

What Qualifies as Creativity

Authorship, Generative AI and Copyright

By Kevin L. Smith

ABSTRACT: Generative artificial intelligence poses a significant challenge to traditional notions of authorship. This article examines some of those issues through the lens of copyright law. By looking at the major input and output challenges issues raised against AI by copyright, we evaluate the ways in which generative AI is forcing authors to reassess how they create, and what it means to be an author. We examine some of the many copyright lawsuits pending against generative AI providers, as well as policy and legislative proposals in the United States and abroad.

It was probably inevitable that, before beginning an article on this topic, I would consult a generative artificial intelligence tool. Actually, though, I turned to AI only after I had developed my own outline and set of references. Only then did I ask Microsoft Copilot to do the same. The result was a competent outline on the topic of AI, copyright, and authorship, but it was one that followed a very standard, one might almost say formulaic, pattern. Almost any topic could have been plugged into the basic structure of that outline. In short, it reflected none of the particular interests or idiosyncrasies of me as an author. I did not find it of much use, except to verify that, in my own outline, I had not missed any major issues. On the other hand, the references, which all appeared to be real, human-authored works, were very useful, and several of them have informed what follows (and are cited in the reference list). It is worth noting that MS Copilot offered to write a complete article for me on this topic. I did not avail myself of this offer.

INTRODUCTION

From its inception, copyright has been interwoven with the historical development of ideas about authorship. Mark Rose (1993), in his *Authors and Owners*, elegantly demonstrates how copyright law arose simultaneously with the Romantic notion of the author as a singular genius, creating *ex nihilo* out of their own personality and imagination. Prior to this development, authors often were seen as divinely inspired, or as working out of an established tradition—the example of Shakespeare reworking the many previous accounts of the legend of King Leir¹ about a century before copyright law was adopted in England is a good illustration (see Shapiro 2016). These traditions were supported in a variety of ways, including public funding, commissions, performances, and patronage. As the individual creativity of the author became the focus, the idea that that creator should own their work in some exclusive sense came to the fore and was embodied in copyright laws. These laws have, over time, been instrumental in defining authorship by focusing on concepts of originality, creativity, and ownership, all of which are now being challenged by the development of generative AI (Gaidartzi and Stamatoudi 2025).

Authorship can be a fraught topic in theological arenas as well. In Judaism, Christianity, and Islam, the so-called “Religions of the Book,” issues of the source and authorship of scriptural texts

Kevin L. Smith is the former Michael and Eugenia Wormser Director of Libraries at Colby College and the former Dean of Libraries at the University of Kansas, where he also taught in the KU Law School.

can be hotly debated. Indeed, we shall discuss a case in which the question of copyright for a work alleged to have been authored by a “celestial being” was adjudicated (Urantia Foundation 1997). Here again, the relationship between our concepts of authorship and the way we protect literary works is complex and contested.

Artificial Intelligence is challenging and further complicating contemporary notions of authorship. Since ChatGPT was released publicly in November 2022, over 50 lawsuits about the impact of generative AI on the creative arts have been filed as of September 12, 2025, according to the website *Chat GPT is Eating the World* (n.d.). Most of these lawsuits are filed on behalf of creators, challenging generative artificial intelligence on several grounds related to the integrity and marketability of copyrighted works. As Sarah Lim notes, these developments “[have] begun reshaping the creative landscape” in which authors work and make their livings (2025, 1). Since copyright has become the fundamental protection for authors for over 300 years, this reshaping will, inevitably, have profound impacts.

Although AI has been around in a variety of forms for quite a while, what is new in the past couple of years is the widespread availability of generative AI—AI tools that actually create outputs that have not existed before. This apparent creativity is where generative AI poses its most fundamental challenge to copyright law and to our concept of authorship. In addition, these outputs of generative AI are genuinely unpredictable; the person who prompts the AI tool cannot know exactly what the output will look like, nor predict what a particular prompt, even one they have used before, will produce. This unpredictability, we will see, has an important impact on legal discussions about the copyrightability of AI outputs.

This article will examine the ways in which that reshaping is being played out in copyright laws and litigation. We will begin by looking at the three major challenges that AI poses for copyright law and then map those challenges to our notions of authorship. At the beginning, it is important to be cognizant of the difference between genuine issues involving AI and copyright and simple AI-assisted plagiarism. Had I accepted the offer to have MS Copilot write this article for me, it would certainly have been a failure of authorship, but it would not have engaged with the deep issues of copyright and authorship in any significant way. As Lemley and Oullette (2025) explain, these are distinct problems, and the distinction is important. If I submitted an article written by an AI tool as my own work, that would be plagiarism, an ethical failing, but not, probably, the legal wrong of copyright infringement. The questions we will examine are issues of how the law is changing the environment for creative authors. AI plagiarism, since it is a failure of authorship, is not really a copyright issue *per se*, but we will discuss it briefly in the conclusion.

THREE CHALLENGES TO COPYRIGHT LAW

It is common, and sensible, to divide the issues involving generative artificial intelligence and copyright into input and output issues. The articles by Sarah Lim (2025) and Carys Craig (2021) are both structured in this way. As we look at the three major copyright questions that are raised by generative AI, it is helpful to keep this distinction in mind, but also to be attentive to the point at which it breaks down.

By raising the distinction between input issues and output issues, we begin to “peek under the hood” of generative AI tools. The process of developing such tools is extremely complex; there are many points at which copyright issues can arise, and copyrighted materials can be introduced. The article by Lee, Cooper and Grimmelmann (2024) about the “generative-AI supply chain” does a nice

job of explaining the process and all the points along that process where copyright is implicated. The metaphor of a supply chain is effective for conveying the complexity we must address.

The major copyright problems that arise for generative AI can be broken out in different ways, but these three questions are one way to pose the issues, and they will serve as an effective outline for our consideration of authorship:

- 1) Does it infringe copyright when previously copyrighted material is used to train a large language model, which is the underlining structure for all generative AI? Such training often involves “scraping” a vast amount of material from the Web, and the process does not usually differentiate between copyrighted and public domain materials. Thus, on the input side, generative AI faces the challenge that it is “founded on the theft of intellectual property.”²
- 2) On the output side, are the products of generative AI eligible for copyright protection? That is, when there is a computer that intervenes between a human person and the output, which is almost never predictable to that person, should the work have a copyright? Who should own that right, if it exists?
- 3) What happens when outputs from generative AI appear to infringe previously existing copyrighted material? Although framed as an output question, this is the point at which the input/output division can break down, because courts have usually been confronted with such allegedly infringing output when it is presented as evidence that the infringed content must necessarily have been used as training data. That is, plaintiffs in many cases have argued that infringing output must prove infringement at the input stage.

