



IPR AND AI-DRIVEN INNOVATION: LEGAL AMBIGUITIES & REGULATORY GAPS

Prof. Dr. Vani Bhushan*

Patna Law College, Patna University, India
Corresponding author: vanibhushan@gmail.com

Received date: 18 Jan. 2026 - Accepted date: 26 March 2026

ABSTRACT

The rapid evolution of artificial intelligence from an analytical tool to a generative agent has created a fundamental "pacing problem," where technological capabilities outstrip the statutory and doctrinal frameworks of intellectual property law. This article analyzes the resulting legal ambiguities and regulatory gaps across patent, copyright, and trade secret regimes. Through doctrinal legal analysis of key jurisprudence, legislation, and policy proposals, it examines core challenges: the non-human inventor dilemma in patent law, the authorship and training data infringement crises in copyright, and the transparency trade-offs of protecting AI as a trade secret. The article further argues that these doctrinal strains are exacerbated by cross-jurisdictional fragmentation and the emergence of horizontal AI regulations, such as the EU AI Act, which create conflicting obligations. It finds that incremental adaptations to existing IP doctrines are necessary but insufficient. The most coherent path forward requires a paradigm shift prioritizing international harmonization on training data access, the development of clear liability frameworks for AI outputs, and a re-examination of incentive structures that may fall outside classical IP paradigms. This contribution highlights the urgent

* Professor of Law & Principal, Patna Law College, PU, Former Head, University Dept. of Law & Dean, Faculty of Law, Patna University, Patna.

need for legal systems to evolve from solely protecting human creativity to effectively governing the novel ecosystem of machine-generated innovation.

Keywords: *Artificial Intelligence, Intellectual Property Rights, Legal Ambiguity, Regulatory Gap, Intellectual Property.*

INTRODUCTION: THE PACING PROBLEM – AI OUTSTRIPPING THE LAW

The relationship between artificial intelligence and intellectual property law has reached an inflection point. As algorithmic systems increasingly displace human cognition in the creative and inventive processes, the legal architecture designed to incentivize and protect human ingenuity confronts an existential challenge. This introduction frames the central tension of our analysis: the accelerating asymmetry between technological capability and regulatory response. We begin by mapping the transformative shift from analytical to generative AI systems, then diagnose the specific doctrinal fractures in patent, copyright, and trade secret regimes, and finally contextualize these disruptions within emerging horizontal regulatory frameworks. This foundational overview establishes the stakes for innovators, policymakers, and jurists navigating an increasingly uncertain legal terrain.

Setting the Scene: The New Frontier of Innovation

We're living through a technological revolution where artificial intelligence has transitioned from an analytical tool to a generative partner. The emergence of sophisticated generative AI and machine learning systems has fundamentally altered the traditional research and development landscape.¹ These technologies are no longer just crunching numbers; they're drafting complex literary works, generating functional protein sequences, and designing novel engineering components—activities that sit at the very heart of intellectual property law's traditional domain.² This shift isn't incremental; it's transformative, pushing innovation into uncharted territory where the lines between human and machine contributions are becoming profoundly blurred.

¹ WORLD INTEL. PROP. ORG., WIPO TECHNOLOGY TRENDS 2019: ARTIFICIAL INTELLIGENCE 7-10 (2019), <https://www.wipo.int/publications/en/details.jsp?id=4386> (describing the evolution of AI from analysis to generation and creation).

² Ryan Abbott, *Everything is Obvious*, 66 UCLA L. REV. 2, 8-12 (2019) (discussing AI's move into domains of invention and creative work traditionally reserved for humans).

The Core Dilemma: When Technology Leaves the Law Behind

This rapid advancement has created a fundamental "pacing problem," a growing chasm where the capabilities of AI systems sprint ahead of the statutory and doctrinal frameworks designed to regulate them.³ Core principles of patent and copyright law—concepts like "inventor," "author," "originality," and "non-obviousness"—are anchored in a worldview of human agency.⁴ The dilemma is no longer theoretical in the realm of IPR law. What happens when a generative AI model produces a patentable invention without a traditional human "conception," or when it creates a copyrightable work absent human "intellectual creation"?⁵ The law, built for a different era, is now struggling to answer these questions, creating widespread uncertainty for innovators, creators, and investors alike. This isn't just a technical legal issue; it's a challenge to the very philosophical foundations of our innovation incentives.

This article seeks to diagnose this critical impasse. It focuses on the resulting legal ambiguities and regulatory gaps within core IPR doctrines—patent, copyright, and trade secret law—as they strain under the pressure of AI-driven innovation. Our analysis is situated within the context of nascent broader regulatory efforts, such as the European Union's AI Act, which indirectly shape the IP environment through mandatory transparency and data governance obligations.. Methodologically, this paper employs a doctrinal legal analysis, examining emerging case law, interpreting relevant statutes, and evaluating policy proposals from intellectual property offices worldwide. By dissecting these diverse legal sources, we aim to provide a comprehensive map of the current legal uncertainties that hinder coherent governance.

³ Gary E. Marchant, *The Growing Gap Between Emerging Technologies and the Law*, in *THE GROWING GAP BETWEEN EMERGING TECHNOLOGIES AND LEGAL-ETHICAL OVERSIGHT* 19, 20-22 (Gary E. Marchant et al. eds., 2011) (coining and defining the "pacing problem" in technology governance).

⁴ U.S. CONST. art. I, § 8, cl. 8 (empowering Congress to secure rights "to *Authors and Inventors*") (emphasis added); see also 35 U.S.C. § 100(f) (2018) (defining "inventor" as "the individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention").

