notecards and told them to copy it. But in order to guide the creation of a three-dimensional

---

210. In *Video Pipeline v. Buena Vista Home Entertainment, Inc.*, 342 F.3d 191 (3d Cir. 2003), the Third Circuit affirmed an earlier preliminary injunction in the case, finding that the accused digital clips did not qualify as fair use because the clips “[would] likely serve as substitutes” for the copyrighted works. See *Video Pipeline v. Buena Vista Home Entertainment, Inc.*, 342 F.3d 191, 199 (3d Cir. 2003).

211. *Id.* at 559.

212. *Id.*

213. *Rogers v. Koons*, 960 F.2d 301, 303 (2d Cir. 1992).

214. *Id.* at 306.

215. *Id.* at 307 (observing that, even without direct evidence, infringement could have been found from access and substantial similarity).

sculptural piece from the two-dimensional photograph, Koons communicated extensively with the Demetz Studio. He visited it once a week during the period the piece was being carved by the workers and gave them written instructions. In his “production notes” Koons stressed that he wanted “Puppies” copied faithfully in the sculpture. For example, he told his artisans the “*work must be just like photo—features of photo must be captured;*” later, “*puppies need detail in fur. Details—Just Like Photo!*”; other notes instruct the artisans to “*keep man in angle of photo—mild lean to side & mildly forward—same for woman,*” to “*keep woman’s big smile,*” and to “*keep [the sculpture] very, very realistic;*” others state, “*Girl’s nose is too small. Please make larger as per photo;*” another reminds the artisans that “*The puppies must have variation in fur as per photo—not just large area of paint—variation as per photo.*” (emphasis supplied).

To paint the polychromed wood “String of Puppies” sculptures, Koons provided a chart with an enlarged photocopy of “Puppies” in the center; painting directions were noted in the margin with arrows drawn to various areas of the photograph. The chart noted, “Puppies, painted in shades of blue. Variation of light-to-dark *as per photo*. Paint realistic *as per photo*, but in blues.” and “Man’s hair, white with shades of grey *as per black and white photo!* (emphasis supplied)<sup>216</sup>

The district court’s original decision in the case<sup>217</sup> characterized the accused sculpture as a derivative work and the plaintiff’s claim as infringement of the derivative-work right.<sup>218</sup> Nothing in the revised district court opinion<sup>219</sup> or in the appellate opinion<sup>220</sup> disagrees with that characterization.

Generative AI systems do this kind of derivation writ large. But they go about it differently, and that difference may defeat the derivative-work argument. Generative AI systems piece together, not verbatim fragments of the works in their learning databases; they extract features from those works and reassemble the features into new works. That seems more like someone who takes a novel or play and critically dissects general characteristics of plot, characters, and setting and writes a screenplay.

Judge Hand’s abstractions test<sup>221</sup> can decide whether the features taken from the underlying works constitute protected expression and thus whether incorporating them into the new work constitutes infringement.<sup>222</sup> Application of the abstractions test in the practical context of a generative AI system requires

---

216. *Id.* at 305.

217. *Rogers v. Koons*, 751 F. Supp. 474, 477 (S.D.N.Y. 1990).

218. *Id.* at 478.

219. *Rogers v. Koons*, 777 F. Supp. 1 (S.D.N.Y. 1991).

220. *Rogers v. Koons*, 960 F.2d 301 (2d Cir. 1992).

221. *See Nichols v. Universal Pictures Corp.*, 45 F.2d 119, 121 (2d Cir. 1930).

222. *See infra* text accompanying notes 245–251.

concrete comparison between features of specific works in the learning database with similar features in the AI systems output.

But there is another difference between generative AI systems and the facts of cases like *Video Pipeline*. In *Video Pipeline*, the accused images could be matched, one for one, with the pre-existing images. Not so with generative AI machine learning. Generative AI systems learn by synthesizing millions of input works and creating new works from the characteristics of that synthesis. The computerized process resembles a musician or an author who absorbs the hallmarks of an entire genre and then creates a new member of that genre.

The plaintiff claiming infringement of the derivative work right is going to have a difficult time. If the scope of derivative work is broad enough to reach the highly transformative output of a generative AI program, the degree of transformativeness likely supports a fair use defense. In the unusual case where the AI program simply spits out something from the learning database more or less intact, an infringement plaintiff not only will easily avoid a fair use defense but also have a strong case of infringement of the reproduction right as well as the derivative work right.

### 3. *Distribution*

The copyright owner has the exclusive right “to distribute copies or phonorecords of the copyrighted work to the public by sale or other transfer of ownership, or by rental, lease, or lending.”<sup>223</sup> The distribution right is implicated in robot creation whenever the output is made available through the Internet.

### 4. *Performance*

The owner of a copyright in a literary, musical, dramatic or choreographic work, pantomime, or motion pictures and other audiovisual works has the exclusive right “to perform the copyrighted work publicly.”<sup>224</sup> Material published on the website or another Internet-accessible source constitutes public display and public distribution, and so a plaintiff unable to show infringement of the reproduction right or of the derivative work right might have an easier time showing infringement of the display and distribution right. But to do so, she must particularize the identification of both underlying and accused works and establish substantial similarity and defeat fair use, just as she would with respect to the reproduction and derivative work rights.

Infringement of the performance right is unlikely with respect to literary works, but more likely with respect to musical, photographic, and video works. But the performance right is infringed only if a performance to the public occurred. As § III.G points out, the plaintiff is going to have to show that her particular copyrighted work, or protected elements of it, was performed by the

---

223. 17 U.S.C. § 106(3) (2018).

224. 17 U.S.C. § 106(4) (2018).



generative AI system. This likely would occur with computer-generated musical or audio-visual works.

### 5. Display

The owner of a copyright in a literary, musical, dramatic, or choreographic work, pantomime, or pictorial, graphic, or sculptural work, including the individual images of a motion picture or other audiovisual work, the exclusive right “to display the copyrighted work publicly.”<sup>225</sup>

In *Perfect 10, Inc. v. Amazon.com, Inc.*,<sup>226</sup> the Ninth Circuit considered a claim that an Internet search engine infringed the display right and the distribution right of the owner of a copyright in nude images by displaying thumbnails of the images.

The image stored in the computer is the “copy” of the work for purposes of copyright law . . . The computer owner shows a copy “by means of a . . . device or process” when the owner uses the computer to fill the computer screen with the photographic image stored on that computer, or by communicating the stored image electronically to another person’s computer.<sup>227</sup> In sum, based on the plain language of the statute, a person displays a photographic image by using a computer to fill a computer screen with a copy of the photographic image fixed in the computer’s memory. There is no dispute that Google’s computers store thumbnail versions of Perfect 10’s copyrighted images and communicate copies of those thumbnails to Google’s users. Therefore, Perfect 10 has made a prima facie case that Google’s communication of its stored thumbnail images directly infringes Perfect 10’s display right.<sup>228</sup>

Because Google did not communicate copies of the full-sized images, but only html code enabling a user to find them, in *Perfect 10, Inc. v. Amazon.com, Inc.*,<sup>229</sup> Google did not infringe either the display or the distribution right with respect to the full images.<sup>230</sup>

Infringement of the display right depends on how the robot-generated information was communicated and presented. If it merely is presented on one user’s computer screen or on one hard copy, the statutory element of “to the public” is not satisfied. On the other hand, if it was put up on a website or made available through various video entertainment channels, it has been displayed publicly. As § III.G points out, the plaintiff must show that her particular work, or protected elements of it, was displayed.

---

225. 17 U.S.C. § 106(5) (2018).

226. *Perfect 10 Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1154 (9th Cir. 2007).

227. 17 U.S.C. § 101 (2018).

228. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d at 1160 (some internal citations omitted).

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

230. *Id.* at 1161–63.

### C. Identification of Particular Works

Regardless of what kind of infringement is claimed, a claim of infringement must begin by particularized identification of the protected work. That must be done even before establishing ownership of copyright—usually identified as the first step in copyright infringement analysis.<sup>231</sup> If a plaintiff claims that someone copied her work, she must define the work that was copied. If she claims the someone prepared a derivative work without her permission, she must define the work from which the accused work was derived. If she claims infringement of the performance right, she must show what was performed. Likewise, if she claims infringement of the display or distribution right she has to show what particular work of hers was distributed or displayed.

