Sterne Kessler
STERNE KESSLER GOLDSTEIN & FOX

2025

AI IP Year in Review

ANALYSIS & TRENDS | 2ND EDITION

Table of Contents

Introduction

This year, intellectual property (IP) law and policy directly addressed the realities of generative artificial intelligence (AI). With a new director at the helm, the U.S. Patent and Trademark Office (USPTO) issued new guidance for AI-assisted inventions. On the litigation side, the Federal Circuit issued its first *Alice* analysis for machine learning patents in *Recentive Analytics, Inc. v. Fox Corp., et al.* 2023-2437, addressing subject matter eligibility under § 101. And several federal district courts weighed in on fair use in the context of training generative AI. AI is also reshaping industries, including the life sciences, with the FDA issuing draft guidance for AI use in medical devices, drugs, and biologics.

Globally, the European Union implemented the world's first comprehensive legal framework for AI through the Artificial Intelligence Act, setting enforceable requirements that already affect how AI systems are developed, deployed, and licensed, including for U.S. companies. Other major jurisdictions, including the United States, China, and Brazil, are advancing parallel AI governance regimes that rely heavily on deployment and procurement controls and that increasingly shape technical design choices around auditability, traceability, provenance, and safety.

We also saw law firms innovate and leverage AI technology in IP practice. In March 2025, Sterne Kessler introduced Patent Assist AI, a proprietary tool designed to streamline patent specification drafting. In addition, we launched our internal R&D team, Sterne Kessler Labs, which is responsible for building, testing, and integrating AI technology into our IP practice.

In our second edition of AI IP Year in Review, we present a comprehensive overview of the legal decisions, policy updates, and technology developments that defined AI and IP in 2025. This report builds on the insights we shared throughout the year, highlighting the trends, practical implications, and strategic considerations for businesses and innovators as they navigate an increasingly complex and rapidly evolving IP and AI landscape in 2026.

The information provided in this review is the result of a collaborative process. We would like to extend our thanks to our contributing authors, whose efforts have enriched this publication.

We appreciate your interest in this report, and we encourage you to read our firm's other recently released publications: 2025 Design Patents Year in Review: Analysis & Trends, 2025 PTAB Year in Review: Analysis & Trends, Federal Circuit IP Appeals: Summaries of Key 2025 Decisions, and 2025 ITC Year in Review: Analysis & Trends, which are available at sternekessler.com or by request. Please contact us if you have questions about this report or would like to discuss AI and IP issues.

Todd Hopfinger
Director, Electronics Practice Group

Lestin Kenton, Jr.
Director, Electronics Practice Group

Editors & Authors

Todd Hopfinger, Director Editor
thopfinger@sternekessler.com

Lestin Kenton, Jr., Director Editor
lkenton@sternekessler.com

Matthew Bodenstein, Director
mbodenstein@sternekessler.com

Dohm Chankong, Director
dchankong@sternekessler.com

Richard Coller III, Director
rcoller@sternekessler.com

Richard Crudo, Director
rcrudo@sternekessler.com

Jason Eisenberg, Director
jasone@sternekessler.com

Ivy Clarice Estoesta, Director
iestoest@sternekessler.com

Roozbeh Gorgin, Counsel
rgorgin@sternekessler.com

Sasha Rao, Counsel
srao@sternekessler.com

Lauren Schleh, Counsel
lschleh@sternekessler.com

Ryan Estatico, Associate*
restatico@sternekessler.com

**Admitted only in Pennsylvania. Supervised by D.C. Bar member Dohm Chankong.*

Benjamin M. Gitzinger, Associate
bgitzinger@sternekessler.com

Brooke N. McLain, Associate
bmclain@sternekessler.com

Jean Selep, Associate
jselep@sternekessler.com

Ananth Josyula, Technical Specialist
ajosyula@sternekessler.com

Regulatory Update: FDA Issues Draft Guidances for AI Use in Medical Devices and for Drugs and Biologics

BY: SASHA RAO AND MATTHEW BODENSTEIN

In January 2025, the U.S. Food and Drug Administration (FDA) issued two draft guidance documents discussing: (1) the use of artificial intelligence (AI) to produce information to support a regulatory decision about a drug or biological product's safety, effectiveness, or quality; and (2) the development and marketing of safe and effective AI-enabled devices.

Drug & Biological Products Guidance

The use of AI to produce data and information supporting regulatory filings regarding safety, effectiveness, and quality of drugs faces challenges such as dataset bias, model opacity, output accuracy verification, and performance drift. To mitigate these, FDA proposes a risk-based credibility assessment framework featuring seven steps: defining the AI model's role and scope (steps 1-2), assessing model risk based on influence and consequences (step 3), creating and executing a credibility assessment plan (steps 4-6), and evaluating the model's adequacy for its specific use (step 7). The guidance emphasizes maintaining AI credibility throughout a product's life cycle and provides hypothetical examples in clinical development and manufacturing.

Medical Device Guidance

The FDA's second draft guidance outlines requirements for marketing submissions for medical devices with AI-enabled software functions. This guidance includes FDA's current thinking on transparency and bias in FDA's Total Product Life Cycle of AI-enabled devices. FDA provides recommendations and rationales for, among other things, risk assessment, data management, and performance validation. The guidance explains that sponsors must provide detailed information covering device uses, inputs, outputs, architecture, development, performance, installation, and maintenance in both submissions and product labels.

Takeaways

- Sponsors implementing AI solutions for drug or biologic safety, effectiveness, or quality information should consider FDA's credibility assessment framework when designing and training their AI model.
- Sponsors seeking approval for medical devices containing AI components should evaluate FDA's requirements over a product's predicted lifetime.
- Early engagement between the company's AI experts and in-house counsel is key to ensure a sponsor is meeting all of FDA's requirements to move a product to market.

Updates

The public comment period during which the FDA requested feedback on: (1) performance metrics and indicators, (2) real-world evaluation methods and infrastructure, (3) post-market data sources and quality management, (4) monitoring triggers and response protocols, (5) human-AI interactions and user experience, and (6) additional considerations and best practices closed on December 1, 2025. Public feedback can be found here: Measuring and Evaluating Artificial Intelligence-enabled Medical Device Performance in the Real-World.

Read our original client alert, FDA Issues Draft Guidance Documents on Artificial Intelligence for Medical Devices, Drugs, and Biological Products.

AI Regulation Meets Innovation: IP Considerations as the EU AI Act Takes Effect

BY: LESTIN KENTON, JR., AND ROOZBEH GORGIN

The EU Artificial Intelligence Act—the world's first comprehensive regulatory framework for AI—will reshape how companies build, document, and protect innovations deployed in Europe. Although the Act was adopted in 2024, its major compliance obligations will phase in through 2026 and 2027, steadily expanding the scope of technical information companies must disclose to regulators. These mandatory disclosures, while designed to ensure safety and accountability, carry profound implications for patent protection, inventorship, and trade secret management.

A central challenge stems from the Act's risk-tiered structure. High-risk AI systems—including those used in healthcare, employment, education, and critical infrastructure—must undergo a conformity assessment before entering the EU market. Providers must supply regulators with detailed technical documentation covering system architecture, datasets, testing methodologies, risk mitigation measures, and human oversight mechanisms. Providers of general-purpose and foundation models face additional obligations, such as preparing dataset summaries and providing systematic evaluations of model behavior.

For patent owners and practitioners, these requirements create a new disclosure risk: regulatory filings may reveal algorithmic or data-related details that overlap with potential patent applications. If submitted before a patent filing, such information may constitute a public disclosure that threatens novelty, undermines claim scope, or narrows enforceability. The Act also impacts trade secret strategies. Mandatory dataset summaries and technical files weaken confidentiality protections, shifting the calculus toward patenting innovations that can no longer be reliably maintained as secrets.

Takeaways

- **Regulatory disclosures may become prior art.** Technical submissions must be sequenced to avoid prefiling exposure of inventive concepts.
- **Foundation-model documentation affects value.** Dataset summaries and risk evaluations may reveal use of copyrighted or unlicensed data, creating litigation or freedom-to-operate risks that impact portfolio valuation.
- **Inventorship disputes become more complex.** The Act's recordkeeping obligations create evidentiary trails that may help delineate human inventive contributions from AI assistance—essential in jurisdictions recognizing only human inventors.
- **EU conformity assessments may constrain claim drafting.** EU-aligned claims may need to be narrower than their U.S. counterparts to match certified system configurations.

Best Practices for Practitioners and IP Owners

- Align patent filings with regulatory timelines; file before submitting documentation.
- Use mixed protection strategies: patent where disclosure is unavoidable; maintain trade secrets only where confidentiality can be preserved.
- Develop claim sets tailored separately for EU-regulated configurations and broader U.S. variants.
- Conduct portfolio audits to identify assets intersecting with EU AI Act compliance obligations.
- Integrate patent counsel into AI governance, compliance planning, and product design reviews.

AI Regulation Meets Innovation: IP Considerations as the EU AI Act Takes Effect *continued*

Significant Developments

On **November 19, 2025**, the European Commission proposed a *Digital Omnibus* package to streamline several tech regulations. Notably, it recommended **delaying certain high-risk AI Act obligations from August 2026 to December 2027**, including requirements affecting biometric identification, road traffic applications, utilities management, job application and exam systems, health services, creditworthiness assessments, and law enforcement uses. The proposal also simplifies cookie consent rules. While still subject to debate and approval, the move reflects pressure from industry and international stakeholders—even as the Commission maintains that the AI Act's overall framework remains robust.

Read the original article that appeared in *Bloomberg Law*, EU AI Act Demands Informed, Disclosure-Aware Patent Strategies.

AI Regulation in 2025 (Around the World): What IP Lawyers and In-House Counsel Need to Know

BY: LESTIN KENTON, JR.

By the end of 2025, artificial intelligence (AI) regulation moved decisively from abstract policy debate into enforceable, jurisdiction-specific compliance regimes. The United States, the European Union, China, and emerging markets such as Brazil have pursued materially different approaches, but the direction is unmistakable. AI governance is no longer speculative or optional. It is operational, enforceable, and increasingly central to intellectual property strategy, licensing, and product deployment.

For IP practitioners, whether in private practice or in-house, 2025 marked the year when AI regulation became deeply intertwined with copyright provenance, data governance, trade secret protection, contractual risk allocation, and enterprise diligence.

A Common Thread Across Jurisdictions

Although regulatory models diverged in 2025, several themes converged across jurisdictions.

