Cahill

AI and the Erosion of Intellectual Property Protections: Copyright, Trade Secret, Privilege, and Patent Vulnerabilities

Monday, August 17, 2026

INTRODUCTION

The explosive adoption of artificial intelligence in coding, engineering, writing, and creative development has triggered a simultaneous erosion of protections across every major category of intellectual property. Organizations using AI tools now face converging legal forces—uncopyrightable AI-generated works, trade secret destruction through careless disclosure to public platforms, privilege waiver upon submission of confidential information to consumer-grade tools, the human-inventorship requirement for patents, and the genuine risk that AI inputs may also constitute public disclosures triggering patent bars. Together, these forces create what may fairly be characterized as a systemic erosion of IP protections across multiple doctrinal categories.

The central finding is significant. The February 2026 decision in *United States v. Heppner* has confirmed the third party-disclosure theory that underlies both the privilege waiver and the patent disclosure risk. Organizations that integrate AI into their workflows without rigorous legal controls may face a scenario in which their most valuable creations lack copyright protection, their secrets are no longer secret, their privileged communications are stripped of protection, their inventions may enter the public domain before they can be patented, and their proprietary software may be contaminated by undisclosed open-source license obligations.

This article discusses AI-induced intellectual property issues through four parts. Part I examines copyright vulnerabilities—the uncopyrightability of AI-generated works, the unresolved fair use questions surrounding AI training, and the risk of open-source license contamination through AI coding tools. Part II addresses trade secret erosion when proprietary information is disclosed to third-party AI platforms. Part III analyzes patent risks from both the human-inventorship requirement and the

possibility that submissions to public AI platforms constitute novelty-destroying disclosures under § 102. Part IV considers the implications for attorney-client privilege and work product doctrine, drawing on the first judicial decisions to address these questions. This article also suggests organizational imperatives for mitigation of AI-induced IP risks.

1. COPYRIGHT

Generative AI systems capable of producing text, images, music, and code—often indistinguishable from human-created work—have placed extraordinary pressure on U.S. copyright law. AI models trained on vast corpora of copyrighted material have moved from research curiosities to commercially deployed tools, forcing courts, regulators, and the U.S. Copyright Office to confront questions existing frameworks were never designed to answer: whether machines can be "authors," whether ingesting copyrighted works to build AI models constitutes fair use, and whether AI-generated outputs resembling protected works give rise to infringement liability. The result is a legal landscape in which core IP protections are being tested simultaneously on multiple fronts.

These questions have generated unprecedented litigation, administrative guidance, and legislative proposals. Major copyright holders have filed dozens of lawsuits against leading AI companies, while AI developers have mounted aggressive fair use defenses that could fundamentally reshape the economics of content creation. The Copyright Office has issued registration decisions and policy reports attempting to draw workable lines between permissible AI-assisted creation and impermissible AI-generated output.

The copyright implications of AI are examined below in three categories: whether AI-generated content can be copyrighted; using copyrighted works to train AI models; and copyright infringement by AI systems. For each category, this Part discusses the statutory framework and leading case law, followed by industry-specific implications.

1.1. COPYRIGHTABILITY OF AI-GENERATED WORKS

1.1.1. STATUTORY FRAMEWORK AND CASE LAW

The Copyright Act of 1976, 17 U.S.C. § 102(a), extends protection to "original works of authorship fixed in any tangible medium of expression." The Supreme Court has long interpreted "author" to require human creative input. *See Burrow-Giles Lithographic Co. v. Sarony*, 111 U.S. 53 (1884). The Copyright Clause, Art. I, § 8, cl. 8, presupposes human intellectual effort, and the Copyright Office's Compendium states it "will not register works produced by a machine or mere mechanical process" without creative input from a human author.

Whether AI-generated works qualify for copyright has been decisively answered for works created autonomously without meaningful human contribution. In *Thaler v. Perlmutter*, the Copyright Office rejected registration for a visual artwork generated entirely by an AI system listed as sole author. On August 18, 2023, Judge Howell of the D.C. District Court affirmed, holding the Copyright Act requires a human "originator with the capacity for intellectual, creative, or artistic labor." *Thaler v. Perlmutter*, No. 1:22-cv-01564, 2023 WL 5333236 (D.D.C. Aug. 18, 2023).

On March 18, 2025, the D.C. Circuit affirmed, holding that "the Copyright Act requires all eligible work to be authored in the first instance by a human being," finding that statutory definitions imply humanity through

references to property ownership, lifespan, nationality, and intent. *Thaler v. Perlmutter*, 130 F.4th 1039 (D.C. Cir. 2025). As of this writing, the Supreme Court has not granted certiorari, leaving this position unchallenged.

In January 2025, the Copyright Office's Report confirmed that a work must have a human author for copyright eligibility, and that merely using AI to generate content—even with extensive prompting and iteration—does not satisfy this requirement. *See* U.S. Copyright Office, Copyright and Artificial Intelligence, Part 1: Digital Replicas (2025). The Office adopted a "separability" framework: if a human contributes original expression "perceptible and separable" from AI-generated elements, that contribution alone is eligible for registration, while AI-generated portions are excluded. Where AI enhancements are so integrated that the original human authorship cannot be identified as distinct, the work as a whole will not qualify. The Office rejected the notion that humans and AI can be joint authors. That guidance, however, leaves something to be desired: artistic works that parody or pastiche other original works, such as "Jane Eyre & Zombies," are clearly copyrightable, but appear to involve the same type of expression that the Copyright Office deems uncopyrightable when generated by artificial intelligence: for example, the prompt "please write a version of Jane Eyre in which the action of the novel occurs against the backdrop of a zombie apocalypse," would create a result very similar to the types of parodies that are routinely granted copyright protection, but would not be copyrightable.

1.1.2. IMPLICATIONS FOR THE SOFTWARE INDUSTRY

The implications for the software industry are significant. A Carlton Fields analysis (April 2026) observed that "copyright protection for software code is being sacrificed, knowingly or not, for the speed and efficiency of AI coding" and that "vibe coding"—where humans describe desired software in natural language and AI writes the code—is now "pervasive." *See* Carlton Fields, AI Makes Securing Copyright Protection for Software Code Tricky, Bloomberg Law (Apr. 15, 2026). Industry estimates indicated that approximately 42% of code was AI-generated or AI-assisted as of October 2025. *See* SonarSource, 2026 State of Code Developer Survey Report (fieldwork conducted Oct. 2025). Purely AI-generated code lacking human authorship is ineligible for copyright protection under the framework established in *Thaler*.