In *ADR International Limited v. Institute for Supply Management Inc.*, the court cited cases holding that copyright infringement plaintiffs had failed to identify infringing works with sufficient particularity, and distinguished them.<sup>232</sup> *Warner Bros. Entertainment Inc. v. RDR Books*,<sup>233</sup> involved an encyclopedia (a “lexicon”) of information about Harry Potter fiction. The court described the relationship between the accused encyclopedia and the copyrighted Harry Potter novels:

The Lexicon entries cull every item and character that appears in the *Harry Potter* works, no matter if it plays a significant or insignificant role in the story. The entries cover every spell (e.g., Expecto Patronum, Expelliarmus, and Incendio), potion (e.g., Love Potion, Felix Felicis, and Draught of Living Death), magical item or device (e.g., Deathly Hallows, Horcrux, Cloak of Invisibility), form of magic (e.g., Legilimency, Occlumency, and the Dark Arts), creature (e.g., Blast-Ended Skrewt, Dementors, and Blood-Sucking Bugbears), character (e.g., Harry Potter, Hagrid, and Lord Voldemort), group or force (e.g., Aurors, Dumbledore’s Army, Death Eaters), invented game (e.g., Quidditch), and imaginary place (e.g., Hogwarts School of Witchcraft and Wizardry, Diagon Alley, and the Ministry of Magic) that appear in the *Harry Potter* works. The Lexicon also contains entries for items that are not explicitly named in the *Harry Potter* works but which Vander Ark has identified, such as medical magic, candle magic, wizard space, wizard clothing, and remorse. Some of the entries describe places or things that exist in the

---

231. *William v. Gaye*, 895 F.3d 1106, 1119 (9th Cir. 2018) (internal citation omitted) (listing ownership as first element of infringement claim).

232. *See ADR Int’l Ltd. v. Inst. for Supply Mgmt. Inc.*, No. 22-CV-1914, 2023 U.S. Dist. LEXIS 58090 (S.D. Tex. Mar. 30, 2023); *see also ADR Int’l Ltd. v. Inst. for Supply Mgmt. Inc.*, No. 4:22-CV-1914, 2023 U.S. Dist. LEXIS 58090, at \*8–11 (S.D. Tex. Feb. 24, 2023) (internal citations omitted).

233. *Warner Bros. Ent. Inc. v. RDR Books*, 575 F. Supp. 2d 513, 520 (S.D.N.Y. 2008).

real world but also have a place in the *Harry Potter* works, such as moors, Greece, and Cornwall.

Each entry, with the exception of the shortest ones, gathers and synthesizes pieces of information relating to its subject that appear scattered across the *Harry Potter* novels, the companion books, The Daily Prophet newsletters, Famous Wizard Cards, and published interviews of Rowling. The types of information contained in the entries include descriptions of the subject's attributes, role in the story, relationship to other characters or things, and events involving the subject. Repositories of such information, the entries seek to give as complete a picture as possible of each item or character in the *Harry Potter* world, many of which appear only sporadically throughout the series or in various sources of *Harry Potter* material.

The snippets of information in the entries are generally followed by citations in parentheses that indicate where they were found within the corpus of the *Harry Potter* works. The thoroughness of the Lexicon's citation, however, is not consistent; some entries contain very few citations in relation to the amount material provided. When the Lexicon cites to one of the seven *Harry Potter* novels, the citation provides only the book and chapter number. Vander Ark explained that page numbers were excluded from the citations because the various editions of the *Harry Potter* books have different pagination, but the chapter numbers remain consistent. The Lexicon neither assigns a letter to each edition nor specifies a standard edition while providing a conversion table for other editions, practices which Plaintiffs' expert Jeri Johnson testified were common for reference guides.

While not its primary purpose, the Lexicon includes commentary and background information from outside knowledge on occasion. For example, the Lexicon contains sporadic etymological references, analogies to characters outside the *Harry Potter* world such as Merlin, and observations of Rowling's allusions to other works of literature such as "the weird sisters" from Shakespeare's *Macbeth*. The Lexicon also points to the very few "flints," or errors in the continuity of the story, that appear in the *Harry Potter* series.<sup>234</sup>

After further analysis of the details of the two works, the court said:

Although it is difficult to quantify how much of the language in the Lexicon is directly lifted from the *Harry Potter* novels and companion books, the Lexicon indeed contains at least a troubling amount of direct quotation or close paraphrasing of Rowling's original language. The Lexicon occasionally uses quotation marks to indicate Rowling's language, but more often the original language is copied

---

234. *Id.* at 525–26 (internal citations omitted).

without quotation marks, often making it difficult to know which words are Rowling's and which are Vander Ark's.<sup>235</sup>

The court found sufficient copying to constitute infringement.<sup>236</sup> Its analysis illustrates particularity in dissecting the elements of both protected and accused works.

#### D. Only in Protected Expression

##### 1. Not Ideas and Facts

The Copyright Act explicitly excludes facts and ideas from copyright protection. Section 102(b) says, “[i]n no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in such work.”<sup>237</sup> *Feist Publications, Inc. v. Rural Telephone Service Co.*<sup>238</sup> explicitly says that facts are not copyrightable.<sup>239</sup>

The robot-generated story about Philander says that the AI system he accuses of infringement “copied some of the key phrases and ideas.”<sup>240</sup> The first phrase describes verbatim copying. That would, indeed, be infringement, but only if what was copied is protectable. If the robot copied only an idea or mere facts, copyright protection does not extend to it.<sup>241</sup>

##### 2. Scenes A Faire, Abstraction, and Filtration

The exclusion of ideas and facts from copyright reinforces the requirement that the expression be “original” to qualify for copyright protection,<sup>242</sup> another statutory requirement mandated by the Constitution.<sup>243</sup> Unprotectable ideas often are referred to as “scenes a faire.”

This idea/expression dichotomy received attention from Judge Learned Hand in *Nichols v. Universal Pictures Corp.*<sup>244</sup> The case involved a claim by the author of a play that a motion picture infringed the copyright in the play. Judge Hand, finding no infringement, began his analysis by noting that infringement

---

235. *Id.* at 527.

236. *Id.* at 535–36.

237. 17 U.S.C. §102(b) (2018).

238. *Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340 (1991).

239. *Id.* at 344.

240. *See supra*, note 1.

241. 17 U.S.C. § 102(b) (2018) (excluding from copyright protection “any idea, procedure, process, system, method of operation, concept, principle, or discovery”).

242. *See Nichols v. Universal Pictures Corp.*, 45 F.2d 119, 122 (2d Cir. 1930) (noting that extending scope of copyright too far into realm of ideas would clash with “originality” requirement).

243. *See* 17 U.S.C. § 102(a) (2018).

244. *Nichols v. Universal Pictures Corp.*, 45 F.2d 119, 120 (2d Cir. 1930).

can be shown without proving literal copying. “[E]lse a plagiarist would escape by immaterial variations,” he said.<sup>245</sup>

But when the plagiarist . . . [takes] an abstract of the whole, decision is more troublesome. Upon any work, and especially upon a play, a great number of patterns of increasing generality will fit equally well, as more and more of the incident is left out. The last may perhaps be no more than the most general statement of what the play is about, and at times might consist only of its title; but there is a point in this series of abstractions where they are no longer protected, since otherwise the playwright could prevent the use of his “ideas,” to which, apart from their expression, his property is never extended.<sup>246</sup>

A copyright, and therefore infringement, might extend to plot or characters, but “the less developed the characters, the less they can be copyrighted.”<sup>247</sup> Therefore, specificity of the candidate work is an important touchstone for copyrightability.

In *Nash v. CBS*,<sup>248</sup> Judge Easterbrook characterized Judge Hand’s formulation as the “abstractions test,” which, he said, dominates this area of copyright law.<sup>249</sup>

[It] all depends on the level of abstraction at which the court conceives the interest protected by the copyright. If the court chooses a low level (say, only the words the first author employed), then a copier may take the plot, exposition, and all other original material, even though these may be the most important ingredients of the first author’s contribution. As a practical matter this would mean that anyone could produce the work in a new medium without compensating the original author, despite the statute’s grant to the author of the privilege to make “derivative works.” If on the other hand the court should select a high level of abstraction, the first author may claim protection for whole genres of work (“the romantic novel” or, more modestly, any story involving doomed young lovers from warring clans, so that a copyright on *Romeo and Juliet* would cover *West Side Story* too). Even a less sweeping degree of abstraction creates a risk of giving copyright protection to “the idea” although the statute protects only “expression.”<sup>250</sup>

Distinguishing ideas from expression is necessary not only in the traditional literary and dramatic context but also in the context of music and computer software.