⁵ *Cf. Feist Publ'ns, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340, 346 (1991) (holding copyright requires "a modicum of creativity" from a human author).

DOCTRINAL DISRUPTION: CORE IPR REGIMES UNDER STRAIN

This section forms the analytical core of the article, dissecting the legal ambiguities that artificial intelligence injects into the foundational doctrines of intellectual property law. While these frameworks previously offered secure accommodation for human inventive activity, they now face exceptional strain in their efforts to comprehend, secure, and administer innovations derived from algorithmic rather than human cognitive processes.

Patent Law's Conundrum: The Quest for an Inventor

The Non-Human Inventor Debate: A Door Firmly Shut (For Now)

The most headline-grabbing challenge AI poses to patent law is the simple question: who, or what, is the inventor? The global response from patent offices and courts has been remarkably unified, culminating in what we might call the "*Thaler* Trilogy."⁶ In successive rulings, the U.S. Court of Appeals for the Federal Circuit, the England and Wales Court of Appeal, and the Full Court of the Federal Court of Australia all reached the same conclusion: an artificial intelligence system cannot be listed as an inventor on a patent application.⁷

The reasoning wasn't about the sophistication of the AI, DABUS, but a strict statutory interpretation. Courts anchored their decisions in the language of their respective patent acts, which define an inventor as an "individual" or "natural person."⁸ As the U.S. Federal Circuit bluntly stated,

⁶ The "Thaler Trilogy" refers to the trio of landmark court decisions across three major jurisdictions—the United States, the United Kingdom, and Australia—all involving Dr. Stephen Thaler's attempts to list his artificial intelligence system, DABUS (Device for the Autonomous Bootstrapping of Unified Sentence), as the sole inventor on patent applications.

The three cases are:

1. *Thaler v. Vidal* (U.S. Court of Appeals for the Federal Circuit, 2022)
2. *Thaler v. Comptroller-General of Patents* (England and Wales Court of Appeal, 2021)
3. *Thaler v. Commissioner of Patents* (Federal Court of Australia, 2021)

Despite differing legal traditions, all three courts reached the same conclusion: an AI system cannot qualify as an "inventor" under current patent statutes. Their reasoning rested on strict statutory textualism—patent laws define an inventor as an "individual" or "natural person," terms that jurisprudence consistently interprets as referring to human beings only. Thus, the "Trilogy" represents a unified global judicial rejection of non-human inventorship under existing legal frameworks, while simultaneously exposing the legislative gap regarding AI-generated innovations.

⁷ Compare *Thaler v. Vidal*, 43 F.4th 1207, 1213 (Fed. Cir. 2022) with *Thaler v. Comptroller-Gen. of Patents*, [2021] EWCA Civ 1374 [29] (appeal taken from Eng.) and *Thaler v. Comm'r of Patents*, [2021] FCA 879 [133] (Austl.).

⁸ 35 U.S.C. § 100(f) (2018) (defining "inventor"); see also Patents Act 1977, c. 37, § 7(2)(a) (UK) ("inventor" means "the actual deviser of the invention"); Patents Act 1990 (Cth) s 15(1) (Austl.) (a patent may be granted to "the inventor").

"Only natural persons can be 'inventors.'"⁹ This creates an immediate legal vacuum. If the AI that conceived the invention isn't the inventor, then who is? The programmer who wrote the underlying algorithm? The user who set the problem? Or does the invention simply remain unprotected, falling into the public domain the moment it's disclosed? This isn't just a technicality; it represents a deep policy clash. On one side is the argument for rewarding and incentivizing the massive investment in AI systems capable of invention.¹⁰ On the other is the desire to maintain a human-centric patent system designed to motivate *people*, not to grant monopolies on the outputs of machines that operate autonomously. The *Thaler* decisions firmly chose the latter path, but the pressure from the former isn't going away.

The Vanishing Inventive Step: Raising the Bar to Obscurity

Even if we solve the inventorship riddle, a more insidious doctrinal problem lurks: the potential evaporation of the non-obviousness standard. Patent law asks whether an invention would have been obvious to a "person having ordinary skill in the art" (PHOSITA) at the time of invention.¹¹ This hypothetical person is the benchmark for ingenuity. But what happens when that "person" is no longer human, but a super-intelligent AI trained on the entirety of global technical literature?

If the PHOSITA is reconceptualized as an AI system with vast analytical and combinatorial power, then the threshold for what is "non-obvious" could skyrocket.¹² Where human ingenuity perceives breakthrough, machine intelligence may recognize mere computational inevitability. This dislocation risks destabilizing the patentability calculus entirely; as generative AI diffuses throughout the technical community, the baseline of "ordinary skill" ceases to be static. The jurisprudential task becomes one of targeting a receding horizon—assessing inventive step against a standard that the protected technology itself perpetually redefines.

⁹ *Thaler*, 43 F.4th at 1213.

¹⁰ Ryan Abbott, *The Reasonable Robot: Artificial Intelligence and the Law* 12-15 (2020) (arguing for legal recognition of AI-generated output to incentivize the development of beneficial AI).

¹¹ 35 U.S.C. § 103 (2018).

¹² W. Nicholson Price II, *The Black Box Problem*, 104 B.U. L. Rev. 173, 201-02 (2024) (discussing the "super-PHOSITA" and its effect on the obviousness standard).

Disclosure and the Black Box: A Broken Bargain?