This creates a practical dilemma: AI-generated code may infringe the copyrights of works in the training data, yet the company shipping the AI-generated code cannot itself claim copyright protection. Competitors can freely copy AI-generated code without infringement liability because there is no copyright to enforce. The AI models generating such code were trained on billions of lines scraped from the internet, much of it under open-source licenses with conditions (attribution, share-alike) that the models do not track. Uncopyrightable work falls into the public domain, and an employer's claim over purely AI-generated code rests entirely on confidentiality obligations—contractual protection that evaporates upon disclosure.

1.2. USE OF COPYRIGHTED WORKS FOR AI TRAINING

1.2.1. STATUTORY FRAMEWORK AND CASE LAW

AI models are trained on enormous datasets scraped from the internet, often including copyrighted text, images, and code. AI companies argue this constitutes fair use under 17 U.S.C. § 107, analogizing to search engines indexing web pages. The fair use doctrine weighs four factors: (1) purpose and character of the use, including whether transformative or commercial; (2) nature of the copyrighted work; (3) amount and substantiality of the portion used; and (4) effect on the potential market. AI developers contend training is "transformative" because the model learns statistical patterns rather than reproducing works. Critics counter that the sheer scale of copying and

commercial nature of AI products weigh against fair use. No appellate court has yet definitively ruled on whether large-scale AI training constitutes fair use, and early trial-level results illustrate the contested nature of these claims.

Factor One **(Purpose and Character).** AI developers' strongest argument is that training is "transformative"—extracting statistical patterns rather than reproducing works—drawing on *Google LLC v. Oracle America, Inc.*, 593 U.S. 1 (2021). *See Thomson Reuters Enter. Ctr. GmbH v. Ross Intelligence Inc.*, 765 F. Supp. 3d 382 (D. Del. 2025). However, the court in *Thomson Reuters v. Ross* expressly rejected this reasoning. Ross argued that converting Westlaw's headnotes into numerical data to train its AI was transformative, but the court distinguished *Google v. Oracle* and *Sony Computer Entertainment, Inc. v. Connectix Corp.*, 203 F.3d 596 (9th Cir. 2000), on the ground that the copied material was creative written text rather than functional computer code—a distinction that subjects the use to more demanding fair use scrutiny.

Factor Two **(Nature** of the **Work).** AI training datasets include both highly creative works and factual works; courts afford greater fair use latitude for the latter. The heterogeneous nature of training datasets makes this factor difficult to apply in the aggregate, potentially requiring assessment at the individual-work level rather than corpus-wide.

Factor Three **(Amount Used).** This factor weighs most heavily against AI developers. Training requires ingestion of entire works. In *Authors Guild v. Google*, the Second Circuit held copying entire books for a search index was permissible because the purpose was limited. *See Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015). AI developers argue analogously that models retain only abstract statistical relationships. Critics respond that models can reproduce substantial portions of training data when prompted, as demonstrated by *The New York Times* complaint showing GPT-4 reproducing near-verbatim passages.

Factor Four **(Market Effect).** This is likely the decisive factor. Where AI outputs serve as substitutes for original works—AI-written news competing with *The Times*, AI images competing with stock photography—the market-harm factor strongly disfavors fair use. In *Thomson Reuters*, the court found market harm on summary judgment where Ross Intelligence built a competing legal research tool trained on Westlaw content. Where the AI product serves a fundamentally different market, the argument is weaker.

Several cases have addressed fair use in the AI training context with varying results. In the most significant authors' class action against an AI company, plaintiffs alleged the defendant scraped millions of copyrighted books to train its model without compensation. The court's analysis drew on *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569 (1994), and *Authors Guild v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015), but the ultimate resolution of the transformative use question in the AI training context remains unsettled. Other major suits remain pending, including *The New York Times Co. v. Microsoft Corp.* (S.D.N.Y.), alleging OpenAI copied millions of articles and that GPT models can reproduce substantial portions. In *Thomson Reuters Enter. Ctr. GmbH v. Ross Intelligence Inc.*, 765 F. Supp. 3d 382 (D. Del. 2025), the court granted summary judgment to Thomson Reuters, holding that an AI startup's use of Westlaw headnotes to train a competing product was not fair use—the first trial-court ruling rejecting fair use as a matter of law in the AI training context.

Two recent decisions from the Northern District of California reached the same ultimate result—finding fair use—through strikingly different reasoning. In *Bartz v. Anthropic PBC*, No. 3:24-cv-05417 (N.D. Cal. June 23, 2025), Judge Alsup issued a split decision on summary judgment, holding that Anthropic's use of copyrighted books to train its Claude LLMs was "spectacularly" transformative and a fair use, while simultaneously holding that

downloading pirated copies from sites like Books3 and Library Genesis was "inherently, irredeemably infringing." The court rejected the argument that AI training displaced an emerging licensing market. Two days later, in *Kadrey v. Meta Platforms, Inc.*, No. 23-cv-03417-VC, 2025 WL 1752484 (N.D. Cal. June 25, 2025), Judge Chhabria granted Meta's cross-motion for summary judgment, finding Meta's use of copyrighted books to train LLaMA "highly transformative," but in far more skeptical terms—criticizing Judge Alsup's market-effect reasoning as "inapt" and signaling that "[i]n cases involving uses like Meta's, it seems like the plaintiffs will often win" where they present meaningful evidence of market dilution.

Together, these decisions underscore that the fair use question in the AI training context remains fundamentally unsettled. Both courts found fair use on the records before them but diverged sharply on market effects and the persuasiveness of the transformative use analogy—and neither addressed output infringement. The judicial disagreement even among courts reaching the same conclusion confirms that this area of law will continue to develop as plaintiffs bring stronger evidentiary records.

1.2.2. IMPLICATIONS FOR INDUSTRY

The unresolved fair use question has significant practical consequences. For publishers and media companies, if courts broadly hold AI training constitutes fair use, content creators will have no legal mechanism to prevent AI companies from ingesting entire catalogs to build competing products without compensation. The music and visual arts sectors face analogous concerns, with AI-generated content trained on copyrighted works flooding platforms in direct competition with human creators. For technology and software sectors, companies relying on proprietary datasets face the prospect that competitors can train AI models on their content without licensing it. The lack of a definitive standard has also created uncertainty for AI companies themselves, many of which have entered voluntary licensing agreements while simultaneously arguing in court that no license is required.

1.3. COPYRIGHT INFRINGEMENT BY AI SYSTEMS

1.3.1. THE MECHANICS OF COPYLEFT INFECTION IN AI-ASSISTED DEVELOPMENT

Traditional software composition analysis (SCA) tools scan declared dependencies but cannot detect GPL-licensed snippets inserted by AI coding assistants. Attribution requirements from original licenses vanish because AI models strip provenance during training and generation. *See Doe 1 v. GitHub, Inc.*, No. 4:22-cv-06823-JST (N.D. Cal.). If a developer's AI assistant suggests code originating from a GPL-3.0 project, the developer may inadvertently trigger copyleft obligations for the entire proprietary product. *See* GNU General Public License, Version 3.0, § 5. While the theoretical remedy includes injunctive relief requiring source code release, in practice GPL enforcement—led by organizations like the Software Freedom Conservancy—typically results in negotiated compliance rather than forced open-sourcing. Nevertheless, litigation risk, reputational damage, and remediation costs remain substantial.

