AI Copyright Debate and the Inevitable Knowledge Transmission Argument for LLMs and Future AI

Artificial intelligence changes the speed and scale of learning. The underlying practice is centuries old. Copyright law must distinguish the acquisition of knowledge from the reproduction of protected expression.

By Basil C. Puglisi, MPA
Human-AI Collaboration Strategist

On November 6, 2026, Judge Sidney Stein approved 56 amicus briefs in *In re OpenAI, Inc. Copyright Infringement Litigation*, the multidistrict copyright litigation against OpenAI and Microsoft in the Southern District of New York. Authors, publishers, technology companies, academics, librarians, unions, and public-interest organizations filed them in the second half of October. The briefs divide sharply: 27 support the plaintiffs, 18 support OpenAI and Microsoft, 7 were filed in support of neither party, and 4 stated no position. Together they press the court on a central question: whether training artificial intelligence on copyrighted material constitutes fair use.

Much of the debate treats artificial intelligence as though learning from copyrighted material represents an entirely new activity. It does not.

Knowledge lawfully acquired by people moves into new minds and new systems, artificial intelligence included, through teaching, writing, and now computation. That movement concerns ideas and capabilities, not the production of an identical model or the legality of every intermediate copy, and the position that copyright law should respect that movement is what I call the Inevitable Knowledge Transmission Argument.

Human civilization develops through the consumption, interpretation, and transmission of existing knowledge. We read books, attend lectures, study published research, consult libraries, and learn from people who previously acquired knowledge from other sources. Artificial intelligence introduces a technological method for accelerating that established process.

The central question becomes whether copyright protects the original expression of an author or extends to controlling the knowledge that others acquire by studying that expression. American copyright law already recognizes this distinction.

The Freedom to Learn From Copyrighted Material

Consider how knowledge moves through an ordinary university. A professor purchases a copyrighted textbook, reads it, and develops an understanding of its contents. The professor then teaches students, discusses the underlying concepts, and applies those ideas to new problems. Those students eventually become researchers, engineers, teachers, and employees who carry the knowledge into laboratories, businesses, and institutions.

The original author receives copyright protection for the expression contained in the book. That protection does not generally prevent others from learning its ideas or independently applying the knowledge they acquire.

Libraries extend the same principle across society. A person can borrow a copyrighted book without purchasing an individual copy, because the library ordinarily acquires the books it lends and allows successive readers to encounter the same intellectual material. Readers discuss what they learn with people who never opened the original work, teachers transmit the knowledge to generations of students, and researchers build new findings upon established scholarship. Knowledge moves beyond the original publication without requiring a new license every time someone learns or teaches an idea.

The United States Copyright Act supports this distinction. Section 102(b) excludes ideas, procedures, processes, systems, methods of operation, concepts, principles, and discoveries from copyright protection, whatever form they take in the work. Section 107 expressly names teaching, scholarship, and research among the purposes that may qualify for fair use.

Specific acts of reproduction, distribution, or public display remain subject to copyright restrictions and applicable exceptions. Library lending itself rests on the first-sale doctrine, which lets the owner of a lawful copy lend or sell that copy without making a new one, and in *Hachette v. Internet Archive* the Second Circuit held in September 2024 that scanning print books and lending the digital copies was not fair use, even one copy at a time. The library analogy therefore holds for the ideas a book carries rather than for copying the book. That limit leaves the core intact: the underlying acquisition and communication of unprotected ideas are ordinary, lawful features of education, and that established practice matters when evaluating AI.

Artificial Intelligence Accelerates an Existing Method

Imagine an organization that wants to build an artificial intelligence system capable of explaining economics. It hires hundreds of economists, educators, and researchers, each of whom reads lawfully obtained books, academic papers, and historical publications. They develop original instructional material, explain economic principles, create examples, and document relationships between concepts. Engineers then use those independently developed educational materials to train a machine-learning system.

The resulting system gains capabilities derived from human knowledge, including knowledge that originated in copyrighted publications. No author possesses a blanket right to prevent that knowledge from being taught simply because it appeared in a copyrighted book.