First, regulators increasingly demanded documentation. Model descriptions, intended-use statements, risk assessments, training data summaries, and governance processes are now expected compliance artifacts, not aspirational best practices. Second, transparency obligations expanded, particularly with respect to training data, synthetic content, and disclosures to downstream users. Third, accountability structures emerged, requiring organizations to designate responsible actors and establish internal oversight mechanisms. Finally, regulators increasingly relied on

AI Regulation in 2025 (Around the World): What IP Lawyers and In-House Counsel Need to Know *continued*

procurement eligibility and deployment conditions as the primary means of enforcing AI governance requirements.[1-2,8-10]

These developments matter to IP lawyers because they shape diligence expectations, commercial contracting, and future litigation narratives. What a company documents today may determine what it can defend tomorrow.

European Union: From Legislation to Compliance Infrastructure

The European Union spent 2025 converting the AI Act from a legislative framework into a functioning compliance regime. Regulation (EU) 2024/1689 entered into force in 2024, but 2025 was the year in which operational obligations began to crystallize.[2]

A central focus was general purpose AI. On July 10, 2025, the European Commission published the *General-Purpose AI Code of Practice*. Although formally voluntary, the Code was expressly designed to provide a structured pathway for demonstrating compliance with the AI Act's obligations applicable to providers of general-purpose AI models, which began to apply on August 2, 2025. The Code addresses transparency, copyright safeguards, safety and security risk mitigation, and internal governance controls.[1]

The publication of the Code followed months of debate regarding feasibility, cost, and timing. Several major technology providers publicly declined to sign on. That resistance, however, does not diminish the Code's practical significance. Regulators, enterprise customers, and commercial counterparties are already using it as a benchmark for reasonable conduct under the AI Act.[2]

For IP practitioners, the implications are immediate. Copyright and training data provenance are no longer peripheral litigation concerns; they are components of regulatory compliance. Licensing negotiations increasingly include representations tied to AI Act readiness, and enterprise customers are beginning to request documentation aligned with the Code's structure. In-house counsel should expect AI governance materials to be requested not only by regulators but also by auditors and commercial partners. Private practitioners should anticipate that AI compliance artifacts will become standard diligence materials in transactions and disputes involving AI-enabled products.

China: Operational Control Through Labeling and Standards

China's approach to AI regulation in 2025 continued to emphasize layered, operational control rather than a single omnibus statute. The regulatory framework combines content labeling requirements, security and filing obligations, and an expanding body of national and industry standards.[3-6]

One visible and continuing obligation concerns labeling and transparency for AI-generated and synthetic content. Requirements under China's generative AI service regulations mandate disclosure and labeling of AI-generated content distributed to the public, reflecting an emphasis on accountability and information control.[3]

At the same time, China has accelerated development of a comprehensive AI standardization system. According to official guidance and state-affiliated reporting, Chinese authorities plan to formulate more than 50 AI-related standards by 2026, spanning technical requirements, governance processes, data annotation practices, cybersecurity, and industrial applications.[4]

AI Regulation in 2025 (Around the World): What IP Lawyers and In-House Counsel Need to Know *continued*

In 2025 alone, multiple national standards addressing generative AI security and data annotation safety were issued, signaling an effort to codify expectations for model development and deployment at a granular level.[5]

These efforts are reflected in the *Guidelines for the Construction of a Comprehensive Standardization System for the National Artificial Intelligence Industry (2024 Edition)*, translated and analyzed in 2025. The document frames AI standardization as a strategic instrument for industrial development, national security, and international influence, emphasizing systematic planning across the AI technology stack and application domains.[6]

China has also paired domestic standardization with international engagement. In 2025, Chinese leadership publicly promoted global cooperation on AI governance and technical standards, culminating in the release of the *Global AI Governance Action Plan at the World Artificial Intelligence Conference*.[7] This initiative underscores China's ambition to shape not only domestic practices but also international norms governing AI safety, documentation, and interoperability.

For IP practitioners, China's regulatory landscape signals both compliance obligations and structural expectations. The growing body of national standards establishes a de facto baseline for technical and process controls that must be reflected in diligence, licensing, and cross-border agreements. The strategic emphasis on standardization further suggests that documentation, traceability, and governance practices will play an increasingly prominent role in China-related AI transactions.

United States: Governance by Directive and Federal–State Tension

Unlike the EU and China, the United States did not enact a comprehensive AI statute in 2025. Instead, governance evolved through executive-branch directives, agency guidance, and an increasingly contested state-level regulatory environment.[8-11]

Federal oversight continued to develop through executive and administrative action. Office of Management and Budget guidance and agency-specific policies embedded AI risk management, transparency, and oversight expectations into federal procurement and deployment practices. While these instruments do not regulate private actors directly, they shape how the federal government purchases and deploys AI systems, effectively exporting governance requirements into the private sector.[8]

In July 2025, the White House released *America's AI Action Plan*, emphasizing innovation, competitiveness, and a streamlined national approach to AI policy. The Plan explicitly warned against fragmented regulatory regimes and signaled a preference for federal uniformity over state-by-state AI regulation.[9] Later in 2025, executive actions reinforced this position, criticizing state-level AI laws as potential obstacles to national AI policy and setting the stage for future preemption disputes.[10]

At the same time, several states continued to advance AI-specific legislation. For example, California enacted transparency and risk documentation requirements for certain frontier AI models, imposing disclosure and recordkeeping obligations with clear implications for IP governance and compliance.[11]

AI Regulation in 2025 (Around the World): What IP Lawyers and In-House Counsel Need to Know *continued*

For IP practitioners, the significance lies less in any single mandate and more in how governance expectations propagate. Federal procurement standards often become templates for private-sector contracting, and state AI laws increasingly require documentation that intersects with training data analysis, model design, and internal governance. Counsel must navigate a fragmented but intensifying compliance environment.

Brazil: A Risk-Based Framework Takes Shape

Brazil's primary AI legislative effort, Bill No. 2338/2023, advanced meaningfully in 2025. After Senate approval in late 2024, the bill moved to the Chamber of Deputies, where a special commission was formed to evaluate, revise, and refine the proposal. Throughout 2025, the commission held hearings and received stakeholder input, marking a shift from high-level principles toward implementation details.[12]

In parallel, Brazil's federal government announced broader plans for a national AI governance system, signaling that institutional oversight and coordinated enforcement will play a central role regardless of the bill's final form.[13-15]

For companies operating in or expanding into Brazil, the key takeaway is that AI compliance expectations are coming into focus. Risk classification, documentation obligations, and accountability mechanisms are likely to resemble European concepts, even if enforcement structures differ.

What IP Lawyers Should Do Now

The regulatory developments of 2025 suggest several practical steps for IP practitioners.

First, AI compliance documentation should be treated as an IP-adjacent asset. Model documentation, risk assessments, training data inventories, and governance processes should be curated with the same care as patent portfolios or trade secret registers.

Second, training data governance and copyright analysis should be elevated from litigation defense to proactive compliance. Regulatory scrutiny is pushing these issues upstream.

Third, contract drafting must adapt to jurisdictional divergence. Modular AI compliance provisions are increasingly preferable to generic "compliance with law" clauses.

Finally, regulation should be viewed as a signal for innovation strategy. Auditability, traceability, provenance tracking, and safety tooling are emerging as concrete technical problems with commercial demand and defensible patent opportunities.

Conclusion

By the end of 2025, AI regulation had become executable. The rules differ across jurisdictions, but the expectations are clear. Documentation, transparency, accountability, and risk management are no longer optional. For IP practitioners, the challenge and opportunity lie in integrating these requirements into licensing, litigation strategy, and innovation planning before they become points of failure.

AI Regulation in 2025 (Around the World): What IP Lawyers and In-House Counsel Need to Know *continued*

European Union

1. **European Commission, *General-Purpose AI Code of Practice* (July 10, 2025)**
   https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai
2. **European Commission, *Artificial Intelligence Act – Implementation and Guidance Materials***
   (Main AI Act portal, including guidance, templates, timelines, and GPAI materials)
   https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

China

3. **Cyberspace Administration of China, *Interim Measures for the Management of Generative AI Services***
   (Official CAC English translation page)
   https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm
4. **China Daily, *Directive on AI standards points the way to 'intelligent' tomorrow,* (July 4, 2024)**
   https://global.chinadaily.com.cn/a/202407/04/WS6685da38a31095c51c50c36e.html
5. **National Information Security Standardization Technical Committee (China), *Generative AI Security and Data Annotation Standards (2024–2025)***
   (Official TC260 standards portal; includes GB/T and draft standards)
   https://www.tc260.org.cn
   Supplementary overview (English):
   https://cset.georgetown.edu/publication/china-gen-ai-data-labeling-safety-standard-draft/
6. **Georgetown Center for Security and Emerging Technology (CSET), *China's AI Standardization Guidelines* (2024 Edition) (2025 analysis & translation)**
   https://cset.georgetown.edu/publication/china-ai-standardization-guidelines-2024/
7. **World Artificial Intelligence Conference, *Global AI Governance Action Plan* (July 2025)**
   (Official WAIC / Chinese government release)
   https://aiii.global/waic-2025/#
   English summary via ANSI:
   https://www.ansi.org/standards-news/all-news/8-1-25-china-announces-action-plan-for-global-ai-governance

United States

8. **Office of Management and Budget, AI Governance and Acquisition Memoranda (2025)**
   (OMB AI policy landing page; includes M-series memoranda)
   https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf
   https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf
9. **White House, *America's AI Action Plan* (July 2025)**
   https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf
10. **Executive Office of the President, *Executive Actions on National AI Policy* (December 2025)**
    (White House Presidential Actions – AI tag)
    https://www.whitehouse.gov/presidential-actions/?s=artificial+intelligence
11. **California Senate Bill 53, *Transparency in Frontier Artificial Intelligence Act***
    (Official California Legislative Information page)
    https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260SB53

Brazil

12. **Brazilian National Congress, *Bill No. 2338/2023* – Legislative History and Proceedings**
    (Official Chamber of Deputies page)
    https://www.camara.leg.br/proposicoesWeb/fichadetramitacao?idProposicao=2487262
13. **Senate History:**
    https://www25.senado.leg.br/web/atividade/materias/-/materia/157233
14. **Government of Brazil, *National Artificial Intelligence Governance Announcements* (2025)**
    (Ministry of Science, Technology and Innovation – AI policy portal)
    https://www.gov.br/mcti/pt-br/acompanhe-o-mcti/transformacaodigital/inteligencia-artificial
15. **Brazilian AI Strategy (ENIA):**
    https://www.gov.br/mcti/pt-br/acompanhe-o-mcti/transformacaodigital/estrategia-brasileira-de-inteligencia-artificial

Key Takeaways from the Federal Circuit's First *Alice* Analysis for Machine Learning Patents

BY: SASHA RAO AND TODD HOPFINGER

The Federal Circuit's decision in *Recentive Analytics, Inc. v. Fox Corp., et al.*, 2023-2437, addressed a question of first impression: whether claims that do no more than apply established methods of machine learning to a new data environment are patent eligible under Section 101. This case involved machine learning used to optimize scheduling and geographical placement of television shows.