---

245. *Id.* at 121.

246. *Id.* (internal citation omitted).

247. *Id.*

248. *Nash v. CBS, Inc.*, 899 F.2d 1537 (7<sup>th</sup> Cir. 1990).

249. *Id.* at 1540.

250. *Id.* (internal citations omitted).

This idea/expression dichotomy achieved constitutional dimensions in *Feist Publications, Inc. v. Rural Telephone Service Co.*<sup>251</sup> The *Feist* Court reasoned that the constitutional justification for the monopolies represented by copyright and patent protection is need for economic incentives for inventor and author activity.<sup>252</sup> Unless someone has added value to what was already there, he has done nothing qualifying for an incentive. Facts and raw ideas are part of the stock of public domain information and are not entitled to copyright protection.<sup>253</sup> Under the Copyrights and Patents Clause of the Constitution, Congress lacks the constitutional power to extend protection to such public domain information.<sup>254</sup>

Moreover, the contribution of a compiler is not entitled to protection under the statute unless it is an “original work of authorship.”<sup>255</sup> The work of a compiler is not original unless there is something about the selection or arrangement of the facts making up the compilation that is creative. In *Feist*, the Supreme Court found that the alphabetical arrangement of a telephone book and the selection of its elements—the name, address, and telephone number—were dictated by its nature and thus constituted neither original arrangement nor selection.<sup>256</sup>

The robot-generated story about Philander says that the AI system he accuses of infringement used “patterns and structures it had learned from the input data.” Under *Feist*, “patterns and structures,” the third key phrase in the story, may constitute selection and arrangement of facts that are protectable by copyright, even if the underlying facts are outside protection.

In *Lexmark International, Inc. v. Static Control Components, Inc.*,<sup>257</sup> the Sixth Circuit explored the scenes-a-faire doctrine in the context of an infringement claim involving laser printer cartridge interfaces. It began by translating the French phrase and explaining its meaning with respect to literary works:

[W]hen external factors constrain the choice of expressive vehicle, the doctrine of “scenes a faire”—“scenes,” in other words, “that must be done”—precludes copyright protection. In the literary context, the doctrine means that certain phrases that are standard, stock, . . . or that necessarily follow from a common theme or setting” may not obtain copyright protection.<sup>258</sup>

It explained that the concept extends to works involving computer technology:

---

251. *Feist Publ’ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 349-350 (1991).

252. *Id.*

253. *Id.* at 358.

254. *Id.* at 361.

255. 17 U.S.C. § 102(a) (2018).

256. *Feist Publ’ns.*, 499 U.S. at 361.

257. *Lexmark Int’l, Inc. v. Static Control Components, Inc.*, 387 F.3d 522, 528 (6th Cir. 2004) (vacating preliminary injunction because of failure to show likelihood of success on merits of copyright infringement claim).

258. *Id.* at 535 (internal citations omitted).

In the computer-software context, the doctrine means that the elements of a program dictated by practical realities—*e.g.*, by hardware standards and mechanical specifications, software standards and compatibility requirements, computer manufacturer design standards, target industry practices, and standard computer programming practices—may not obtain protection. As “an industry-wide goal,” programming “[e]fficiency” represents an external constraint that figures prominently in the copyrightability of computer programs.<sup>259</sup>

In *Shaw v. Lindheim*,<sup>260</sup> the district court on remand after appeal entered judgment for the defendant, finding that similarities between the accused work and the underlying copyrighted works constitutes “ideas, concepts, and other unprotected elements.”<sup>261</sup> Similar issues in the music context were litigated in the copyright infringement trial between Ed Sheeran and Marvin Gaye in the United States District Court in New York.

The idea/expression dichotomy and the scenes-a-faire doctrine go to the heart of generative AI system functionality. These systems work, not by verbatim copying of phrases, imagery, or sounds from their learning databases, but by synthesizing semantic concepts from those underlying works. It would not be too far off the mark to characterize deep learning systems as extractors of scenes-a-faire or corallers of ideas.

Such a characterization obviously points to the conclusion that the systems should not be liable for copyright infringement, because they have not copied protectable expression (or displayed it, performed, or distributed it). That does not exclude the possibility that their use of the copyrighted material in the learning databases might be characterized as preparation of derivative works, but even the most expansive understanding of the concept of derivative work cannot result in infringement being attached to anything that is not protectable expression in the derivative work.

The possibility remains, of course, that a plaintiff might show, in a particular case, that a system has done more than appropriate her ideas and scenes-a-faire. If she can show appropriation of specific protectable expression from specific works in specific outputs of the system, she has established a *prima facie* case of infringement.

#### *E. Vicarious Infringement*

Individuals are entities with control over intermediaries may be contributorily or vicariously liable for intermediary infringement. This possibility is suggested by the phrase, “the AI was an independent agent, and that they had no control over what it generated” in the robot generated story about Philander.<sup>262</sup>

---

259. *Id.*

260. *Shaw v. Lindheim*, 809 F. Supp. 1393 (C.D. Cal. 1992).

261. *Id.* at 1402.

262. *See supra*, note 1.

The Supreme Court's decision in *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*<sup>263</sup> explained that secondary liability for copyright infringement comprises two branches: *contributory* infringement and *vicarious* infringement. Contributory infringement arises when the copyright owner can prove intent that an intermediary engaged in direct infringement.<sup>264</sup> Vicarious infringement arises when a copyright owner can show that the defendant profited from intermediary conduct that he knew involved direct infringement.<sup>265</sup>

The robot-generated story about Philander had him overcoming the defense that the defendant had no control over what his robot did. That would not necessarily be enough to avoid secondary liability. To recover for contributory infringement, the fictional Philander will have the burden of presenting evidence that the operator of the generative AI system intended that it infringe copyright.

To recover for vicarious infringement, Philander must show that the user of the system knew that it was likely to infringe copyrights like Philander's and that the user profited from use of the system.

Under *Sony Corporation v. Universal City Studios, Inc.*,<sup>266</sup> "the sale of copying equipment, like the sale of other articles of commerce, does not constitute contributory infringement if the product is widely used for legitimate, unobjectionable purposes. Indeed, it need merely be capable of substantial noninfringing uses."<sup>267</sup> The Sony court applied this standard to a determination of vicarious liability.<sup>268</sup>

So, the user of a generative AI system can avoid vicarious liability by showing that the intermediary is capable of noninfringing uses. In *Grokster*, however, the Supreme Court held that those noninfringing uses must be more than trivial and that the defendant must not have intended infringement.<sup>269</sup> *Grokster*'s argument failed because the evidence showed that it actively promoted use of the computer system to infringe copyrights and music, and that that's mostly what the system did.

The *Grokster* Court understood *Sony* as standing for the proposition that when a system is "'good for nothing else' but infringement," one can infer an intent to

---

263. *MGM Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 929–930 (2005) (internal citations omitted).

264. *Id.* at 930 ("One infringes contributorily by intentionally inducing or encouraging direct infringement") (internal citations omitted).

265. *Id.* (One "infringes vicariously by profiting from direct infringement while declining to exercise a right to stop or limit it").

266. *Sony Corp of Am. v. Universal City Studios, Inc.*, 464 U.S. 417 (1984) (holding that manufacturers of video recorders were not secondarily liable for infringement of copyrights in movies).

267. *Id.* at 442.

268. *Id.* at 438–39 (characterizing the case as involving vicarious, not contributory, infringement claim).

269. *MGM Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 933 (discussing "substantial" noninfringing use; i.e. non-trivial); *id.* at 936–937 (discussing intent).



infringe on its user.<sup>270</sup> On the other hand, “the doctrine absolves the equivocal conduct of selling an item with substantial lawful as well as unlawful uses, and limits liability to instances of more acute fault than the mere understanding that some of one’s products will be misused.”<sup>271</sup> “Thus, where evidence goes beyond a product’s characteristics or the knowledge that it may be put to infringing uses, and shows statements or actions directed to promoting infringement, *Sony’s* staple-article rule will not preclude liability.”<sup>272</sup> It found the record of the *Grokster* case itself established that “the unlawful objective is unmistakable.”<sup>273</sup> If the user of the generative AI system profits from its use and knows that it infringes copyrights like those belonging to Philander, *Grokster* says it is liable for vicarious infringement. But *Sony* says that if the system is capable of substantial noninfringing use, and there is no intent, vicarious liability does not arise.