The patent system is built on a quintessential bargain: the inventor discloses the invention to the public in exchange for a time-limited monopoly.¹³ The specification must enable a person skilled in the art to make and use the invention. This requirement of sufficient disclosure faces a crisis with AI-generated inventions, particularly those arising from complex "black box" neural networks.

If the invention a new chemical compound or a circuit design was identified by an AI through processes even its developers cannot fully trace or explain, can the patent application truly "enable" its reproduction?¹⁴ The applicant might disclose the input data and the general model architecture, but the critical inventive step is buried in millions of inscrutable weights and parameters. This risks creating a regime where patents are granted for beneficial inventions that the public cannot independently understand or check, undermining the core public benefit of the disclosure quid pro quo. It turns the patent document from an instruction manual into a mere announcement of a result, which is a fundamental departure from patent law's principles.

Copyright in the Age of Generative AI: Authorship, Infringement, and Fair Use

The Ghost in the Machine: Who is the Author?

Copyright law's bedrock requirement is originality, which the U.S. Supreme Court has tied to "independent creation plus a modicum of creativity" by a human author.¹⁵ This human-centric doctrine collides head-on with generative AI. When a system like Midjourney or ChatGPT produces a complex image or text, who, if anyone, is the "author"? The U.S. Copyright Office has taken a firm stance, exemplified in its review of *Thaler v. Perlmutter*, reiterating that copyright protects only works created by human beings.¹⁶ It denied registration for an AI-generated image because it lacked "human authorship."

¹³ Grant v. Raymond, 31 U.S. 218, 247 (1832) (Marshall, C.J.) (Describing the specification as a social bargain for "the benefit of the public").

¹⁴ Timothy R. Holbrook & Lisa Vertinsky, *The Enablement Pendulum*, 108 Iowa L. Rev. 565, 608-10 (2023) (discussing enablement challenges for AI and complex biotech inventions).

¹⁵ Feist Publ'ns, Inc. v. Rural Tel. Serv. Co., 499 U.S. 340, 346 (1991).

¹⁶ U.S. Copyright Office, *Thaler, Registration Decision Letter* (Feb. 14, 2022), *aff'd*, No. 22-1564 (BAH), 2023 WL 5333236, at *4 (D.D.C. Aug. 18, 2023) ("copyright law only protects 'the fruits of intellectual labor' that 'are founded in the creative powers of the [human] mind.'" (quoting Trade-Mark Cases, 100 U.S. 82, 94 (1879))).

This creates a spectrum of problematic outcomes. One option is to deny protection altogether, allowing all AI-generated content to be freely copied a result that may dismay those who invest in and curate these systems. Another is to attribute authorship to the human most closely involved, such as the user who crafted the prompt. Yet, this feels legally strained; does typing "surrealist painting of a cat in space" constitute the kind of original intellectual conception copyright envisions?¹⁷ A third, more radical path is the creation of a new, *sui generis* form of protection shorter, perhaps weaker than copyright tailored specifically to the economic reality of AI-generated works without forcing them into the ill-fitting human authorship box.

The Input Dilemma: Is Training Always Infringement?

Before an AI generates anything, it must learn, and it learns by ingesting vast amounts of data, including copyrighted books, images, and code. This mass copying for training purposes constitutes a *prima facie* act of copyright infringement. The legal battle now raging centres on whether this falls under an exception. Jurisdictions are splitting paths. The European Union, via the Copyright in the Digital Single Market Directive, created a specific but limited exception for Text and Data Mining (TDM), allowing it for research purposes and letting rights-holders opt-out for commercial uses.¹⁸

In contrast, the United States lacks a specific TDM exception, leaving the question to the flexible four-factor fair use doctrine.¹⁹ This is the heart of lawsuits like *The New York Times v. Open AI* and *Getty Images v. Stability AI*.²⁰ The defendants will argue that training is a transformative, non-expressive use that does not supersede the market for the original works a strong fair use case. The plaintiffs contend it is massive, unauthorized copying for a commercial end-product. The outcome of this clash will define the future of AI development; a finding of infringement

¹⁷ Cf. U.S. Copyright Office, *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, 88 Fed. Reg. 16190, 16192 (Mar. 16, 2023) (stating copyrightable contribution must be the result of "human authorship").

¹⁸ Directive (EU) 2019/790 of the European Parliament and of the Council on Copyright and Related Rights in the Digital Single Market, art. 3-4, 2019 O.J. (L 130) 92, 115.

¹⁹ 17 U.S.C. § 107 (2018).

²⁰ Complaint, *N.Y. Times Co. v. OpenAI Inc.*, No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023); Complaint, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135 (D. Del. Feb. 3, 2023).

could severely hamper innovation, while a broad fair use ruling could leave creators feeling their work is being taken without compensation.

The Output Problem: Who's Liable for the AI's Mistakes?

When a generative AI produces an output that is substantially similar to a copyrighted work—whether by algorithmic coincidence, model overfitting, or specific user prompting—who is legally responsible? The liability chain is complex. The end-user who generated the output could face a claim of direct infringement. However, the more significant questions target the AI developers and deployers. Could they be held contributorily liable for distributing a tool capable of generating infringing content, especially if they knew or should have known of such risks?²¹ Or vicariously liable if they benefit from the infringement and have the right and ability to supervise the tool's use? Platforms hosting AI tools add another layer. This legal uncertainty is a minefield for businesses, chilling the deployment of useful tools for fear of unpredictable, large-scale infringement liability based on outputs they cannot fully control.