Applying the two-step *Alice* framework familiar in software patent cases, the Federal Circuit affirmed the lower court's holding of subject matter ineligibility "because the patents are directed to the abstract idea of using a generic machine learning technique in a particular environment, with no inventive concept."

For *Alice* Step 1, the court found that "[b]oth sets of patents rely on the use of generic machine learning technology in carrying out the claimed methods for generating event schedules and network maps," and the claimed machine learning technology was "conventional." The court rejected the argument that the claimed methods' application of machine learning to a new field of use conferred eligibility. The greater speed and efficiency gained through machine learning could not confer patent eligibility. And the court found the claims did not delineate steps that improve the machine learning technology.

For *Alice* Step 2, the court found nothing in the claims, individually or in their ordered combination, that would transform the patents into something significantly more than "the abstract idea of generating event schedules and network maps through the application of machine learning."

Takeaways

- Patents that do no more than claim the application of generic machine learning to new data environments, without disclosing improvements to the machine learning models to be applied, claim patent ineligible subject matter under Section 101.
- The Federal Circuit's *Alice* analysis for machine language models differs little from its applications in more traditional software contexts. The court's prior holdings that "the application of existing technology to a novel database does not create patent eligibility" are apropos to the machine learning context. Practitioners may want to consider these prior decisions in composing their Section 101 arguments.

Read the original alert, Federal Circuit's First *Alice*-Analysis for Machine Learning Patents.

USPTO Issues New Inventorship Guidance for AI-Assisted Inventions

BY: LESTIN KENTON, JR., AND TODD HOPFINGER

What happened. On November 26, 2025, the U.S. Patent and Trademark Office (USPTO) issued new inventorship guidance (formally published on November 28) clarifying how inventorship should be evaluated when AI tools are used during the development of an invention. The Office withdrew its February 2024 guidance, which had suggested applying the *Pannu* joint inventorship factors to AI-assisted inventions and clarified that those factors apply only to situations involving multiple human contributors. Inventorship continues to require a natural person who formed a definite and complete idea of the claimed invention.

Takeaway. The USPTO is not creating any special rules for AI-assisted inventions. The Office is relying entirely on existing inventorship law, including the long-established conception standard. AI systems are treated the same way as other research tools, and only natural persons may be named as inventors.

Key Points to Know

AI tools are not inventors under U.S. law.

The USPTO reaffirmed that AI systems cannot be listed as inventors or joint inventors. Inventorship requires that a natural person form a definite and complete idea of the claimed invention, consistent with long-standing Federal Circuit precedent.

The inventorship standard remains conception.

The Office emphasized that the same legal standard applies to all inventions regardless of how they were created. The Office will continue to apply the traditional conception analysis. No separate test or exception has been adopted for AI-assisted work.

Joint inventorship applies only to human contributors.

If several individuals contribute to the conception of an invention, joint inventorship may be appropriate and the traditional *Pannu* factors continue to apply in that context. The USPTO made clear, however, that the *Pannu* framework does not apply when an invention is developed by a single human using AI assistance, because AI systems cannot be joint inventors. In other words, the presence of AI tools does not trigger a joint inventorship analysis, and the familiar joint inventorship principles remain limited to human contributors.

Priority claims require alignment with U.S. inventorship rules.

A U.S. application cannot validly claim priority to a non-U.S. filing that names an AI system as the sole inventor. If a non-U.S. application lists both a human and an AI system, only the human inventor may be listed on the U.S. filing. Applicants should ensure consistency across jurisdictions.

Practical Implications

The central theme of the new guidance is that AI-assisted inventions are governed by the same rules that have always applied. The USPTO is not creating new standards, modifying the inventorship test, or imposing additional evidentiary burdens. Instead, the Office is reiterating that inventorship turns solely on human conception of the claimed subject matter.

Although the legal framework is unchanged, AI-assisted research and development (R&D) can raise practical questions about who conceived particular elements of an invention. Companies should be prepared to identify the human inventor who formed the complete and oper-

USPTO Issues New Inventorship Guidance for AI-Assisted Inventions *continued*

ative idea reflected in the claims. AI output, even if sophisticated, does not replace the need for human conception.

This clarification is also important for global filing strategies. Naming AI systems as inventors in non-U.S. patent applications can create inconsistencies that may jeopardize priority rights in the United States. Applicants should review their non-U.S. filing practices with an eye toward how different jurisdictions treat AI-assisted inventions. Although no major patent office currently permits non-human inventorship, several jurisdictions have examined or consulted on the issue. Companies should understand these differences, account for local procedural requirements, and ensure that non-U.S. filings remain compatible with U.S. inventorship rules to preserve priority rights.

Recommended Steps

To align with the USPTO's guidance, companies should consider the following steps:

1. Review current and pending filings to ensure that inventorship determinations accurately reflect human conception in AI-assisted projects.
2. Implement forward-looking documentation practices that clearly record the inventive contributions of human inventors in AI-assisted R&D.
3. Educate R&D teams on how AI tools should be used and how their output should be incorporated into invention disclosures.
4. Coordinate global filing strategies by reviewing inventorship rules and filing requirements in each jurisdiction and ensuring that non-U.S. patent applications remain compatible with U.S. inventorship law.

Conclusion

The USPTO's new guidance confirms that AI does not alter the fundamental standards of inventorship. Human conception remains the core requirement. For companies integrating AI into research and development, this is an appropriate time to review inventorship practices and maintain clear documentation of the human contributions that support the claimed subject matter.

Navigating the Crossroads of AI Patent Eligibility: Translating USPTO Optimism into Long-Lasting Protection

BY: BEN GITZINGER AND LESTIN KENTON, JR.

Artificial intelligence (AI) has become an engine of growth in the modern economy. Innovations in AI-based technologies are reshaping how industries operate by integrating speed and efficiency into fields such as healthcare, financial services, and content creation. As these technologies increasingly define market competition, protecting intellectual property (IP) is becoming an ever more urgent priority for AI companies. Yet, securing robust patent protection for AI inventions has proven to be a persistent area of difficulty given the nature of the technology and the industry.

Trade secret protection, which has traditionally served as the fallback method for guarding proprietary algorithms, source code, and training data, requires the implementation of ongoing confidentiality precautions. However, these measures have become increasingly difficult to maintain for AI companies given the collaborative work patterns and high employee turnover that are typical of the tech industry. Compounding the problem, advanced AI systems themselves can sometimes derive or "reverse engineer" proprietary methods by analyzing model outputs, raising concerns that core system behavior may be partially inferable despite conventional barriers.

Patents often serve as the typical alternative to trade secrets for protecting scientific inventions, but they also face practical problems in tech. AI and machine learning (ML) innovations frequently encounter rejections by examiners at the United States Patent and Trademark Office (USPTO) and federal judges under 35 U.S.C. § 101. This is primarily due to AI and ML's close relationship with algorithms and computational processes, which courts often treat as unpatentable abstract ideas under § 101.

However, recent actions by the USPTO's new leadership suggest a potential broadening of what qualifies as patent-eligible subject matter at the agency. The USPTO has indicated a more permissive stance toward AI-related inventions under the management of Director John Squires. This optimism regarding the patentability prospects for AI applicants is expressed most clearly in the Director's recent *sua sponte* intervention in *Ex Parte Desjardins*, where he touted the eligibility of a patent application that claimed a method of training an ML model.

Still, optimism alone cannot substitute for enforceability. The value of a patent depends not only on whether it can be obtained but also on whether it will withstand validity challenges years down the line. Because federal courts have the final word on patent validity, AI stakeholders must approach drafting and prosecution with caution, even if the USPTO appears more welcoming on eligibility. This article examines (1) the USPTO's shifting posture on AI patentability, (2) the continuing obstacles that remain for AI patents, and (3) strategies for drafting AI applications that can endure both USPTO and Federal Circuit scrutiny.

New Directions at the USPTO

Upon taking office, USPTO Director John Squires made clear that patent eligibility reform would be a priority. His early actions signal a desire to broaden protection for emerging technologies at the agency.

During his first week in the position, Director Squires ceremonially signed two patents in technology areas traditionally thought to be unpatentable (diagnostics and cryptocurrency) in an effort to underscore a policy

Navigating the Crossroads of AI Patent Eligibility: Translating USPTO Optimism into Long-Lasting Protection

continued

interest in expanding eligibility.[1] He also provided testimony before the Senate in support of the Patent Eligibility Restoration Act (PERA), where he framed the expansion of patent eligibility as a matter of national security, global leadership, and economic competitiveness.[2]

The Director's most notable effort to elucidate his view on AI patentability was his decision in *Ex Parte Desjardins.*[3] Exercising the Appeals Review Panel in only its second use to date, the Director authored a panel opinion overturning a decision by the Patent Trial and Appeal Board (PTAB) that had found an ML-related patent ineligible.

The application was directed to an ML model designed to address "catastrophic forgetting," a degradation phenomenon in continual learning systems where models lose performance on previously learned tasks when trained on new data. After finding the claims obvious under 35 U.S.C. § 103, the PTAB issued a new ground of rejection under § 101, concluding the claims were directed to abstract mathematical processes lacking a "practical application."[4]

Director Squires disagreed, finding that adjusting model parameters to "optimize performance… while protecting performance on the first task" reflected a technical improvement to the operation of the ML model itself.[5] Because he found the claims integrated the abstract idea into a practical application and provided a concrete technological solution, he reversed the § 101 rejection.

In doing so, the Director cautioned against "categorically excluding AI innovations from patent protection," warning that such an approach "jeopardizes America's leadership in this critical emerging technology."[6] These remarks signal a USPTO policy preference for treating § 101 as a coarse filter rather than a strict gatekeeping doctrine.

For AI innovators, this suggests a warmer reception at the examination stage. Yet even a favorable agency posture does not eliminate the broader legal headwinds facing AI patents.