Each of these questions asks us how to find a balance between innovation and the rights of creators. They pose the dilemma of how we can maintain the incentive for creativity which is part of the justification for copyright law in the first place. In what follows, we will look at each of these questions both as legal matters and for what they tell us about evolving concepts of authorship.

AI TRAINING AND FAIR USE

Nearly all the more than fifty lawsuits that have been filed against companies that have developed generative AI tools are focused on the issue of how AI models are trained, claiming that copyrighted works were infringed when they were included in training data. Most of these cases have been brought by authors and other creators (*ChatGPT is Eating the World* provides a running update on this litigation). For authors, the issues raised by these cases regarding innovation and fair use are central to establishing the limits of authorial control. What kinds of reuse of any work must be authorized by the creator? As new potential uses develop, should authors always be consulted before their works are used? Should they be compensated for all such new and previously unimagined uses?

For some authors, the issues raised by generative AI go beyond the concern over an allegedly unauthorized use of their works. In many of the lawsuits that have been filed over the past couple of years, the plaintiffs have asserted that their works have been copied and used to train an AI tool so that it can produce similar works that will compete with them in their particular creative market. This idea that their works are copied and somehow stored inside the AI tool, and that the products of the tool will damage the market for their own books, songs, etc. can have strong emotional impact on authors and has led some plaintiffs to assert that all output from a generative AI tool should be considered infringing derivative works of each item in its training data. Although we have very few final legal rulings about generative AI, the one court that has considered this claim about derivative works has rejected it, requiring that plaintiffs provide evidence of a specific link between a known item of training data and some particular output (Andersen 2023). This is a nearly impossible evidentiary

hurdle for the claimants. We will discuss the other part of this claim, the idea of memorization—that AI tools store copies of at least some of their training data—later in this section.

Apart from the ruling about derivative works that was addressed in the above-cited preliminary ruling in the *Andersen v. Stability AI Ltd.* case, we have three final decisions, all at the lower District Court level, that rule on whether the use of copyrighted materials in AI training data is infringement or fair use. These rulings provide us with only a cloudy picture of the situation.

The earliest “final” ruling (meaning a decision on the merits from a lower court that is susceptible to appeal) that we saw in a case involving training generative AI on copyrighted works held that such use of prior works was not fair use. That was in the case of *Thomson Reuters v. Ross Intelligence*, where the District Court in Delaware found that the fact that the product created by the AI tool would be used to directly compete against the products from Thomson Reuters that were used to train the AI precluded fair use (Thomson Reuters 2025). The evidence of direct, commercial competition was central to this case, and it likely means that this decision, as it stands, will not have a significant impact on subsequent cases. The decision is also on appeal to the Third Circuit Court of Appeals.

Just a few months after the Thomson Reuters case, other district courts reached the opposite conclusion about AI training and fair use. In June 2025 two different district judges in the Northern District of California both held that training generative AI was a transformative use of previously copyrighted materials and was, therefore, fair use (Bartz 2025; Kadrey 2025). In these cases, the expected arguments that what was used was information *about* the previously copyrighted works, such as the relationships between words and grammatical structures that were stored in the AI tools neural network, and that the products were not used in direct commercial competition prevailed. These arguments were expected because they harkened back to the *Google* and *HathiTrust* cases about book digitization for the purposes of indexing and preservation. In those cases, the Second Circuit found clearly transformative purposes for the copying (see BLTJ Blog 2015), and in *Bartz* and *Kadrey*, the judges likewise thought that the transformative use was clear; Judge Alsup in the *Bartz* case called it “exceedingly transformative” (Bartz 2025, 9). In addition to that finding, Judge Alsup ruled that the amount used by Anthropic, the defendant in the case, was necessary for its purpose and was not used to make any of the original works publicly available. In finding that “training the LLM did not result in any exact copies nor even infringing knockoffs of their works being provided to the public,” the judge also held that training this AI tool (called Claude) would not usurp demand for the original works (Bartz 2025, 28).

Although these decisions were both strong and, in my opinion, correct analyses of fair use for AI training, both rulings contain reasoning that compromised what might have otherwise appeared as clear victories for fair use. In the *Kadrey* case, Judge Chhabria raised an entirely new possibility for the fair use analysis by claiming that his analysis might have been different if the plaintiffs had argued that the products of a generative AI tool would “dilute” the market for the previously copyrighted works (Kadrey 2025, 28-36). This issue of dilution has never been a part of the four-factor analysis before. If it were, it could demolish fair use because each new work produced could, arguably, dilute the market for previous works. That Judge Chhabria devotes so much space to this idea, even though he ultimately does not employ the theory in his ruling, is troubling, and it prevents this case from being as clear as it ought to be.

The *Bartz* case also has to be qualified in terms of its fair use implications because Judge Alsup held that some of the materials used, specifically those works that were taken for AI training from so-called “pirate libraries,” were not fairly used (Bartz 2025). This idea that the source of the work that is used has an impact on whether or not the subsequent use is fair represents an inquiry that is

not common in fair use cases, and it has caused considerable debate (Lim 2025). Because the books sourced from infringing works were not included in the fair use ruling, the plaintiffs were able to force a massive settlement for authors in this case; Anthropic agreed to a \$1.2 billion payout (see Hansen 2025).

Internationally, these legal challenges involve a different approach than in the United States; they evaluate the issue of training data using the copyright exceptions for text and data mining (TDM), which have been incorporated into the copyright laws of most European Union countries. These exceptions, intended to facilitate text and data mining for research purposes, have been held by EU courts to encompass the training activities for LLMs (Gaidartzi and Stamatoudi 2025). Although the specific application of the TDM exceptions has been debated, the fact that most jurisdictions have such a copyright exception generally points to training an LLM being permissible. However, as we turn to the vexing issue of memorization, we must note a very recent case. In *GEMA v. OpenAI*, decided in November 2025, a court in Munich found infringement because it believed that the AI model in question, ChatGPT, had retained complete, or nearly complete, copies of copyrighted works. The exceptions for TDM do not apply in a case where the original work has been copied *in toto*, so the court ruled that, based on the evidence it had of “memorization of training data,” Open AI could be held liable (Guadamuz 2025).