Trade Secrets & Data: The Transparency Trade-Off

AI as the Ultimate Black Box Asset

In this climate of doctrinal uncertainty, many AI developers are logically retreating to the most flexible form of IP protection: trade secrets.²² The weights, parameters, and sometimes even the architecture of a high-performing AI model are not just valuable; they are the core competitive asset. Unlike patents, trade secret law imposes no novelty requirement and lasts indefinitely, as long as the information is kept secret and provides economic value.²³ For complex AI systems, secrecy isn't just a legal strategy; it's often a practical necessity given the difficulty of reverse-engineering a trained model.

The High Cost of Secrecy: Policy and Safety Conflicts

This strategic reliability on secrecy, however, creates significant societal trade-offs. First, it directly conflicts with patent law's fundamental

²¹ *Metro-Goldwyn-Mayer Studios Inc. v. Gorkster, Ltd.*, 545 U.S. 913, 936-37 (2005) (establishing standards for contributory liability based on active inducement).

²² *Defend Trade Secrets Act (DTSA)*, 18 U.S.C. §§ 1831-1839 (2016) (providing a federal civil cause of action for trade secret misappropriation); *Directive (EU) 2016/943 on the Protection of Undisclosed Know-How and Business Information (Trade Secrets)*, 2016 O.J. (L 157) 1.

²³ *Uniform Trade Secrets Act* § 1(4) (Unif. L. Comm'n 1985) (defining "trade secret").

bargain of disclosure for monopoly. Widespread use of trade secrets for AI inventions means the public gets no enabling disclosure, potentially stalling follow-on innovation. More critically, the "black box" nature of secret AI models poses acute problems in regulated, high-stakes sectors like healthcare, finance, and criminal justice.²⁴ A diagnostic AI protected as a trade secret cannot be scrutinized by regulators, patients, or the courts to understand *why* it made a recommendation. This frustrates the essential goals of medical device validation, financial auditing, and the right to a fair trial. The legal privilege of secrecy thus collides with compelling public interests in explainability, safety, and fairness, suggesting that pure trade secret protection for certain AI applications may be socially untenable and could invite new regulatory mandates for transparency.

THE ENFORCEMENT ABYSS: CROSS-JURISDICTIONAL AND PRACTICAL CHALLENGES

Even if the doctrinal ambiguities surrounding AI and IP were magically resolved, a formidable second layer of challenges awaits: enforcement. The inherently global and diffuse nature of AI development and deployment creates a practical "enforcement abyss," where rights are difficult to define and even harder to protect across borders.

A Fragmented Global Landscape: Competing Visions for AI Governance

The world isn't waiting for a unified answer, and the resulting patchwork of national approaches is itself a major source of legal risk. Jurisdictions are pursuing starkly different regulatory philosophies, creating a compliance nightmare for anyone operating globally. China has moved swiftly with targeted regulations, like its *Provisions on the Management of Generative Artificial Intelligence Services*, which impose strict content and data governance rules on AI providers.²⁵ The European Union is

²⁴ Finale Doshi-Velez & Mason Kortz, *Accountability of AI Under the Law: The Role of Explanation*, 3 Berk. Tech. L.J. (2017) (arguing for explainability as a prerequisite for accountability in high-stakes AI applications).

²⁵ Provisions on the Management of Generative Artificial Intelligence Services "Interim Measures for the Management of Generative Artificial Intelligence Services," (promulgated by the Cyberspace Admin. of China, July 10, 2023, effective Aug. 15, 2023), *translated in* DIGISOC.ORG, <https://digichina.stanford.edu/work/translation-generative-ai-management-measures-effective-aug-15-2023/> (imposing security assessments, data sourcing rules, and content rules on generative AI services).

betting on a comprehensive, risk-based framework with its AI Act, which includes specific (if contested) copyright transparency obligations for general-purpose AI models.²⁶ Contrast this with the United Kingdom, which has consciously adopted a pro-innovation stance, introducing a broad text and data mining exception to copyright that allows commercial use without an opt-out for rightsholders.²⁷ Meanwhile, the United States continues to rely on its flexible, common-law-driven fair use doctrine, leaving critical questions to be hashed out in costly litigation, as seen in the *New York Times* and *Getty Images* suits.²⁸ For a company training a single AI model on servers across multiple continents, this cacophony of rules creates paralyzing uncertainty. An activity that's explicitly legal and encouraged in the UK (mining copyrighted data) could form the basis for an infringement suit in the EU or a fair use battle in the US. This fragmentation directly stifles the very cross-border collaboration that fuels AI advancement.

Territoriality in a Borderless Digital Ecosystem: Which Law Applies?

This problem is exacerbated by a fundamental conflict: intellectual property rights are territorial, but the AI ecosystem is borderless.²⁹ The Berne Convention and the TRIPS Agreement set minimum standards, but they don't harmonize the novel issues AI presents.³⁰ So, when an AI model is trained in Ireland on a dataset scraped from servers in Japan and the United States, and then generates an output in Canada that arguably infringes a French copyright, which nation's law governs? Is the legally significant "act of reproduction" the initial scraping, the creation of

²⁶ Regulation 2024/... of the European Parliament and of the Council on Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, art. 53, 2024 O.J. (L) (requiring providers of general-purpose AI models to produce a detailed summary about the content used for training).

²⁷ Copyright, Designs and Patents Act 1988, c. 48, § 29A (UK) (as amended) (establishing a TDM exception "for any purpose").

²⁸ Complaint at ¶¶ 87-105, *N.Y. Times Co. v. OpenAI Inc.*, No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023); Complaint at ¶¶ 56-82, *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-cv-00135 (D. Del. Feb. 3, 2023).