Remaining Obstacles for AI Patents

AI patent applicants will continue to face significant hurdles despite a more receptive USPTO. Even while § 101 rejections may decrease in some art units, applicants will continue to face challenges under §§ 102, 103, and 112. Indeed, in *Desjardins*, although Director Squires reversed the § 101 rejection, he left the § 103 rejection in place. If the agency adopts the Director's view that the remaining statutory bars to patentability are "the proper statutory tools for limiting the scope of patents,"[7] then the overall number of §§ 102, 103, and 112 rejections may actually increase, becoming the means by which examiners may invalidate claims that would traditionally have been viewed as patent ineligible under § 101.

More important, regardless of the Director's views on eligibility, practitioners might not experience much change in practice if the Federal Circuit does not share the agency's perspective. Following the Supreme Court's decision in *Loper Bright Enterprises v. Raimondo*, courts are no longer obligated to defer to agency

1 Dani Kass, *Squires Jumps Right Into Patent Eligibility Reform*, Law360 (Oct. 1, 2025)

2 The Patent Eligibility Restoration Act–Restoring Clarity, Certainty, and Predictability to the U.S. Patent System: Hearing on S. 1546 Before the S. Subcomm. on Intell. Prop. of the S. Comm. on the Judiciary, 119th Cong. (2025) (statement of John A. Squires, Director of the United States Patent and Trademark Office).

3 *Ex Parte Desjardins*, 16/319,040, 2025 WL 3095778 (P.T.A.B. Sept. 26, 2025).

4 *Id.* at 4.

5 *Id.* at 9.

6 *Id.*

7 The Patent Eligibility Restoration Act–Restoring Clarity, Certainty, and Predictability to the U.S. Patent System: Hearing on S. 1546 Before the S. Subcomm. on Intell. Prop. of the S. Comm. on the Judiciary, 119th Cong. (2025) (statement of John A. Squires, Director of the United States Patent and Trademark Office).

Navigating the Crossroads of AI Patent Eligibility: Translating USPTO Optimism into Long-Lasting Protection

continued

interpretations of law. As a result, the Federal Circuit will continue applying its own framework for eligibility, irrespective of USPTO policy preferences.

In April 2025, the Federal Circuit decided *Recentive Analytics, Inc. v. Fox Corp.*, the court's first § 101 decision focused squarely on machine-learning claims.[8] The court affirmed the ineligibility of two patents directed to ML-based TV scheduling. After characterizing the claims as directed to mathematical algorithms, the court held that requiring the model to be "iteratively trained or dynamically adjusted" did not convert the abstract idea into a patent-eligible technical improvement.[9] In the court's view, this behavior was merely "incident to the very nature of machine learning."[10] The court further clarified that computer-implemented speed improvements or efficiency gains must stem from enhancements to computational techniques or system operation, not simply from automating a process that could be performed mentally.

Although *Recentive Analytics* may temper enthusiasm generated by *Desjardins*, the decision also clarified that not all ML models are ineligible. Judge Dyk emphasized that the ruling addressed only claims that "apply established methods of machine learning to a new data environment," signaling that claims describing concrete improvements to model architecture, system performance, or computational behavior remain viable under the Federal Circuit's application of § 101.[11]

This distinction creates a meaningful, but navigable, tension. While the USPTO may be willing to recognize the practical applications of AI models as sufficient for eligibility, the Federal Circuit requires that those applications reflect concrete technical improvements. Practitioners must therefore draft with both audiences in mind.

Drafting Strategies for AI Patent Applications

The juxtaposition between *Desjardins* and *Recentive Analytics* provides guidance on how applicants can strengthen their AI claims. The following drafting principles can help bridge the gap between a more permissive USPTO and a more demanding Federal Circuit.

1. Start with a Technical Narrative

A strong AI patent begins with a clear technical story. From the outset, applicants should articulate the specific technical problem addressed by the invention, whether tied to data structures, training behavior, inference latency, model stability, resource utilization, or deployment constraints. In finding the ML model from *Desjardins* to be patent eligible, Director Squires relied heavily on the application's description of reduced storage usage and decreased system complexity, treating these as indicators of a genuine technological improvement.

The specification should contextualize the improvement and avoid describing the invention solely at a high level. Clear articulation of how the model modifies gradient updates, alters feature-embedding behavior, reduces computational cost, or improves memory access patterns can strengthen § 101 positions and reinforce inventive concepts under §§ 102 and 103.

8 *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025).

9 *Id.* at 1212.

10 *Id.*

11 *Id.* at 1211.

Navigating the Crossroads of AI Patent Eligibility: Translating USPTO Optimism into Long-Lasting Protection

continued

2. Embed Technical Detail and Avoid Result-Oriented Claiming

The claims in *Recentive* failed largely because they implemented generic machine learning processes in a new environment without altering the underlying technology. Applicants should instead make AI central to the claimed invention by describing:

- Model architectures (e.g., encoder-decoder path, attention mechanisms, gating functions);
- Training pipelines and parameter update techniques;
- Feature-engineering strategies;
- Deployment-specific constraints (e.g., on-device inference, memory bounds, latency targets); and
- Real-world operational challenges (e.g., model drift, catastrophic forgetting, noisy labels).

Grounding claims in the "how," not merely the "what," helps differentiate AI inventions from abstract algorithmic processing.

3. Address § 112 Head-on While Balancing Trade Secret Sensitivities

AI inventions commonly face written-description and enablement challenges, particularly after *Amgen, Inc. v. Sanofi*.[12] Applicants must provide enough technical detail to show possession of the invention and to enable others to make and use it, but they must also be mindful that disclosures become public and may reveal aspects of the system that a company would prefer to maintain as a trade secret.

A balanced approach requires articulating the invention's core technical mechanisms without exposing proprietary implementation details that provide competitive advantage. Applicants should consider describing, at an appropriate level of abstraction:

- Inputs, intermediate representations, and outputs;
- Training datasets and preprocessing pipelines (e.g., categories or characteristics, not necessarily raw data);
- Iterative training behaviors and update logic;
- Inference workflows and data flow sequencing; and
- System-level interactions and deployment constraints.

Providing this detail not only satisfies § 112 but also helps demonstrate technical improvements under § 101, while thoughtful scoping of the disclosure helps preserve sensitive information that may be better protected through trade secrets.

4. Layer and Diversify Claims

Applicants should draft claims across multiple statutory classes, including method, system, non-transitory computer-readable medium, and component-level claims. Dependent claims can capture:

- Specific model components;
- Data preparation steps;
- Training environment features;
- Memory management techniques;
- Optimization strategies; and
- Deployment configurations.

Layered claiming increases the likelihood that at least one claim survives eligibility and prior art challenges.

***

The USPTO's recent actions under Director Squires point toward a more flexible interpretation of § 101 for

12 *Amgen Inc. v. Sanofi*, 598 U.S. 594 (2023).

Navigating the Crossroads of AI Patent Eligibility: Translating USPTO Optimism into Long-Lasting Protection

continued

AI-related inventions. But the Federal Circuit's decision in *Recentive* underscores that judicial scrutiny remains a critical barrier. AI inventors should be encouraged by the USPTO's signals, but they must draft with the understanding that a patent examined today may be litigated years later before a court that applies § 101 more narrowly.

Long-term protection requires careful alignment of specification detail, technical storytelling, and layered claiming. Applicants who articulate concrete technological improvements, rather than functional or result-driven descriptions, will be best positioned to secure durable protection for their most important AI assets.

Patent Protection for AI Technologies: Navigating § 101 After *Recentive Analytics*

BY: SASHA RAO AND TODD HOPFINGER

As companies lean into AI innovation to revolutionize a wide range of industries, they will encounter opportunities to secure patents on their advances. This article provides a practical roadmap to securing valid, enforceable patents and avoiding prosecution pitfalls in the evolving AI patent landscape.

The most authoritative guidance from the Federal Circuit on AI patentability under § 101 comes from *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025). The court affirmed that AI-based patents are subject to the two-step framework from *Alice Corp. v. CLS Bank International* (2014). The USPTO's published guidance also underscores that many AI claims will implicate abstract ideas that may require deeper scrutiny under Step 2 of the *Alice/Mayo* framework. The guidance instructs examiners to look for features that confine or transform the recited abstract idea into a particular technological environment.

This direction given by the Federal Circuit and the USPTO can guide practitioners in drafting strong claims for their AI inventions.

Build a Compelling Technical Narrative

The patent should have a narrative, beginning with the specification and flowing through the claims, that ties together the technical problem, the inventive mechanism, and the measurable improvement.

- **Define a Concrete Technical Problem.** A strong patent application identifies a specific technical limitation in the prior art.
- **Describe the Technological Solution in Detail.** The specification should disclose how the invention changes the computational landscape through architecture, process, or resource optimization techniques.

Patent Protection for AI Technologies: Navigating § 101 After *Recentive Analytics* *continued*

Several recurring errors continue to derail AI patent applications:

- **Avoid Concessions of Conventionality** – Counsel should probe for the non-obvious components (e.g., architectural changes to the AI model, the feature extraction method, or algorithmic optimization technique) that differentiate the invention from a generic application of AI.
- **Overly Abstract Functional Language** – Claims written at a high functional level without linking operations to technical context invite § 101 rejections.
- **Outcome-Focused Drafting and Overreliance on Performance Metrics** – A related error is claiming an invention by its end result rather than its technological mechanism.
- **Disconnected Claims and Specification** – A lack of narrative continuity between the specification and claims weakens the patent's technical foundation outline.

Takeaways

- Courts and the USPTO demand specificity, substance, and a demonstrable link between algorithmic techniques and technological advancement.
- For legal practitioners, success depends on technical storytelling as much as legal drafting.
- Practitioners should consider trade secrets, copyright, and contracts/NDAs to protect other aspects of AI technology.

Read the full article, Turning AI Innovation into Patent Protection: Key Considerations, which first appeared in *Westlaw Today*.

Navigating Patent and IP Risks in AI-Powered Blockchain and Crypto Technologies

BY: ROOZBEH GORGIN

As blockchain and cryptocurrency technologies continue to evolve—from decentralized finance (DeFi) and smart contract platforms to AI-powered trading and compliance tools—the intellectual property landscape surrounding these innovations has grown significantly more complex. Companies navigating this rapidly shifting ecosystem must contend not only with fast-moving technical developments but also with heightened scrutiny from the USPTO and courts, particularly in fintech and AI-driven applications.

