AI and IP: what companies need to know to avoid legal landmines

David Kappos and Callum Sproule
Cravath Swaine & Moore

31 July 2025

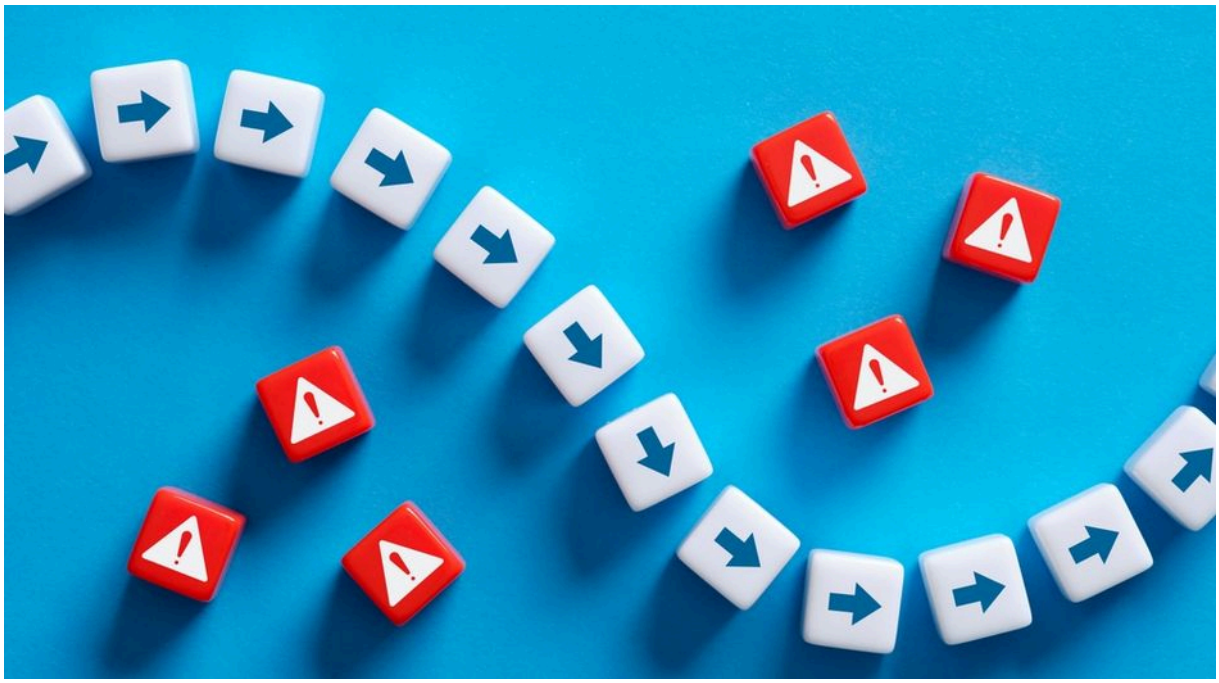

Shutterstock/Cagkan Sayin

- **Until user prompts can sufficiently control expressive elements in AI-generated outputs to reflect human authorship, generative AI outputs may be uncopyrightable in the US**
- **A natural person must have significantly contributed to every claim in a patent or patent application or risk the patent being cancelled**
- **There is a real question as to whether any agentic AI-generated output can be protected under current US copyright and patent frameworks**

AI in business is no longer theoretical. For companies across sectors, it is already reshaping core business practices and surfacing complex IP-related questions. A recent **McKinsey report** found that 78% of respondents say that their organisations use AI in at least one business function as of late 2024 (up from 55% in 2023), including 71% of respondents reporting the use of generative AI (as opposed to merely analytical AI) for a growing share of business functions, including the creation of text, images, code and video.

While the use of generative AI presents exciting opportunities, it also introduces unprecedented IP risk. AI-generated content is showing up at the centre of transactions, in internal company platforms and as part of marketing campaigns. Such content is often treated as protectable IP when it may not be. Unresolved questions of authorship, ownership and infringement now sit at the core of many companies' most valuable assets.

Companies and investors need to address these risks now. Delay invites downstream legal uncertainty and IP disputes, particularly where AI-generated content is being treated as valuable proprietary IP.