Now consider another organization pursuing a similar objective through computational training. Its engineers obtain a collection of works and use software to identify linguistic relationships, conceptual patterns, and other statistical features. The computational process makes intermediate copies and adjusts numerical model parameters. Its intended result is a system capable of generating new responses.

The technological process differs substantially from human reading, yet its intended intellectual objective remains recognizable: acquire patterns and knowledge from existing materials and apply them to new situations. One method relies on human intermediaries, while the other relies on computational analysis. That distinction in method requires legal examination of any copying involved, but it does not automatically establish that the underlying purpose of learning becomes infringement. The human-learning analogy informs the purpose of training without settling the legality of each copy made along the way.

We already accept technological acceleration in research, data analysis, and communication, and copyright law has accommodated computational uses of protected works before. In *Authors Guild v. Google*, the Second Circuit in 2015 recognized fair use when Google scanned entire copyrighted books to create a searchable index that did not provide substantial substitutes for the originals. The court assessed the purpose and effects of the use, including the copying necessary to make the technology function.

The precedent does not automatically resolve generative AI training. It shows that making complete copies for a technical purpose can be fair use when the resulting use serves a sufficiently different purpose and does not impermissibly substitute for protected expression.

The Fruit of the Poisonous Tree and Inevitable Discovery

The copyright debate also invites an analogy from criminal law. The doctrine commonly described as fruit of the poisonous tree addresses evidence derived from unlawful government conduct, and the Supreme Court recognized an inevitable discovery exception in *Nix v. Williams* on June 11, 1984. Under that doctrine, evidence that would inevitably have been discovered through lawful means may remain admissible despite an earlier constitutional violation, provided the prosecution shows by a preponderance of the evidence that lawful means would have found it.

That doctrine governs evidence in criminal proceedings and does not provide a copyright infringement defense. Its underlying logic nevertheless offers a useful conceptual question: what happens when a result could have emerged through an established lawful process, but technology produces it more efficiently?

An AI system can acquire concepts through instructional material developed by people who lawfully study copyrighted works. A company can employ researchers who read licensed journals, consult purchased books, access libraries, and produce original educational explanations. Those explanations can become part of a separately constructed training process. Over time, a substantial body of knowledge can reach artificial intelligence through ordinary human education and lawful transmission.

The pathway is slower, more expensive, and less comprehensive. It may not reproduce the precise capabilities or parameters of a model trained directly on an enormous source collection. Yet the underlying point remains: the knowledge itself is not exclusively controlled by the original publisher.

Knowledge lawfully acquired by human beings can be taught, reorganized, and transmitted into new systems, including artificial intelligence. Technological acceleration changes the process through which that knowledge is acquired without establishing ownership over the underlying ideas. That is the Inevitable Knowledge Transmission Argument introduced at the outset, and the Nix analogy marks its limit: it is an argument about the foreseeable transmission of ideas and capabilities, not the inevitability of an identical model, and not a claim that criminal evidence law excuses unlawful copying.

The existence of an alternative lawful learning pathway cannot erase an infringement that already occurred. It can, however, challenge the premise that the intellectual capability emerging from training must itself be treated as unlawfully appropriated property. The principle asks the law to separate its treatment of the method from its treatment of the knowledge produced, examining the conduct involved rather than declaring knowledge contaminated simply because a computational system acquired it. Read that way, the analogy speaks to remedy: where acquisition was unlawful, as with the pirated library copies in *Bartz*, the penalty belongs on that step rather than on the trained capability, unless the model itself retains protected expression.

The Same Pathway, Traveled by One Practitioner

The lawful pathway described above is not hypothetical for me. In September 2025, I built a grid of 22 thought leaders whose published work shaped artificial intelligence, from Stuart Russell and Yoshua Bengio to Joy Buolamwini, Kate Crawford, Daron Acemoglu, and Lina Khan, as a structured selection tool. Each received a score from 1 to 5 on three dimensions, Depth, Immediate Critique, and Future Partner, and those scores went through the same multi-AI methodology I use for everything else. I read their books, papers, and public testimony, most of it copyrighted, and tested my own frameworks against what they had published.