1.3.2. STATUTORY FRAMEWORK AND CASE LAW

AI coding assistants such as GitHub Copilot, powered by OpenAI's Codex model, were trained on billions of lines of source code from publicly available sources, including code in public GitHub repositories. *See Doe 1 et al. v. GitHub, Inc.*, No. 4:22-cv-06823-JST (N.D. Cal.), Compl. ¶¶ 82-83. Much of this training data was published under open-source licenses ranging from permissive (MIT, Apache) to strong copyleft (GPL-2.0, GPL-3.0, AGPL-3.0). These models do not track provenance—when Copilot suggests code, it does not inform the user whether it derives from GPL-licensed, MIT-licensed, or proprietary material. GitHub acknowledged Copilot "reproduces code

from training data about 1% of the time," which at enterprise scale represents substantial exposure for license violations. *Id*., ¶¶ 56, 71, 74, 87-90

All eleven open-source licenses on GitHub require derivative works or copies to include: (i) attribution; (ii) a copyright notice; and (iii) the license terms. See, e.g., MIT License; Apache License, Version 2.0. Copyleft licenses (GPL, AGPL) go further: any combined or modified work incorporating GPL code must be released under the same terms if distributed—the mechanism of "license infection." *See* GNU GPL v.3.0, § 5; GNU AGPL v.3.0, § 13. If a developer unknowingly incorporates a GPL snippet suggested by an AI tool into a proprietary codebase, the entire codebase may become subject to GPL terms, potentially requiring public release of proprietary source code.

The class action *Doe 1 et al. v. GitHub, Inc.*, No. 4:22-cv-06823-JST (N.D. Cal.), directly addresses these concerns. Anonymous programmers sued GitHub, Microsoft, and OpenAI, alleging Copilot reproduces open-source code without adhering to licensing terms—stripping required attribution and license information. *See* Developers Sue GitHub, Microsoft, and OpenAI Over Copyright in Creating AI Tool Copilot, Lexology (2022).

In a May 2023 order, Judge Tigar issued a mixed ruling. *See Doe 1 v. GitHub, Inc.*, No. 4:22-cv-06823-JST, 2023 WL 3449131 (N.D. Cal. May 11, 2023). Most claims were dismissed but claims under DMCA §§ 1202(b)(1) and (b)(3) (prohibiting intentional removal of copyright management information) and breach of open-source licenses survived. Plaintiffs cited a GitHub blog post acknowledging that Copilot suggested the GNU GPL "more than a whopping 700,000 different times during training." *See* Albert Ziegler, "GitHub Copilot Research Recitation," The GitHub Blog (June 30, 2021, updated Aug. 16, 2022), available at https://github.blog/ai-and-ml/github-copilot/github-copilot-research-recitation. In a June 2024 order, DMCA claims were dismissed because generated code was not identical enough to plaintiffs' work, but license violation and breach of contract claims survived. *See Doe 1 v. GitHub, Inc.*, No. 4:22-cv-06823-JST (N.D. Cal. June 24, 2024). The DMCA dismissal is on appeal to the Ninth Circuit. Plaintiffs estimated statutory damages exceeding $9 billion. Id. Comp. ¶¶ 2–4.

1.3.3. IMPLICATIONS

When AI-generated code is committed to public repositories, uncopyrightable code enters the public domain by default because there is no human author to claim copyright. *See Thaler v. Perlmutter*, No. 1:22-cv-01564 (D.D.C. 2023). Any third party can freely copy, use, and redistribute it. The company may well not have a legal mechanism to prevent competitors from using it. This creates a feedback loop: as more AI-generated code enters public repositories, AI models ingest increasingly uncopyrightable material, generating outputs derivative of other AI outputs—a recursive cycle in which protectable code steadily declines. An employer's claim over AI-generated code amounts to a contractual restraint rather than a property right—and that restraint disappears the moment the code becomes public.

When proprietary codebases containing AI-generated code are released publicly, the uncopyrightable portions cannot be reclaimed through copyright enforcement. Companies engaged in heavy "vibe coding" derive less strategic advantage from copyright claims and more from maintaining trade secret protection over prompting strategies and architectural decisions—though trade secret protection is itself precarious when AI tools are involved, given risks of inadvertent disclosure through AI system logs, training data extraction attacks, and employee use of external AI platforms.

AI companies have responded with contractual indemnity programs. Microsoft's Copyright Commitment defends Copilot Business and Enterprise customers against IP infringement claims, though subject to conditions (use of

content filters) and standard liability caps. *See* Microsoft, Copilot Copyright Commitment (Sept. 2023). OpenAI offers tools "as-is" with no IP indemnity. *See* OpenAI, Terms of Use (effective Mar. 14, 2023, as amended). More fundamentally, no terms of service can override copyright or license law. When a provider's terms state "we assign all rights in the output to you," they may be assigning something that does not exist if the output is predominantly AI-generated. An indemnity does not retroactively remove a GPL obligation once triggered. *See* GNU GPL v.3.0, § 5.

The copyright risks examined above—uncopyrightable outputs, unresolved fair use, and open-source contamination—represent one dimension of AI-related IP vulnerabilities. But these risks concern what comes *out* of AI systems. The next dimension concerns what organizations feed *into* them—and the trade secret protections that may be destroyed in the process.

2. TRADE SECRET

Trade secret law presents a distinct—and in some ways more insidious—set of AI-related vulnerabilities. Unlike copyright, where the risk materializes in the *output* of AI systems, trade secret erosion occurs at the *input* stage. Engineers, developers, and in-house counsel routinely feed proprietary source code, internal technical documents, and other confidential materials into AI-powered code-completion tools and large language model chatbots to obtain debugging assistance, code optimization, or technical analysis. The critical question is whether such submissions—particularly to consumer-grade platforms that retain and train on user inputs—constitute a waiver of trade secret protection. The answer turns on two doctrinal pillars: whether the owner maintained "reasonable efforts" to preserve secrecy, and whether the information has become "generally known" or "readily ascertainable" by others.

2.1. TRADE SECRET PROTECTION AND THE RISK OF WAIVER

2.1.1. STATUTORY FRAMEWORK AND CASE LAW

Trade secret protection in the United States is governed primarily by state law—with the vast majority of jurisdictions having adopted some version of the Uniform Trade Secrets Act ("UTSA")—and at the federal level by the Defend Trade Secrets Act of 2016 ("DTSA"), 18 U.S.C. § 1836 *et seq*. Both adopt substantially equivalent definitions of a "trade secret," describing it as information that (1) has independent economic value precisely because it is neither generally known to, nor easily discoverable by proper means by, those who could benefit from its disclosure, and (2) has been protected through measures that are reasonable given the circumstances to preserve its confidentiality.

