AI on Trial: How Copyright Litigation is Reshaping AI Risk

California courts are emerging as the focal point of high-stakes litigation over generative artificial intelligence (AI) and intellectual property (IP). In the three-and-a-half years since ChatGPT's public release, generative AI tools have become embedded in daily tasks, from drafting emails to creating meal plans. But large language models (LLMs) must be trained on vast amounts of data, much of which belongs to rights holders, raising questions about how that data is sourced, regulated and used.

There are currently no concrete rules specifying what qualifies as "fair use" when an AI developer draws upon protected works to train and operate their AI systems. That has triggered a wave of AI copyright litigation testing how existing laws apply to the evolving technologies. From the widely watched *Bartz v. Anthropic* case to pending lawsuits brought by major media organizations, including the New York Times, plaintiffs are advancing claims that challenge both training methodology and output risk.

> As the risk landscape surrounding AI takes shape in real time and friction remains between traditional IP doctrine and the new technology, defensibility for California developers, for now, largely depends on data provenance and operational oversight.

Where AI Litigation Is Playing Out

California courts, particularly the Northern District, emerged as an early epicenter of AI copyright litigation. However, the center of gravity is shifting. Courts in the Southern District of New York (SDNY) are now handling some of the most consequential disputes brought by major news organizations against AI developers. Delaware is also seeing an uptick in AI-based copyright claims due to its popularity as a corporate legal home. However, many of the Delaware-initiated cases ultimately migrate to alternative jurisdictions as litigation progresses.

These cases reflect the increasing involvement of media, publishing and entertainment companies whose business models depend on protecting and monetizing proprietary content.

The practical implication is clear: While California remains central to AI IP litigation, it is no longer the default venue expected to shape the future of IP law. Practitioners should monitor developments across jurisdictions, particularly in New York, where key fair use questions are being tested.

Training Versus Output: Where the Real Fight Is Emerging

At the core of the disputes is a fundamental question: When does the use of copyrighted material to train an AI system constitute infringement?

Early rulings suggest judicial openness to the argument that training AI models on protected works may qualify as "transformative" and therefore is fair use of the works, particularly when systems learn patterns rather than reproducing specific works. This line of reasoning supports AI developers, particularly when training data is lawfully obtained, drawing on analogies to cases such as *Authors Guild v. Google*. But those precedents did not address systems that generate new expressive outputs, leaving key questions unresolved.

However, counsel should not over-rely on this trend, as litigation is increasingly focusing on the outputs of AI systems. Copyright holders, content creators, authors, artists and news organizations are highlighting instances in which AI-generated content closely resembles or replicates protected works, arguing that such outputs undermine claims of transformation and instead constitute actionable infringement. Through discovery and technical analysis, litigants are attempting to demonstrate that models can reproduce substantial portions of copyrighted material under certain conditions.

The ongoing ***New York Times Co. v. Microsoft Corp. (OpenAI)*** lawsuit in the SDNY is a primary test for whether "transformative use" holds when AI systems generate content, other than training-based content, that mirrors protected works. In April 2025, the court denied the defendants' motions to dismiss secondary infringement claims, signaling that developers and potentially downstream users face ongoing exposure for infringing outputs. That risk is heightened by evidence that models can reproduce memorized portions of their training data, including near-verbatim passages from news articles, books and lyrics.

This distinction between input and output is likely to become a defining fault line in AI IP litigation and should be a focal point of risk assessment.

Pirated vs. Lawful Data: A $1.5 Billion Distinction

Not all training data is treated equally, and courts appear less receptive to claims involving unlawfully obtained data.

The **$1.5 billion settlement in *Bartz v. Anthropic*** established critical persuasive authority. Senior U.S. District Judge William Alsup in the Northern District of California suggested in the ruling that while training on legally acquired data may be fair use, the act of downloading pirated content from shadow libraries such as LibGen is likely infringing. The settlement included authors of the roughly seven million works stolen by the shadow libraries, which amounted to about $3,000 per registered work.

Recently, **BMG and other major publishers filed a $3 billion suit against Anthropic** in the Northern District, alleging the unauthorized use of over 20,000 song lyrics and citing the same "shadow library" theories used in Bartz.

In contrast, a case brought by the **Chicago Tribune against Perplexity**, filed in the SDNY, focuses on Perplexity's use of nonpirated, paywalled content obtained through Retrieval-Augmented Generation (RAG). This technique combines large language models with external knowledge sources — such as databases, document sets or search engines — enabling systems to generate responses grounded in retrieved content rather than solely on learned patterns.

That distinction raises questions about secondary liability. Because RAG systems rely on retrieved material, their outputs are more likely to include verbatim or near-verbatim content. The Tribune argues that Perplexity's use of RAG reproduces its original reporting, while Perplexity maintains that the resulting outputs remain transformative. Therefore, the case presents a direct test of how courts will treat lawfully accessed but protected material.

As these cases evolve, they are testing not only how courts evaluate training data and outputs but also how responsibility is assigned when AI systems generate infringing content.