My ideas and solutions found alignment where their published positions already pointed the same way, and they improved where those positions pressed against them. That grid became the backbone of *The Minds That Bend the Machine*, released November 6, 2026, which adds Yuval Noah Harari and Yann LeCun to the original 22 and profiles 24 minds across 28 chapters, closing with my own practice, the governance voices the grid did not reach, and the AI industrialist. The grid's three dimensions determined selection, chapter sequence, and the intensity of the simulated critique applied to each figure.

Using AI accelerated that learning without replacing it. I drafted and reviewed the profile chapters across five AI platforms and used them to build simulated critiques, reconstructions of what each figure would argue against my frameworks drawn from that person's published record. My own direct corrections decided what stayed. The book states plainly that those critiques do not represent the views, words, opinions, or endorsements of the people named. Reviews of the 22 chapters drafted at that point, conducted with five platforms, never flagged that Yann LeCun was missing, and I caught the gap myself.

Those critiques changed what I built. Lina Khan's antitrust argument gave structural vocabulary to the policy of keeping several independent AI providers in play, which I call AI Provider Plurality, before that policy had a name. Timnit Gebru's documentation work shaped the transparency requirements in every specification of the governance ecosystem I call HAIA, and Stuart Russell's preference uncertainty principle mirrored the rule that no single AI should be treated as authoritative.

My work shows the human pathway that the Inevitable Knowledge Transmission Argument describes. Copyrighted materials informed my understanding, and that understanding became original frameworks, published analysis, and structured material that other people or AI systems can study in turn. The knowledge moves across human and technological boundaries without transferring ownership of the original authors' protected expression, and the distinction that matters is between learning from those works and reproducing them.

Copyright Protects Expression, Not the Result of Education

Consider the implications of treating every AI capability learned from copyrighted materials as an infringement. A person studies a copyrighted programming manual and subsequently develops new software. A researcher studies published medical literature and discovers a new relationship. An engineer studies a copyrighted textbook and applies its principles to a commercial product. A writer studies dozens of novels and develops an original story informed by literary techniques encountered throughout a lifetime.

Each may generate economic value from knowledge acquired through copyrighted materials, yet commercial benefit alone does not transform the original act of learning into copyright infringement. The same distinction should guide AI training.

A model that develops general capabilities from statistical analysis of protected works presents different legal questions from a model or retrieval system that supplies unauthorized reproductions of those works. An AI system capable of explaining a concept encountered in a copyrighted book is not necessarily replacing that book's protected expression. An AI system delivering substantial passages from the same book, or producing outputs that unlawfully replicate protected expression, presents a different problem.

The difference matters, because both activities may involve the same technology while their relationship to the copyright holder's exclusive rights differs. The conduct therefore deserves examination in stages: acquisition, copying and retention, training use, memorization inside the model, and outputs and their distribution. A finding at one stage does not decide the others.



Figure 1. Five stages of AI training conduct, and how the record treats each one.

Commercial Use Is a Legitimate Question, but Not the Entire Question

Copyright holders raise a substantial economic objection. Companies invest considerable resources in producing books, research, journalism, music, and other creative works, and AI developers can create commercial products using computational training based on those materials. Publishers argue that

unlicensed training reduces the market for licensing their works and may support competing products that divert readers, customers, and advertising revenue.

Those concerns deserve evidence-based examination, and the statute provides for it: commercial purpose is explicitly included in the four-factor fair use analysis, and so is the effect on existing and potential markets. Fair use also examines the nature of the copyrighted work and the amount and substantiality of the material copied. These factors matter when evaluating computational training, particularly when a model requires complete copies to extract statistical relationships. Because the amount copied is weighed in relation to the purpose of the use, the balance cannot be determined by commercial intent alone.

However, the existence of a commercial product does not independently establish that the underlying learning process infringes copyright. Universities develop commercially valuable technologies through publicly available research. Employees contribute knowledge gained from previous education, and businesses build products using ideas developed and published by others. Copyright law has never guaranteed an author exclusive control over every subsequent economic opportunity arising from exposure to a work's unprotected ideas.