Two elements of this definition are central to the AI-disclosure question. First, the information must not be "generally known" or "readily ascertainable" by others. Second, the trade secret owner must have taken "reasonable efforts" to maintain secrecy. A failure to satisfy either element defeats trade secret status, or a voluntary disclosure to a third party without adequate confidentiality protections may undermine both. The following discussion examines each element in turn before applying the framework to AI platform disclosures.

2.1.1.1. THE REASONABLE EFFORTS REQUIREMENT

Courts have long held that the "reasonable efforts" requirement is not a demand for perfection but rather an inquiry into whether the trade secret owner has taken steps that are proportionate to the value of the information and the circumstances of the case. In *Rockwell Graphic Systems, Inc. v. DEV Industries, Inc.*, 925 F.2d 174 (7th Cir. 1991), Judge Posner emphasized that the requirement serves an evidentiary function: it tends to confirm that the information is in fact secret and that the owner values it as such. Similarly, in *Learning Curve Toys, Inc. v. PlayWood Toys, Inc.*, 342 F.3d 714 (7th Cir. 2003), the Seventh Circuit reiterated that the standard is one of reasonableness, not absolute secrecy, and that the owner need not take every conceivable precaution.

Reasonable efforts typically include measures such as restricting access to information on a need-to-know basis, employing physical and electronic security measures, marking documents as confidential, and requiring employees and third parties to sign nondisclosure agreements ("NDAs"). The critical point for present purposes is that sharing trade secret information with a third party is not inherently fatal to trade secret status, provided that the disclosure is accompanied by adequate confidentiality protections. As the Restatement (Third) of Unfair Competition § 39, comment f, explains, a trade secret owner may disclose information to employees, licensees, and other third parties without forfeiting protection so long as the disclosure is made subject to a duty of confidentiality, whether express or implied.

2.1.1.2. THE "GENERALLY KNOWN" AND "READILY ASCERTAINABLE" PRONGS

Even if a trade secret owner can demonstrate reasonable efforts, trade secret protection may be lost if the information becomes "generally known" to persons who can derive economic value from it. The question is whether submission to an AI platform causes information to enter the public domain in this sense. If the platform's terms of service permit the use of inputs for model training, and if the trained model is then made available to millions of users, there is a plausible argument that the information has been injected into a system from which it is, at least in theory, extractable by competitors. However, the counterargument is also strong: a model's internal parameters are not human-readable, and the mere possibility that fragments of the information might be reproduced does not mean that the information is "generally known" in the trade secret sense. The information remains, in most cases, buried in billions of numerical weights and is not practically accessible through normal interaction with the model. Similarly, although large language models can have their chatbots circumvented to return inputs they were trained upon verbatim, by exfiltrating those inputs through an algorithm known as "Best-of-N jailbreaking," such methods would seem to qualify more as hacking or the unauthorized taking of trade secrets by former employees, neither of which in practice destroy trade secret protection and in fact are categorized as forms of trade secret theft. *See* Ahmed, "Extracting books from production language models," Jan 6, 2026, https://arxiv.org/abs/2601.02671 ("In some cases, jailbroken Claude 3.7 Sonnet outputs entire books near verbatim").

The better view is that whether trade secret status is lost depends on the specific facts of each case, including the platform's data handling practices, the sensitivity and distinctiveness of the information, the likelihood of memorization and reproduction, and the contractual protections in place. A blanket rule that any submission to any AI platform destroys trade secret protection would be inconsistent with the flexible, fact-intensive approach that courts have historically applied to the reasonable effort inquiry. Conversely, a blanket rule that trade secret protection survives any such submission would ignore the genuine risks that these platforms pose.

2.1.2. IMPLICATIONS

The act of feeding proprietary information to a third-party AI platform raises a distinct analytical challenge. Unlike a disclosure to a human counterpart bound by social and legal norms of confidentiality, the "third party" here is a software system operated by a technology company whose terms of service may or may not impose confidentiality obligations. The analysis must therefore proceed on two levels: the contractual terms governing the platform, and the technical realities of how it processes user input.

At the contractual level, the terms of service and privacy policies of major AI platforms vary considerably. Some platforms, such as OpenAI's ChatGPT Enterprise and Microsoft's Azure OpenAI Service, offer contractual commitments that user inputs will not be used for model training and will be handled in accordance with enterprise-grade confidentiality obligations. Other platforms, particularly consumer-facing versions, reserve the right to use user inputs to improve and train their models, and their terms of service may include broad licenses to user-submitted content. If user inputs are used for model training, the information may become encoded—however abstractly—in the model's parameters, raising the question of whether the information has been disclosed in a manner that renders it "generally known" or "readily ascertainable" by others who interact with the model.

At the technical level, the concern is whether information submitted to an AI platform can be reproduced, regurgitated, or otherwise extracted by third parties. Research in the machine learning community has demonstrated that large language models are capable of memorizing and reproducing portions of their training data, particularly when those data are repeated or distinctive. If a model is trained on a user's proprietary source code and subsequently reproduces that code in response to another user's prompt, the trade secret may have been effectively disclosed to the public, regardless of the contractual terms governing the original submission.

The more difficult case arises when the AI platform does not use inputs for training and contractually commits to confidentiality, but the user has not executed a standalone NDA or other confidentiality agreement with the platform provider. Courts evaluating the "reasonable efforts" requirement generally look to the totality of the circumstances, and a click-through terms-of-service agreement may or may not be deemed a sufficient confidentiality measure depending on the jurisdiction and the specifics of the provision. In *Compulife Software Inc. v. Newman*, 959 F.3d 1288 (11th Cir. 2020), the Eleventh Circuit considered the effect of disclosing information through a publicly accessible interface, and the court's analysis suggests that the absence of a formal confidentiality agreement, combined with the broad accessibility of the platform, would weigh against a finding of reasonable efforts.

The prevailing view is that disclosure to a service provider does not waive trade secret protection per se, but that the trade secret owner must demonstrate that adequate confidentiality measures accompany the disclosure. The absence of such measures—particularly when the service provider's terms of service affirmatively disclaim confidentiality obligations or reserve the right to use customer data—may be treated as a failure to maintain reasonable efforts.