This question of memorization—the idea that the LLMs somehow retain complete copies of items from their training data—is hotly debated in both the courts and in the literature about AI. The most compelling case that seems to assert memorization is *New York Times v. Microsoft*, filed in 2023 in the Southern District of New York. In their complaint, *The New York Times* asserts that it has been able to prompt MS Copilot and ChatGPT to produce verbatim, or nearly verbatim, copies of *The New York Times* columns when carefully prompted to do so. This is evidence, the plaintiffs claim, that copies of their copyrighted works are being retained and distributed by the companies that market these AI tools.

In his article “Copyright Safety for Generative AI,” Matthew Sag (2023) asserts that LLM models do not retain wholesale copies. The size of the models is enough to prove that they do not store a large amount of their training data. At most, he argues, it must be a very small amount of training data detail that is memorized in any meaningful way. Sag also describes what is sometimes called “the Snoopy problem,” where a model seems to retain particular aspects of a previous work when it is trained on multiple copies of an identical element. Thus, if there are 10,000 images of Snoopy in a model’s training data, it is much more likely to memorize that figure, as evidenced, allegedly, by that fact that a prompt seeking a beagle sleeping on a doghouse will very often produce an output that looks a lot like Snoopy.

Courts have also taken notice that the *New York Times* case and others often use elaborate prompting that is specifically designed to produce verbatim copies; Sag calls these “extraction attacks” (2023, 326). In the *New York Times* case, the prompts specifically asked for verbatim copies, and they gave the models significant chunks of the desired text in order to prompt the model to output more of the same text (New York Times 2023). In a similar kind of prompting in a music-generation case, *UMG Records et al. v. Suno*, the plaintiffs asserted that the defendants’ AI model could produce nearly exact copies of the plaintiff’s song “Johnny B. Goode.” The defendants noted that the prompting provided a large amount of very specific extracts and information about the song, while also asserting that there are many versions of “Johnny B. Goode” on the Internet, so that similar output is not evidence that one specific version was memorized (see UMG Records 2024). Indeed, it is possible that, in some cases, the AI model simply acts like a search engine and finds the material it has been so specifically requested to find on the Web.

A recent article on memorization by A. Feder Cooper and James Grimmelman (2025) explores the phenomenon in considerable depth. Like Sag, they suggest that the ability to extract verbatim copies most likely is the result of the highly detailed information about previous works stored in a model's neural network, rather than copies of those works, but that does not prevent a model from sometimes reproducing exact items from its training data. To understand this phenomenon, Cooper and Grimmelman distinguish between memorization, which would happen during training; extraction, a kind of prompting; and the regurgitation of a work as output. They opine that, if the details retained in a neural network are sufficiently complete to permit extraction and regurgitation of a verbatim copy of training data, then, for the purpose of copyright law, a copy has been made even if strict memorization did not happen during training. Interestingly, this appears to be the logic used by the Munich court in the *GEMA* case.

These issues of training, fair use, and memorization clearly have important implications for how we understand authorship. Whenever a work is shared in any way, authorship involves a loss of control, an opening up of space in which the reader can make use of the work as they wish. Fair use is arguably a legal definition of this space where the user has freedom. But the idea that a machine can make use of this same space, while it may be vital for technological innovation, is frightening and offensive to many authors. The idea of memorization, especially when not well understood, makes the use of an author's work in training an AI model even more frightening, because it raises the fear that that author's own work could literally be turned against them to compete with, or even substitute for, their own creations.

WHO OWNS THE OUTPUTS OF GENERATIVE AI?

When we turn to the question of copyright and ownership of the output of generative AI, we probe at the very heart of what it means to be an author. In many jurisdictions, the answer to the question of who owns AI outputs is that no one does, which is simultaneously an affirmation of the humanity that is required for authorship and a threat to the identity of those authors who want to use generative AI as their creative medium.

The United States Copyright Office has been consistent in its insistence that only works of human authorship are eligible for copyright protection. The recent report from the Copyright Office on this topic (part two of its study of issues raised by generative AI) details this history quite thoroughly (United States Copyright Office 2025). As the report states, "the incentives authorized by the Copyright Clause are to be provided to human authors as the means to promote progress" (36). While the report firmly maintains the principle that human authorship is a prerequisite for copyright protection, it also details some of the circumstances in which a work can have protection for the human-authored elements it contains, while AI-generated portions of the work would be in the public domain.

Most jurisdictions take a similar approach to that of the United States, at least insofar as they prioritize human authorship. Canadian copyright scholar Carys Craig (2021) explains that the language of the law in her country, referring to the author's place of residence and their lifespan to define the scope of protection, creates "the obvious implication that an 'author' is a natural person" (5). Nevertheless, some countries do take a more nuanced approach. In their comparative study on copyrightability, Gaidartzi and Stamatoudi (2025) note that in the UK, the law stipulates that a copyright in a computer-generated work can be held by "the person by whom the arrangements necessary for the creation of the work are undertaken" (4, quoting from the UK Copyright, Designs, and Patents Act 1988). The EU takes a somewhat different approach, focusing on the idea of originality, which requires that a work can have protection if "it is its author's own intellectual creation" (4). Although

this formulation ducks the outright question of human personhood, one could clearly argue that Carys Craig's reasoning about the implication of a natural person as the author in Canadian law applies to this EU approach as well.

In the United States, the development of the position that only human-authored works are entitled to copyright protection can be illustrated by the history of three images. In the 2023 case of *Thaler v. Perlmutter*, the Circuit Court for the District of Columbia upheld the Copyright Office's denial of copyright registration for an image called *A Recent Entrance to Paradise* (Figure 1), which was submitted with the assertion that the image was solely the work of a generative AI tool (Thaler 2023). To arrive at this conclusion, the Circuit Court looked back at several earlier cases, a couple of which are worth examining here.



Figure 1: Stephen Thaler's *A Recent Entrance to Paradise* is in the public domain because of the rulings of the Copyright Office and the D.C. District Court.

The Supreme Court encountered a case with clear relevance to modern disputes about generative AI all the way back in 1884, in *Sarony v. Burrow-Giles Lithography*. Napoleon Sarony took photographic portraits of famous people out of his New York City studio. His photographs of Oscar Wilde (see Figure 2) were widely used to promote Wilde's speaking tour of the United States, but they were also used without permission by Burrow-Giles to create visiting cards. Sarony sued for copyright infringement, and Burrow-Giles defended the suit by claiming that photography was a mechanical process in which a machine simply captured what was before it. In short, they said that photographs were not a product of human authorship and therefore could not have a copyright.

This argument is very similar to the one being used today to deny copyright to the products of generative AI. Since the argument was rejected by the Supreme Court in the *Sarony* case, it is necessary to distinguish the two situations clearly. In *Sarony*, the Supreme Court held that the photographer had made a series of decisions about his photographic process, including lighting, background and camera setting, that, when taken together, embody a creative vision that undergirded the photographs. Thus, they found that Sarony's copyright was valid and had been infringed (Sarony 1884).



