

PUBLICATION

The Impact of the Supreme Court's *Goldsmith* Decision on Copyright Enforcement Against AI Tools

Authors: Edward D. Lanquist, Jr, Dominic A. Rota

July 27, 2023

Introduction

The U.S. Supreme Court's opinion in *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith* sent ripples through the legal and artistic communities. Months later, legal scholars and art journalists continue to debate whether the decision opens the door for federal courts to act as "art critics." Many, however, downplay how the Supreme Court's decision impacts the ways in which copyright owners may enforce their rights against generative AI tools.

Conversations around generative artificial intelligence (generative AI) are dominating the social stratosphere, as generative AI is regularly atop the headlines within the context of OpenAI's ChatGPT or Google's Bard. By simulating human cognitive thinking, generative AI can produce new types of text, imagery, audio, and synthetic data by using patterns and informational elements obtained from prior works.

Because generative AI often relies on pools of data and third-party creations to create new content, the community at large is concerned that generative AI may, whether intentionally or inadvertently, exploit copyright-protected content to develop purportedly original content. Although not currently being perceived as such, the Supreme Court's decision in *Goldsmith* just might provide insights on enforcement of copyright against generative AI's misuse of protected works.

The Supreme Court Clarifies Fair Use

In *Goldsmith*, the Supreme Court held that certain Andy Warhol silkscreen portraits of the musical artist Prince, which were derived from third-party photographs, constituted impermissible fair use. Provided below is a factual and procedural background of the case, followed by a discussion of the doctrine of fair use and the Supreme Court's application of it.

Factual Background and Procedural History

In 1981, Lynn Goldsmith, an acclaimed professional photographer, captured Prince in concert and in studio. Later, Goldsmith licensed to *Vanity Fair* the rights to a black and white photographic portrait of Prince for the purpose of serving as an "artist reference for an illustration." In turn, *Vanity Fair* hired Andy Warhol to create, for publication, a silkscreen portrait of Prince derived from Goldsmith's image. Andy Warhol created fifteen additional works derived from Goldsmith's photograph.

Goldsmith learned of Warhol's silkscreen portraits after *Vanity Fair's* parent company, Condé Nast, used one of Warhol's previously unpublished silkscreen portraits (Orange Prince) on the cover of a commemorative magazine following Prince's death in 2016. After Goldsmith notified the Andy Warhol Foundation (AWF) of her belief that AWF infringed her copyright, AWF sued Goldsmith and her agency for a declaratory judgment of noninfringement or, in the alternative, fair use. Goldsmith counterclaimed for infringement. For reference, the allegedly infringing Orange Prince is juxtaposed against Goldsmith's black-and-white photographic portrait.



The district court granted summary judgment in favor of AWF, finding that Warhol's silkscreen portraits made fair use of Goldsmith's photography. However, the Court of Appeals for the Second Circuit reversed and remanded, holding that all four fair use factors favored Goldsmith. The United States Supreme Court granted certiorari on a "narrow issue" decided by the lower courts – whether the purpose of Warhol sufficiently *transformed* Goldsmith's photograph to constitute fair use.

The Doctrine of Fair Use

The fair use doctrine is an objective inquiry into what a user does with an original work. Fair use, a common-law doctrine now codified under 35 U.S.C. § 107, balances the value of copying against the value of the original work. There are four factors to consider in determining whether the use made of a work in any particular case is a *fair use*. The first (and most important) of these factors explores "the purpose and character of the use," including whether such use is commercial or is for non-profit educational purposes.

The four statutory factors are typically considered together; however, in *Goldsmith*, AWF only challenged the Second Circuit's decision as to the first factor, having conceded the commercial nature of the license to Condé Nast. Consequently, the sole question before the Court was whether the purpose of Warhol's use of Goldsmith's photograph was sufficiently transformative to establish fair use.

Court's Prior Precedent and Reasoning

Prior to its decision in *Goldsmith*, the Court held that a use is "transformative" when it alters the first work with new expression, meaning, or message. In the landmark decision *Campbell v. Acuff-Rose Music, Inc.*, the Supreme Court determined 2 LIVE CREW's parody of Roy Orbison's "Oh, Pretty Woman" constituted fair use because 2 LIVE CREW's version went beyond a mere "derivative" of Orbison's original. The *Campbell* Court

held that the purpose and character of the new use was distinctly different than the original, notwithstanding the commercial nature of the parody.

However, in *Goldsmith*, the Court narrowed the standard from *Campbell*, finding that a use which merely adds some new expression, meaning, or message is **alone** insufficient to satisfy the first fair use factor. The Court declared that a contrary rule would swallow the copyright owner's exclusive right to prepare derivative works. Consequently, when the commercial nature of a new use is undisputed, and the purpose of the new use is similar to the original, additional justification is needed to satisfy the first fair use factor.