2.1.2.1. TRADE SECRET LOSS

Trade secret protection is fundamentally conditioned on the owner maintaining secrecy. Commentators have warned that data entered into consumer AI platforms may no longer be considered private, as it can be used by the platform to respond to future queries or for other purposes by the platform operator. Trade secret protections can be lost permanently if the information is made public, even if the disclosure is accidental or unintentional. The

Second Circuit established this principle in *Defiance Button Machine Co. v. C & C Metal Products Corp.*, 759 F.2d 1053, 1063 (2d Cir. 1985), holding that "upon disclosure, even if inadvertent or accidental, the information ceases to be a trade secret and will no longer be protected."

An entity claiming trade secret status bears the burden of identifying and demonstrating that the material is included in categories of protected information under the statute and must additionally take active steps to maintain its secrecy. *See* 18 U.S.C. § 1839(3); Uniform Trade Secrets Act § 1(4). One frequently raised defense to a trade secret misappropriation claim is that the original owner effectively abandoned the secret by not taking adequate steps to maintain its confidentiality. This defense takes on new significance in the context of consumer-facing AI platforms, where inputting proprietary information into a system that may incorporate it into responses provided to other users could constitute abandonment of the trade secret.

2.1.2.2. THE SAMSUNG INCIDENT AS A CAUTIONARY PARADIGM

In March 2023, Samsung Electronics employees in its semiconductor business unit inadvertently leaked sensitive corporate data through ChatGPT in three separate incidents. *See* Samsung Employees Leaked Corporate Data in ChatGPT: Report, CIO Dive (Apr. 6, 2023); Incident 768: ChatGPT Implicated in Samsung Data Leak of Source Code and Meeting Notes, AI Incident Database (2023). One employee entered faulty source code related to a facility measurement database to find a solution; another entered program code for identifying defective equipment seeking code optimization; and a third converted a smartphone recording of a company meeting into a document file and entered it into ChatGPT to generate meeting minutes. *See* Samsung Electronics Co., Internal Memo on Generative AI Use Restrictions (May 2023), as reported in Bloomberg News (May 2, 2023). The AI retained the input data, leading to a breach of confidentiality, and Samsung subsequently banned employee use of generative AI tools entirely. *See* Samsung Bans Generative AI Use by Staff After ChatGPT Data Leak, Bloomberg (May 2, 2023).

The Samsung incident illustrates a pervasive threat. ChatGPT's web-based interface, at the time of the incident, used input data to train and improve the tool, meaning that if employees input company data as part of their prompt, the information could become part of the model's knowledge base and be available to other users in some form. *See* OpenAI, Terms of Use (effective Mar. 14, 2023). Research indicated that 3.1% of workers had pasted confidential company data into ChatGPT, according to a Cyberhaven report, while 43% of professionals reported using AI tools at work. *See* Cyberhaven Labs, Data Trends Report: AI Use in the Enterprise (Mar. 2023).

2.1.2.3. ENTERPRISE VERSUS CONSUMER AI AND THE CONFIDENTIALITY DIVIDE

The Berkeley Technology Law Journal published a comprehensive analysis by Professor Camilla Hrdy titled "Keeping ChatGPT a Trade Secret While Selling It Too," examining how generative AI companies like OpenAI maintain trade secret protection through closed-source formats that hide algorithms, code, training data, and underlying model architecture, while attaching contractual terms of use that limit reverse engineering. Camilla A. Hrdy, *Keeping ChatGPT a Trade Secret While Selling It Too*, 39 Berkeley Tech. L.J. 1 (2024). The article emphasized the critical distinction between consumer-grade AI tools and enterprise solutions. OpenAI's Enterprise product includes mutual confidentiality obligations and guarantees that customer data will not be used for model training. *Id*. By contrast, the free consumer versions retain broad rights to use input data.

Trade secret owners who use public, consumer-grade AI tools risk creating what may amount to an irremediable waiver of trade secret protection. The implication for organizations is that internal policies restricting employee

use of consumer AI tools are not merely best practices but essential elements of maintaining the "reasonable measures" required under the Defend Trade Secrets Act, 18 U.S.C. § 1839(3)(A), and the Uniform Trade Secrets Act, § 1(4)(ii), to preserve trade secret status.

Trade secret erosion through AI disclosure is serious, but patent law presents an even more unforgiving framework. While trade secret protection turns on a flexible reasonableness inquiry that examines the totality of circumstances, patent novelty can be defeated by a single unrestricted disclosure—making the stakes of AI-related information sharing particularly acute in the inventorship and prior art contexts.

3. PATENT

Patent law's novelty requirement presents a distinct vulnerability in the AI era. Under 35 U.S.C. § 102(a)(1), as amended by the America Invents Act, an invention is unpatentable if it was "in public use, on sale, or otherwise available to the public" before the filing date. When engineers submit technical information describing patentable inventions to AI platforms—whether for debugging, optimization, or analysis—the question arises whether such submission constitutes a "public use" or a disclosure rendering the invention "otherwise available to the public." Unlike the trade secret context, where the inquiry focuses on whether the owner maintained reasonable secrecy efforts, the patent inquiry turns on whether the information has been made accessible to the relevant public in a manner that defeats novelty—a lower threshold that can be triggered even by a single unrestricted disclosure.

3.1. PATENT LAW AND THE QUESTION OF PUBLIC USE OR DISCLOSURE UNDER 35 U.S.C. § 102(a)

3.1.1. THE NOVELTY FRAMEWORK UNDER THE AIA

Under 35 U.S.C. § 102(a)(1), as amended by the AIA, a person is not entitled to a patent if the claimed invention was "patented, described in a printed publication, or in public use, on sale, or otherwise available to the public before the effective filing date of the claimed invention." The AIA's addition of the catch-all phrase "otherwise available to the public" was intended, according to its legislative history, to emphasize that the statutory bars in § 102(a)(1) are concerned with whether the invention has been made available to the public. This has led to a vigorous debate over whether the AIA narrowed the scope of the "public use" and "on sale" bars to require actual public accessibility, or whether those bars retain their pre-AIA breadth.

The Federal Circuit addressed this question in *Helsinn Healthcare S.A. v. Teva Pharmaceuticals USA, Inc.*, 855 F.3d 1356 (Fed. Cir. 2017), aff'd, 139 S. Ct. 628 (2019), holding that the "on sale" bar under the AIA is not limited to public sales and that a sale or offer for sale can trigger § 102(a)(1) even if the details of the invention are not publicly disclosed. The Supreme Court affirmed, but the decision was narrowly focused on the on-sale bar and did not squarely address the scope of the "public use" bar. Accordingly, the question of whether a confidential or semi-confidential use of an invention can constitute a "public use" under the AIA remains subject to some uncertainty.