It is conceivable that a generative AI system could be designed to extract protectable expression from works in its learning database and to reassemble them and present them as its own. Proof of such a design would set up a very strong case of vicarious liability, because it would seem to exclude the possibility of non-infringing uses, and imply an intent for the system to be used for infringement.

But not that is not the way most generative AI systems work, and most of them generate non-infringing expression far more often than they generate infringing expression, that characteristic would trigger the *Sony* defense and spare them from vicarious liability—unless there is proof of infringing intent by the user. *Grokster* makes it clear that intent, in addition to mere awareness of the possibility of infringing uses is necessary, “where evidence goes beyond a product’s characteristics or the knowledge that it may be put to infringing uses, and shows statements or actions directed to promoting infringement, *Sony’s* staple-article rule will not preclude liability.”<sup>274</sup>

The Stability AI lawsuit includes a claim for vicarious copyright infringement.<sup>275</sup> It alleges that the ability of the accused system “to respond to prompts containing specific artists’ names was designed by defendants with either the knowledge of or reckless disregard for the fact that this functionality could easily be used to create fakes.”<sup>276</sup> It alleges that the defendants “have directly and indirectly profited from acts of infringement by imposters.”<sup>277</sup> It thus touches the bases required by *Grokster*.

---

270. *MGM Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 932 (2005) (internal citation omitted).

271. *Id.* at 932–33.

272. *Id.* at 935.

273. *Id.* at 937–40.

274. *Id.* at 935.

275. Stability AI Complaint at ¶¶ 169–177.

276. *Id.* at ¶ 172.

277. *Id.* at ¶ 176.

### F. Fair Use: Four Factors

Even if a copyright-infringement plaintiff can establish violation of one of her exclusive rights under section 106<sup>278</sup> the defendant may nevertheless escape liability if he can establish *Fair Use*. The statute codifies the Fair Use privilege, and mandates consideration of four factors:

- (1) the purpose and character of the use, *including* whether such use is of a commercial nature or is for nonprofit educational purposes;
- (2) the nature of the copyrighted work;
- (3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
- (4) the effect of the use upon the potential market for or value of the copyrighted work.<sup>279</sup>

The legislative history of section 107 and its post-enactment judicial interpretations make it clear that the four fair-use factors are neither collectively exhaustive nor mutually exclusive.<sup>280</sup> The nature of the accused work in the nature of the copyrighted work impact market effect. Likewise, the amount of copyrighted material taken for the accused work impacts market effect. The *Warner Bros. Entertainment Inc. v. RDR Books*,<sup>281</sup> court found that evaluation of the four factors failed to support a fair use defense for the Harry Potter lexicon.<sup>282</sup>

In *Google LLC v. Oracle America, Inc.*, the Supreme Court extended the fair-use privilege to Google for copying copyrighted lines of code from Oracle's Java SE programming platform.<sup>283</sup> The copying was necessary to allow Google users to employ familiar Java calls for Google's Android cellphone operating system.

The Court summarized its holding as follows: "[W]here Google . . . [took] only what was needed to allow users to put their accrued talents to work in a new and transformative program, Google's copying of the Sun Java API was a fair use of that material as a matter of law."<sup>284</sup>

#### 1. Purpose and Character of the Use & Transformativeness

Outputs from generative AI programs are highly transformative in both content and purpose, and so the first factor tilts in favor of the generative AI output. Outputs of generative AI programs are already being used to earn

278. 17 U.S.C. § 106 (2018).

279. 17 U.S.C. § 107 (2018).

280. See *N. Jersey Media Grp. Inc. v. Pirro*, 74 F. Supp. 3d 605, 614 (S.D.N.Y. 2015) (explaining how factors are to be considered and weighed together); see also *Sony Corp. v. Universal City Studios, Inc.*, 464 U.S. 417, 476 (1984) (Blackmun, J., dissenting) (noting the list of factors is not intended to be exclusive).

281. *Warner Bros. Ent. Inc. v. RDR Books*, 575 F. Supp. 2d 513 (S.D.N.Y. 2008).

282. *Id.* at 551.

283. *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183, 1190 (2021).

284. *Id.* at 1209.

revenue and make a profit. But commercial, as opposed to private use, is not fatal to a fair use defense.

The caselaw supports this conclusion. On the first of the four fair use factors, the purpose and character of the use, the *Google*,<sup>285</sup> Court found Google's copying to be transformative:

Here Google's use of the Sun Java API seeks to create new products. It seeks to expand the use and usefulness of Android-based smartphones. Its new product offers programmers a highly creative and innovative tool for a smartphone environment. To the extent that Google used parts of the Sun Java API to create a new platform that could be readily used by programmers, its use was consistent with that creative "progress" that is the basic constitutional objective of copyright itself.<sup>286</sup>

\* \* \*

These and related facts convince us that the "purpose and character" of Google's copying was transformative—to the point where this factor too weighs in favor of fair use.<sup>287</sup>

In *Fox News Network, LLC v. TVEyes, Inc.*, the district court found that a news aggregator's scraping and organization of copyrighted news stories constituted fair use.<sup>288</sup> The defendant admitted verbatim copying of plaintiff's copyrighted content.

The court distinguished *Associated Press v. Meltwater U.S. Holdings, Inc.*,<sup>289</sup>:

TVEyes distinguishes itself from those cases by the different character of its database. Print is fixed in form, and regularly available from publishing sources and archives. A service that provides clipping of news articles and columns provides essentially the same service as could be provided by the content provider itself. TVEyes, however, is not a clipping service for print. TVEyes' search results show the combination of visual images and text in a medium that raises the commentator to have the qualities of news itself. The focus of certain programs and talk shows on President Obama's recent golf vacation, for example, was as much the news as the beheading of an American reporter. The actual images and sounds depicted on television are as important as the news information itself—the tone of voice, arch of an eyebrow, or upturn of a lip can color the entire story, powerfully

---

285. *Id.* at 1183.

286. *Id.* at 1203 (quoting *Feist Publ'ns Inc. v. Rural Tel. Serv. Co.*, 499 U. S. 340, 349 (1991) ("The primary objective of copyright is not to reward the labor of authors, but 'to promote the Progress of Science and useful Arts.'" (quoting U. S. Const., Art. I, § 8, cl. 8))).

287. *Id.* at 1204.

288. *Fox News Network, LLC v. TVEyes, Inc.*, 43 F. Supp. 3d 379, 400 (S.D.N.Y. 2014).

289. *AP v. Meltwater U.S. Holdings, Inc.*, 931 F. Supp. 2d 537 (S.D.N.Y. 2013). The following two paragraphs in the text, *infra*, distinguish *Meltwater*.

modifying the content. The service provided by TVEyes, indexing and collecting visual and audio images, allows subscribers to categorize, not only content in the response to key search words, but also “information [that] may be just as valuable to [subscribers] as the [content], since a speaker’s demeanor, tone, and cadence can often elucidate his or her true beliefs far beyond what a stale transcript or summary can show.”<sup>290</sup> Unlike the indexing and excerpting of news articles, where the printed word conveys the same meaning no matter the forum or medium in which it is viewed, the service provided by TVEyes is transformative. By indexing and excerpting all content appearing in television, every hour of the day and every day of the week, month, and year, TVEyes provides a service that no content provider provides. Subscribers to TVEyes gain access, not only to the news that is presented, but to the presentations themselves, as colored, processed, and criticized by commentators, and as abridged, modified, and enlarged by news broadcasts.

There also is a second relevant distinction that makes the district judge’s opinion in *Meltwater* less helpful to deciding the disposition here. *Meltwater* aggregated content already available to the individual user who was willing to perform enough searches and cull enough results on the Internet. The service provided simply “crawled” the Internet, gathering extant content. TVEyes, however, creates a database of otherwise unavailable content. TVEyes is the only service that creates a database of *everything* that television channels broadcast, twenty-four hours a day, seven days a week. The Internet does not and cannot house the entirety of this content because Fox News, for example, does not provide all of its content online.