Having narrowed the standard, the Court determined that the purposes for Goldsmith's photograph and Warhol's silkscreen were largely the same, as both works were intended to depict "portraits of Prince used in magazines to illustrate stories about Prince." Because AWF provided no further compelling justification, the Court determined Warhol's new use failed the transformative test.

Generative AI and its Applications

The *Goldsmith* decision is sure to have an impact beyond the propriety of Warhol's works. Although the full implications of this decision remain uncertain, the Court's rationale will likely have a significant impact on policing infringing content created using generative AI tools. So, how do generative AI tools present a threat to copyrighted material?

Generative AI consists of algorithms, or neural networks, that use training data to create new content in the form of text, images, or audio (See IBM's "[What is Generative AI?](#)"). These tools are driven by various mechanisms, such as deep learning models, large language models, natural language processing, and diffusion models, that scour over training data to generate new content. The training data often consists of webpages, books, articles, and other publicly available resources. Of course, many of these resources are copyrighted material, and the generative AI tools' use of this material, to create new content opens the door to claims of copyright infringement against the developers or end users of the AI tools.

Copyright Infringement and the Fair Use Defense

To prove copyright infringement, a copyright-owner must show that the alleged infringer had access to the copyrighted work and that the allegedly infringing work is *substantially similar* with the copyrighted work (See 3 Paul Goldstein, *Goldstein on Contracts* § 9.2.1 [2023]). As with the case of several open-market generative AI tools, such as OpenAI's suite of platforms, generative AI tools have direct access to copyrighted works. These AI tools are often driven by crawling (or "scraping") mass quantities of information, as made publicly available through the internet. Because these tools pull information from openly accessible resources, the training data for these generative AI tools relies on information that is otherwise protected by copyrights, trademarks, and other intellectual property regimes (and combinations thereof).

By ingesting training data containing copyright-protected content, generative AI run the risk of producing outputs that are *substantially similar* to copyrighted material owned by third parties. Therefore, generative AI outputs present an imminent risk for infringement of copyrighted material (and other intellectual property) by way of ingesting the training data and generating content based on the training data.

If there is a path forward for the use of generative AI, there must be an intellectual property defense, or justification, for the use of AI tools. Otherwise, without the requisite guardrails, this groundbreaking technology may become the crux of never-ending intellectual property litigation.

Under the current framework, fair use likely presents the best defense. Unlike other defenses, fair use implicitly acknowledges that a work copies the protected expression of a copyrighted material without authorization. But the question of fair use almost always points to the first factor: whether the use of the copyrighted expression

was sufficiently "transformative." In the context of AI, the question is whether the outputs of generative AI tools are sufficiently transformative of original works to justify the tools' fair use of the copyrighted material.


After the Supreme Court's decision in *Goldsmith*, the fair use defense may be a non-starter. Unless the outputs are used for a different purpose than the original works, it may be difficult to show that the AI-generated work did not otherwise misappropriate the protected expression of the copyrighted work.

Risks and Opportunities for Generative AI Tools

Application and evolution of the fair use defense, in the context of generative AI, continue to play out in real-time as AI platforms face infringement litigation from all areas of the creative community.

OpenAI and Stability AI, both popular generative AI platforms, have recently faced lawsuits alleging copyright infringement.¹ These lawsuits claim that the generative AI models' use of copyrighted material, by and through its ingestion (and/or use) of the training data, constitutes infringement of the copyright owner's protected copyrights.

In a lawsuit filed in federal court in Delaware, Getty alleged that Stable Diffusion has copied more than twelve million copyrighted images belonging to Getty. Getty provided one such example of the alleged infringement of its images:



Given that Stable Diffusion had access to Getty's online portfolio images and seeing at least some *substantial similarity* between the AI-generated work and the original Getty image, Stable Diffusion faces a tall task in making a case for fair use under *Goldsmith*. Making this more difficult, Stable Diffusion's AI tool also produced a distorted version of Getty's watermark. It is unlikely that the distorted watermark supports any proposition that the generative AI output sufficiently transformed Getty's copyrighted material. In fact, the watermark very likely evidences the tool's use of the image for a similar purpose – to produce a commercially viable image that may be licensed in the same way Getty images are licensed.

This case poses an illuminating quagmire for AI platforms after the Supreme Court's decision in *Goldsmith*: can Stable Diffusion successfully argue that the two photos were for different purposes under the first factor of fair use? If not, the tool's use of the copyrighted image may constitute infringement.