Figure 2: Although the copyright in Sarony's photos of Oscar Wilde was upheld by the U.S. Supreme Court, these images are now in the public domain because those rights have expired.

At first blush, the issue in *Thaler* seems very similar. A human's creative vision, arguably, is at the root of the prompting that causes an AI tool to create an image like *A Recent Entrance to Paradise*. Indeed, a human may modify and iterate those prompts until they get an image that satisfies them.

So why was the *Thaler* case different? According to the D.C. Circuit Court, the principal difference is in the *unpredictability* of the output of generative AI. Regardless of what the person prompting the AI tool may imagine, the output is never going to conform to a specific vision the way a photograph does. And if that person submits the same prompt multiple times, in most cases the output is going to be different each time. This unpredictability distinguishes AI output from photographs, according to the D.C. Circuit Court, and means that such output does not have the same intimate relationship to a human author's vision for the work (Thaler 2023). Thus, Thaler's work was not eligible for copyright even though the *Sarony* case had seemed like a relevant precedent.

One other case that the Copyright Office considered in deciding that a human author was required for copyright protection was one that is familiar to many people, the case of the monkey selfie. That case involved photographs taken by a group of macaque monkeys using camera equipment left unattended by wildlife photographer David Slater (see Figure 3). Although People for the Ethical Treatment of Animals sued on the monkey's behalf (they named the monkey Naruda), the Ninth Circuit held that copyright is a benefit conveyed by human laws for the encouragement of human authors. As such, it cannot be held by a monkey (Naruda 2018).



Figure 3: The Ninth Circuit Court of Appeals held that this picture is in the public domain because copyright is intended to benefit human beings, and therefore cannot be held by a monkey, no matter how photogenic.

This case is interesting for several reasons. First, it defines authorship as necessarily human in distinction not only from machines, but also from other creatures. The decision seems to recall, at

least implicitly, the idea of a creative vision. Naruda, of course, did not have a vision for the photos; they merely saw themselves in the lens and pushed the buttons. In that sense, the idea of predictability is also implicit since a non-human creator like Naruda could not understand the consequence of the actions they took. These cases help us to understand the aspects of humanity—our ability to plan and imagine—that underlie the requirement that an author must be human to receive copyright protection for their work.

On what we might call the other end of the spectrum, we need to consider a case in which, according to the court, “the parties agreed that the original authors of the work were celestial beings” (Urantia Foundation 1997, 956). In this dispute, the issue was an alleged infringement by an enthusiastic believer in the revelations supposedly made by these beings, after she published a guide to the revelatory documents that included their full text. The founding group of the Urantia Foundation, known as the Contact Commission, filed suit, and one of the defenses that was raised was that these works could not have copyright because there was not a human author. In these specific circumstances, the Ninth Circuit rejected this defense. As Christina Rhee (1998) notes, “because the members of the Contact Commission were, evidently, the first humans to arrange the information gathered from the celestial beings, the court found that they were entitled to the copyright in the book as a compilation” (72).

It is interesting to note here that the court is granting copyright because it found enough human involvement to recognize a copyright held by those humans. Yet the court itself realizes that their decision might have consequences for questions about computer-generated works:

Maaherra claims that there can be no valid copyright because the Book lacks the requisite ingredient of human creativity. The copyright laws, of course, do not expressly require “human” authorship, and considerable controversy has arisen in recent years over the copyrightability of computer-generated works. (Urantia Foundation 1997, 958)

In this 1997 decision, the court seems optimistic that it will be possible and desirable to grant copyright to humans over computer-generated work. Of course, that was before generative AI introduced a level of disconnection between what the human who prompts an AI engine expects and the output they receive. As Rhee (1998) observes in her note about the case, “If courts strictly apply the *Urantia* decision to future copyright claims over computer-generated works, they will grant the copyright to a human, most likely the computer program user” (76). She rightly observes, however, that a minimum level of creative input by humans will be necessary, and that is precisely what the Copyright Office and the *Thaler* court found lacking; the unpredictability of the output of generative AI seems to cast doubt on how much human creativity is really expressed in such output.

WORKS THAT COMBINE HUMAN AUTHORSHIP AND AI-GENERATED CONTENT

These cases stand firmly for the principle that authorship is a characteristic of humans. The precedential bar for changing this principle in U.S. law is very high, although, as we shall see shortly, there are some pressures to do so. While it is affirming to authors to know that they do not face competition from computers on the copyright front, it is also the case that many authors want to use generative artificial intelligence as part of their creative process. We need to briefly examine their situation.

The rule that the Copyright Office has always used for works that include public domain content is to grant copyright registration only for new material; a new copyright does not reawaken protection for public domain material that is included in a new work. Thus, a new edition of Dickens’ *Oliver*

Twist will get copyright protection for any new material—an introduction and notes, for example—but the new publication does not change the public domain status of the text of *Oliver Twist* itself.

This rule is reaffirmed and extended to works of generative AI in the Copyright Office’s report on the copyright status of AI products (United States Copyright Office 2025). It is most fully articulated, however, in a 2023 letter that the Copyright Office addressed to author Kristina Kashtanova (through her lawyer) about her graphic novel *Zarya of the Dawn* (United States Copyright Office 2023). As the letter makes clear, Ms. Kashtanova was the author of the text of the novel, but used an AI tool, Midjourney, to create the images. The letter summarizes the situation:

We conclude that Ms. Kashtanova is the author of the text as well as the selection, coordination and arrangement of the work’s written and visual elements. That authorship is protected by copyright. However, as we discuss below, the images in the work that were generated by Midjourney technology are not the product of human authorship. (United States Copyright Office 2023, 1)

After an extensive examination of how Midjourney works, “The [Copyright] Office concludes that the images generated by Midjourney contained within the work are not original works of authorship protected by copyright” (United States Copyright Office 2023, 8). The result of this painstaking examination is similar to the example we provided above about *Oliver Twist*; the human-authored parts of the work are protected, just like the new material in our imagined new edition of Dickens, but significant elements within it, the images, are not.

It is interesting to consider how far Ms. Kashtanova’s protection would extend and what that might mean for other situations in which generative AI is used to supplement or support human authorship. For *Zarya of the Dawn*, it is clear that Kashtanova could sue for infringement if her entire work were reproduced and/or distributed without her consent. She would also have protection against an appropriation of the text into another context, subject to the possibility of fair use. But another author could simply take the illustrations from *Zarya* and reuse them in a different work without liability because they are in the public domain; they are outside the scope of copyright protection.

