POLSINELLI

Publications

November 18, 2026 • Updates

AI Copyright Decisions Raise Broader Questions About Institutional Governance

Why It Matters

- Institutions adopting AI must consider more than the legality of the technology itself. Decisions about data use, human oversight and accountability can create legal and governance risks that remain the institution's responsibility.

Next Steps

- Institutional leaders should evaluate how AI is used in consequential decisions, establish meaningful human oversight and assess the rights and restrictions governing information incorporated into AI workflows.

Artificial intelligence is quickly moving from a tool that helps employees write emails or summarize documents into something much more consequential. Institutions are beginning to use AI to analyze information, recommend courses of action, evaluate risks, allocate resources and inform decisions that previously depended almost entirely on human judgment.

That change creates enormous opportunities. It also creates difficult choices.

Recent copyright decisions underscore why institutions must look beyond the capabilities of AI technology to consider the legal and governance responsibilities associated with its use. Courts are beginning to address when copyrighted material may be used to train AI systems and when AI-generated output may constitute infringement. Those decisions also highlight a broader question facing companies, universities, nonprofit organizations, government bodies and other institutions:

How much decision-making authority should an institution give to artificial intelligence, and what responsibilities remain with the people who govern it?

Copyright Provides an Early Warning

Recent cases demonstrate why institutional leaders cannot simply assume that the

Related People

- Michael D. Pegues
- Clement A. Asante

Related Capabilities

- Intellectual Property
- Artificial Intelligence & Machine Learning
- Intellectual Property Litigation
- Trademark, Copyright & Brand Protection

technology provider has resolved the legal questions surrounding AI.

In *Thomson Reuters Enterprise Centre GmbH v. ROSS Intelligence Inc.*[1], the Third Circuit rejected a fair use defense involving Westlaw headnotes used to develop an AI-powered legal research product. The court emphasized the commercial nature of the use and the fact that the resulting product competed with Westlaw.

By contrast, federal courts in *Bartz v. Anthropic* [2] and *Kadrey v. Meta* [3] found that the use of copyrighted books for large language model training could qualify as fair use on the records before them. Those decisions, however, did not establish a blanket rule that AI training is lawful.

Bartz is particularly instructive. The court separated the question of using books for model training from the question of how Anthropic obtained some of those books. Training could be transformative while the acquisition and retention of unauthorized copies presented a different copyright problem.

The Ninth Circuit added another distinction in *Doe v. GitHub, Inc.*[4] The court rejected a claim under Section 1202 of the Digital Millennium Copyright Act based on AI-generated code that did not contain copyright management information. Yet the court left open whether substantially similar AI-generated code might violate traditional copyright law.

Together, these decisions point toward an emerging principle. There may be no single answer to the question, “Is our use of AI lawful?” The better questions are more specific: Where did the information come from? What rights govern its use? What does the system do with it? What does the system ultimately produce?

Those same questions should inform AI governance more broadly.

The Harder Question Is Who Makes the Decision

Suppose an institution uses AI to review hundreds of grant applications, evaluate job candidates, analyze disciplinary matters, recommend investments or identify which programs should receive additional resources.

The AI system may provide useful analysis. It may identify patterns that human decision makers would miss.

But the institution still has to decide what role the machine should play.

Should AI simply organize information for human review? Should it recommend a decision? Should a human reviewer be permitted to accept an AI recommendation without independently examining the underlying information?

Those are governance decisions, not technology decisions.

The risk is that institutions may gradually delegate judgment without consciously deciding to do so. A system begins as an efficiency tool. Its recommendations prove useful. Over time, decision makers increasingly rely on its output. Eventually, human review can become little more than approval of a decision made by the system.

That is precisely where governance becomes important.

The National Institute of Standards and Technology’s (NIST) AI risk management guidance emphasizes clearly defining human roles and responsibilities for AI oversight. The European Union’s AI Act similarly requires meaningful human oversight for certain

high-risk AI systems and contemplates that the people performing that oversight must have appropriate competence and authority.

The principle is straightforward: human oversight should mean more than having a person available to press “approve.”

Institutions Must Decide What Information AI Should See

Copyright cases also highlight another governance issue.

AI systems become more useful when they receive more information. That creates a natural temptation to provide them with internal reports, proprietary research, personnel information, confidential communications, copyrighted publications or other sensitive material.

The relevant question should not simply be whether the technology permits the upload.

Institutional leaders should ask whether the institution has the right to use the information for that purpose and whether providing it to an AI system is consistent with contractual obligations, confidentiality requirements and internal policies.

An employee using an AI system to summarize a licensed publication may implicate different rights from an institution using thousands of those publications to create its own internal knowledge system.

The recent copyright cases demonstrate why provenance matters. Knowing where information originated and what rights accompany it should become part of AI governance rather than an inquiry that begins only after litigation arises.

Efficiency Cannot Replace Accountability

One of AI’s greatest attractions is efficiency. Institutions can analyze more information and reach decisions more quickly.

But efficiency is not necessarily the same thing as sound governance.

Some institutional decisions involve competing values that cannot be resolved solely by identifying the most statistically efficient result. A university deciding how to allocate scholarships, a board deciding whether to discontinue a community program or a company evaluating employees may need to consider context that an algorithm cannot fully appreciate.

AI can inform those decisions. It should not allow the institution to avoid responsibility for them.

That also means institutions should be able to explain consequential decisions. “The AI recommended it” is unlikely to satisfy an employee, customer, regulator, donor or court seeking to understand why an important decision was made.

The EU AI Act reflects a similar concern by requiring human oversight and, for certain uses, information to individuals affected by AI-assisted decisions.

What Institutional Leaders Should Do Now

Institutions do not need to choose between embracing AI and avoiding it. They need governance structures that allow them to use the technology without surrendering

judgment to it.

That starts with identifying the decisions for which AI is being used and determining whether the system is providing information or effectively making the decision.

Institutions should also identify who has authority to approve particular AI uses. Responsibility should not disappear simply because a third-party vendor provides the technology.

For consequential decisions, meaningful human review should be built into the process. The reviewer should understand what the AI system considers, recognize its limitations and have genuine authority to reject its recommendation.

Finally, institutions should understand the information entering the system. Copyright ownership, license restrictions, confidentiality obligations and data provenance should be considered before valuable information is incorporated into an AI workflow.

The Larger Lesson

Developing copyright cases are important because they tell us something larger about artificial intelligence.

AI rarely eliminates a difficult legal or institutional question. More often, it moves the question somewhere else.

A dispute over training data becomes a question about where the information came from. An efficiency tool becomes a question about who actually made the decision. An algorithmic recommendation becomes a question about who will accept responsibility when the recommendation is wrong.

Institutions should take advantage of what artificial intelligence can do. But the most important decisions about AI will not be made by AI.

They will be made by the people entrusted with governing the institution.

[1] 139 F.4th 271

[2] 787 F. Supp. 3d 1007

[3] 788 F. Supp. 3d 1026

[4] No. 24-7700