3.1.1.1. THE "PUBLIC USE" BAR

Under pre-AIA law, the public use bar was construed broadly. In *Egbert v. Lippmann*, 104 U.S. 333 (1881), the Supreme Court held that a single use of an invention by a person other than the inventor, without restriction or obligation of secrecy, could constitute a public use. The Federal Circuit refined this doctrine in subsequent

decisions, holding that the public use inquiry turns on whether the inventor maintained control over the invention and limited access to it, and whether the use was primarily for the inventor's commercial benefit. In *Lough v. Brunswick Corp.*, 86 F.3d 1113 (Fed. Cir. 1996), the court emphasized that even a use that is not visible to the public may constitute a "public use" if the inventor fails to impose limitations, such as a confidentiality agreement, on those who have access to the invention.

Under the AIA, the Federal Circuit has not definitively resolved whether the "public use" bar continues to encompass confidential or restricted uses, or whether the "otherwise available to the public" residual clause imposes a general public-accessibility requirement on all of § 102(a)(1)'s categories. In *Electro Scientific Industries, Inc. v. Dynamic Details, Inc.*, and other post-AIA decisions, the court has generally continued to apply pre-AIA precedent to the public use bar, suggesting that a use need not be literally "public" to trigger the bar so long as the inventor has not taken adequate steps to preserve confidentiality.

3.1.2. IMPLICATIONS

The submission of source code or technical documents describing a patentable invention to an AI platform could potentially trigger the "public use" bar, the "printed publication" bar, or the "otherwise available to the public" catch-all. It needs to be emphasized, however, that our analysis remains an as yet untested extension of existing patent law. No court has squarely held that submission of information to an AI platform constitutes a public disclosure under § 102(a)(1), and the application of the traditional public use and printed publication doctrines to AI platform interactions represents a novel doctrinal question that courts have not yet had occasion to resolve.

Public Use. The strongest argument for applying the public use bar is that the submission constitutes a use of the invention (or at least of information describing the invention) by a third party—the AI platform operator—without adequate limitations on further use or disclosure. Under *Lough* and its progeny, the absence of a confidentiality restriction on the recipient of the invention has been a critical factor in finding public use. If the platform's terms of service do not impose confidentiality obligations on the platform operator, or if they affirmatively permit the operator to use submitted content for model training (and thus to make the information accessible, in some form, to other users), the submission may be analogized to an unrestricted disclosure. However, the counterargument is also substantial: the user is not "using" the invention in the traditional sense of deploying it for its intended purpose. Submitting a description of an invention to an AI chatbot for debugging assistance is categorically different from, for example, selling a product embodying the invention or allowing a third party to use the invention in the field. Courts have generally required that the "use" in question involve the invention being applied for its intended purpose or a commercial purpose, and a submission for informational or analytical assistance may not satisfy this requirement.

Printed Publication. Under well-established Federal Circuit precedent, a "printed publication" is a document that has been made sufficiently accessible to the public interested in the art. In *In re Klopfenstein*, 380 F.3d 1345 (Fed. Cir. 2004), the Federal Circuit established a framework with several factors to assess whether a displayed reference constitutes a printed publication. These factors include: the duration of the display, the level of expertise possessed by the intended audience, whether there were reasonable expectations that viewers would refrain from copying the displayed material, and how easily the material could have been reproduced. Under this framework, a document submitted to an AI platform is unlikely to qualify as a "printed publication" in most circumstances. The submission is directed to a machine, not to the "public interested in the art," and the submitter typically does not intend for the information to be disseminated. However, if the AI platform's model is trained on the submission

and subsequently reproduces the technical information in response to queries from third parties—particularly competitors or persons skilled in the art—a more difficult question arises. The information may have effectively been "published" through the medium of the AI model, even if no traditional document was ever made available. That will certainly be the case when large language models output text they were trained upon verbatim, which can be exfiltrated from such models through an algorithm known as "Best-of-N jailbreaking," which iterates, randomizes and shuffles nested queries to AI chatbots until a harmful response is elicited. *See* Ahmed, "Extracting books from production language models," Jan 6, 2026, https://arxiv.org/abs/2601.02671 (" In some cases, jailbroken Claude 3.7 Sonnet outputs entire books near-verbatim").

Otherwise Available to the Public. The AIA's catch-all provision—"otherwise available to the public"—was designed to capture forms of disclosure that do not fit well into the traditional categories. If a submission to an AI platform causes the submitted information to become available to the public through the platform's outputs, this provision may be the most natural statutory hook. However, the mere submission of information to the platform, without evidence that the information has actually become accessible to others, is unlikely to trigger this bar. The critical question is whether the platform's processing of the information—including potential model training—constitutes the information being made "available to the public."

The USPTO has formally acknowledged this risk. Its guidance on practitioner duties warns that patent attorneys are bound to maintain the confidentiality of all client information under 37 C.F.R. § 11.106, and that inputting novel, enabling data into a public AI whose Terms of Service grant the vendor the right to store and use the data for model training constitutes a breach of the duty of confidentiality, as the information is disclosed to an unauthorized third party who is not bound by any obligation of confidentiality. *See* USPTO, Guidance on Use of Artificial Intelligence-Based Tools in Practice Before the United States Patent and Trademark Office (2024). The USPTO's position signals that it views information shared with a public AI provider as no longer confidential, potentially satisfying the "publicly available" requirement under § 102.

Beyond the risk of direct repetition of confidential inputs, modern AI models introduce additional disclosure vectors. A non-enabling fragment of confidential input can be combined by the AI with its internal database of public prior art to generate an output that is fully enabling—a risk described as the "Black Box Problem" or extrapolation risk. Patent examiners may combine multiple public references under 35 U.S.C. § 103 to demonstrate obviousness, meaning that sharing pieces of an invention across different public AI platforms can provide the components necessary to render claims unpatentable. *See* KSR Int'l Co. v. Teleflex Inc., 550 U.S. 398 (2007) (obviousness standard).

3.2. PATENT INVENTORSHIP AND THE AI TOOL PARADIGM

3.2.1. STATUTORY FRAMEWORK AND CASE LAW

The Federal Circuit's 2022 decision in *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022), definitively established that only natural persons can be named as inventors on patent applications under U.S. law. Stephen Thaler had attempted to list his AI system, DABUS, as the inventor on patent applications for two inventions he claimed were autonomously conceived by the AI without human assistance. The Federal Circuit held that the statutory term "inventor" in 35 U.S.C. § 100(f) means "the individual . . . who invented or discovered the subject matter of the invention," and that "individual" unambiguously refers to a natural person. *Id.* at 1210-12. The Supreme Court denied certiorari, leaving this interpretation as settled law. *Thaler v. Vidal, 143 S. Ct. 1783 (2023) (cert. denied).*

The USPTO has issued two rounds of guidance on AI-assisted inventions, reflecting the evolving policy

landscape. In February 2024, under Director Kathi Vidal, the USPTO issued its initial "Inventorship Guidance for AI-Assisted Inventions," which adapted the *Pannu v. Iolab Corp.*, 155 F.3d 1344 (Fed. Cir. 1998), factors—originally for assessing joint human inventorship—to evaluate whether human contributions in AI-assisted scenarios were sufficiently "significant." The 2024 guidance established that while AI-assisted inventions are not categorically unpatentable, each claim must have been invented by at least one named natural person who made a significant contribution.