While that situation may be fairly clear, it is increasingly the case that authors use generative AI as a tool to support their creative work in diverse ways. As I indicated at the start of this essay, I used an AI tool to generate an outline (which I did not use) and some references (a few of which I did use). These products of AI would not undermine in any way the fact that this essay is entirely a work of my original expression (i.e., authorship). But what if I had incorporated AI summaries of court cases? Those elements would not be protected, although the other parts would be. And what if I used AI as an editor? That would probably not undermine my copyright, although it would depend on how extensive the editing was. If the AI editor contributed any significant original expression to the work, it could restrict the scope of the protection the work would enjoy.

THE REQUIREMENT OF DISCLOSURE

As this situation grows increasingly complex, the Copyright Office depends on detailed and truthful disclosure in copyright registration applications. In its letter regarding *Zarya of the Dawn*, the Copyright Office notes that it was actually canceling the original copyright registration that was granted “because the application had not disclosed the use of artificial intelligence.” Such application, they said, “was incorrect, or at a minimum, substantively incomplete” (United States Copyright Office 2023, 2-3). In this early case, the Copyright Office did its own investigation of the circumstances surrounding the creation of the novel and ultimately granted copyright registration for the human-made elements, as described above. The problem is that more and more authors are using generative AI in more

and more diverse ways. Authors may not know how to disclose the role of AI on their registration applications, or they may dissemble in a misguided attempt to gain more protection over their work.

It is important to note why an inaccurate or dishonest registration would not ultimately benefit an author. Copyright, of course, covers work as soon as it is created. But the scope of that right would only really be defined if there was a conflict, where the judgement of the Copyright Office and the courts would determine, for the purpose of preventing infringement, if the original work did have protection in the first place. Registration is required to bring legal action for infringement, so if the registration application was inaccurate in some way, the adversarial nature of a lawsuit for infringement would probably bring those inaccuracies to light. Thus, the scope of the “automatic” copyright protection could be defined retroactively.

PRESSURES ON THE HUMAN-AUTHORSHIP STANCE

Although the position that copyright, as a body of law that governs human society, should be reserved for works of human expression, may seem logical, there are a lot of pressures on that apparently intuitive stance. In conversations about the fact that, at this moment, works generated by AI do not enjoy copyright, I have been surprised at the number of people who have immediately asserted that “they will have to change that!” Although it is by no means clear that the legal bodies involved can or will change this policy, that reaction is an instinctive recognition of the several pressures that complicate the current situation regarding copyright for the output of generative AI. Indeed, a recent online article in the *Harvard International Law Journal* wonders, “Why the obsession with human creativity?” and suggests that the US policy will indeed have to change (Ambartsumain and Cannon 2025).

One such pressure has already been discussed—the reliance of the Copyright Office on voluntary disclosure about the use of AI in order to process copyright registrations. As AI becomes increasingly integrated into the daily work of many kinds of creators, it is very likely that there will be inconsistency in the levels of disclosure. Such inconsistency is more likely to be the result of confusion, or even the natural attitude that AI is just another tool in a personal act of creation, than of intentional deception. Regardless of the reason, reliance on voluntary disclosure will inevitably make it difficult to apply the rule that only human authorship deserves copyright.

Related to this difficulty is the fact that it is becoming increasingly hard to tell works of human creativity and those generated by AI apart. Many commentators have noted this problem; Sarah Lu (2025), for example, has a nice discussion. In the case of *Zarya of the Dawn*, the Copyright Office had the luxury of doing its own investigation. But we are already at the point where that would be impossible because of the amount of time and labor involved, and because it would be nearly impossible to know, apart from what is disclosed, where the boundary between human and AI actually lies. As Lu expresses it, “this may inevitably make it more difficult to administer justice and result in differential treatment to creations of similar quality based on authorship” (87).

Differential treatment can result, as Lu (2025) notes, in unequal incentives. Copyright law is intended to create an incentive for creators so that a steady stream of art, writings and music will continue to “promote the progress of Science and the useful Arts” (U.S. Constitution, art. I, § 8). As generative AI becomes a bigger part of such creation, there is a danger that by denying copyright to the products of generative AI, we may create a negative incentive. The policy might “have an unfairly negative effect on the value of AIGC [artificial intelligence generated content]” (Lu 2025, 87). The result could be a disincentive to create, where the rewards for creative use of AI are diminished because

copyright is not available. This would be exactly the opposite impact on authorship that copyright law was intended to have. This pressure on the policy of granting copyright only for human authorship may be the most worrying of all.

Finally, although we have noted that most countries have some form of the policy that copyright is reserved for human authors, a lack of consistency internationally can also create pressure on the US policy. In several jurisdictions, it seems to be the case that an AI-generated work can be subject to copyright, where the right is held by the human person responsible for causing the computer to produce the work (see Gaidartzi and Stamatoudi 2025). In a similar way, an Internet Court in China has held that the person who prompted an AI tool to produce the outputted work can hold a copyright of that work (Wang, 2024; Ambartsumain and Cannon 2025). These differences may well lead to concerns over competitiveness, the argument being that those jurisdictions where copyright over AI works is available will have a competitive advantage over the United States in fostering investment in AI. This possibility creates an additional reason why the United States might feel pressured to modify or at least introduce more nuance into its policy about copyright and generative AI (see Gaidartz 2025).

INFRINGEMENT LIABILITY AND COMPETITION

The final broad question with which generative AI confronts copyright law, the issue of infringing output, really rehashes some of the input and output problems we have already discussed, so we can deal with it relatively quickly. It is simply not the case, based on experience so far, that generative AI tools produce a great deal of infringing content simply in response to ordinary prompting; their neural networks provide too many possible combinations for that to happen (see Lee 2024). If that did happen, the courts have been clear that such infringement would be evaluated by the same standard of “substantial similarity” that is used for other claims of infringement (see, for example, Andersen 2023, 11). What we do need to examine, however, regarding authors and authorship, are some ancillary questions raised by the potential of infringing output, no matter how unlikely that is. And even with non-infringing output, there are concerns about competition and dilution that legitimately worry authors.

In our discussion of the *Andersen* case above, we already noted the claim that if an author’s copyrighted work was used to train an AI model, all of the output from that model would potentially infringe that original work as unauthorized derivatives (Andersen 2023). Judges so far have dismissed that claim, asking that plaintiffs provide more concrete evidence both that their works were included in the training data and that some specific output from the generative AI tool actual shows substantial similarity to such works. Presenting such evidence would be almost impossible for most plaintiffs, since training data is usually only known by large groups of materials, and outputs are seldom actually substantially similar to prior work. The *New York Times* case suggests that such evidence could be presented, since *The New York Times* has been able to produce apparently verbatim copies of their articles (New York Times 2023). In that case, however, the claim is for direct infringement, so it will not test the theory that all outputs are derivative works.