Shifting Away from 'AI Did It'

An internal IBM training document in 1979 said, "A computer can never be held accountable." And for years, the claim that the system acted independently was a popular defense in autonomous AI cases.

Recent legislative efforts in California suggest a more constrained view. AB 316, signed into law by Governor Gavin Newsom in October 2025, bars the "AI did it" defense, prohibiting claims that AI acted independently. Instead, the law shifts focus to the developer's or operator's role in designing, deploying and overseeing those systems.

Courts and lawmakers are following suit by increasingly treating AI outputs as extensions of the organizations that create and use them. For companies, the shift raises the stakes around governance, oversight and risk management, regardless of how complex or opaque the underlying technology may be.

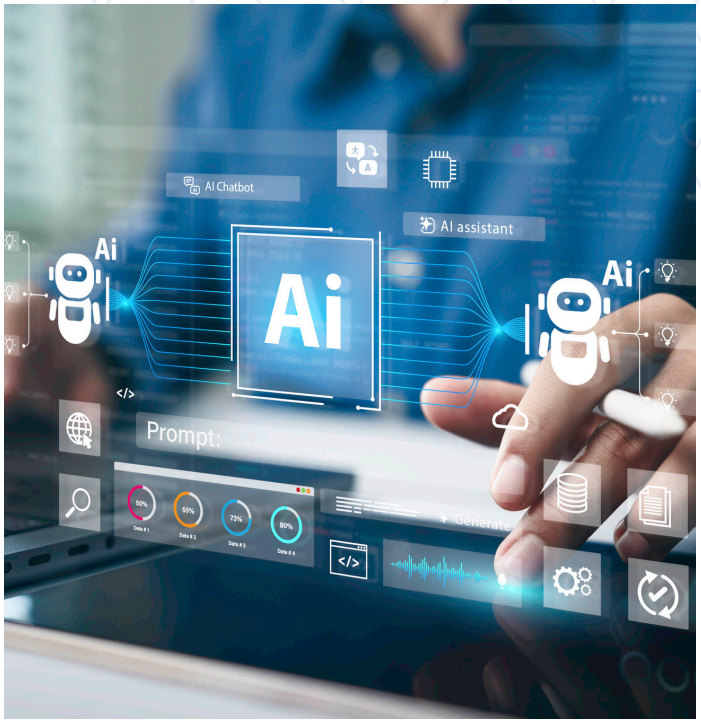

California's 2026 Regulatory Guardrails

Federal guidance on AI remains limited. Much like earlier waves of technological disruption, such as the rise of peer-to-peer file sharing and search engine indexing, courts and litigants are operating within traditional statutory frameworks to scrutinize AI. The Copyright Act of 1976 remains the governing authority over copyright. The Patent Act remains the authority over patented works, and the Lanham Act governs trademarks.

Still, California state lawmakers continue to take an active role by introducing a range of bills addressing transparency, training data disclosure and measures to support AI governance and compliance. New laws taking effect in 2026 include AB 316, SB 53 and AB 853. SB 53, the Transparency in Frontier Artificial Intelligence Act, requires developers of advanced "frontier" models, such as OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, Perplexity and X's Grok, to publish safety frameworks and disclose catastrophic risk assessments. AB 853, the California AI Transparency Act, requires providers to offer free AI-detection tools and include "latent disclosures" (watermarking) in AI-generated content by Aug. 2.

Other states, such as Colorado, have enacted landmark initiatives. The state passed and signed into law the controversial Artificial Intelligence Act, which is now undergoing amendment. As enacted, it required developers and deployers of certain "high-risk" AI systems to implement measures to prevent algorithmic discrimination and support AI governance.

> As of March 2026, 45 states have introduced more than **1,500 AI-related bills**, including 90 in just California.

The state developments do not replace federal copyright law. Instead, companies must navigate an increasingly layered environment in which evolving state-level requirements intersect with existing federal statutes.

CEB

What Corporate Counsel Should Be

The lack of clear legal standards is no reason to take a wait-and-see approach. Retrofitting compliance will be more costly than taking measured steps today.

Counsel should consider:

- Establishing clear internal policies instructing how AI tools may be used and what data they can access
- Reviewing where training data comes from and whether it is properly licensed
- Monitoring AI system outputs to reduce the risk of infringement
- Establishing "kill switches" and preauthorized transaction limits for autonomous AI agents that comply with AB 316
- Monitoring discovery as disclosures required under AB 853 may create a paper trail for future copyright plaintiffs
- Monitoring AI legislation and rulings to ensure AI practices align

Companies that build thoughtful guardrails now will be better positioned as courts begin to draw clearer lines. Waiting for that clarity may seem practical, but in reality, it risks shifting liability downstream. Retrofitting compliance later is likely to be far more costly and disruptive than taking measured steps today.

For a deeper look at how intellectual property law is evolving alongside generative AI and what it means for practitioners, check out **CEB's IP in the AI Era white paper**.

CEB

Contact us at 1-800-232-3444 or visit us online to learn more.