On November 28, 2025, under the new administration, the USPTO issued revised guidance that rescinded the February 2024 guidance in its entirety. *See* Revised Inventorship Guidance, Eliminating Separate Standard for AI-Assisted Inventions (2025). The revised guidance fundamentally changed the framework by withdrawing the application of *Pannu* factors in the AI context. *Id.* It establishes that "the same legal standard for determining inventorship applies to all inventions, regardless of whether AI systems were used in the inventive process," that AI systems are tools analogous to laboratory equipment, and that the traditional "conception" test—whether a natural person formed the "definite and permanent idea of the complete and operative invention"—governs uniformly. *See* Revised Inventorship Guidance for AI-Assisted Inventions, USPTO (Nov. 28, 2025); *see also* USPTO Issues Revised Inventorship Guidance, Eliminating Separate Standard for AI-Assisted Inventions (2025).

3.2.2. IMPLICATIONS

The revised framework creates a potential gap: where AI generates solutions with limited human contribution, and no human can demonstrate that they conceived the invention, the invention may be unpatentable because there is no valid human inventor. As one commentary noted, "overreliance on AI, where no single person can demonstrate a significant contribution to conception, may lead to inventions being dedicated to the public, as no person could make a valid claim of inventorship." *See* USPTO, Revised Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54636 (Nov. 28, 2025). This creates the perverse result that the more powerful AI becomes as an inventive tool, the more likely it is that genuinely novel inventions cannot be patented, expanding the public domain involuntarily.

The patent vulnerabilities discussed above share a common thread with the final category of risk: both turn on whether information submitted to an AI platform retains its confidential character. In the privilege context, however, the consequences are felt not at the patent office but in active litigation—where waiver can result in the forced disclosure of defense strategy, legal analysis, and communications that attorneys and clients believed were protected.

4. ATTORNEY-CLIENT PRIVILEGE AND WORK PRODUCT

The integration of AI tools into legal practice has created acute tensions with the attorney-client privilege and work product doctrine—two foundational protections that depend on the maintenance of confidentiality. Lawyers and their clients increasingly use AI platforms to draft briefs, analyze legal questions, and develop litigation strategy. Yet submitting privileged or work-product-protected information to a third-party AI system—particularly a consumer-grade platform whose terms of service permit data retention and model training—raises the specter of waiver. Courts have now begun to address these questions directly, and the resulting framework turns on familiar doctrinal distinctions: the nature of the communication, the identity of the recipient, the purpose of the disclosure, and the steps taken to preserve confidentiality.

4.1. ATTORNEY-CLIENT PRIVILEGE AND WORK PRODUCT DOCTRINE IN THE AI ERA

4.1.1. STATUTORY FRAMEWORK AND CASE LAW

On February 17, 2026, U.S. District Judge Jed S. Rakoff of the Southern District of New York issued the first major judicial decision directly addressing whether communications with a publicly available AI platform are protected by attorney-client privilege or the work product doctrine. In *United States v. Heppner*, No. 1:25-cr-00503 (S.D.N.Y. Feb. 17, 2026), the defendant Bradley Heppner—charged with securities fraud, wire fraud, conspiracy, and related offenses—had used the publicly available version of Anthropic's Claude to prepare approximately thirty-one documents outlining his defense strategy after learning he was the target of a federal investigation. When the FBI raided his home and seized these documents pursuant to a search warrant, his lawyers asserted attorney-client privilege and work product protection.

Judge Rakoff identified at least two, and likely three, independently dispositive deficiencies in the privilege claim. First, because Claude is not an attorney, the communications were not between a client and their attorney, failing the most basic element of attorney-client privilege. Second, the communications were not confidential because the platform's privacy policy authorizes data collection on user inputs and outputs, use of such data to train the model, and disclosure to third parties including "governmental regulatory authorities." Third, Heppner did not communicate with the AI platform for the purpose of obtaining legal advice from the platform, which itself disclaims providing legal advice. *See* A Tale of Two Cases—Attorney Client Privilege and the Use of AI Tools (2026); AI Privilege and Waiver: What Courts Are Actually Saying (And What They're Not), Carpe Datum Law (2026).

Judge Rakoff's ruling on the confidentiality point is particularly consequential because it establishes that sharing information with a public AI platform is legally equivalent to disclosing it to a third party—the classic waiver fact pattern. *See* Public AI Exposure Waives Privilege, Confidentiality, and Privacy, US Judge Rules, TransPerfect Legal, JDSupra (2026). The court's reasoning was "unequivocal": non-privileged communications are not "somehow alchemically changed into privileged ones upon being shared with counsel." *See* Attorney-Client Privilege and Work Product in the Age of Generative AI, White & Case LLP (2026). However, the court left open that had counsel directed Heppner to use the platform, "the platform might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege." *See* Cahill Gordon & Reindel LLP, SDNY Judge Finds Criminal Defendant Case Strategy Communications with Gen AI Tool Not Privileged or Work Product, Client Alert (Mar. 18, 2026), https://www.cahill.com/publications/client-alerts/2026-03-18-sdny-judge-finds-criminal-defendant-case-strategy-communications-with-gen-ai-tool-not-privileged-or-work-product; *See also* AI Tools May End Attorney-Client Privilege, Harris Beach Murtha (2026).

Just one week before *Heppner*, on February 10, 2026, the U.S. District Court for the Eastern District of Michigan reached a different result in *Warner v. Gilbarco, Inc.*, No. 2:24-CV-12333, 2026 WL 373043 (E.D. Mich. Feb. 10, 2026). In that civil case, a *pro se* plaintiff had used ChatGPT to prepare legal briefs in anticipation of litigation, and the defendant moved to compel production of all documents concerning the plaintiff's use of AI tools. The court denied the motion, holding that the materials were protected under the work product doctrine.

The critical distinction was that work product protection is waived only by "disclosure to an adversary or in a way likely to get in an adversary's hand," not by mere disclosure to a third party. *See* A Tale of Two Cases—Attorney Client Privilege and the Use of AI Tools (2026). The court characterized the defendant's theory as one that, "if accepted, would nullify work-product protection in nearly every modern drafting environment, a result no court

has endorsed." The court also characterized the AI tools as "tools, not persons," recognizing that AI-assisted drafting is functionally analogous to using any other research or writing tool.