Another oblique challenge regarding output that could make the claim for infringement is the theory advanced by Judge Chhabria in the *Kadrey* case regarding dilution. In his ruling on cross motions seeking summary judgment, the judge was clear that, on the record before him, using copyrighted works as training data was transformative fair use. Nevertheless, Judge Chhabria went to some length discussing a theory that the plaintiffs did *not* argue—that using copyrighted works for training input might be infringing if the output of the generative AI tool diluted the market for

those inputted works. The judge suggests that such dilution might impact the market harm factor in the fair use analysis, undermining his conclusion about fair use for training (Kadrey 2025). Because this argument was not raised by the plaintiffs, the judge does not base any conclusion or ruling on it, but he does seem to suggest that such a line of reasoning might be a way for evidence about output to impact the fair use analysis of the input in an AI model.

This dilution theory is quite new, and, if it were accepted, it could radically alter the fair use analysis in many cases. All new publications, of course, dilute the market for previously existing works. A new mystery novel, by adding another choice, dilutes the market for earlier mystery novels. A new commentary on Ephesians likewise dilutes the market for prior commentaries. Such dilution has never been considered the type of market harm that should be considered when analyzing fair use. Imagine, for example, that that new Ephesians commentary wants to show why a prior commentary misconstrued some aspects of the text. To do that, the new author would have to quote some part of the earlier work. Could the prior author bring an infringement action for those quotations and then argue that, because of market dilution, the fourth fair use factor on market harm should count against fair use? This argument runs counter to the constitutional purpose of copyright law to encourage creative production, and it has been rejected many times. Judge Chhabria's introduction of it in the early days of fair use analysis around generative AI is therefore very concerning. We do not know if other courts will follow this line of thinking—it seems unlikely—but if they did, it would radically redefine how authorship is conceived in the AI world.

Of course, fear of a dramatic increase in competition must surely underlie many authors' antipathy to generative artificial intelligence. The fear that inferior works (inferior, certainly, from the point of view of those human authors) could flood the market for commentaries on Ephesians must worry and discourage authors of potential new commentaries. One can sense that fear when reading the complaints in various cases against generative AI tools. A lot of space in those complaints is spent describing the uniqueness of the plaintiffs' works, their value, and the amount of labor that went in to creating them. To allow "easy" competition from generative AI is surely not fair. This fear of unlimited competition is clear in the *New York Times* case, which is deeply worried that people could access articles from *The New York Times* without paying the fees they have established (New York Times 2023). Likewise, the lawsuit against the AI music generation tool Suno expresses a significant fear that people could listen to songs like "Johnny B. Goode" without purchasing them from the rights owner (UMG Records 2024). In fact, the defendant's answer in this case deliberately calls the court's attention to that fear, arguing that the record companies were less concerned about potential copyright infringement than they were interested in preventing competition (UMG Records 2023).

A final way in which the issues of infringement and competition could impact authors comes from the settlement in the *Bartz* case, where the court found that works that were obtained for training purposes from unauthorized sources were not fairly used in the training. That is, most training of a large language model would be considered fair use, but not if those works were taken from an infringing source in the first place, like a so-called "pirate library" (Bartz 2025, 18–25). This ruling resulted in a large group of authors, those whose works became input data for Anthropic's AI tool from pirate sources, winning a significant infringement judgment. That win is the reason for a \$1.5 billion settlement between those authors and Anthropic (see Hansen 2025). So, a potential source of income for authors could arise from licensing their works for AI training. Authors cannot rely, of course, on having their works scraped up for AI training from a pirate library, but opportunities to proactively license their works are already beginning to arise, and some legislative proposals may codify this licensing opportunity for authors.

WHERE DO AUTHORS (AND LAWMAKERS) STAND TODAY?

Amidst the continuous change, the conflicts, and the ongoing uncertainty about how the technology will develop next, what is the current situation for authors regarding their copyrights and generative artificial intelligence? It is still early days, so all conclusions must be regarded as provisional. We have seen only one court decision that can be called final, and that is because of the settlement in the *Bartz* case, which really is an anomaly. Other than that, all of the lawsuits around generative AI have more to come, either to reach a final decision in the lower courts in which they were filed, or on appeal from lower court rulings. In this situation, it is hard to state any firm conclusions, but there are some broad trends emerging, which we will try to summarize here.

Although the *Bartz* case may be an anomaly, it certainly does remind us that authors could look for a new revenue stream from generative AI. The huge settlement in *Bartz*, the largest in the history of copyright law (see Hansen 2025), is not likely to be repeated, but the plaintiff authors in a great many of the lawsuits are seeking some kind of compensation from AI companies, as are their legislative advocates around the globe. Some publishers are already signing deals to sell content over which they own the rights for the purpose of AI training. Wiley is one such example, and they have said publicly that authors of works published by Wiley will be compensated “according to their contracts” (Battersby 2024). That is a very uncertain commitment, because in most cases it depends on how various publishers will interpret the royalty provisions in their contracts with authors, contracts which did not imagine or allow for generative AI when they were executed. Still, as new publication agreements are drafted, licensing revenue for AI training may be the most reliable form of authorial profit from AI.

One thing the progress of the cases so far does tell us is that generative AI is not a copyright-free zone of technological development. Throughout its history, copyright law has struggled, with more or less success, to keep up with and adapt to new technologies. Generative AI is no exception. We know several very basic things about the current situation. First, authors can enforce their rights against generative AI developers. Second, those rights will be subject to fair use, just as they are in all other contexts. And finally, we are going to see increasing reliance on contractual agreements, and even collective rights organizations, as we move ahead. Because of the scale of reuse in generative AI, it seems almost certain that schemes to compensate authors by requiring contributions to collective rights organizations will be developed. Such compulsory licenses have had mixed results in the past, but they are about the only way to compensate creators on the massive scale of generative AI.

There are, of course, many proposals about how to “fix” the copyright system so that authors receive legitimate recognition and compensation. The Congressional Research Service produced a good summary of the various options earlier this year (Harris 2025). We can identify three broad categories of concern in the legislative proposals so far—safety, transparency, and compensation.