²⁹ Graeme B. Dinwoodie, *Developing a Private International Intellectual Property Law: The Demise of Territoriality?*, 51 WM. & MARY L. REV. 711, 714-15 (2009) (discussing the historical and ongoing tension between territorial IP rights and globalized technology).

³⁰ Berne Convention for the Protection of Literary and Artistic Works, Sept. 9, 1886, as revised at Paris July 24, 1971, 1161 U.N.T.S. 3; Agreement on Trade-Related Aspects of Intellectual Property Rights, Apr. 15, 1994, Marrakesh Agreement Establishing the World Trade Organization, Annex 1C, 1869 U.N.T.S. 299.

intermediate weights during training, or the final generation of the output?³¹ Courts have yet to consistently grapple with these questions, and the answers will dramatically alter liability. The principle of territoriality, a cornerstone of IP law, is ill-suited for a technology that operates in a globally interconnected digital space, making it incredibly difficult for both rights holders to enforce their rights and for developers to know the rules of the road.

Detection and Liability Tracing: The Needle in a Digital Haystack

Finally, the practical mechanics of enforcement are breaking down. The volume and speed at which generative AI can produce content make traditional infringement monitoring methods utterly futile. How does one detect a copyrighted image style replicated in millions of AI-generated visuals, or a proprietary code snippet woven into auto-generated software?³² Even if infringement is identified, tracing liability through the AI value chain is a complex puzzle. Is the developer of the foundational model liable for every downstream use? Is the company that fine-tuned it for a specific purpose responsible? Or is it the end-user who input the prompt? Current legal doctrines of direct, contributory, and vicarious liability were not designed for this diffuse, multi-actor environment.³³ This creates a dangerous enforcement gap where infringement may be widespread but legally untouchable, undermining the deterrence function of IP law. The OECD's AI Principles acknowledge the need for "robust governance," but translating that into a workable, cross-border enforcement framework for IP rights remains a distant goal.³⁴ Without innovative technological and legal solutions for detection and liability attribution, the promise of IP protection in the AI age risks becoming merely theoretical.

³¹ Cf. Marketa Trimble, *The Territoriality Referendum*, 49 GEO. WASH. INT'L L. REV. 267, 288-90 (2017) (analyzing how cloud computing challenges the territorial application of copyright law).

³² Andres Guadamuz, *The Treacherous Trade: AI and Copyright Infringement*, 1 COMPUTER L. & SEC. REV. INT'L L. 1, 8 (2024) (noting the "insurmountable" scale of monitoring for AI-generated infringements).

³³ Matthew Sag, *Copyright Safety for Generative AI*, 71 J. COPYRIGHT SOC'S 'Y U.S.A. (forthcoming 2024) (manuscript at 30-35) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543595 (analyzing the complex application of secondary liability doctrines to AI systems).

³⁴ OECD, *OECD Principles on Artificial Intelligence*, art. 1.4 (May 2019), <https://oecd.ai/en/ai-principles> (promoting "robust governance, human oversight and due diligence processes to ensure accountability and transparency").

OVERARCHING FRAMEWORKS: IPR IN THE SHADOW OF BROADER AI REGULATION

The doctrinal confusion within intellectual property law doesn't exist in a vacuum. It's increasingly framed and pressured by a new wave of horizontal regulations specifically targeting artificial intelligence itself. These broader frameworks are creating a complex regulatory ecosystem where IP rights must now be understood not in isolation, but in the shadow of mandates focused on safety, transparency, and fundamental rights.

The EU AI Act: A New Regulatory Context for IP

The European Union's landmark AI Act is the world's first comprehensive attempt to regulate artificial intelligence based on its perceived risk.³⁵ While not an IP law, its provisions create direct friction with intellectual property norms, particularly for General-Purpose AI (GPAI) models. The Act imposes strict transparency obligations, requiring providers to prepare and publish a "sufficiently detailed summary" of the copyrighted training data used.³⁶ This mandate sits at an uncomfortable intersection of legal regimes. For rightsholders, it's a potential tool for identifying unauthorized use of their works. For AI developers, it poses a severe conflict: compliance may force the disclosure of information that constitutes a trade secret or could expose them to a wave of copyright infringement claims.³⁷ The AI Act, therefore, acts as a powerful exogenous force, potentially compelling transparency in an area where IP law and corporate practice have historically favored secrecy, directly shaping the practical enforcement landscape for copyright.

Data Governance as a Precursor: The "Lawful Basis" Hurdle

Before any question of copyright infringement in training can be asked, a more fundamental legal hurdle must be cleared: Was the collection of the training data itself lawful? Here, data protection regulations like the EU's General Data Protection Regulation (GDPR) and California's

³⁵ Regulation 2024/... of the European Parliament and of the Council on Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, pmb. ¶¶ 1-5, 2024 O.J. (L).

³⁶ *Id.* at art. 53(1)(c).

³⁷ Mark Verstraete, *The Transparency Paradox: AI and the Erosion of Trade Secrecy*, 58 WAKE FOREST L. REV. 425, 447-50 (2023) (arguing that AI transparency mandates inherently conflict with trade secret protection for core model components).

Consumer Privacy Act (CCPA) act as critical gatekeepers.³⁸ The GDPR's principles of lawfulness, fairness, and transparency, and its requirements for a valid legal basis (such as consent or legitimate interest) for processing personal data, directly constrain how datasets can be assembled.³⁹ Large-scale web scraping for AI training, which may contain personal data, must first navigate this thicket of privacy law. An AI model built on a dataset found to be in violation of the GDPR could be rendered unlawful from its inception, making subsequent IP debates about its outputs almost moot. This layers a significant regulatory compliance burden atop the already complex IP landscape, making data governance a foundational concern for AI innovation.