The sooner these risks are identified and addressed, the stronger and more future-proof a company's IP position will be. This article outlines where AI risks often arise, and what companies and investors can do now to avoid IP landmines in the future.

Copyright

In the United States, it is well established that copyright will only protect human-created works. The US Copyright Office (USCO) has spoken clearly: AI-generated content generally does not qualify, even if that content is curated or modified by a human.

In one February 2023 **decision**, the USCO found that artwork created through the use of the Midjourney generative AI tool was "not the product of human authorship" and therefore not copyrightable. Guiding image creation through prompt engineering, and curating even hundreds of iterations of output, is insufficient to deem users to be authors for copyright purposes, as users "lack sufficient control" over generated images to be treated as the "master mind" behind them.

In another **decision**, in September 2023, the USCO found that where a work contains "more than a *de minimis* amount of AI-generated content", even with extensive manual modification by the user, the AI-generated content that "remain[ed] in substantial form in the final [w]ork" is not the product of human authorship and therefore not copyrightable.

These requirements are further explained in the USCO's 2023 **AI Registration Guidance**, which requires that applicants disclose and describe, for works containing more than a *de minimis* amount of AI-generated material, the human author contributions to such works. The position set forth in these USCO decisions and the *AI Registration Guidance* was recently affirmed by the US Court of Appeals for the DC Circuit, which unanimously **found** that the USCO's denial of a copyright application for an image generated by a generative AI tool was appropriate because copyright can only protect works of human authorship. The court also rejected arguments that AI-generated content could constitute "work for hire" on behalf of a human being.

A newly issued USCO report entitled "**Copyright and Artificial Intelligence**" recently clarified what kinds of content created or modified by AI-generated tools can be copyrightable. The report distinguishes between use of AI tools to “assist in the creation of works” or as a “brainstorming tool” on the one hand, and to serve as “a stand-in for human creativity” on the other.

In the latter case, even highly detailed or repeatedly revised prompt engineering is akin to issuing instructions to an author, which the courts have long established as insufficient to grant a right of copyright to the delegating party.

Conversely, where a user “select[s] or arrange[s] AI-generated material in a sufficiently creative way” the work as a whole may constitute an original work of authorship (eg, as a compilation, though the AI-generated material contained within such compilation would not be protectable outside of such work), or a user may modify AI-generated output “to such a degree that **the modifications** meet the standard for copyright protection” (*emphasis added*).

In either case, copyright will only protect the human-authored aspects of the work, which are independent of and do not affect the copyright status of the AI-generated material itself.

On balance, until user prompts can sufficiently control expressive elements in AI-generated outputs to reflect human authorship, generative AI outputs may be uncopyrightable under current law. For businesses using generative AI tools for the creation of content, this restriction causes uncertainty around ownership and enforceability.

In general, companies should avoid overreliance on generative AI output being protected as proprietary IP unless human authorship is documented, dominant and can be proven.

Companies that utilise AI-generated content as part of broader works (eg, generating song lyrics for a human singer) must carefully parse elements of such works that do or do not qualify for copyright protection. Regardless of use case, content created by generative AI tools should be adequately documented and special consideration should be given to the positioning of AI-generated content in a company’s broader IP strategy, including whether the use of AI output undermines the potential strength of copyrights on derivative works.

A second-order question is whether the use of AI outputs generated from training data containing copyrighted material without authorisation of the copyright holder further impacts the copyrightability of generative AI outputs. The USCO has not taken a definitive stance on this issue, but, in May 2025, took the unusual step of releasing a “pre-publication” version of its Part III of the **Copyright and Artificial Intelligence Report**.

Part III of the report found, among other things, that large-scale data scraping likely does not fully satisfy the requirements for “fair use” under copyright doctrine. Lawsuits brought against numerous generative AI developers are ongoing. Copyright claims are frequently surviving motions to dismiss.

Companies utilising generative AI tools should monitor developments in this area closely. In addition, there is a latent risk of AI-generated code resembling or replicating open-source software and thereby triggering copyleft obligations (eg, prohibitions on monetisation or distribution rights). Consequently, AI-generated code should be made subject to appropriate copyleft **compliance monitoring**.