Visual-based AI tools are not the only AI platforms that have come under legal scrutiny, and AI platforms built on text-based inputs and outputs are also facing challenges. A group of authors, including comedian Sarah Silverman, recently filed class-action lawsuits against OpenAI and Meta, alleging, among other things, that the AI tools' use of their copyrighted books in training data constitutes copyright infringement. In an exhibit to the complaint, the plaintiffs provided interrogations of OpenAI's tool ChatGPT, which responded by providing the user with an accurate summary of the authors' copyrighted material:

> Summarize in detail the first part of "The Bedwetter" by Sarah Silverman

"The Bedwetter: Stories of Courage, Redemption, and Pee" is a memoir by comedian Sarah Silverman, published in 2010. The book is divided into various chapters and covers Silverman's life from her early childhood to her career in comedy.

The first part of the book largely deals with Silverman's childhood and formative years, which were rife with several notable incidents that shaped her life and career.

Born in 1970, Silverman grew up in Bedford, New Hampshire, in a Jewish family. Her father, Donald Silverman, owned a clothing store, while her mother, Beth Ann, was a photographer and theater enthusiast. Silverman has three sisters, one of whom is a rabbi, and her upbringing in a somewhat unconventional family influenced her sense of humor and perspective on life.

As a child, Silverman had a rebellious streak and a proclivity for the unconventional, which would later come to define her stand-up persona. For instance, at a young age, she was already fond of pushing boundaries with her humor and was not averse to being the center of attention.

While the plaintiffs admitted that the tool did not reproduce the exact contents of the copyrighted material, the ChatGPT prompt demonstrates evidence that the tools had access to the copyrighted material and used the material to generate a *substantially similar* output.

Like the dispute over Stability AI's photo-generated outputs, if OpenAI fails to demonstrate that the tool's use of the copyrighted material is for purposes different than that of the authors' initial creation, the fair use defense for infringement may be a non-starter.

However, the *Goldsmith* decision is not all bad news for the future of generative AI. Although the Court's rationale provides fertile soil for new infringement claims against producers and uses of generative AI technologies and generated material, the Court signaled a potential silver lining. If the *purpose* of including copyrighted material in training data is substantially different than the *purpose* of the original works, then developers of AI tools (and the users thereof) may have a pathway to establishing fair use.

The future trajectory of generative AI will likely hinge on the key question, "what is the *purpose* of training data?" Because *Goldsmith* did not answer this exact question, lower courts are left to grapple with whether the *purpose* of training data is different *enough* from that of the original copyrighted works to justify a fair use defense, or whether using copyrighted works as training data is *ipso facto* copyright infringement.

Technical Fixes for Generative AI Tools

Until courts refine the analysis regarding infringement based on AI tools' use of copyrighted material, copyright holders are left with conventional mechanisms to enforce their protected copyright against generative AI platforms (and the users thereof).

As with other technological innovations, the response may interweave legal analysis and frameworks with technological systems designed to limit misappropriation of third-party intellectual property. For example, Google reports that it is developing, and has developed, its Bard system in such a way that it will be more controlled in what is used to teach the large-language model.² As another example, the University of Chicago has released free software, titled "Glaze," which is designed to thwart copying of a visual work by generative AI.³ Similarly, other companies are developing tools to *proof* protected copyright works from ingestion by AI tools.

As the courts continue to grapple with these queries, protecting copyrighted works from misuse by AI platforms will require a concerted effort from the courts, Congress, the creative community, and technology companies. In addition to the courts' development of precedent centered around AI and the fair use defense, Congress, or the U.S. Copyright Office, will be instrumental in revisiting currently enforceable statutes and regulations in view of AI engines. Further, the creative community and technology companies will need to work together and with Congress to prioritize copyright protection and empower creative contributions from humans, rather than generative AI tools.

Conclusion

The Supreme Court's decision in *Goldsmith* poses significant implications for the future of generative AI and highlights the risk of infringement litigation for producers and users of generative AI. Consequently, organizations that leverage generative AI tools should be mindful of how the tools are used in a commercial context to mitigate the risk of infringing uses. Likewise, owners of intellectual property should be aware of how their works are used by generative AI models and the users of these tools, and timely action should be taken to defend intellectual property against infringement.

For more information on this or other intellectual property matters please contact [Edward D. Lanquist](#), [Dominic Rota](#), or a member of Baker Donelson's [Intellectual Property](#) and [AI](#) teams.

Tyler Dove and Rebecca Villanueva, summer associates at Baker Donelson, contributed to this alert.

¹ Amended Complaint, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135-GBW (D. Del. Mar. 29, 2023), ECF No. 13; Complaint, *Silverman et al v. OpenAI, Inc. et al*, No. 4:23-cv-03416 (N.D. Cal. Jul. 7, 2023), ECF No. 1.

² Joe Toscano, *ChatGPT or Google Bard? Privacy or Performance? Outstanding Questions Answered*, FORBES (June 24, 2023), available at <https://www.forbes.com/sites/joetoscano1/2023/06/24/chatgpt-or-google-bard-privacy-or-performance-outstanding-questions-answered/>.

³ *What is Glaze*, Glaze, available at <https://glaze.cs.uchicago.edu/> (last visited July 17, 2023).