The September 29, 2026, appellate decision in *Thomson Reuters v. ROSS Intelligence* illustrates the important boundary. The Third Circuit rejected fair use where a competing legal-research service copied protected editorial materials for a highly similar commercial purpose, and it distinguished that service from general-purpose generative AI systems capable of producing original expression. The relevant questions therefore include the nature of the copying, the purpose of the resulting system, and whether the challenged use threatens legally recognized markets for protected expression.

This distinction is central to the Inevitable Knowledge Transmission Argument. The court examined the function and commercial purpose of the technology and recognized that different AI systems may present different legal circumstances. The ruling does not establish that generative AI training is always lawful, but it shows why a categorical judgment against every technological form of learning would overlook distinctions federal courts already draw. Economic competition deserves scrutiny, but the mere existence of competing knowledge-based products does not establish infringement.

Lawful Acquisition and Training Must Remain Distinct

One more issue deserves direct treatment. AI developers may obtain materials through purchases, licenses, publicly accessible websites, libraries, or unauthorized repositories, and these sources do not necessarily confer identical legal rights. Purchasing a book permits reading and many ordinary subsequent uses, but it does not automatically authorize every form of digital duplication or redistribution. Likewise, publicly accessible material remains subject to copyright.

On June 23, 2025, the federal court in *Bartz v. Anthropic* found that the challenged use of books to train large language models qualified as fair use while declining to excuse the separate acquisition and retention of pirated library copies. The litigation then concluded in a $1.5 billion class-action settlement

that Judge Araceli Martínez-Olguín approved on July 20, 2026, at approximately $3,000 per work before attorneys' fees, which the final order cut to approximately $101.56 million. The settlement did not overturn the court's distinction between training and the separate conduct involving pirated library copies.

That distinction supports a more precise legal framework, in which a developer can face legitimate liability for unauthorized acquisition or reproduction while maintaining a separate fair use argument concerning the purpose of model training. Courts should evaluate these acts individually. A system trained to reproduce protected works may implicate copyright differently from a general-purpose model trained to identify statistical patterns and produce independently generated responses, and this distinction preserves accountability without presuming that every educational or computational use of protected material requires a separate license.

The Risk of Making Knowledge Exclusive

The long-term implications extend beyond OpenAI, Microsoft, or any individual copyright dispute. A legal requirement demanding individual licensing arrangements for every copyrighted work used in AI training could substantially increase development costs. Large companies possessing extensive capital and existing content agreements would be positioned to absorb those costs, while smaller developers, universities, independent researchers, and public-interest institutions could face additional barriers. The result could concentrate advanced AI capabilities among the organizations already holding the greatest financial and technological advantages. How large that effect would be remains an empirical question, and the Open Markets Institute research discussed below reads the evidence differently.

Copyright law exists to encourage creation and advance public knowledge. Its protections should not unintentionally convert accumulated human knowledge into an exclusive resource available only to organizations capable of purchasing comprehensive licensing portfolios. The economic interests of original creators remain legitimate, and so does society's interest in learning from their contributions.

Paying Creators Without Owning Learning

A balanced framework would protect both interests, and it would need a mechanism as well as a principle. Fair use addresses whether particular uses constitute copyright infringement; it does not answer how creators share in the value their work helps produce. Two paths can carry that payment without turning learning into property. Both are prospective legislative options for sharing value, not a statement that training infringes today and not a concession that learning requires a license.

The first borrows from music. Under Section 115 of the Copyright Act, once a nondramatic musical work has been released with the owner's authority, anyone may record and distribute a new version without the owner's consent by paying the statutory royalty. Since January 1, 2021, digital music services have obtained that license through a blanket license administered by the Mechanical Licensing Collective,

while individual notices of intention remain for physical formats such as CDs and vinyl. Payment without permission already exists in copyright law, and for digital services it already runs through a collective.