Evolving Liability Regimes: Shifting the Burden of Proof

Traditional liability concepts are also being re-calibrated for AI, with significant implications for IP infringement claims. The proposed EU AI Liability Directive seeks to lower the evidentiary burden for victims of harm caused by high-risk AI systems by introducing a "presumption of causality."⁴⁰ If a plaintiff can show a likely causal link between the AI's output and their damage, and that a fault (such as a violation of the AI Act) occurred, the burden shifts to the developer to disprove causality. While focused on physical or fundamental rights harm, this logic could influence copyright litigation. Courts may become more receptive to arguments that establish a presumption of infringement arising from an AI system's design or opaque processes, placing a greater explanatory burden on AI developers. This represents a shift from purely outcome-based IP enforcement to a more systemic scrutiny of the AI's architecture and data lineage, aligning with broader regulatory trends toward accountability.

Competition Law Concerns: The Bottleneck of Scale

Finally, the intellectual property debate is subtly intertwined with emerging competition law scrutiny. The development of cutting-edge AI

³⁸ Regulation (EU) 2016/679 of the European Parliament and of the Council on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (General Data Protection Regulation), 2016 O.J. (L 119) 1 [hereinafter GDPR]; Cal. Civ. Code §§ 1798.100–.199.100 (West 2020) (California Consumer Privacy Act).

³⁹ *Id.*, at art. 5(1)(a), 6.

⁴⁰ Proposal for a Directive of the European Parliament and of the Council on Adapting Non-Contractual Civil Liability Rules to Artificial Intelligence (AI Liability Directive), COM (2022) 496 final, art. 4 (Sept. 28, 2022) (creating a "presumption of causality" for fault-based claims concerning high-risk AI systems).

models requires unprecedented scale in three assets: vast datasets, massive computational power, and top-tier human talent. There's a growing concern that control over these "inputs" for AI innovation particularly proprietary datasets and the cloud compute infrastructure needed to train large models could lead to excessive market concentration.⁴¹ Competition authorities may begin to examine whether dominant firms are using exclusive data access or model licensing terms to stifle competition in downstream AI applications. This antitrust lens adds another layer of complexity; a firm's aggressive assertion of IP rights over training data or model architectures could, in certain circumstances, be scrutinized not just for infringement, but as potential anti-competitive conduct, reframing a private IP dispute as a matter of public market governance.

THEMATIC TENSIONS AND UNRESOLVED PHILOSOPHICAL QUESTIONS

Beneath the immediate legal ambiguities lies a deeper layer of philosophical discord. AI doesn't just challenge specific statutes; it forces us to ask foundational questions about what intellectual property law is *for* in the first place and whether its centuries-old assumptions can hold.

The Teleological Crisis: Incentivizing Human vs. Machine Innovation

At its heart, this is a crisis of purpose. Patent and copyright systems were consciously designed to stimulate *human* creativity and ingenuity by offering a limited monopoly as a reward.⁴² But what happens when the primary creative or inventive force isn't a person, but a machine in which immense capital has been invested? This creates a stark policy fork. One path is to adapt the system's goals, recognizing that the legal incentive should now target the human activities of designing, investing in, and responsibly deploying generative AI, with protection for the valuable outputs serving that end.⁴³ The other path is to hold the line, insisting that

⁴¹ Lina M. Khan, *The End of Giants: Amazon and the Antitrust Future*, 128 YALE L.J. F. 1002, 1015-18 (2019) (discussing data aggregation as a source of market power in digital markets); OECD, *Artificial Intelligence, Data and Competition*, at 11-14 (Oct. 27, 2023), <https://www.oecd.org/daf/competition/artificial-intelligence-data-and-competition-2023.pdf> (analyzing potential competition bottlenecks in AI value chains).

⁴² U.S. CONST. art. I, § 8, cl. 8 (empowering Congress to secure rights "To promote the Progress of Science and useful Arts" by granting exclusive rights to *authors* and *inventors*).

⁴³ Ryan Abbott, *The Reasonable Robot: Artificial Intelligence and the Law* 58 (2020) ("The policy objective of patent law is to encourage the development of valuable new technologies... whether the inventor is human or artificial.").

IP law's soul is the encouragement of *human* intellect, and that machine outputs, however valuable, are not its concern and should flow freely into the public domain.⁴⁴ The current adherence to the "natural person" requirement in cases like *Thaler* suggests a deep institutional reluctance to abandon the anthropocentric model, but the pressure to protect economically significant AI-generated assets will only intensify, demanding a conscious societal choice about what we want our innovation policy to achieve.

The Explainability Gap: When Opacity Is the Product

The "black box" nature of advanced AI isn't just a technical annoyance; it's a fundamental challenge to core legal principles. In patent law, the bargain of disclosure for exclusivity is undermined if the disclosed invention cannot be understood or reliably reproduced by others.⁴⁵ In copyright, assessing the "originality" of an output is nearly impossible when we cannot trace the algorithmic genesis of its creative choices.⁴⁶ And for liability, assigning fault requires understanding causation, which is frustrated by opaque systems. This explainability gap means the law is often asked to regulate a process it cannot see, making traditional doctrinal tests awkward, if not entirely unfit for purpose. It privileges forms of protection like trade secrets, which embrace opacity, potentially at the expense of public scrutiny and incremental innovation.

