



Digital Article / Intellectual Property

Generative AI Has an Intellectual Property Problem

Strategies to help companies mitigate the legal risk and ensure they're in compliance with the law. *by Gil Appel, Juliana Neelbauer, and David A. Schweidel*

Published on HBR.org / April 7, 2023 / Reprint [H07K15](#)



HBR Staff/Pexels

Generative AI can seem like magic. Image generators such as Stable Diffusion, Midjourney, or DALL·E 2 can produce remarkable visuals in styles from aged photographs and water colors to pencil drawings and Pointillism. The resulting products can be fascinating — both quality and speed of creation are elevated compared to average human performance. The Museum of Modern Art in New York hosted an AI-

generated installation generated from the museum's own collection, and the Mauritshuis in The Hague hung an AI variant of Vermeer's Girl with a Pearl Earring while the original was away on loan.

The capabilities of text generators are perhaps even more striking, as they write essays, poems, and summaries, and are proving adept mimics of style and form (though they can take creative license with facts).

While it may seem like these new AI tools can conjure new material from the ether, that's not quite the case. Generative AI platforms are trained on data lakes and question snippets — billions of parameters that are constructed by software processing huge archives of images and text. The AI platforms recover patterns and relationships, which they then use to create rules, and then make judgments and predictions, when responding to a prompt.

This process comes with legal risks, including intellectual property infringement. In many cases, it also poses legal questions that are still being resolved. For example, does copyright, patent, trademark infringement apply to AI creations? Is it clear who owns the content that generative AI platforms create for you, or your customers? Before businesses can embrace the benefits of generative AI, they need to understand the risks — and how to protect themselves.

Where Generative AI Fits into Today's Legal Landscape

Though generative AI may be new to the market, existing laws have significant implications for its use. Now, courts are sorting out how the laws on the books should be applied. There are infringement and rights of use issues, uncertainty about ownership of AI-generated works, and questions about unlicensed content in training data and whether users should be able to prompt these tools with direct reference other

creators' copyrighted and trademarked works by name without their permission.

These claims are already being litigated. In a case filed in late 2022, *Andersen v. Stability AI et al.*, three artists formed a class to sue multiple generative AI platforms on the basis of the AI using their original works without license to train their AI in their styles, allowing users to generate works that may be insufficiently transformative from their existing, protected works, and, as a result, would be unauthorized derivative works. If a court finds that the AI's works are unauthorized and derivative, substantial infringement penalties can apply.

Similar cases filed in 2023 bring claims that companies trained AI tools using data lakes with thousands — or even many millions — of unlicensed works. Getty, an image licensing service, filed a lawsuit against the creators of Stable Diffusion alleging the improper use of its photos, both violating copyright and trademark rights it has in its watermarked photograph collection.

In each of these cases, the legal system is being asked to clarify the bounds of what is a “derivative work” under intellectual property laws — and depending upon the jurisdiction, different federal circuit courts may respond with different interpretations. The outcome of these cases is expected to hinge on the interpretation of the fair use doctrine, which allows copyrighted work to be used without the owner's permission “for purposes such as criticism (including satire), comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research,” and for a transformative use of the copyrighted material in a manner for which it was not intended.

This isn't the first time technology and copyright law have crashed into each other. Google successfully defended itself against a lawsuit

by arguing that transformative use allowed for the scraping of text from books to create its search engine, and for the time being, this decision remains precedential.

But there are other, non-technological cases that could shape how the products of generative AI are treated. A case before the U.S. Supreme Court against the Andy Warhol Foundation — brought by photographer Lynn Goldsmith, who had licensed an image of the late musician, [Prince](#) — could refine U.S. copyright law on the issue of when a piece of art is sufficiently different from its source material to become unequivocally “transformative,” and whether a court can consider the meaning of the derivative work when it evaluates that transformation. If the court finds that the Warhol piece is not a fair use, it could mean trouble for AI-generated works.

All this uncertainty presents a slew of challenges for companies that use generative AI. There are risks regarding infringement — direct or unintentional — in contracts that are silent on generative AI usage by their vendors and customers. If a business user is aware that training data might include unlicensed works or that an AI can generate unauthorized derivative works not covered by fair use, a business could be on the hook for willful infringement, which can include damages up to \$150,000 for each instance of knowing use. There’s also the risk of accidentally sharing confidential trade secrets or business information by inputting data into generative AI tools.

Mitigating Risk and Building a Way Forward

This new paradigm means that companies need to take new steps to protect themselves for both the short and long term.

AI developers, for one, should ensure that they are in compliance with the law in regards to their acquisition of data being used to train

their models. This should involve licensing and compensating those individuals who own the IP that developers seek to add to their training data, whether by licensing it or sharing in revenue generated by the AI tool. Customers of AI tools should ask providers whether their models were trained with any protected content, review the terms of service and privacy policies, and avoid generative AI tools that cannot confirm that their training data is properly licensed from content creators or subject to open-source licenses with which the AI companies comply.

Developers

In the long run, AI developers will need to take initiative about the ways they source their data — and investors need to know the origin of the data. Stable Diffusion, Midjourney and others have created their models based on the LAION-5B dataset, which contains almost six billion tagged images compiled from scraping the web indiscriminately, and is known to include substantial number of copyrighted creations.