Patents

The US Patent Act unambiguously **requires** that inventors be natural persons; AI systems cannot be inventors for the purposes of patent eligibility. The USPTO’s 2024 **Inventorship Guidance for AI-Assisted Inventions** states that AI-assisted inventions are not “categorically unpatentable”, but a natural person must have provided a significant contribution to the invention. The guidance relies on the three-factor *Pannu* test developed by the Federal Circuit in 1998 and originally applied to joint inventors.

The *Pannu* test for determining inventorship states that, to be deemed to have made a “significant contribution” and thus be eligible for patent protection, an inventor must:

- “contribute in some significant manner to the conception or reduction to practice of the invention";
- "make a contribution to the claimed invention that is not insignificant in quality, when that contribution is measured against the dimension of the full invention"; and
- "do more than merely explain to the real inventors well-known concepts and/or the current state of the art”.

Where an inventor uses AI-assistance in making an invention, that inventor must satisfy the *Pannu* factors for the invention to be patent-eligible, and only natural persons may be named on the patent application (ie, AI systems cannot be named inventors).

The guidance sets forth the following principles for the application of the *Pannu* factors to an AI-assisted invention:

- Use of AI does not negate inventorship where a natural person makes a significant contribution under the *Pannu* test.
- Prompting an AI system with a problem and receiving an answer does not constitute a “significant contribution”.

- Reduction to practice of an invention, standing alone, is insufficient; a natural person must have contributed to the invention's conception or otherwise made significant contributions to the AI output.
- A natural person who "develops an essential building block from which the claimed invention is derived" may be deemed to have provided a significant contribution to the conception of the claimed invention (thus, the natural person who designs, builds or trains an AI-system in view of a specific problem to elicit a particular solution could be an inventor).
- Maintaining "intellectual dominion" over an AI system does not, on its own, make a person an inventor of any inventions created through the use of the AI system.

The guidance further notes that patent examiners may request the submission of information reasonably necessary to properly examine a pending application or reexamination request, including where there is a reasonable basis to believe that one or more named inventors may not have contributed significantly to the claimed subject matter (eg, where an invention is the product of AI output).

A natural person must have significantly contributed to every claim in a patent or patent application. Where it is determined that no natural person significantly contributed to a claimed invention, the patent claim must be cancelled or amended.

Consequently, companies should carefully evaluate whether AI-assisted inventions are patent-protectible, and, where applicable, ensure that human inventorship can be clearly demonstrated under the *Pannu* factors.

Companies that utilise AI-generated content as part of their inventive processes (eg, incorporating AI in drug discovery to identify potential drug candidates) must carefully parse elements of such works that do or do not qualify for patent protection. Regardless of use case, and whether in the context of business acquisitions, patent assignments or exclusive licences, buyers should seek appropriate representations that the underlying inventions satisfy the requirements for patent eligibility, including human authorship.

Trademarks

Unlike in the case of copyrights and patents, no agency guidance has been issued relating to the use of generative AI tools in creating registered trademarks.

Among the open issues regarding AI and trademarks, two stand out prominently:

- whether a human authorship requirement should be imposed; and
- whether the creation of a new trademark using generative AI trained on existing trademarks can become evidence of infringement of a pre-existing trademark by an applicable AI-generated trademark.

As for the first issue, US courts have not yet examined whether, in light of recent AI capabilities, trademarks should have a human authorship requirement, as in the case of copyrights and patents.

Given the statutory basis and function of trademarks in protecting against consumer confusion, it seems less likely that a human authorship requirement would be imposed.

Regarding evidence of infringement, generative AI output is reliant on underlying training data, which in many cases includes trademark-protected brands and trade dress.

Future challenges to AI-generated marks may argue that the use of generative AI tools trained on data sets that include protected third-party marks demonstrates infringement where AI-generated marks include similar design elements as those found in the protected marks.

Companies should exercise care to ensure that AI-generated outputs do not infringe existing third-party trademarks. Companies leveraging generative AI should monitor legal developments closely to get ahead of any developments that threaten to undermine registered trademarks comprised of AI-generated outputs.