Safety, in this context, refers to measures to prevent harm from AI tools and to control how sensitive information is handled by these tools. As Harris notes, most of the action in the US Congress so far has focused on administrative use and control of AI, so the goal has really been safety. In Europe, the EU’s AI Act also adopts a “risk-based approach,” requiring that AI systems be classified accord to the risks they pose in several distinct areas (Harris 2025).

Transparency focuses on the ability of authors to know if their work has been used in AI training. As we have seen, this is a major issue for authors in most of the pending lawsuits. In the same EU AI Act, specific transparency requirements are imposed on AI developers (Gaidartz 2025). This is a classic example where the devil will be in the details, however. It seems unlikely that AI developers will be able to produce lists of every specific work used in training an AI model, so transparency will

be accomplished on a larger scale, where collections are identified. Thus authors may still struggle to know if their specific works are part of a set of training data.

As mentioned above, compensation will almost certainly rely on collective rights organizations, if it is to happen at all. Several legislative and policy proposals have suggested this approach. There is one court case out of China, however, which has taken a different approach. In *Li v. Liu* (2023), the Beijing Internet Court not only granted copyright ownership to the person who developed the prompts that led Stable Diffusion to create a specific image, a ruling clearly at odds with US policy, it also ordered a subsequent user to pay direct compensation to that copyright owner (Gaidartz 2025). But such direct liability will certainly remain the exception, so any proposals for widespread compensation will have to rely on organizations built on the model of ASCAP and other existing collective rights organizations.

As the legislative approaches develop, one tension that is becoming clear is that between consent before a work is used in training and an “opt out” approach. Here the conflict between innovation and individual rights is especially clear. Although some wish to see a consent regime, where authors must give prior permission before their works can be used to train large language models, such a mandate would likely hamstring AI development. On the other hand, the EU’s 2019 Directive on Copyright already has an “opt out” approach for using materials in text and data mining, and, for now, that seems to also be how AI training is going to be handled (Lim 2025). As with other proposals, this effort to make the circumstances fairer for authors would also seem to increase the burden on those authors to pay careful attention to where their works end up and how they are being used.

CONCLUSION – AI, AUTHORSHIP AND CREATIVITY

All of these complicated scenarios focus on how an author’s work gets used by AI models. The situation for use and compensation will undoubtedly change, or at least become clearer, as some of these proposals are adopted. Of course, the other big issue is how authors use AI in their own creative process. The use of generative AI is already widespread in the creation of content, so we need to look carefully at the benefits and the risks.

The benefits of using generative AI are probably obvious. AI can produce combinations and possibilities much faster, and in greater variety, than a human brain can. But it also lacks the judgment that comes standard with the human brain. So, the best uses of generative AI will combine its speed and versatility with reasoned judgment. At the beginning of this article, I explained how I used a generative AI tool before I began writing. That process used the AI to fill gaps and make suggestions and was controlled because I retained the final decisions over everything the AI tool suggested. Most importantly, I informed all of you, the readers, about what I had done. Transparency is, I believe, the key to responsible AI use. In the legislative proposals about AI, transparency usually refers to disclosing the materials used to train a particular large language model. In the context of using AI, however, transparency means treating AI like any other source for a work by disclosing its use to those who read or view or use the creation. AI tools can offer support and even editorial assistance, but the notion of authorship should preclude using AI to create a complete work, and it should always include being transparent about what the AI was used to accomplish.

Using AI to produce a work for which the human user of the AI tool takes credit or even using the tool to do significant work on the final product without disclosure, is what we mean by AI plagiarism. The essay by Lemley and Oulette (2025) does a nice job of defining and distinguishing AI plagiarism from copyright infringement. The point is that plagiarism in this context is the same as in other

situations; it is not copyright infringement, but it is a violation of the community norms that exist in many academic or creative communities. This is at the core of what Olivia Guest and her coauthors (2025) are arguing in “Against the Uncritical Adoption of ‘AI’ Technologies in Academia.” Academic norms are strict and usually fairly clear. Taking credit for work the alleged author did not create is a violation, and even a betrayal, of the fundamental educational mission. In many contexts, academic and otherwise, such uncredited use of AI can lead to sanctions and professional disrepute. There are, for example, already many instances of attorneys getting into trouble for using AI to create court documents (see, for example, Yang 2025). Not only is this a deception practiced on the court, but, as in the case Yang describes, it can lead to providing false or misleading information to a court, which is always a cause for discipline.

To the list of rules for responsible use of AI for authors that we have already identified, namely providing careful human oversight and always disclosing the use to readers, we can add one more. Consider carefully the context for which you are creating and the expectations of that community of potential users; avoid using generative AI in a way that would disappoint or undermine those expectations. Your community probably does not care if you create a flyer for an event or a description of the bicycle that you have for sale using Claude or Copilot. But it will care very much if you are writing a column for the local newspaper or a legal brief to file in court and use these tools without appropriate acknowledgement.

At its most fundamental level, authorship is about communication within a community. A successful author recognizes, adopts, and also challenges the norms of the community for which they write. Authorship is always constructive and communal; we build on work that has gone before us, and we write as part of a conversation with a specific, albeit sometimes very broad, community. If authors keep in mind the responsibilities imposed on them by those who went before and those who will use their work subsequently, those considerations will guide them toward responsible authorship that uses, and limits, the role of generative artificial intelligence.