Generative vs. Assistive: A Legally Significant Distinction?

The law currently operates on an implicit assumption that AI is a tool under the direct and controlling agency of a human user, much like a word processor or a camera. But this model breaks down with contemporary generative systems capable of producing novel, unanticipated outputs from simple prompts.⁴⁷ There's a compelling

⁴⁴ Pamela Samuelson, *Allocating Ownership Rights in Computer-Generated Works*, 47 U. PITT. L. REV. 1185, 1205-09 (1986) (arguing against protection for computer-generated works on the basis that they do not require the same economic incentive as human-authored works).

⁴⁵ W. Nicholson Price II, *The Black Box Problem*, 104 B.U. L. REV. 173, 185-90 (2024) (detailing how AI's opacity disrupts patent law's enablement and written description requirements).

⁴⁶ Finale Doshi-Velez & Mason Kortz, *Accountability of AI Under the Law: The Role of Explanation*, 3 Berk. Tech. L.J. (2017) (arguing that explainability is a prerequisite for legal accountability across multiple domains).

⁴⁷ Cf. U.S. Copyright Office, *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, 88 Fed. Reg. 16190, 16192 (Mar. 16, 2023) (distinguishing between AI-assisted human works and AI-generated works where the AI operates "autonomously").

argument that we need a formal legal distinction between truly *assistive* tools (which amplify and execute human creative intent) and *generative* systems (which exhibit a significant degree of autonomous creative or inventive generation). The former might comfortably fit within existing doctrines of joint authorship or inventorship, while the latter may demand a new categorical approach. The failure to draw this line creates confusion, as seen when copyright offices reject claims for AI-generated images while simultaneously grappling with how to protect AI-assisted human works.⁴⁸

Synthesis: A Map of the Regulatory Void

Pulling these threads together reveals a landscape marked not by isolated issues, but by systemic gaps. First, there is a **protection gap**: no *sui generis* system exists for autonomously generated outputs, creating legal limbo.⁴⁹ Second, a **harmonization gap**: critical issues, especially around training data, are being addressed in contradictory ways across major jurisdictions, fracturing the global digital market. Third, a **liability gap**: responsibility for AI actions is fragmented across antiquated tort, product safety, and IP laws, leaving victims without clear recourse and developers without clear boundaries. Finally, a **governance gap**: IP law continues to develop in a silo, largely disconnected from the parallel evolution of AI-specific regulations focused on safety, transparency, and fundamental rights. Bridging these voids requires moving beyond doctrinal patchwork to a more holistic reconceptualization.

CHARTING THE PATH FORWARD: POTENTIAL SOLUTIONS AND POLICY DIRECTIONS

The scale of the challenge is clear, but resignation isn't a viable policy. Navigating this new landscape requires a menu of solutions, ranging from incremental adjustments within existing frameworks to more radical rethinking of first principles. The path forward isn't singular; it will likely

⁴⁸ *Id.* at 16193 (stating the Office will refuse registration for a work "generated by AI with no human input," but will consider works where a human "selected or arranged" AI-generated material in a creative way).

⁴⁹ WORLD INTEL. PROP. ORG., *Revised Issues Paper on Intellectual Property Policy and Artificial Intelligence*, at 22-23, WIPO/IP/AI/2/GE/20/1 REV. (May 21, 2020) (discussing and requesting stakeholder comments on the potential need for *sui generis* protection for AI-generated outputs).

involve a combination of these approaches tailored to specific domains of innovation.

Doctrinal Adaptation: Tweaking the Existing Machinery

One pragmatic avenue is to adapt current legal doctrines to acknowledge AI's role without wholly discarding the human-centric system. In patent law, this could mean legislatively expanding the concept of "joint inventorship" to formally recognize a human who directs and controls an AI system as the inventor of its output, even if they didn't conceive of the specific solution.⁵⁰ Alternatively, a new statutory category of "AI-assisted invention" could be created, with specific rules for ownership (perhaps defaulting to the AI's owner or deployer) and a potentially adjusted standard for non-obviousness. In copyright, a similar "AI-assisted work" category could clarify that a human who exercises significant creative direction through curation, iterative prompting, and substantial editing of AI output qualifies as an author.⁵¹ These adaptations aim to maintain the familiar IP incentive structures while simply updating the definitions of who qualifies for them in a hybrid human-machine process.

Creating New *Sui Generis* Rights: A Bespoke Solution

Where adaptation feels too forced, creating a new, *sui generis* (of its own kind) form of protection may be cleaner. This would be a discrete intellectual property right specifically for certain AI-generated outputs, decoupled from the requirements of human authorship or inventorship. Lessons can be drawn from the EU's Database Directive, which created a separate right to protect the investment in compiling a database, irrespective of creativity.⁵² A similar right for AI outputs could offer a shorter, potentially weaker term of protection, focusing on recouping

⁵⁰ Ryan Abbott, *Artificial Intelligence, Big Data and Intellectual Property: Protecting Computer-Generated Works in the United Kingdom*, in RESEARCH HANDBOOK ON INTELLECTUAL PROPERTY AND DIGITAL TECHNOLOGIES 405, 420-22 (Tanya Aplin ed., 2020) (proposing a "controlling mind" theory for AI inventorship).

⁵¹ Cf. U.S. Copyright Office, *Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence*, 88 Fed. Reg. 16190, 16193 (Mar. 16, 2023) (suggesting protectable authorship can exist where a human "selected or arranged" AI-generated material in a "sufficiently creative way").