Thus, without TVEyes, this information cannot otherwise be gathered and searched. That, in and of itself, makes TVEyes’ purpose transformative and different in kind from *Meltwater*’s, which simply amalgamated extant content that a dedicated researcher could piece together with enough time, effort, and Internet searches. These differences further reduce the persuasive value of the district court opinion in *Meltwater*.<sup>291</sup>

Generative AI systems are highly transformative in both content and purpose, as § III.F points out. They are much more than search and retrieval tools providing hooks to particular artifacts of content. The most sophisticated generative AI applications just began appearing in the market in 2022, and is premature to say with any confidence how they will be used and how markets for them will develop after the excitement dies down and prices for their use stabilize.

---

290. *Fox News Network*, 43 F. Supp. 3d at 392–93 (quoting *The Swatch Grp. Mgmt. Ltd. v. Bloomberg L.P.*, 756 F.3d 73, 84 (2d Cir. 2014)).

291. *Fox News Network*, 43 F. Supp.3d at 392–393.

It appears likely, however, based on the areas of most intense exploration and discussion, that the systems will be used to generate routine textual material such as sports and financial stories based on factual data from sports contests and mandated corporate financial statements. They may also play a role in synthesizing other kinds of news stories for popular news aggregation sites. Computer programmers may use them to come up with first drafts of computer code for certain procedures, and musicians may use them to come up with ideas for new musical works.<sup>292</sup>

It also is highly likely that the availability of more capable generative AI systems will accelerate the trend toward greater automation of video production, especially with regard to animation and special effects.

In all of these examples, the technology enables activities that heretofore have required significant human labor; that alone is transformative, and may or may not involve the requisite originality of expression to qualify for copyright protection.

So, in almost all cases in which a creative robot is accused of copyright infringement, its defenders will have strong arguments as to its transformative character, influencing the application of the derivative work right to it and strongly influencing the availability of the fair use defense.

## 2. Nature of the Copyrighted Work

On the second factor, nature of the underlying work, *Perfect 10, Inc. v. Amazon.com, Inc.*,<sup>293</sup> the court said: “[o]nce Perfect 10 has exploited this commercially valuable right of first publication by putting its images on the Internet for paid subscribers, Perfect 10 is no longer entitled to the enhanced protection available for an unpublished work.”<sup>294</sup>

The character of the copyrighted work included in the learning database for a generative AI system can affect fair use analysis in a number of ways. If the copyrighted work is nonfiction and analyzes natural phenomena, ideas, or facts, the Supreme Court’s decision in *Feist*<sup>295</sup> counsels thin protection for the work. Its use to teach a generative AI system is likely to involve mostly extraction of un-protectable facts or ideas from the copyrighted work and thus not to constitute infringement at all, and that connection with the idea/expression dichotomy warrants similar impact in fair use analysis. On the other hand, copyrighted fiction may support a stronger argument both of appropriation of protectable elements, and of adverse market impact.

Debbie’s work in the robot-generated story at the beginning of this article is described as an “essay” on artificial intelligence technology. That suggests that

---

292. See discussion *infra* Section I.C (assembling press accounts of uses).

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

294. *Id.* at 1167.

295. *Feist Publ’ns, Inc. v. Rural Tel. Servs. Co.*, 499 U.S. 340, 363–64 (1991).

it is a technical paper explaining and synthesizing facts, mathematical equations and algorithms, thus entitling it to relatively thin protection.

Also, the robot generated story does not make it clear whether Debbie expects to monetize her essay. If she does not, any adverse market impact from use in a learning database for a generative AI system is muted.

### 3. Amount and Character of Appropriation

On the third factor, the machine learning process for generative AI takes only discrete pieces of the works included in a learning database of millions of pre-existing works. Accordingly, the third factor militates in favor of fair use.

The Court in *Google*<sup>296</sup> held that the amount and substantiality of the portion used factor tilted in Google's favor:

If one considers the declaring code in isolation, the quantitative amount of what Google copied was large. Google copied the declaring code for 37 packages of the Sun Java API, totaling approximately 11,500 lines of code. Those lines of code amount to virtually all the declaring code needed to call up hundreds of different tasks. On the other hand, if one considers the entire set of software material in the Sun Java API, the quantitative amount copied was small. The total set of Sun Java API computer code, including implementing code, amounted to 2.86 million lines, of which the copied 11,500 lines were only 0.4 percent.<sup>297</sup>

As to the third factor, the amount of the underlying work taken, the New York district court in *Fox News Network*<sup>298</sup> observed: "where copying the entire work is necessary to accomplish the transformative function or purpose, as is the case, here, this factor, like the second factor, bows to the importance and priority of the first factor's finding of transformative use."<sup>299</sup> The amount and nature of the material taken from the copyrighted work will, in most generative AI systems be small; the system learns by taking little pieces of millions of different works in the database. So, absent an unusual design in the learning algorithms or the learning database, this factor militates strongly in favor of fair use for the generative AI system.

### 4. Impact on Market

The Court in *Google*<sup>300</sup> analyzed market effects:

The fourth statutory factor focuses upon the "effect" of the copying in the "market for or value of the copyrighted work."<sup>301</sup> Consideration

---

296. *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183 (2021).

297. *Id.* at 1204-05.

298. *Fox News Network, LLC v. TVEyes, Inc.*, 43 F. Supp.3d 379 (S.D.N.Y. 2014).

299. *Id.* at 394.

300. *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183 (2021).

301. 17 U.S.C. § 107(4) (2018).

of this factor, at least where computer programs are at issue, can prove more complex than at first it may seem. It can require a court to consider the amount of money that the copyright owner might lose. As we pointed out in *Campbell*, “verbatim copying of the original in its entirety for commercial purposes” may well produce a market substitute for an author’s work.<sup>302</sup> Making a film of an author’s book may similarly mean potential or presumed losses to the copyright owner. Those losses normally conflict with copyright’s basic objective: providing authors with exclusive rights that will spur creative expression.

But a potential loss of revenue is not the whole story. We here must consider not just the amount but also the source of the loss. As we pointed out in *Campbell*, a “lethal parody, like a scathing theatre review,” may “kill[] demand for the original.”<sup>303</sup> Yet this kind of harm, even if directly translated into foregone dollars, is not “cognizable under the Copyright Act.”<sup>304</sup>

Further, we must take into account the public benefits the copying will likely produce. Are those benefits, for example, related to copyright’s concern for the creative production of new expression? Are they comparatively important, or unimportant, when compared with dollar amounts likely lost (taking into account as well the nature of the source of the loss)?<sup>305</sup>

\* \* \*

As to the likely amount of loss, the jury could have found that Android did not harm the actual or potential markets for Java SE. And it could have found that Sun itself (now Oracle) would not have been able to enter those markets successfully whether Google did, or did not, copy a part of its API.

\* \* \*

Looking to these important differences, Google’s economic expert told the jury that Android was not a market substitute for Java’s software. As he explained, “the two products are on very different devices,” and the Android platform, which offers “an entire mobile operating stack,” is a “very different typ[e] of produc[t]” than Java SE, which is “just an applications programming framework.”

\* \* \*

---

302. *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 591 (1994).

303. *Id.* at 591–92.

304. *Id.* at 592.

305. *Cf. MCA, Inc. v. Wilson*, 677 F.2d 180, 183 (2d Cir. 1981) (calling for a balancing of public benefits and losses to copyright owner under this factor).