Stability.AI, which developed Stable Diffusion, has announced that artists will be able to opt out of the next generation of the image generator. But this puts the onus on content creators to actively protect their IP, rather than requiring the AI developers to secure the IP to the work prior to using it — and even when artists opt out, that decision will only be reflected in the next iteration of the platform. Instead, companies should require the creator's opt-in rather opt-out.

Developers should also work on ways to maintain the provenance of AI-generated content, which would increase transparency about the works included in the training data. This would include recording the platform that was used to develop the content, details on the settings that were employed, tracking of seed-data's metadata, and tags to facilitate AI reporting, including the generative seed, and the specific prompt that was used to create the content. Such information would not only allow for the reproduction of the image, allowing its

veracity to be verified easily, but it would also speak to the user's intent, thereby protecting business users that may need to overcome intellectual property infringement claims, as well as demonstrate that the output was not due to willful intent to copy or steal.

Developing these audit trails would assure companies are prepared if (or, more likely, when) customers start including demands for them in contracts as a form of insurance that the vendor's works aren't willfully, or unintentionally, derivative without authorization. Looking further into the future, insurance companies may require these reports in order to extend traditional insurance coverages to business users whose assets include AI-generated works. Breaking down the contributions of individual artists who were included in the training data to produce an image would further support efforts to appropriately compensate contributors, and even embed the copyright of the original artist in the new creation.

Creators

Both individual content creators and brands that create content should take steps to examine risk to their intellectual property portfolios and protect them. This involves proactively looking for their work in compiled datasets or large-scale data lakes, including visual elements such as logos and artwork and textual elements, such as image tags. Obviously, this could not be done manually through terabytes or petabytes of content data, but existing search tools should allow the cost-effective automation of this task. New tools can even promise obfuscation from these algorithms.

Content creators actively should monitor digital and social channels for the appearance of works that may be derived from their own. For brands with valuable trademarks to protect, it's not simply a matter of looking for specific elements such as the Nike Swoosh or Tiffany Blue. Rather, there may be a need for trademark and trade dress monitoring

to evolve in order to examine the style of derivative works, which may have arisen from being trained on a specific set of a brand's images. Even though critical elements such as a logo or specific color may not be present in an AI-generated image, other stylistic elements may suggest that salient elements of a brand's content were used to produce a derivative work. Such similarities may suggest the intent to appropriate the average consumer's goodwill for the brand by using recognizable visual or auditory elements. Mimicry may be seen as the sincerest form of flattery, but it also can suggest the purposeful misuse of a brand.

The good news regarding trademark infringement for business owners is that trademark attorneys have well-established how to notify and enforce trademark rights against an infringer, such as by sending strongly worded cease-and-desist notice or licensing demand letter, or moving directly to filing a trademark infringement claim, regardless of whether an AI platform generated the unauthorized branding, or a human did.

Businesses

Businesses should evaluate their transaction terms to write protections into contracts. As a starting point, they should demand terms of service from generative AI platforms that confirm proper licensure of the training data that feed their AI. They should also demand broad indemnification for potential intellectual property infringement caused by a failure of the AI companies to properly license data input or self-reporting by the AI itself of its outputs to flag for potential infringement.

At minimum, businesses should add disclosures in their vendor and customer agreements (for custom services and products delivery), if either party is using generative AI to ensure that intellectual property rights are understood and protected on both sides of the table as well as how each party will support registration of authorship and ownership of those works. Vendor and customer contracts can include AI-related

language added to confidentiality provisions in order to bar receiving parties from inputting confidential information of the information-disclosing parties into text prompts of AI tools.

Some leading firms have created generative AI check lists for contract modifications for their clients that assess each clause for AI implications in order to reduce unintended risks of use. Organizations that use generative AI, or work with vendors that do, should keep their legal counsel abreast of the scope and nature of that use as the law will continue to evolve rapidly.

...

Going forward, content creators that have a sufficient library of their own intellectual property upon which to draw may consider building their own datasets to train and mature AI platforms. The resulting generative AI models need not be trained from scratch but can build upon open-source generative AI that has used lawfully sourced content. This would enable content creators to produce content in the same style as their own work with an audit trail to their own data lake, or to license the use of such tools to interested parties with cleared title in both the AI's training data and its outputs. In this same spirit, content creators that have developed an online following may consider co-creation with followers as another means by which to source training data, recognizing that these co-creators should be asked for their permission to make use of their content in terms of service and privacy policies that are updated as the law changes.

Generative AI will change the nature of content creation, enabling many to do what, until now, only a few had the skills or advanced technology to accomplish at high speed. As this burgeoning technology develops, users must respect the rights of those who have enabled its creation – those very content creators who may be displaced by it. And while we

understand the real threat of generative AI to part of the livelihood of members of the creative class, it also poses a risk to brands that have used visuals to meticulously craft their identity. At the same time both creatives and corporate interests have a dramatic opportunity to build portfolios of their works and branded materials, meta-tag them, and train their own generative-AI platforms that can produce authorized, proprietary, (paid-up or royalty-bearing) goods as sources of instant revenue streams.

This article was originally published online on April 7, 2023.



Gil Appel is an assistant professor of marketing at the George Washington University School of Business. His research focuses on digital innovations and its impact on consumers.



Juliana Neelbauer is a partner at Fox Rothschild LLP in the corporate, intellectual property, emerging markets, and entertainment and sports law groups. She is a lecturer at the University of Maryland and Georgetown University regarding securities law, negotiations, digital assets, and business law.



David A. Schweidel is the Goizueta Chair in Business Technology and a professor of marketing at Emory University's Goizueta Business School.