4.1.2. IMPLICATIONS

The apparent conflict between *Heppner* and *Warner* dissolves upon careful analysis of the factual distinctions. In *Heppner*, a represented criminal defendant used a public AI platform on his own initiative without counsel's direction, the materials were seized via search warrant, and the claim was for attorney-client privilege—which requires confidential communication with a lawyer. *See* AI Privilege and Waiver: What Courts Are Actually Saying (And What They're Not), Carpe Datum Law (2026). In *Warner*, a *pro se* litigant used AI as part of her own litigation preparation, opposing counsel sought production through a discovery request, and the claim was for work product protection—which requires only preparation in anticipation of litigation and non-disclosure to an adversary. *Id.*

The overarching principle that emerges is that the AI tool itself is neutral—privilege and work product analyses turn on the same factors they always have: confidential communication with a lawyer for legal advice, materials prepared in anticipation of litigation, and maintenance of confidentiality. *Id.* Using an enterprise AI tool with contractual confidentiality guarantees under the direction of counsel stands on far stronger footing than using consumer-grade public platforms without attorney supervision.

5. PRACTICAL RISK MITIGATION

The foregoing analysis yields a clear set of organizational imperatives. Companies should adopt comprehensive policies governing AI platform use by employees and contractors, specifying approved platforms and permissible information types. At a minimum, these policies should prohibit submitting trade secret information or patentable inventions to consumer-tier platforms that reserve the right to use inputs for model training. The following measures represent essential components of a robust AI governance framework:

Terms of Service Review. The contractual terms governing AI platforms are among the most important variables in determining whether a submission forfeits IP protection. OpenAI's consumer ChatGPT product permits use of inputs to improve models unless the user opts out, while ChatGPT Enterprise and the OpenAI API expressly prohibit using customer inputs for model training. Similar distinctions exist across Google (Gemini vs. Vertex AI), Anthropic, Microsoft, and Amazon Web Services. A submission to a consumer-tier platform whose terms permit model training is far more likely to be treated as a waiver of trade secret protection and a potential public use than a submission to an enterprise platform with contractual confidentiality commitments and data segregation.

Technical Architecture and Data Handling. Beyond contractual terms, the platform's technical architecture matters. Key questions include whether inputs are stored persistently, whether they are used for real-time fine-tuning, whether the platform employs data isolation for enterprise customers, and whether safeguards prevent memorization and reproduction of user data. A platform that processes inputs in memory without persistent storage presents a categorically different risk profile than one that retains inputs indefinitely and uses them to train a publicly available model.

Nondisclosure Agreements. Companies seeking to preserve trade secret protection while using AI platforms should execute standalone NDAs with the platform provider, rather than relying solely on standard terms of service. A standalone NDA creates a clear record of confidentiality obligations and provides a stronger basis for

arguing "reasonable efforts" under the DTSA and UTSA. In the patent context, a confidentiality agreement is one of the most important factors in rebutting a finding of public use, as courts have consistently held that a disclosure made subject to such an obligation is not "public" for purposes of § 102.

Documentation of Human Contribution. For both copyright and patent purposes, meticulous documentation of human creative and inventive contribution has become essential. The Copyright Office requires applicants to disclose AI-generated content in works submitted for registration, and has registered hundreds of works incorporating AI-generated material where human authorship was identifiable and separable. *See* U.S. Copyright Office, Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence Systems, 88 Fed. Reg. 16,190 (Mar. 16, 2023). For patents, the revised 2025 USPTO guidance requires that applicants demonstrate a human conceived the "definite and permanent idea of the complete and operative invention," making contemporaneous records of human inventive thought critical to establishing valid inventorship. *See* USPTO, Revised Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54636 (Nov. 28, 2025).

Open-Source License Governance in AI Workflows. Organizations using AI coding tools must implement automated scanning of all AI-generated code for license compliance before deployment. Traditional software composition analysis tools that scan only declared dependencies are insufficient because they miss GPL-licensed snippets inlined by AI assistants. Emerging tools from providers such as FOSSA and Codacy are beginning to address this gap, but the effectiveness of these tools as evidence of compliance is untested in litigation. Any code committed to a public repository should be reviewed for potential copyleft contamination, and organizations should maintain a dynamic Software Bill of Materials (SBOM) that tracks the license type for every dependency and AI-generated component. *See* NTIA, Software Bill of Materials (SBOM) Minimum Elements (July 2021).

The File-First-Disclose-Later Doctrine. For patent-conscious organizations, the safest strategy is to file patent applications before any interaction with AI tools that might involve novel, enabling information. This "file first, disclose later" approach is essential for international patent protection under absolute novelty regimes and provides maximum safety under the U.S. grace period as well. *See* 35 U.S.C. § 102(b)(1)(A) (grace period for inventor's own disclosures). Even in the U.S., the grace period merely provides a one-year window from the date of disclosure—a window that begins running the moment information enters a public AI system and that many organizations may not realize has started.

6. CONCLUSION

The convergence of legal principles governing AI and intellectual property reveals a pattern of accelerating erosion across multiple doctrinal categories—a dynamic in which the adoption of AI tools simultaneously undermines copyright, trade secret, privilege, and patent protections faster than legal frameworks have adapted. The case law confirms the trajectory: the *Thaler* decisions establish that purely AI-generated works enter the public domain; *Heppner* confirms that public AI tools are third parties for confidentiality purposes; the Samsung incident demonstrates how easily trade secrets are destroyed; and the USPTO's revised inventorship guidance confirms that only humans can be inventors while AI grows increasingly responsible for actual conception. The analytical framework under § 102 suggests that inputs to public AI tools may constitute novelty-destroying disclosures—a conclusion no court has yet reached but one that follows from *Heppner*'s reasoning about the third-party character of public AI platforms.

The practical effect is that organizations racing to adopt AI are unknowingly feeding their most valuable intellectual property into systems that simultaneously strip it of legal protection while potentially encumbering it with undisclosed open-source obligations. Legislative action—redefining authorship, creating sui generis protection for AI-generated works, clarifying the patent disclosure status of AI inputs, or establishing provenance-tracking requirements—remains possible but has not materialized. Until it does, the organizations that navigate these risks successfully will be those that implement rigorous AI governance today: enterprise-grade tools with contractual confidentiality, documented human contributions at every stage, patent applications filed before AI interactions, automated license scanning of all AI-generated code, and institutional controls that prevent the casual destruction of irreplaceable intellectual property rights.

* * *

If you have any questions about the issues addressed in this memorandum, or if you would like a copy of any of the materials mentioned in it, please do not hesitate to call or email authors: Britton Davis (partner) at 202.862.8932 or bdavis@cahill.com; or Rahul Sarkar (counsel) at 212.701.3062 or rsarkar@cahill.com; or Nick Lee (associate) at 212.701.3257 or nhlee@cahill.com.