Finally, given programmers' investment in learning the Sun Java API, to allow enforcement of Oracle's copyright here would risk harm to the public. Given the costs and difficulties of producing alternative APIs with similar appeal to programmers, allowing enforcement here would make of the Sun Java API's declaring code a lock limiting the future creativity of new programs. Oracle alone would hold the key. The result could well prove highly profitable to Oracle (or other firms holding a copyright in computer interfaces). But those profits could well flow from creative improvements, new applications, and new uses developed by users who have learned to work with that interface. To that extent, the lock would interfere with, not further, copyright's basic creativity objectives.<sup>306</sup> . . . After all, "copyright supplies the economic incentive to [both] create and disseminate ideas,"<sup>307</sup> and the reimplementing of a user interface allows creative new computer code to more easily enter the market.<sup>308</sup>

In *Authors Guild v. Google, Inc.*,<sup>309</sup> the court of appeals affirmed summary judgment for Google, holding that Google's display of textual excerpts of copyrighted books was privileged by fair use. It held that the alleged infringing activity was transformative,<sup>310</sup> and that Google's commercial purpose did not override the transformative character of its conduct;<sup>311</sup> that the second factor weakly supported Google's defense;<sup>312</sup> and that, under the third factor, Google's copying of the entire work did not militate against fair use because it did not make the complete copy available to the public.<sup>313</sup> The fourth factor, likewise, supported fair use:

Especially in view of the fact that the normal purchase price of a book is relatively low in relation to the cost of manpower needed to secure an arbitrary assortment of randomly scattered snippets, we conclude that the snippet function does not give searchers access to effectively competing substitutes. Snippet view, at best and after a large commitment of manpower, produces discontinuous, tiny fragments, amounting in the aggregate to no more than 16% of a book. This does

---

306. See *Sony Comput. Ent., Inc. v. Connectix Corp.*, 203 F.3d 596, 607 (9th Cir. 2000); see also *Sega Enters. v. Accolade, Inc.*, 977 F.2d 1510, 1523–24 (9th Cir. 1992) ("[a]n attempt to monopolize the market by making it impossible for others to compete runs counter to the statutory purpose of promoting creative expression"); *Lexmark Int'l, Inc. v. Static Control Components, Inc.*, 387 F.3d 522, 544 (6th Cir. 2004) (noting that where a subsequent user copied a computer program to foster functionality, it was not exploiting the programs "commercial value as a copyrighted work" (emphasis in original)).

307. *Harper & Row Publishers v. Nation Enters.*, 471 U.S. 539, 558 (1985).

308. *Google LLC v. Oracle Am., Inc.*, 141 S. Ct. 1183, 1206–208 (2021).

309. *Authors Guild v. Google, Inc.*, 804 F.3d 202, 229 (2d Cir. 2015).

310. *Id.* at 207–08, 214–18.

311. *Id.* at 218–19.

312. *Id.* at 220.

313. *Id.* at 221–22.



not threaten the rights holders with any significant harm to the value of their copyrights or diminish their harvest of copyright revenue.<sup>314</sup>

In *Perfect 10, Inc.*,<sup>315</sup> the Ninth Circuit found a *prima facie* case of infringement of the display right. It found, however, that Google's display and distribution of the thumbnail images qualified as fair use. It was highly transformative:

We conclude that the significantly transformative nature of Google's search engine, particularly in light of its public benefit, outweighs Google's superseding and commercial uses of the thumbnails in this case. In reaching this conclusion, we note the importance of analyzing fair use flexibly in light of new circumstances. We are also mindful of the Supreme Court's direction that "the more transformative the new work, the less will be the significance of other factors, like commercialism, that may weigh against a finding of fair use."<sup>316</sup>

In sum:

Google has put Perfect 10's thumbnail images (along with millions of other thumbnail images) to a use fundamentally different than the use intended by Perfect 10. In doing so, Google has provided a significant benefit to the public. Weighing this significant transformative use against the unproven use of Google's thumbnails for cell phone downloads, and considering the other fair use factors, all in light of the purpose of copyright, we conclude that Google's use of Perfect 10's thumbnails is a fair use.<sup>317</sup>

As to the fourth factor, impact on the market for the copyrighted work, the court in *Fox News Network*<sup>318</sup> said:

Crucially, this factor "is concerned with only one type of economic injury to a copyright holder: the harm that results because the secondary use serves as a substitute for the original work."<sup>319</sup> Thus any economic harm caused by transformative uses does not factor into this analysis, "because such uses, by definition do not serve as substitutes for the original work."<sup>320</sup> This factor also requires a "balancing of the

---

314. *Id.* at 224.

315. *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1160 (9th Cir. 2007).

316. *Id.* at 1166; *United States v. Campbell*, 510 U.S. 569, 579 (1994) (some internal quotations and citations omitted).

317. *Perfect 10, Inc.*, 508 F.3d at 1168.

318. *Fox News Network, LLC v. TVEyes, Inc.*, 43 F. Supp. 3d 379 (S.D.N.Y. 2014).

319. *Id.* at 395 (quoting *Authors Guild Inc v. HathiTrust*, 755 F.3d 87, 99 (2d Cir 2014)).

320. *Id.*

benefit the public will derive if the use is permitted and the personal gain the copyright owner will receive if the use is denied.”<sup>321</sup>

\*\*\*

No reasonable juror could find that people are using TVEyes as a substitute for watching Fox News broadcasts on television. There is no history of any such use, and there is no realistic danger of any potential harm to the overall market of television watching from an “unrestricted and widespread conduct of the sort engaged in by defendant.”<sup>322</sup>

In most cases, the highly transformative character of the generative AI system means that the market for its output will be quite distinct from the market for the underlying copyrighted works in its learning database. There may be unusual situations where this is not so, but the plaintiff should be burdened to construct a persuasive argument on why her pre-existing work and the generative AI system serve the same market. This might be established in the news aggregator case, or maybe in the robot-generated music case at least when the new musical work in the pre-existing work are in the same genre.

### G. Transformativeness

Professor Reese usefully distinguishes two types of transformativeness.<sup>323</sup> Transformative *content* modifies the content of the original work but seeks to serve the same purpose. Transformative *purpose* may involve content exactly the same or very similar to the original content but organized, formatted, and communicated to serve a different market and a different purpose. Cleaning up the Disney movies in *Clean Flicks v. Soderbergh*<sup>324</sup> did not transform content. The Google communication of thumbnails in *Perfect 10*,<sup>325</sup> served a transformative purpose.

The output of generative AI programs is transformative in both senses. The content of the underlying works has been significantly modified by synthesizing among millions of them. Additionally, the use to which the output of a transformative AI program is to be put is likely to be quite different from the use intended by the authors of the underlying works.

That high degree of both kinds of transformativeness militates against imposing infringement liability on the output of these systems, either because

---

321. *Id.* at 395 (quoting *Bill Graham Archives v. Dorling Kindersley Ltd.*, 448 F.3d 605, 610 (2d Cir. 2006) (internal quotations omitted)).

322. *Fox News Network*, 43 F. Supp. 3d at 396 (quoting *United States v. Campbell*, 510 U.S. 569, 590 (1994) (internal citations and quotations omitted)).

323. R. Anthony Reese, *Transformativeness and the Derivative Work Right*, 31 COLUM. J. L. & ARTS 467, 484–86 (2008).

324. *Clean Flicks of Colo., LLC v. Soderbergh*, 433 F. Supp. 2d 1236, 1241–42 (D. Colo. 2006).

325. *Perfect 10, Inc.*, 508 F.3d, 1165–67.

they are sufficiently different from the original that they do not infringe any of the exclusive rights, or because they constitute fair use. It does not much matter which pathway leads to a finding of no liability.

#### H. The Plaintiff's Proof Challenge

To recover on a copyright infringement claim, the claimant must prove:

- That her work was included in the learning database;<sup>326</sup>
- That the learning algorithms consulted her work;<sup>327</sup>
- What parts of her work the algorithms copied or otherwise manipulated;<sup>328</sup> and
- How those parts were incorporated into the accused work.<sup>329</sup>

Is not surprising that the infringement test cases involve images and program code. In those applications of generative AI, the evidentiary challenges associated with text-based systems like chatGPT are fewer. The process of machine learning from images is easier to understand than the process of machine learning from massive amounts of text, and so a plaintiff has a simpler task in making allegations necessary to establish infringement.

In the computer-coding application, the likelihood is high that the generative AI system will come up with code sentences and sections that are identical to or close to those covered by the plaintiff's copyright. The problem the plaintiff faces, of course, is scenes a faire: there are only so many ways to code a sorting algorithm.

The two high profile lawsuits alleging copyright infringement by generative AI programs implicitly admit the weakness of any claim that the ways in which the program learn constitute copyright infringement. The Open AI plaintiffs do not allege copyright infringement at all.<sup>330</sup>

The Stable Diffusion plaintiffs admit that, "none of the Stable Diffusion output images provided in response to a particular Text Prompt is likely to be a close match for any specific image in the training data."<sup>331</sup> So, instead of alleging facts that constitute copying, or in addition to such allegations, they assert more imaginative theories.

---

326. This reflects the requirement to prove access in order to prove infringement. See discussion *supra* Section II.B.1.c.