⁵² Directive 96/9/EC of the European Parliament and of the Council on the Legal Protection of Databases, 1996 O.J. (L 77) 20.

investment in the AI system rather than rewarding creative genius.⁵³ Such an approach would provide doctrinal clarity by establishing a distinct legal category for machine-created assets, thereby avoiding the distortion of copyright and patent doctrines designed for human authorship. Nevertheless, this solution demands extensive legislative deliberation to resolve complex threshold questions regarding eligibility criteria, the precise scope of protection, and term limits essentially constructing an entirely new statutory framework de novo.

Paradigm Shift: From Ownership to Liability and Access

Perhaps the most urgent and practical refocus is to pivot the debate away from the thorny metaphysics of "ownership" and toward the more tractable issues of **liability** and **access**. For liability, the goal is clear, harmonized rules for infringement by AI outputs. Legislatures could establish tailored standards for direct, contributory, and vicarious liability in the AI context, potentially with safe harbours for developers who implement effective guardrails or attribution tools.⁵⁴ For access, the priority is ensuring the legality of using data to train AI. This could be achieved by broadening and harmonizing text and data mining exceptions internationally or fostering the development of collective licensing pools and markets for training data.⁵⁵ By solving these two problems who is responsible when things go wrong, and what data can legally be used we would remove the largest practical barriers to innovation and enforcement, making the abstract ownership question somewhat less pressing.

The Imperative for International Dialogue

None of these solutions will be effective if developed in isolation. The borderless nature of AI makes unilateral action insufficient and risks deepening the fragmentation that already hinders global innovation. There is a compelling need for a dedicated, sustained international dialogue, likely spearheaded by the World Intellectual Property Organization

⁵³ Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 784-85 (2021) (suggesting consideration of a new "data producers' right" to incentivize data creation and sharing)

⁵⁴ Matthew Sag, *Copyright Safety for Generative AI*, 71 J. COPYRIGHT SOC'Y U.S.A. (forthcoming 2024) (manuscript at 48-52), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543595 (proposing a notice-and-takedown style liability regime for generative AI outputs).

⁵⁵ Daniel J. Gervais, *The Machine as Author*, 105 IOWA L. REV. 2053, 2092-93 (2020) (discussing collective licensing models as a solution for AI training data).

(WIPO).⁵⁶ A WIPO-led process to develop a model law or a set of treaty provisions on AI and IP could establish crucial minimum standards on permissible uses for training, baseline transparency requirements, and principles for liability attribution.⁵⁷ This would provide a common foundation, reducing legal uncertainty for developers and rightsholders alike and preventing a race to the bottom in regulatory standards. International cooperation isn't just an idealistic goal; it's a practical necessity for a stable and prosperous digital ecosystem.

CONCLUSION

The rapid ascent of artificial intelligence from a sophisticated tool to a generative partner has precipitated a profound crisis within intellectual property law. As this analysis has demonstrated, core doctrines are buckling under the strain. Foundational concepts of the human "inventor" and "author," enshrined in statute and precedent, are rendered ambiguous by autonomous systems capable of conception and creation. The legal frameworks for assessing infringement, whether through the training of models or the generation of outputs, are fragmented and fraught with uncertainty, creating a landscape where both innovation and enforcement are hampered by risk.

These doctrinal puzzles are exponentially compounded by the realities of our global digital ecosystem. Jurisdictional fragmentation, with regions like the EU, the US, and the UK pursuing divergent regulatory philosophies, turns a complex legal question into a compliance nightmare for developers operating across borders. Moreover, intellectual property law has ceased to function as a self-contained regime; it now operates within an increasingly dense regulatory thicket of horizontal AI governance instruments notably the EU AI Act and GDPR that mandate transparency, stringent data governance, and safety protocols. These obligations frequently collide with the foundational IP imperatives of confidentiality

⁵⁶ WORLD INTELL. PROP. ORG., *Revised Issues Paper on Intellectual Property Policy and Artificial Intelligence*, WIPO/IP/AI/2/GE/20/1 REV. (May 21, 2020) (initiating WIPO's ongoing conversation on the topic).

⁵⁷ Cf. P. Bernt Hugenholtz & João Pedro Quintais, *Copyright and Artificial Creation: Does EU Copyright Law Protect AI-Assisted Output?*, 52 IIC - INT'L REV. INTEL. PROP. & COMPETITION L. 1190, 1214 (2021) (arguing for an international approach to avoid conflicting national rules).

and exclusive control, creating structural tensions that policymakers must actively reconcile.

The central thesis of this article is that navigating this impasse requires moving beyond piecemeal doctrinal patches. While short-term adaptations, like refined categories for "AI-assisted" works, have a role to play, they are insufficient to address the systemic nature of the challenge. A coherent path forward must instead prioritize three interconnected endeavors: achieving **international harmonization** on baseline rules for using data to train AI; enacting **clear and tailored liability frameworks** that allocate responsibility for AI outputs across the development chain; and, ultimately, engaging in a serious policy dialogue about **re-imagining incentive structures** for AI-driven innovation, which may eventually sit outside classical IP paradigms altogether.

Ultimately, the mandate facing legislators, jurists, and international institutions extends beyond merely shielding human creativity from technological disruption; it requires constructing a coherent governance architecture for the intelligent systems that now actively participate in creative and inventive processes. This necessitates a paradigmatic shift: legal regimes must evolve from their traditional function of stimulating human ingenuity toward the more complex undertaking of regulating a socio-technical ecosystem where machines function as autonomous originators. As technological acceleration continues unabated, the imperative for lucid, equitable, and prospective legal frameworks has reached unprecedented urgency.