Some core strategies for protecting blockchain and AI-enabled crypto innovations while mitigating risks posed by prior art, open-source environments, and potential validity challenges are provided:

Prioritize Technical Implementation Over High-Level Concepts

Patent claims in the crypto domain face substantial §101 scrutiny. Successful applications highlight concrete technical improvements—such as novel consensus mechanisms, secure key management architectures, or system-level enhancements to network performance—rather than abstract business logic.

Capture Innovations Early and Use Strategic Claim Layering

Rapid iteration is common in blockchain systems. Identifying patentable features during product development and building modular portfolios (system, method, and component-level claims) ensures coverage of both core infrastructure and application-specific use cases.

Patenting AI-Blockchain Integrations Requires Added Specificity

For AI-powered crypto tools—including trading models, fraud detection, and automated compliance—patent protection depends on detailing model architectures, training techniques, and how AI interfaces with deterministic blockchain systems. Hybrid designs combining probabilistic predictions with verifiable on-chain processes can yield strong patentable subject matter.

Draft with Post-Grant Challenges in Mind

Crypto patents face no special presumption in post-grant proceedings. Robust specifications that explain system architecture, algorithms, and data flows create the strongest foundation for defending validity.

Best Practices for Practitioners and IP Owners

- Emphasize tangible technological improvements when framing claims, especially for fintech-related inventions.
- Use continuation applications to maintain strategic flexibility as blockchain and AI markets shift.
- Develop IP-aware open-source governance, particularly for teams contributing to public repositories.
- Balance global filings with realistic enforcement prospects in key jurisdictions.
- Proactively address prior art, particularly common open-source frameworks or academic disclosures.
- Align patent strategy with business objectives, whether protocol adoption, enterprise deployment, or long-term licensing.

Read the original Q&A with Roozbeh Gorgin that appeared in *DataBird Business Journal*, Navigating the Patent Landscape: Blockchain, AI, and Cryptocurrency IP Strategy.

Key Takeaways—AI-Equipped Medical Devices Require Speculative Patent Approach

BY: DOHM CHANKONG AND RICHARD COLLER, III

Traditional hardware-driven patent strategies long used by the medical device industry should be updated for the AI era. Because AI and software evolve much faster than physical hardware or regulatory (FDA) timelines, companies should adopt a patent approach that is more speculative—one that looks past current product versions to anticipate future software capabilities and technical milestones.

Takeaways

- **Move Beyond Hardware**: Historically, medical device patents focused on physical designs and circuitry. In the AI era, value has shifted to the "code-based" innovations, such as how a device makes decisions, adapts to patient data, or transforms clinical workflows.
- **Adapt to AI Roadmaps**: Companies should use their internal product roadmaps to identify and "extrapolate" future features early in the lifecycle. This allows them to file initial applications that contain a range of anticipated technical variations before they are even built.
- **Adopt a Nimble Mindset for Dynamic Timelines**: AI features are often only a few coding sessions away from reality and move much faster than the slow iterations of regulated hardware. Patents must be filed early to capture these rapid software milestones before competitors do.
- **Focus on Technical Novelty**: Teams should prioritize features that improve clinical performance in a technically novel way that would be difficult for a competitor to replicate or work around and ensure that those features are properly claimed.
  - *Weak*: Claiming functional features that merely use AI, such as to "detect arrhythmias faster."
  - *Stronger*: Claiming a specific implementation or architecture, such as a trained neural network, that performs the functions, such as classifying waveforms by dynamically adjusting feature extraction parameters based on historical patient-specific data.

Read the full article, AI-Equipped Medical Devices Require Speculative Patent Approach, which first appeared in *Bloomberg Law*.

The Market Harm Dilemma: What *Bartz* and *Kadrey* Reveal About Fair Use's Future in Generative AI

BY: SASHA RAO AND IVY ESTOESTA

In June 2025, two Northern District of California decisions—*Bartz et al. v. Anthropic PBC* and *Kadrey et al. v. Meta Platforms, Inc.*—revealed fault lines in the forming fair use terrain for generative artificial intelligence (GenAI) copyright infringement actions. While similar in facts and outcome, the *Bartz* and *Kadrey* courts reached their findings of fair use through different analyses.

In *Bartz*, Anthropic had placed pirated or purchased-and-destructively-scanned copies of books into a central research library and used some of them to train Anthropic's large language model (LLM), Claude. Judge Alsup found the fair use doctrine covered training LLMs and converting purchased print library copies into digital library copies but not building a central library using downloaded pirated copies. The court's fair use analysis relied more heavily on the purpose and character of using copyrighted works to train LLMs to generate new text and less on market harm.

A few days later, in *Kadrey*, Judge Chhabria gave greater consideration to the market effect of GenAI infringement. Meta had trained its LLM, Llama, on both publicly available data and potentially pirated information. The court found the use of plaintiffs' books to train its LLM to be highly transformative and therefore fair. But the court emphasized that, in the fair use analysis, harm to the market for the copyrighted work is more important than the purpose for which the copies are made. While *Kadrey*, like *Bartz*, also granted summary judgment on fair use in favor of the GenAI defendant, the court noted it might have found otherwise but for these plaintiffs making "the wrong arguments and fail[ing] to develop a record in support of the right one," particularly as it relates to market harm.

Update

In August, the parties in *Bartz* settled. The court's holding that Anthropic's collection and retention of a centralized library of pirated books did not qualify as fair use set the stage for a damages trial. But Anthropic and the authors jointly filed for a classwide settlement after mediation. Three compelling factors may have driven both parties to settle:

1. **Damages Exposure:** With Anthropic having pirated at least 7 million books, statutory damages for those works had the potential to run close to a trillion dollars.
2. **Damaging Publicity:** Anthropic may have been eager to avoid fact discovery and the uncomfortable disclosures and publicity that could have resulted.
3. **Damaging Precedent:** The court's summary judgment decision foretold the potential for precedent neither party wanted.

Takeaways

- Transformative use is strongly favored when GenAI systems employ mitigating software to prevent infringing or overly presumptuous borrowing in their outputs.
- Market effect can be a necessary counterweight to transformative use to defeat a fair use defense.
- Market harm may be shown more effectively through market dilution than hypothetical loss of licensing fees, which may not be cognizable in certain courts.
- The use of pirated copies, even if immediately used for the transformative use and thereafter discarded, might not qualify as fair use.

Read our original IP Hot Topics on these issues:

A Clash in California: Judicial Tug of War in GenAI Fair-Use Cases

From Trillion-Dollar Risk to Resolution: Settlement in the Anthropic Authors/Class Action

2025 Summer Associate Wade Marshall contributed to this article.

First Federal Ruling Rejects Fair Use Defense for AI Training Data in Copyright Dispute

BY: SASHA RAO AND IVY ESTOESTA

In February 2025, a Delaware district court issued a much-awaited summary judgment decision in *Thomson Reuters Enterprise Centre GmbH et al. v. Ross Intelligence Inc.*, No. 1:20-cv-613. This was the first of a series of district court decisions last year that began determining the contours of fair use analysis in copyright actions involving copyrighted works used to train artificial intelligence (AI) models.

As background, defendant Ross Intelligence developed AI-based legal research software using Westlaw's copyrighted legal materials for training data. The court found Thomson Reuter's copyright in its headnotes and Key Number system to be valid and granted partial summary judgment to Thomson Reuters on direct copyright infringement.

In its analysis, the court found that the fair use doctrine did not apply. For the first fair use factor (the purpose/character of the use), the court found Ross's use was not transformative because it did not have a "further purpose or different character" from that of Thomson Reuters' copyrighted work. Instead, Ross aimed to develop a competing legal research tool. The court also found the fourth fair use factor (the effect on the copyrighted work's market) favored Thomson Reuters because Ross "meant to compete with Westlaw by developing a market substitute." The court dismissed claims of public benefit, as legal opinions are freely available but Westlaw's curated content is not.

Update

The Thomson Reuters decision rapidly became central to generative AI copyright litigation. The decision became highly cited, including by the court in *Bartz v. Anthropic PBC*. And although the court was careful to exclude generative AI from its analysis by distinguishing Ross's AI tool, the court's application of the *Warhol* framework could nevertheless apply to generative AI. Indeed, this case has since been applied in a matter involving an artificial intelligence tool that generates study materials. *Barkley & Assocs., Inc. v. Quizlet, Inc.*, No. 2:24-CV-05964-WLH-E, 2025 WL 2946993, at *6 (C.D. Cal. Sept. 11, 2025)

The ruling strengthened plaintiffs' arguments that AI training isn't transformative use and emphasized market harm, particularly the impact on emerging AI training data licensing markets. However, the court certified this case for interlocutory appeal, noting that "[b]ecause appellate resolution of these questions would change the shape of trial—and possibly avoid a copyright trial altogether—I stay this case pending the Third Circuit's response."

Takeaways

- Transformative use—the go-to defense for AI defendants—may not hold much water when the defendant's product competes with the copyright holder. And intermediate copying may be relegated to specific factual scenarios.
- The fourth fair use factor could be met by a potential market for AI training data. This increases an AI defendant's potential to fail the fourth factor when copyrighted material is part of its training data.

Read the original client alert, Fair-Use Whiplash: *Thomson Reuters v. Ross Intelligence* Revisits Fair Use for Artificial Intelligence.

The (Media) Empire Strikes Back—An Update on *Disney et al. v. Midjourney Inc.*

BY: SASHA RAO AND IVY ESTOESTA

In June 2025, some of the largest movie studios sued the image-GenAI company Midjourney Inc. in the Central District of California, accusing it of large-scale, willful copyright infringement. *Disney et al. v. Midjourney Inc.*, No. 2:25-05275 (C.D. Cal. June 11, 2025).

Plaintiffs allege that Midjourney trains its GenAI model on unauthorized copies of plaintiffs' films, television shows, and character artwork, then monetized that data by selling a tiered subscription service. In use, a subscriber of Midjourney's image service can enter a bare-bones prompt without specifically identifying the copyrighted character and yet the model can still output a copyrighted character. Further, Midjourney publicly uses many of these outputs to entice new customers and upgrade existing ones.

Legal Claims

The suit pleads direct and secondary infringement under the Copyright Act.

Direct infringement. Plaintiffs contend that Midjourney itself makes the infringing reproductions by (1) copying plaintiffs' works into its training set and (2) generating, displaying, and distributing derivative images in response to user prompts, all without license. The infringement is willful because plaintiffs sent Midjourney letters, which it ignored.

Secondary infringement. Plaintiffs allege Midjourney knowingly provides the tools and inducement that make the violations possible. The complaint stresses that Midjourney could have prevented—or at least curbed—copying. And the availability of infringing outputs is a key draw that drives Midjourney's revenue.