327. This can be proved through direct evidence or inferred from access. See discussion *supra* Section II.B.1.c.

328. Copying would prove infringement of the reproduction right under 17 U.S.C. § 106(1). Other manipulation would show infringement of the derivative work right under 17 U.S.C. § 106(2).

329. Copyright infringement claims relate to a particular work that is alleged to infringe. See discussion *supra* Section II.

330. Stability AI Complaint at ¶ 21.

331. *Id.* at ¶ 93.

The GitHub claimants do not plead copyright infringement at all. The GitHub complaint claims:

- Violation of the DMCA;<sup>332</sup>
- Common-law breach of contract for violating licenses;<sup>333</sup>
- Tortious interference with contract relationships;<sup>334</sup>
- Common-law fraud;<sup>335</sup>
- False designation of origin/false passing off in violation of 15 U.S.C. § 1125;<sup>336</sup>
- Unjust enrichment;<sup>337</sup>
- Unfair competition;<sup>338</sup>
- Breach of contract - violation of GitHub privacy policy and terms of service;<sup>339</sup>
- Violation of California Consumer Privacy Act;<sup>340</sup>
- Negligence in handling personal data;<sup>341</sup> and
- Civil conspiracy.<sup>342</sup>

In addition to direct copyright infringement,<sup>343</sup> the Stability AI complaint claims:

- Vicarious Copyright Infringement;<sup>344</sup>
- Violation of the DMCA;<sup>345</sup>
- Violation of the Statutory Right of Publicity;<sup>346</sup>
- Violation of the Common Law right of publicity;<sup>347</sup>
- Unfair competition;<sup>348</sup>
- Breach of contract.<sup>349</sup>

---

332. GitHub Complaint at ¶¶ 138–67.

333. *Id.* at ¶¶ 168–83.

334. *Id.* at ¶¶ 184–189.

335. *Id.* at ¶¶ 190–95.

336. *Id.* at ¶¶ 196–99.

337. *Id.* at ¶¶ 200–06.

338. *Id.* at ¶¶ 207–10.

339. *Id.* at ¶¶ 211–20.

340. *Id.* at ¶¶ 221–29.

341. *Id.* at ¶¶ 230–34.

342. *Id.* at ¶¶ 235–39.

343. See discussion *supra* Section II.B (summarizing copyright infringement claims).

344. Stability AI Complaint at ¶¶ 169–77.

345. *Id.* at ¶¶ 178–200.

346. *Id.* at ¶¶ 201–12.

347. *Id.* at ¶¶ 213–22.

348. *Id.* at ¶¶ 223–26.

349. *Id.* at ¶¶ 227–36.

### III. REMOVAL OF DMCA DIGITAL RIGHTS MANAGEMENT INFORMATION

The Digital Millennium Copyright Act (“DMCA”) prohibits circumvention of a technological measure to protect copyright: “No person shall, without the authority of the copyright owner or the law—

(1) intentionally remove or alter any copyright management information . . . .”<sup>350</sup> Both the GitHub and the Stability AI complaints allege violation of this DMCA technological protection right.<sup>351</sup>

A claim of violation of the DMCA’s technological protection provisions<sup>352</sup> avoids having to prove copying or derivative work; success on such a claim depends upon showing that metadata associated with copy protection or assertion of copyright has been removed from a copyrighted work. But proof of such a claim still must identify a particular work from which technological protection information has been removed, and given the way that inputs to learning databases are mixed into a large and diverse stew of characteristics, that burden is likely to be impossible.

The deep learning systems do indeed remove or at least disregard the copy protection metadata from the entries in their learning databases but then they simply process the content rather than copying it verbatim, publishing, distributing, or displaying it verbatim without the copy protection. So, the copyrighted work never exists in the learning database stripped of its copyright management information; the deep-learning system simply looks at other parts of it as it learns.

In *Murphy v. Millennium Radio Group LLC*,<sup>353</sup> the Third Circuit held that DMCA protection extends beyond an automated copy protection system to a mere incorporation of a photographer’s name as credit on a photograph.<sup>354</sup> The plaintiff photographer alleged that the defendant radio station posted a copyrighted photograph on its website without the credit, of those distributing a work from which copyright management information had been removed in violation of section 1202(b)(3).<sup>355</sup>

In *O’Neal v. Sideshow, Inc.*,<sup>356</sup> the district court dismissed a claim for violation of the DMCA copyright management information provisions. Citing other cases, the court held that no DMCA violation occurs unless the accused work is identical to the copyrighted work from which copyright information was

---

350. 17 U.S.C. § 1202(b)(1) (2018).

351. Github Complaint at ¶¶ 138–67; Stability AI Complaint ¶¶ 179–200.

352. 17 U.S.C. § 1202(b)(1)–(c)(1).

353. *Murphy v. Millennium Radio Grp. LLC*, 650 F.3d 295 (3d Cir. 2011).

354. *Id.* at 305.

355. *Id.* at 301 (characterizing plaintiff’s argument).

356. *O’Neal v. Sideshow, Inc.*, 583 F. Supp. 3d 1282, 1288 (C.D. Cal. 2022); *see also* *O’Neal v. Sideshow, Inc.*, No. 21-cv-07735, 2023 U.S. Dist. LEXIS 42984 (C.D. Cal. Mar. 14, 2023) (granting summary judgment to defendant on copyright infringement claims).

removed.<sup>357</sup> Moreover, the court held, the plaintiff failed to allege the requisite intent. “In fact, O’Neal has not alleged anything with respect to Sideshow and Anzalone’s knowledge that removal or alteration of O’Neal’s CMI would likely result in future infringement.”<sup>358</sup>

In *Kirk Kara Corp. v. Western Stone and Metal Corps.*,<sup>359</sup> cited by the O’Neal court, the district court dismissed a DMCA claim, while denying a motion to dismiss a copyright infringement claim based on substantial similarity.<sup>360</sup>

Defendant argues that Plaintiff’s claim should be dismissed because the images displaying the Subject Designs and the images displaying the Subject Products are not exact copies.<sup>361</sup> The Court agrees. Based on a review of the side-by-side images included in the Complaint, the Court can determine that, while the works may be *substantially similar*, Defendant did not make *identical* copies of Plaintiff’s works and then remove the engraved CMI. In such cases, even where the underlying works are similar, courts have found that no DMCA violation exists where the works are not identical.<sup>362</sup>

But in *ADR International Limited v. Institute for Supply Management Inc.*,<sup>363</sup> the district court adopted a magistrate judge’s recommendation to deny a motion to dismiss a DMCA claim, holding that the accused work from which copyright information was removed need not be identical to the pre-existing copyrighted work. It found *Kirk Kara* unpersuasive.<sup>364</sup>

On the other hand, derivative works from which copyright protection information has been removed would be outside the protection of DMCA section

357. See *O’Neal*, 583 F. Supp. 3d at 1287.

358. *Id.* at 1288.

359. *Kirk Kara Corp. v. W. Stone & Metal Corp.*, No. CV-20-1931, 2020 U.S. Dist. LEXIS 189425 (C.D. Cal. Aug. 14, 2020).

360. *Id.* at \*7 (denying defendant’s motion to dismiss copyright infringement claims based on substantial similarity due to determination being better suited to adjudication).

361. *Id.* at \*18.

362. *Id.* at \*6. The court notes the following cases,

*See Kelly v. Arriba Soft Corp.*, 77 F.Supp.2d 1116, 1122 (C.D. Cal. 1999), *aff’d and rev’d in part on other grounds*, 336 F.3d 811 (9th Cir. 2003) (explaining that displaying thumbnails of the plaintiff’s images without CMI did not violate the DMCA because CMI was not removed from the plaintiff’s original, full-sized images); *Fischer v. Forrest*, 286 F. Supp. 3d 590, 609 (S.D.N.Y. 2018) (finding that CMI was not removed where advertisements were not entirely copied, despite containing the same four phrases); *Frost-Tsuji Architects v. Highway Inn, Inc.*, No. CIV. 13-00496 SOM, 2015 WL 263556, at \*2 (D. Haw. Jan. 21, 2015), *aff’d*, 700 F. App’x 674 (9th Cir. 2017) (stating that because architectural drawings were “not identical . . . this court cannot say that [the defendant] removed or altered [the plaintiff]’s copyright management information”).

*Id.* at \*18–19.