Companies should proactively document human contributions to AI-generated outputs that are intended for trademark use or registration and, where practicable, utilise generative AI tools that are not trained on data that includes registered third-party marks.

Trade secrets and ownership of GenAI outputs

In general, under the US Defend Trade Secrets Act, any valuable information that is not known outside of a business entity (ie, is known only by employees and others involved in the business), is subject to reasonable measures to guard the secrecy of the information, and is difficult for others to acquire or duplicate independently may be protectable under law as a trade secret.

In most cases, reliance on trade secret protections is an alternative to registered copyright or patent protection (each of which requires detailed public disclosure). And, unlike in the case of copyright or patent protection, there is no requirement that trade secrets be originally developed by a human author.

Consequently, many companies are **increasingly protecting** AI-generated outputs as trade secrets as an alternative to registering that content with copyrights or patents. However, protecting AI-generated outputs as trade secrets is not without its own set of risks.

First, trade secrets can generally only be leveraged to protect internal company information, so AI-generated content created for public distribution is not protectable.

In addition, owing to the mechanics of generative AI technology, it is possible for third parties to obtain confidential company information by use of prompt injection attacks that cause generative AI tools to improperly disclose company information.

This risk is particularly heightened for companies that have developed their own generative AI tools as public-facing services, in which case attackers can use prompt injection attacks (ie, using manipulative prompts to trick a chatbot into revealing confidential information) to cause disclosure of the internals of those platforms.

Generative AI customers should carefully assess all generative AI tool terms of service to understand whether and how generative AI inputs or outputs are retained by service providers (or incorporated into training data) and are therefore susceptible to such attacks. Cases involving the alleged theft of proprietary information using prompt injection attacks are beginning to **surface**, accompanied by litigation.

Generative AI customers should review the terms of service of all generative AI platforms used to create content intended as proprietary IP and carefully assess whether those terms assign ownership of outputs to users and whether outputs might be retained by the platforms for re-use (eg, as training data), as such practices vary from platform to platform.

Company policies should include employee training, which prohibits employees from inputting trade secrets into generative AI tools, particularly where inputs are utilised as training data or where platforms may be susceptible to injection prompt attacks.

Looking forward

The development and adoption of agentic AI will mark a significant departure from the status quo, gradually forming around prompt-based generative AI. Whereas prompt-based generative AI intrinsically involves human interaction (whether in designing prompts or selecting or refining outputs), AI agents facilitate content creation without the intervention of human controllers.

Consequently, there is a real question as to whether any agentic AI-generated output could be protected under the United States' current copyright and patent frameworks.

Policymakers currently focus primarily on copyright issues endemic in training data sets. The **Generative AI Copyright Disclosure Act** bill, introduced in April 2024, would require developers of generative AI systems to register copyrights for their training materials with the USCO prior to release, which would curtail the practice of data scraping to generate training data sets.

Currently, the USPTO and the USCO must rely on compliance with applicant obligations to disclose the use of AI-generated tools in copyright and patent applications.

Developments in “watermarking” to **track the provenance** of AI-generated content and other techniques to **detect** potential infringement in AI-generated works may become viable means to identify and regulate AI-generated content.

Indeed, the **California AI Transparency Act**, adopted in late 2024 and effective 1 January 2026, requires that AI platform developers create “AI detection tools” to determine whether certain content was created or altered by the developer’s AI systems. The statute further requires that developers embed metadata into generated content showing AI generation provenance and, in certain cases, give users the option to permanently watermark AI-generated content with disclosures that the content was created by generative AI.

And, of course, changes in priorities under the Trump Administration, including **changes** in leadership, may further impact whether and how the government regulates AI-generated content.

Regardless, companies that proactively document their AI usage, adapt their IP policies and practices, and continuously monitor the fast-moving landscape will be best positioned to successfully navigate the accelerating AI revolution in business.

*A **version** of this article was also published on WTR sister platform IAM on 31 July 2025*

David Kappos

Partner
Cravath, Swaine & Moore

Callum Sproule

Corporate Associate
Cravath, Swaine & Moore