REFERENCES

- Ambartsumain, Yalena, and Maria T. Cannon. 2025. "Why the Obsession with Human Creativity? A Comparative Analysis on Copyright Registration of AI-Generated Works." *Harvard International Law Journal*, February 21, 2025. <https://journals.law.harvard.edu/ilj/2025/02/why-the-obsession-with-human-creativity-a-comparative-analysis-on-copyright-registration-of-ai-generated-works/>.
- Andersen v. Stability AI Ltd.* 2023. N.D. Cal. Order on Motions to Dismiss and to Strike. Filed October 30, 2023. https://storage.courtlistener.com/recap/gov.uscourts.cand.407208/gov.uscourts.cand.407208.117.0_1.pdf.
- Bartz et al. v. Anthropic.* 2025. N.D. Cal. Order on Fair Use. Filed June 23, 2025. <https://admin.bakerlaw.com/wp-content/uploads/2025/07/ECF-231-Order-on-Fair-Use.pdf>.
- Battersby, Mildred. 2024. "Wiley Set to Earn \$44m from AI Rights Deals, Confirms 'No Opt-Out' for Authors." *The Bookseller*, August 30. <https://www.thebookseller.com/news/wiley-set-to-earn-44m-from-ai-rights-deals-confirms-no-opt-out-for-authors>.
- BTLJ Blog. 2015. "Fair Use's Endless Defense: An Update on the Google Books Litigation." *Berkeley Technology Law Journal*, November 3. <https://btlj.org/2015/11/fair-uses-endless-defense-an-update-on-the-google-books-litigation/>.
- Burrow-Giles Lithography Co. v. Sarony.* 1884. 111 U.S. 53.
- Caldwell, Mackenzie. 2023. "What Is an 'Author'? Copyright Authorship of AI Art Through a Philosophical Lens." *Houston Law Review* 61 (2): 411-442.
- ChatGP is Eating the World.* n.d. Accessed September 16, 2025. <https://chatgptiseatingtheworld.com>.
- Craig, Carys. 2021. "AI and Copyright." In *Artificial Intelligence and the Law in Canada*, edited by Florian Martin-Bariteau and Teresa Scassa. Toronto: LexisNexis Canada.
- Cooper, A. Feder, and James Grimmelman. 2025. "The Files Are in the Computer: On Copyright, Memorization, and Generative AI." Cornell Legal Studies Research Paper. *Chicago-Kent Law Review* 100 (1): 141-219. <https://ssrn.com/abstract=4803118>.
- Gaidartzi, Anthi, and Irimi Stamatoudi. 2025. "Authorship and Ownership Issues Raised by AI-Generated Works: A Comparative Analysis." *Laws* 14 (4): 57-75. <https://doi.org/10.3390/laws14040057>.
- Guadamuz, Andres. 2025. "First Thought on GEMA v. OpenAI." LinkedIn (blog). https://www.linkedin.com/posts/andres-guadamuz_first-thoughts-on-gema-v-openai-i-agree-activity-7394119682972237824-XMc2/.
- Guest, Olivia, Marcela Suarez, Barbara Muller, Edwin van Meerkerk, et al. 2025. "Against the Uncritical Adoption of 'AI' Technologies in Academia." *PhilArchive*. <https://philarchive.org/rec/GUEATU>.
- Hansen, Dave. 2025. "The Bartz v. Anthropic Settlement: Understanding America's Largest Copyright Settlement." *Kluwer Copyright Blog*, November 10. https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/?utm_medium=email&utm_source=substack.

- Harris, Laurie. 2025. "Regulating Artificial Intelligence: U.S. and International Approaches and Considerations for Congress." *Congressional Research Service Report*, June 4. <https://www.congress.gov/crs-product/R48555>.
- Kadrey v. Meta Platforms*. 2025. N.D. Cal. Order Denying the Plaintiffs' Motion for Partial Summary Judgment and Granting Meta's Cross-Motion for Partial Summary Judgment. Filed June 25. <https://admin.bakerlaw.com/wp-content/uploads/2025/07/ECF-598-Order-on-Partial-Summary-Judgment.pdf>.
- Lee, Katherine, A. Feder Cooper, and James Grimmelman. 2024. "Talkin' 'Bout AI Generation: Copyright and the Generative-AI Supply Chain." *Journal of the Copyright Society* 72 (3): 251-409. <https://ssrn.com/abstract=4523551>.
- Lemley, Mark and Lisa Larrimore Oullette. 2025. "Plagiarism, Copyright and AI." *University of Chicago Law Review Online*. <https://lawreview.uchicago.edu/online-archive/plagiarism-copyright-and-ai>.
- Lim, Sarah (Sarang). 2025. "Copyright Law in the Age of AI: Navigating Authorship, Infringement, and Creative Rights." New York State Bar Association Advance Publication (June 20). <https://nysba.org/copyright-law-in-the-age-of-ai-navigating-authorship-infringement-and-creative-rights/>.
- Lu, Yiheng. 2025. "Reforming Copyright Law for AI-Generated Content: Copyright Protection, Authorship and Ownership." *Technology and Regulation* 2025: 81-95. <https://techreg.org/article/view/23039>.
- Naruda v. Slater*. 2018. No.16-15469 (9th Cir. 2018).
- New York Times Company v. Microsoft Corp et al.* 2023. S.D.N.Y. Complaint. Filed December 27. https://nytco-assets.nytimes.com/2023/12/NYT_Complaint_Dec2023.pdf.
- Rhee, Christina. 1998. "Urantia Foundation v. Maaherra." *Berkeley Technology Law Journal* 13: 69-82.
- Rose, Mark. 1993. *Authors and Owners: The Invention of Copyright*. Harvard University Press.
- Sag, Matthew. 2023. "Copyright Safety for Generative AI." *Houston Law Review* 61 (2): 295-348.
- Shapiro, James. 2016. *The Year of Lear: Shakespeare in 1606*. Simon & Schuster.
- Thaler v. Perlmutter*. 2023 687 F. Supp. 3d 140 (D.D.C. 2023). *Thomson Reuters v. Ross Intelligence*. 2025. D.Del. Memorandum Opinion. Filed February 11, 2025. https://www.ded.uscourts.gov/sites/ded/files/opinions/20-613_5.pdf.
- UMG Records et al. v. Suno, Inc.* 2024. D. Mass. Answer of Defendant Suno, Inc. to Complaint. Filed August 1. <https://www.musicbusinessworldwide.com/files/2024/08/SUNO-response-to-copyright-suit.pdf>.
- United States Copyright Office. 2023. *Letter re: Zarya of the Dawn (Registration # Vau001480196)*. February 21, 2023. <https://copyright.gov/docs/zarya-of-the-dawn.pdf>.
- United States Copyright Office. 2025. *Copyright and Artificial Intelligence: Part 2: Copyrightability*. <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf>.
- Urantia Foundation v. Maaherra*. 1997. 114 F. 3d 955 (9th Cir. 1997).
- U.S. Constitution, art. I, § 8.

Wang, Yuqian and Jessie Zhang. 2024. "Beijing Internet Court Grants Copyright to AI-Generated Image for First Time." *Kluwer Copyright Blog*. February 2. <https://legalblogs.wolterskluwer.com/copyright-blog/beijing-internet-court-grants-copyright-to-ai-generated-image-for-the-first-time/>.

Yang, Maya. 2025. "US Lawyer Sanctioned after Being Caught Using ChatGPT for Court Brief." *The Guardian*, May 31. <https://www.theguardian.com/us-news/2025/may/31/utah-lawyer-chatgpt-ai-court-brief>.

NOTES

- 1 As part of his reworking, Shakespeare apparently changed the spelling of Lear, perhaps because another play, *The Tragedy of King Leir and his Three Daughters*, was in contemporary circulation.
- 2 This is the way that Amanda Stent, the former director of the Davis Institute for AI at Colby College, framed the situation in an open letter to the campus in 2023.