Damages. Plaintiffs seek statutory damages of up to $150,000 per work or full actual damages and profits, plus injunctive relief halting further use of their content.

Update

Midjourney has since filed its answer. In addition to denying most of plaintiffs' allegations, Midjourney has raised several affirmative defenses including fair use, unclean hands, *de minimis* copying, and failure to comply with the Digital Millennium Copyright Act (DMCA). The unclean hands defense is based on plaintiffs' own use of and benefit from using Midjourney's AI model and plaintiffs' own training of GenAI models using copyrighted materials.

Takeaways

- **Moderation Backfiring.** Midjourney's efforts to ensure improper content does not get generated and to enforce compliance with its community standards backfired.
- **DMCA.** Midjourney's affirmative defense under the DMCA is a notable use of the DMCA as a defensive tactic rather than the more typical use of the DMCA as an offensive tool by copyright holders.
- **Infringement by the Unwary.** GenAI users creating public-facing content should be aware that GenAI can output copyrighted materials even when users' prompts do not specifically call for it.

Read the original article IP Hot Topic: The (Media) Empire Strikes Back.

Challenges with Protocols for Training Data Discovery in AI Litigation

BY: SASHA RAO AND RICHARD CRUDO

As artificial intelligence (AI) litigation increases in frequency, parties may face discovery requests for their training data and must consider how best to protect it. Litigants have often relied on traditional source code inspection protocols. But the unprecedented scale and complexity of AI data require adaptation of such protocols.

The Value of Training Data

Training data is a highly valuable commercial asset. AI companies invest significant resources and expenses to collect, curate, and structure this data to create their models. Gathering this data can also be challenging given the time required and potential inconsistency or bias in available data. Thus, AI companies often treat the data as a valuable asset whose disclosure could cause competitive harm.

This training data can also be highly valuable in AI-based copyright litigation. Training data could include copyrighted works, and LLMs could generate outputs substantially similar to protected works. Access to this training data thus provides litigants with an opportunity to see firsthand the potential unauthorized use of copyrighted materials.

Training Data Inspection Protocols: Challenges and Opportunities

Where disclosure of training data seems likely, parties have stipulated to inspection protocols that govern a plaintiff's ability to request and review a defendant's data. But the unprecedented scale and complexity of AI data will likely require adaptation of the source code review protocols that litigants are familiar with.

Scale and volume. A major difference between source code and AI training data is scale and volume. The scale of AI training data makes it difficult to simply "hand over" in any human-reviewable form. Defendants can use their training data's size to their advantage by making proportionality arguments. Further, both parties benefit from negotiating shortcuts for reviewing raw training data: defendants avoid opening their entire training datasets to review while plaintiffs can review the data more efficiently.

Heterogeneity and accessibility. Training data can use multiple languages and might require different types of indices and tools for facilitating review of different data types and formats. Given these numerous ways to organize training data, most protocols require technical guidance for efficient review.

Takeaways

Training data's importance, both as a commercial asset and in litigation, demands an approach that balances protection with access during discovery. Achieving this balance requires accounting for the unique size, formats, and sources of the training data. Creative solutions beyond source code review protocols are needed by litigants to ensure efficient review while minimizing overproduction and misuse.

Read the full article, Discovery of Training Data in AI Litigation, which first appeared in *Corporate Counsel*.

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

BY: JASON EISENBERG, RYAN ESTATICO, BROOKE MCLAIN, LAUREN SCHLEH, JEAN SELEP, AND ANANTH JOSYULA

The rapid integration of generative artificial intelligence (GenAI) into the professional practice of the legal industry has brought forth several instances of attorneys, experts, and judges knowingly or unknowingly submitting AI hallucinations in their work. This review focuses on a curated selection of such instances from across the United States that have "made the news," and discusses how the courts are proposing to curtail the submission of AI hallucinations.

Currently, there is no explicit guidance for federal courts on how to address the misuse of GenAI, let alone address the submission of AI hallucinations in court filings. Broadly speaking, courts have applied existing rules and laws to attorney, expert, and judge misuse of AI. However, individual judges have taken their own approaches to punishing attorneys who submit AI hallucinations and have inconsistently applied punishments. This approach indicates a necessity for courts to respond but an uncertainty in how to respond. In other words, courts know they must act but lack consensus on how to appropriately do so in response to the submission of AI hallucinations in court filings.

What Do the Federal Rules of Civil Procedure (FRCP) Say?

Presently, the courts are enforcing violations under FRCP 11(b), which in relevant part states, "[B]y presenting to the court a pleading, written motion, or other paper—whether by signing, filing, submitting, or later advocating it—an attorney or unrepresented party certifies that to the best of the person's knowledge, information, and belief, formed after an inquiry reasonable under the circumstances… (2) the claims, defenses, and other legal contentions are warranted by existing law."

What Did We Find?

Our review uncovered the following themes:

1. Governance Through Existing Rules

The courts have mostly shoehorned the submission of AI hallucinations into other rules or laws, largely stemming from the attorney's professional responsibility. The best example is FRCP 11(b).

2. Honesty Is the Best Policy

Courts seem to respond more harshly to the submission of AI hallucinations when use of GenAI is evident but such use is denied by the attorney. In contrast, when up front and apologetic, attorneys receive lesser punishments.

3. Proportionality of Actions

A brief with a few erroneous citations or arguments will receive reduced punishment compared to many erroneous citations in a single brief or several errant briefs.

4. Most Common Hallucinations

There appear to be three categories of AI hallucinations (although others can occur, they are less frequent). The categories are:

- Citations to fictitious cases.
- Fabricated citations to real cases or documents.
- Citations to real quotes from real cases that fail to support or that directly contradict a proposed legal proposition.

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

continued

Exemplary Sanctions for Filings from Plaintiffs, Defendants, Experts, and Judges/Clerks

Plaintiff's Motions

In *Dehghani v. Castro* (U.S. District Court for the District of New Mexico), the petitioner's attorney purchased a brief from a freelance attorney. The freelance attorney likely used GenAI and then destroyed all notes. The purchasing attorney did not review the purchased brief before submitting it to the court. The court issued two show cause orders, but the petitioner's attorney did not adequately respond to either request to clarify and provide copies of the cases. The court issued a sanctions order requiring a fine to be paid to the court, mandatory CLE training, and self-reporting to all appropriate state bars of this misconduct for the petitioner's attorney and the freelance attorney.

In *In re Loletha Hale* (U.S. District Court for the Northern District of Georgia, Atlanta Division), the court found that "[a]n overwhelming majority of the cases cited by [the plaintiff's attorney] either did not exist, did not support the proposition for which they were cited, or misquoted the authority." The attorney blamed her daughter for drafting the brief, who was neither an attorney nor a paralegal. The court issued a sanctions order requiring her to notify all existing clients of the court's decision and file a notice with the court along with a copy of the order in all pending and future cases in the Northern District of Georgia in which she appears for five years from the date of the order.

In *Dubinin v. Papazian* (U.S. District Court for the Southern District of Florida), the plaintiff's attorney filed a response brief with fabricated quotes and relied on a "decision from this Court that does not exist." The attorney blamed a subordinate who was not a member of the bar in the state, who in turn blamed her legal assistant. The court issued a sanctions order dismissing the case without prejudice, requiring that the attorney pay the other parties' attorney fees and referring all attorneys involved to their respective state bars for discipline.

In *Tercero v. Sacramento Logistics, LLC, et al.* (U.S. District Court for the Eastern District of California), the plaintiff's attorney filed a motion for reconsideration that included "non-existent, patently manufactured, or entirely inapplicable" case citations. Specifically, plaintiff cited to "two cases that simply do not exist," "[t]en other cases... [that] do not contain the language quoted by Plaintiff," and 12 cases that "do not support the propositions for which [plaintiff] proffers them." Instead, the plaintiff's counsel argued to the court that they did not use AI. The court found that the attorney was lying about how she did her work. Despite each of the lies being exposed and three opportunities to come clean being offered, the attorney continued to give excuses, even going so far as telling the opposing counsel they were wrong for noting the incorrect citations. The court issued a sanctions order requiring a $1,500 sanction to the court; mandating service by the attorney of a copy of the order on her client; and directing the clerk of the court to serve a copy of the order on the State Bar of California, of which the attorney was a member.

In *Idehen v. Stoute-Phillip* (Civil Court of the City of New York, Queens County), the petitioner's attorney submitted an affirmation that "contained citation to seven fake cases." In response, the attorney submitted a second filing with an appendix described as "an incoherent document that is eighty-eight pages long, has no structure, contains the full text of most of the

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

continued

cases cited, shows distinct indications that parts of the discussion/analysis of the cited cases were written by artificial intelligence, cites multiple cases that do not have anything to do with the issues presented in Respondent's motion to dismiss, and cites to multiple cases for propositions that are not remotely supported by the case cited." The court found that "[the attorney's] conduct at the hearing and the Appendix to the June 2, 2025 Affirmation escalated his conduct from merely frivolous to egregious misconduct that implicates his honesty, trustworthiness, and fitness to practice law", and that "had [the attorney] simply come to the hearing acknowledging that he had used Copilot to conduct his research, was unaware that the software could produce fake cases, and apologized for his mistake, the Court would have likely determined that a sanction was sufficient to address the frivolous conduct." The attorney is under a show cause order.

Defendant's Motions

In *Johnson v. Dunn, et al.* (U.S. District Court for Northern District of Alabama, Southern Division), the defendant cited AI hallucinations in two motions relating to seeking discovery from an incarcerated person. At least one of the defendant's attorneys used ChatGPT to prepare the motions. The court issued a sanctions order disqualifying defendants' attorneys from the case and referring the matter to the state bar. The judge noted that "[i]f fines and public embarrassment were effective deterrents, there would not be so many cases to cite. And in any event, fines do not account for the extreme dereliction of professional responsibility that fabricating citations reflects, nor for the many harms it causes. In any event, a fine would not rectify the egregious misconduct in this case."

In *Mid Cent. Operating Eng'rs Health and Welfare Fund v. Hoosiervac LLC* (U.S. District Court for the Southern District of Indiana), the defendant's attorney submitted three briefs that contained citations to cases that did not exist. The attorney admitted to using GenAI and that he made no effort to verify the cases cited by the AI. Instead of arguing that he did not violate Rule 11 of the FRCP, the attorney argued that monetary sanctions were "moot" because of the reputational damage he incurred as a result of his erroneous citations coming to light. The judge was unsympathetic, noting that it is up to the court to determine the punishment and that "sanctions payable to the Court cannot become moot because of external consequences." However, the judge did sanction the attorney for only $6,000 rather than the full $15,000 recommended by a magistrate judge.