363. *ADR Int’l Ltd. v. Inst. for Supply Mgmt. Inc.*, No. 4:22-CV-1914, 2023 U.S. Dist. LEXIS 58090 (S.D. Tex. Feb. 24, 2023).

364. *Id.* at \*23.

1202.<sup>365</sup> “‘Removal’ of CMI from a copyrighted work is not the same as the failure to add CMI to a nonidentical rendition or a derivative of the protected work.”<sup>366</sup>

The biggest hurdle for a plaintiff claiming violation of the DMCA copyright protection information right is under the cases requiring that the accused work from which protection was removed must be identical to the underlying work. The case law is mixed on that question, however, and so there might be a glimmer of hope for plaintiffs on this claim.

#### IV. CONSIDERATION OF THE PURPOSES OF COPYRIGHT

Judge Posner pointed out that copyright protection cuts both ways in shaping creative expression:

[E]xtending copyright protection is a two-edged sword from the standpoint of inducing creativity, as it would reduce the incentive of subsequent authors to create derivative works (such as new versions of popular fictional characters like Holmes and Watson) by shrinking the public domain. For the longer the copyright term is, the less public-domain material there will be and so the greater will be the cost of authorship, because authors will have to obtain licenses from copyright holders for more material—as illustrated by the estate’s demand in this case for a license fee from Pegasus.

Most copyrighted works include some, and often a great deal of, public domain material—words, phrases, data, entire sentences, quoted material, and so forth. The smaller the public domain, the more work is involved in the creation of a new work. The defendant’s proposed rule would also encourage authors to continue to write stories involving old characters in an effort to prolong copyright protection, rather than encouraging them to create stories with entirely new characters. The effect would be to discourage creativity.<sup>367</sup>

What Judge Posner had to say about extending the term of copyright applies with equal force to the extension of the scope of copyright. As the Supreme Court put it, “[t]he more artistic protection is favored, the more technological innovation may be discouraged; the administration of copyright law is an exercise in managing the tradeoff.”<sup>368</sup>

Charles Cronin said it this way:

Popular songs today are akin to Lego block or Tinker Toy assemblages in which the constituent components may contain greater

---

365. *Id.* at \*23–26 (citing cases).

366. *Id.* at \*24–25 (quoting *Kipp Flores Architects, LLC v. Pradera SFR, LLC*, No. SA-21-CV-00673-XR, 2022 U.S. Dist. LEXIS 68040, at \*7 (W.D. Tex. Apr. 13, 2022)).

367. *Klinger v. Conan Doyle Est., Ltd.*, 755 F.3d 496, 501 (7th Cir. 2014) (affirming summary judgment and rejecting extension of protection for characters after copyright expired).

368. *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 928 (2005).

inventiveness than their combination. Or, the finger paintings of toddlers whose doting parents—like music arrangers, audio engineers, and videographers—will transform them into attractive works by using skillful framing and presentation techniques unknown to their creators. Regardless of the potential appeal or marketability of such creations, however, the more nebulous their authorship, the more charily courts should view the legitimacy of infringement claims based upon them.<sup>369</sup>

Mr. Cronin argues that the law should “cultivate a return to more permissible attitudes toward copying of musical expression through less hospitable reception of infringement disputes involving anything other than replication of substantial musical expression—essentially the entire work—that threatens to supplant it in the marketplace.”<sup>370</sup>

Fair use analysis, particularly the fourth market impact factor provides a legal framework for pragmatic law and economic analysis to get the threshold of incentives right, balancing incentives for creation against possible roadblocks to creation.

The circumstances will be unusual where the market for generative AI output is the same as the market for the protected works constituting the learning database for the AI system. Any infringement plaintiff seeking to overcome a fair use defense should be burdened by showing market similarity and explaining how the existence of the generative AI output diminishes the market for the underlying work.

As the sections of this article analyzing the statute, the case law, and the commentary make clear, the owner of a copyrighted work included in the learning database for a generative AI system will have difficulty establishing infringement and avoiding fair use. No one should bend over backwards to interpret the law to make things easier for these infringement claims.

The need to balance copyright’s incentive effect against its role as a barrier means that courts and lawmakers should think hard about the effect of exposing generative AI innovators to copyright infringement liability for the millions of possibly copyrighted works that their machine learning systems consult in the development process. If a robot entrepreneur is exposed to a proliferation of copyright infringement lawsuits, he is going to have a hard time attracting investment and, once he gets ensnared in the litigation, have little time for product development and entrepreneurship.

Opposing this is the claim of the owners of the pre-existing rights not to be deprived of reasonable returns on the sunk cost represented by their creative investment. Sound policy requires the copyright owner should have a heavy burden to demonstrate exactly how the availability and use of a particular

---

369. Charles Cronin, *I Hear America Sing: Music Copyright Infringement in the Era of Electronic Sound*, 66 HASTINGS L. J. 1187, 1254 (2015).

370. Cronin, 66 HASTINGS L. J. at 1255.



generative AI system interferes with their investment-backed expectations of return on their copyrighted works.<sup>371</sup>

It is not likely that the politics will evolve so as to provide broad immunity from copyright-infringement liability to generative AI systems. On the contrary government agencies in Europe<sup>372</sup> and the United States<sup>373</sup> seem to be in a race to figure out how to impose regulations on artificial intelligence. This implies that exposure to liability will haunt generative AI innovators for the foreseeable future. One way to deal with this is to develop schemes of collective licensing, modeled on what was developed in the early days of recorded music with the major copyright collectives,<sup>374</sup> and in the early days of digital music available through the Internet, 50 years later.<sup>375</sup>

Authors of copyrighted or copyrightable literary, musical, photographic, and video works would agree with new licensing collectives to allow their works to be placed in massive learning databases for generative AI systems. AI system developers would contract with these new collectives to license access to all of the material in the databases. If the terms of the licensing agreements are drafted appropriately, the generative AI system developers would be protected from copyright infringement liability when their systems learn from the databases and when they produce output based on that learning.

Standard licenses already are employed for large collections of computer code and images under the open-source system.<sup>376</sup> These standard licenses could be amended to cover the use of the licensed works in generative AI training databases. That approach would not require the creation of new copyright collectives, and might be a useful place to start working out licensing details in the rapidly changing real-world of AI exploitation.

#### CONCLUSION

Copyright owners will have a difficult time imposing liability for copyright infringement on most generative AI systems, because they will have difficulty showing that protected elements of their particular were copied, turned into

---

371. *Penn Cent. Transp. Co. v. City of New York*, 438 U.S. 104 (1978) (identifying the extent to which regulation interferes with investment-backed expectations as a factor in determining whether a Fifth Amendment taking has occurred).

372. Supantha Mukherjee et al, *EU proposes new copyright rules for generative AI*, REUTERS (Apr. 28, 2023, 2:51 AM), <https://www.reuters.com/technology/eu-lawmakers-committee-reaches-deal-artificial-intelligence-act-2023-04-27/> (reporting on European Parliament committee draft of requirement for disclosure of copyrighted items used to train generative-AI systems).

373. NTIA, AI Accountability Policy Request for Comment, 88 Fed. Reg. 22433 (Apr. 13, 2023) (seeking comments on systems to “audit” AI systems to ensure that they are “accountable”).

374. Eric Priest, *The Future of Music Copyright Collectives in the Digital Streaming Age*, 45 COLUM. J.L. LAW & ARTS 1, 14–16 (2021) (summarizing history of licensing collectives).

375. *Id.* at 11–14 (discussing movement to digital streaming compulsory licenses).

376. See GitHub Complaint at ¶ 34 n.4 (describing license options for when a GitHub user creates a new repository).

derivative works, distributed to the public, displayed to the public, or performed for the public. Even if they establish *prima facie* infringement, defenders of the AI systems are likely to establish the elements of fair use excusing robot-generated material from liability because it is so transformative.

Exposure of robot entrepreneurs to constant threats of copyright infringement lawsuits, will chill innovation. Even risk-taker entrepreneurs will have trouble attracting investment. Even if they get funded, once litigation starts, it will derail plans for product development.

Lawmakers in Europe and the United States are falling all over themselves in a stampede to regulate AI, intensifying the legal threats.

Avoiding the cloud of legal uncertainty may drive the private AI sector to develop mechanisms for collective licensing of content for training databases that shield generative AI engines from copyright infringement liability.