Expert Reports

In *Concord Music Group, Inc., et al. v. Anthropic PBC* (U.S. District Court for the Northern District of California), music publishers, including Concord Music Group, brought suit against Anthropic for infringing copyrights of the plaintiffs. To support a brief on the plaintiffs' requests for prompts and outputs from Claude that relate to song lyrics, Anthropic's expert submitted a report that included at least one citation to an article that was identified as an AI hallucination. The court struck the part of her expert report that included this erroneous citation and took note of her credibility. Anthropic explained that it was an "honest mistake" and that the citation must have been accidentally added when it ran her declaration through Claude for "formatting." The court was displeased with the expert for citing an AI hallucination and with the attorney for submitting the expert's report without checking for AI

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

continued

hallucinations. In other words, the court did not care about the severity of the mistake and noted that the expert, and Anthropic's attorneys, still had a duty to check the declaration for accuracy.

In *Kohls v. Ellison* (U.S. District Court for the District of Minnesota), an expert, testifying about AI and deepfakes, used AI to write his expert report, which was filled with AI hallucinations. To support a motion for preliminary injunction, the plaintiff's attorney submitted two expert reports giving background on AI and deepfake technology. One of the expert reports cited two fictitious articles and a third article that was misattributed. The expert admitted to using ChatGPT to draft his expert report, but said that the substantive parts were still accurate. The court pointed out the irony of an expert relying heavily on AI to comment on the dangers of doing so. The court further commented that regardless of whether the substantive parts of the report were accurate, the fact that the expert had submitted false information (the citations) under penalty of perjury was problematic, especially since "[the expert] typically validates citations with a reference software when he writes academic articles but did not do so when submitting the… legal filing." The court further commented that "[the expert's] citation to fake, AI-generated sources in his declaration—even with his helpful, thorough, and plausible explanation (ECF No. 39)—shatters his credibility with this Court." The court commended the honesty that the expert exhibited in acknowledging his mistake and explaining how the mistake occurred, but the court could not ignore the fake citations as simple mistakes. The court excluded the problematic expert report and did not allow the party to correct or replace the excluded expert report prior to ruling on the preliminary injunction.

Judges and Law Clerks

Unfortunately, violations related to AI are not limited to parties, as some judges and their clerks have also been reprimanded for submitting orders and decisions with AI hallucinations in them. For example, Henry T. Wingate of the U.S. District Court for the Southern District of Mississippi and Julien Xavier Neals of the U.S. District Court for the District of New Jersey recently admitted that their offices used AI when preparing opinions that were found to contain AI hallucinations.

In *Mississippi Association of Educators, et al. v. Board of Trustees of State Institutions of Higher Learning, et al.* (U.S. District Court for the Southern District of Mississippi), Federal Judge Henry T. Wingate issued a temporary restraining order in a civil rights case in Mississippi on July 20, 2025. The defendants, led by the state attorney general, filed a motion raising concerns about the accuracy of the order, pointing out that the order misnamed the parties in the case, misquoted state law, and made factually inaccurate statements that were not supported by the record. A few days later, Judge Wingate replaced his ruling with a corrected version and refused to clarify what happened, stating that the issues were merely "clerical errors referencing improper parties and factual allegations." He refused to allow the original ruling to remain available on the public docket. A few months later, Senator Chuck Grassley of the U.S. Senate Committee on the Judiciary sent an inquiry to Judge Wingate requesting clarification on what occurred with the original order and speculated on the use of AI. In response, Judge Wingate acknowledged that his law clerk had used Perplexity, an AI program, to write the order. Wingate further admitted to inadequately reviewing the order

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

continued

before issuing it. While Judge Wingate claims that he will take steps to ensure this does not happen again, there appear to be no additional corrective measures for Judge Wingate's AI hallucination-laden order.

In *In re CorMedix*, Judge Neals issued an opinion on a motion to dismiss in a biopharma securities case on June 30, 2025. In July, the defendants filed a letter raising their concerns about errors in the opinion, including fictitious quotations and incorrect case outcomes cited. In August, Judge Neals withdrew the previous opinion and issued a new opinion without explanation. In response to a letter similar to that written by Senator Grassley to Judge Wingate, Judge Neals admitted that a law school intern in his office had used ChatGPT to draft the opinion and that he had inadequately reviewed the opinion. Similar to Judge Wingate, Judge Neals also committed to implementing preventive measures to ensure this does not happen again, with no additional corrective measures taken.

New and Proposed Rules to Curtail the Misuse of AI

The courts and the federal bar are taking action to deter submission of AI hallucinations and the misuse of AI in filings. For example, the bar has proposed changes to Federal Rule of Evidence (FRE) 707, which establishes uniform standards for introducing evidence generated or significantly modified by AI systems, such as outputs from machine learning models or deepfakes. This is a response to the current lack of specific rules for AI evidence, which has led to confusion and inconsistency in court. Proposed FRE 707 would subject machine-generated evidence to the expert witness admissibility standards under the existing FRE 702 (which includes the *Daubert* standard for reliability). To admit AI evidence under the proposed rule, the parties need to show:

- The evidence is helpful to the trier of fact (judge or jury).
- The AI tool relied on sufficient facts or data.
- The AI tool used reliable principles and methods.
- These principles and methods were reliably applied to the specific facts of the case.

The proposed rule relies on the court's role, requiring judges to scrutinize underlying technology, data inputs, and validation studies to ensure evidence is reliable before it can be presented to a judge or jury. The proposed rule was approved for publication for public comment in June 2025 by the U.S. Judicial Conference's Advisory Committee on Evidence Rules, such that it is under review and has not been finalized or implemented.

We also looked into how the top patent courts are addressing use of AI through local rules, some of which are being tracked.

Some specific courts have taken action to change local rules.

The Eastern District of Texas modified Local Rule CV-11(g): Use of Generative Artificial Intelligence Technology. Local Rule CV-11(g) now states, "All litigants remain responsible for the accuracy and quality of legal documents produced with the assistance of generative artificial intelligence technology. Litigants are cautioned that certain generative artificial intelligence technologies may produce factually or legally inaccurate content. If a litigant chooses to employ generative artificial intelligence technology, the litigant

AI Hallucinations in Court Filings and Orders: A 2025 Review of Sanctions Across the Courts and Rule Proposals

continued

continues to be bound by the requirements of Fed. R. Civ. P. 11 and must review and verify all content to ensure that it complies with all such standards. See also Local Rule AT-3(m)."

In the Northern District of California, no new rules or procedures have been issued, but individual judges have standing orders requiring disclosure of when GenAI is used in court filings.

The Delaware Supreme Court issued interim guidelines for the use of GenAI in 2024.

The authors will continue to monitor this important issue to see how the U.S. Patent and Trademark Office, federal courts, the U.S. International Trade Commission, and other patent venues will handle these circumstances.

AI Tools Are Hard at Work—and So Are We

BY: RYAN ESTATICO, TODD HOPFINGER, AND LESTIN KENTON, JR.

The era of experimental artificial intelligence (AI) pilots is ending, and enterprise-scale adoption has begun. Generative and foundation models now power everyday copilots that draft documents, HR agents that answer policy questions, and automated workflows that accelerate decision-making.

The promise is speed, scale, and smarter insights. Yet without disciplined governance, these same systems can create operational, legal, and reputational risk. Treating AI as a plug-and-play productivity tool rather than a managed system can invite costly mistakes.

As Bill Gates noted in AI Is About to Completely Change How You Use Computers, "Soon after the first automobiles were on the road, there was the first car crash. But we didn't ban cars—we adopted speed limits, safety standards, licensing requirements, and other rules of the road."[1] The same principle applies here: AI progress must be matched by policy, design, and oversight.

This article explores three common risk patterns in enterprise AI adoption: how they arise, why they matter, and what organizations might consider doing to guard against them.

1. Model Performance, Hallucination, and Version Drift

AI's promise of speed and scale is only as strong as the accuracy of its outputs. When generative systems

1 Bill Gates, AI Is About to Completely Change How You Use Computers, Gates Notes, https://www.gatesnotes.com/meet-bill/tech-thinking/reader/ai-agents

AI Tools are Hard at Work—And So Are We *continued*

sound convincing but are wrong or outdated, trust erodes and the damage can be immense and spread quickly. Errors at the model level can quickly become governance issues at the enterprise level. A single flawed output can cascade through automated systems, policy documents, or customer interactions, amplifying the original mistake.

Two patterns drive this breakdown: hallucination and version drift.

Hallucination: Confidently Wrong

Hallucination occurs when an AI system fabricates information or sources that never existed. Hallucinations are a result of AI models' probabilistic nature and their training to produce fluent, humanlike text even when uncertain. Real-life errors have ranged from complete fabrications of authority to small inaccuracies within citations.

These mistakes often sound authoritative, which makes them dangerous. An invented citation in a legal memo or a misquoted policy in an HR document can quietly ripple through an organization before anyone notices. Once institutional trust is compromised, even verified outputs face skepticism. The cost is not just factual error but loss of confidence in the system itself.

Techniques such as retrieval-augmented generation (RAG) and other fine-tuning of the AI system can reduce hallucination risk but cannot eliminate it altogether. The only reliable safeguard remains human verification, especially in functions like legal, compliance, finance, and policy, where precision determines accountability.

Version Drift: Outdated but Believable

Version drift is more subtle but can be just as costly. Version drift occurs when AI retrieves information that was once correct but no longer is (e.g., outdated policies, contracts, or internal guidance that have fallen out of sync with reality). Unlike hallucination, version drift involves accurate retrieval of obsolete truth.

Version drift, just like hallucination, has caused havoc for companies, both in the media and in court. For example, in *Moffatt v. Air Canada*, a customer seeking bereavement fares asked the airline's chatbot about discount eligibility. The bot confirmed the policy's availability and directed the passenger to purchase tickets and claim the discount later. However, Air Canada had discontinued bereavement fares months earlier. The chatbot, although wrong, did not hallucinate; it cited outdated internal guidance. The British Columbia Civil Resolution Tribunal held the airline liable, finding that a company bears responsibility for what its AI says on its behalf.[2]

This case demonstrates that version drift can turn outdated knowledge into a liability. In a regulated environment, "technically correct but no longer true" can still be grounds for accountability.

Version drift doesn't just threaten accuracy; it can undermine compliance, audit integrity, and customer trust. In global enterprises, local versions of policies or datasets may age at different rates, creating inconsistencies that quietly spread. In the age of generative systems, data freshness has become a form of compliance.

2 *Moffatt v. Air Canada*, 2024 BCCRT 149 (CanLII).

AI Tools are Hard at Work—and So Are We *continued*

Mitigation Practices

- **Timestamping and Versioning**: Many organizations might consider maintaining timestamps and version markers for retrieved content so that references to policies or documents reflect their most recent update dates.
- **Regular ReIndexing of Knowledge Bases**: It might be wise for content repositories to be refreshed on defined cycles—such as compliance or audit intervals—to keep indexed information aligned with current standards.
- **Human-in-the-Loop Review**: In areas involving higher sensitivity (e.g., HR, legal, compliance), it might be wise to incorporate review checkpoints to maintain oversight of AI-generated outputs.
- **Monitoring Correction Latency**: Organizations might consider tracking the time between the detection, correction, and redistribution of updated information as an indicator of information management responsiveness.
- **Ownership for Data Freshness**: It might be wise to designate data or knowledge steward roles to manage content life cycle processes, including version control and update governance.
- **Automated Freshness Indicators**: Metadata tracking, API monitoring, or similar mechanisms may be used to signal when referenced sources approach or exceed a defined staleness threshold.

Accuracy is not merely technical; it is the foundation of trust. Organizations that treat data quality and model oversight as continuous disciplines, not one-time controls, are often best positioned to scale AI confidently and responsibly.

2. Integrating AI Responsibly: Governance, Accountability, and Oversight Gaps

AI adoption is outpacing most enterprise governance systems, creating a dangerous blind spot between perception and reality. Enterprises are often blind to how individuals are using AI to boost their performance. In McKinsey's State of AI 2025 survey, nearly all executives and employees reported familiarity with generative AI—but leaders believed only 4% of employees used it regularly. However, the real number was likely three times higher.[3] Enterprises cannot manage what they cannot see, and many underestimate the extent of unsanctioned AI experimentation happening within their own walls.

When employees use unapproved AI tools to accelerate their work, sensitive data may leave secure environments, and unvetted models may influence decisions without proper oversight. Essentially, employees are running an invisible experimentation of AI tools, creating an enterprise liability. Leaders who lack visibility into which models are influencing operations or what data those models touch cannot credibly manage compliance, privacy, or reputational exposure.

Accountability and Regulatory Risk

Weak oversight is not just a technical risk—there are tangible business and regulatory consequences. Algorithmic decision tools have drawn external scrutiny when their outputs affect consumers without adequate transparency or review.

3 Sheryl Estrada, "MIT report: 95% of generative AI pilots at companies are failing," Aug. 18, 2025. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

AI Tools are Hard at Work—and So Are We *continued*

Recent cases involving AI-enabled hiring systems illustrate the risk. In 2025, a federal court allowed a nationwide collective action to proceed against Workday after plaintiffs alleged that its AI-driven applicant-screening tools disproportionately rejected applicants over the age of 40 across multiple employers using the platform. The court found the claims plausible under the Age Discrimination in Employment Act and authorized the case to move forward, reflecting increasing judicial willingness to scrutinize automated hiring pipelines.[4] In another prominent example, the Equal Employment Opportunity Commission reached a settlement in 2023 with iTutorGroup after alleging that its automated hiring software was programmed to automatically reject older applicants based on age thresholds embedded in the system—resulting in more than 200 qualified candidates being screened out solely due to age.[5]

AI enforcement actions underscore a broader lesson: automation does not absolve companies of accountability. Enterprises deploying AI-driven decision systems are increasingly expected to explain how decisions are generated, what data sources or proxies are used, and which safeguards keep outcomes within legal and ethical boundaries. Unexpected errors can lead to reputational damage and regulatory exposure. Governance gaps, therefore, might include not only missing oversight structures but also the absence of clarity and explanations that regulators increasingly expect from organizations relying on automated decision tools.

Internal Development and Rollout of AI Tools

The following are considerations at a generic enterprise level, but approaches to AI governance should be tailored to match the enterprise's own structure, risk appetite, and technological maturity.[6]

- **Leadership** bodies often play a central role in shaping the organization's AI risk posture by defining elements such as risk appetite, ethical boundaries, and oversight expectations. Regular visibility into AI deployments, identified risks, and compliance status is commonly part of governance reporting practices.
- **Management** functions typically translate governance expectations into operational processes. This can include assigning AI system ownership, establishing access and data controls, and incorporating responsible-use training across departments. Management structures often emphasize traceability, documentation, and review mechanisms across deployed models.
- **Internal audit and assurance groups** provide validation and oversight by examining the effectiveness of controls, assessing transparency mechanisms, and reviewing the explainability and reproducibility of AI-supported decisions.

Effective governance shields organizations from regulatory surprise and maintains trust when AI operates in public contexts. Systems that are owned, reviewed, and documented by accountable humans generally operate predictably. Organizations that institutionalize this discipline often adapt faster, avoid regulatory shocks, and strengthen confidence among clients, employees, and regulators alike.

4 *Mobley v. Workday, Inc.*, 3:23-cv-00770 (N.D. Cal. Feb. 21, 2023).

5 *Equal Employment Opportunity Commission v. iTutorGroup, Inc.*, 1:22-cv-02565 (E.D.N.Y.).

6 The Institute of Internal Auditors, The IIA's Artificial Intelligence Auditing Framework, "Part 3 – AI Auditing Framework," updated 2024, https://www.theiia.org/globalassets/site/content/tools/professional/aiframework-sept-2024-update.pdf

AI Tools are Hard at Work—and So Are We *continued*

3. Human Versus AI-Generated Code

Essentially, copyright laws exist to protect your original work from being copied or claimed by others. But what if that work was never fully yours to begin with? That question sits at the center of the legal uncertainty surrounding AI-generated code.

Copyright protection provides the legal protection of most software assets. But when code is generated, modified, or integrated using artificial intelligence, that foundation becomes uncertain. If no human author can be clearly identified, the resulting code may not qualify for copyright protection.[7] Therefore, without a human-backed system, an enterprise could lose the ability to defend, license, or monetize its own intellectual property.

At the same time, rapid advances in AI coding agents are accelerating this uncertainty. These agents have made significant progress on software engineering benchmarks—standardized tests that measure model performance. Developers increasingly describe a new working style—sometimes called "vibe coding"—where engineers articulate intent in natural language and rely on AI systems to write, refine, and debug entire software components.

Legal and Technical Risk

U.S. copyright law requires that creative works originate from a human being. When AI tools such as GitHub Copilot or Amazon CodeWhisperer assist in writing or refactoring code, the line between human and machine authorship blurs. Adoption is accelerating too: 65% of developers now use AI coding tools at least weekly.[8] Large portions of modern software may already include code that cannot be definitively attributed and therefore are not definitively owned.

This ambiguity also compounds legal risk. AI-generated code can inadvertently reproduce, remix, or approximate existing open-source snippets under restrictive or incompatible licenses, introducing license contamination. For startups seeking investment or enterprises developing regulated products, these issues can trigger costly audits, force re-architecture efforts, delay product launches, or create downstream legal exposure. Companies in regulated industries or those distributing commercial software are especially at risk.

Actionable Mitigation

To reduce ambiguity and still capture the benefits, it is often helpful to adopt structured documentation and governance practices around AI-assisted development:

- **Marking AI Contributions**: Some development teams may consider annotating portions of the codebase that were generated, modified, or influenced by AI tools to maintain visibility into the origin of specific code segments.
- **Provenance Auditing**: Version control platforms such as Git may be leveraged to trace authorship, track changes, and record human review activity associated with AI-assisted contributions.
- **Documenting Human Oversight**: Teams might consider maintaining records of prompts, review notes, and approval checkpoints to provide transparency into where and how human authorship and supervision occurred.

7 U.S. Copyright Office, Copyright and Artificial Intelligence Part 2: Copyrightability 10 (Jan. 2025), https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf

8 Stack Overflow, 2025 Developer Survey: AI, https://survey.stackoverflow.co/2025/ai#ai-agents-ai-agents

AI Tools are Hard at Work—and So Are We *continued*

- **Reviewing Open Source Dependencies.** Open-source license reviews may be performed to understand whether AI-generated code introduces dependencies with restrictive or incompatible licensing terms.
- **Developer Training and Supervision.** Many organizations might consider training on authorship standards and oversight expectations to reinforce the role of meaningful human contribution within AI-supported development processes.

The legal framework around AI-generated code is still evolving. U.S. copyright authorities have reaffirmed that protection requires meaningful human authorship. Until clearer international standards develop, documenting the human contribution in AI-assisted development remains a practical safeguard. Transparent attribution of human and AI contributions can allow teams to navigate the evolving landscape responsibly while continuing to leverage AI's productivity benefits.

Conclusion

The enterprise AI landscape of 2025 mirrors the early stages of other tech revolutions: transformative, promising, and still taking shape. Enterprise structures must evolve as quickly as the technologies themselves. The organizations that succeed will be those that treat AI not as a shortcut but as an engineering discipline. They will build strong guardrails early, so progress can accelerate safely and sustainably.

Sterne Kessler is uniquely situated to counsel clients interested in growing an AI program. Our own AI initiatives reflect the integration of law, engineering, and ethics. Our approach allows us to scale responsibly while preserving the quality, accuracy, and trust that define our practice.

About Sterne, Kessler, Goldstein & Fox

Based in Washington, D.C., Sterne Kessler is one of the world's leading intellectual property law firms, specializing in the full range of IP services globally. We are passionate about IP law, with a unique combination of legal acumen and technical experience across both prosecution and litigation. With more than 200 attorneys, patent agents, and technical specialists across the firm, we are committed to delivering practical, business-driven solutions that support our clients' innovation and long-term objectives. For over 45 years, we have been a trusted partner to the world's most innovative companies and inventors, helping them protect, enforce, and maximize the value of their IP around the globe.

Our highly integrated IP practice focuses on patent prosecution and strategic counseling, post-grant proceedings, litigation in federal district and appellate courts and before the International Trade Commission, and trademark prosecution and enforcement. Our clients include Fortune 500 companies, entrepreneurs, start-ups, inventors, venture capital firms, and universities that are making discoveries, building brands, and creating inventions that impact our daily lives.

For more information about the firm and our services, visit sternekessler.com or follow us on LinkedIn to get the latest news and updates in IP.

Sterne, Kessler, Goldstein & Fox PLLC

1101 K Street NW, 10th Floor

Washington, DC 20005