The same structure could apply to machine learning: a royalty generated when an AI system ingests a work, tracked the way streaming platforms track plays, and paid without granting the owner a veto over what the system learns. That approach would require infrastructure that does not yet exist, along with agreement among publishers, institutions, AI developers, and authors. It would also need a defensible allocation formula, because recording that a work entered a training set does not measure what that work contributed to the model.

The second path is collective, and music supplies that precedent too. When a song plays on U.S. AM or FM radio, performing rights organizations such as ASCAP and BMI license the public performance under blanket licenses and pay the songwriters and publishers, while the station pays nothing to the performers or the owner of the recording. Those organizations have operated since 1941 under Justice Department consent decrees, last amended in 2001 for ASCAP and 1994 for BMI, that keep their licenses nonexclusive and require them to license any applicant on terms equal to those of similarly situated licensees. In its March 20, 2026, legislative recommendations, the White House stated its view that training on copyrighted material does not violate copyright law, left the question to the courts, and asked Congress to consider licensing frameworks or collective rights systems that let rights holders negotiate compensation together without antitrust liability, while leaving open when or whether licensing is required.

Collective negotiation would speak to the cost concern raised above, since it would replace millions of individual agreements with shared channels, and a royalty tied to use would pay creators without granting any of them a veto over what others may learn.

National Security and the Government's Own Path

The executive branch has already tied artificial intelligence to national security. On June 5, 2026, President Donald Trump signed National Security Presidential Memorandum 11, Artificial Intelligence in the National Security Enterprise, directing agencies to speed the adoption and procurement of frontier AI. In the OpenAI litigation, the Justice Department argued in a September 1, 2026, statement of interest that ruling training outside fair use could threaten national security and hand an advantage to foreign adversaries, a filing the Third Circuit later cited as an instance of the Department raising national security concerns. The memorandum says nothing about copyright or training data, and it requires AI use consistent with applicable law.

That raises a question beyond ordinary commercial competition: when the federal government requires advanced AI capability for national defense, must every necessary use of protected intellectual property wait on individual licensing negotiations? Congress answered a version of that question long before AI. Under 28 U.S.C. § 1498(b), when the United States, or a contractor acting for it with its authorization or consent, infringes a copyright, the owner's exclusive remedy is a suit against the government in the

Court of Federal Claims for reasonable and entire compensation. That remedy is compensation from the government, not an order stopping the work.

Together those authorities hold the ground. The executive has made AI a national security directive, and Congress has already provided that when the government uses protected work, the owner is paid rather than handed a veto, the same separation Section 115 makes for recorded music. This rests on a different foundation from the Inevitable Knowledge Transmission Argument and does not establish an unrestricted exemption for private AI development, because Section 1498(b) is a remedy for infringement rather than a finding that training is lawful, and it reaches only use by or for the government. A model trained on a company's own initiative and only later licensed to an agency was not trained for the government, so the statute reaches government-authorized work rather than earlier commercial training. The argument also needs evidence. In ROSS, discussed below, the Third Circuit rejected a national security defense because nothing in the record connected that legal-research platform to national security.

Where the Argument Meets Resistance

The argument has serious opposition, and the record should carry it.

The 27-to-18 amicus split noted at the outset leans against this position, though a count of briefs does not measure the strength of the legal arguments inside them. The plaintiffs' side, according to Tech Policy Press, includes publishers, news groups, music organizations, unions, authors, and a group of copyright law professors. It also includes the Open Markets Institute, which has published research arguing that AI is becoming dominated by Big Tech, and its presence on that side challenges the claim that licensing requirements would serve only incumbents.

The Justice Department filed a statement backing OpenAI on September 1 and rejected the publishers' market-dilution theory. The U.S. Copyright Office's May 2025 prepublication report takes a more cautious position on generative AI training, particularly on market effects and licensing. A June 30, 2026, Supreme Court order, denying the government's application for a stay, allowed Register of Copyrights Shira Perlmutter to remain in office while litigation over her removal continues, which may clear the way for formal publication of the report, although the order resolved nothing about the report's conclusions or the copyright status of AI training.

Intermediate copying draws the sharpest objection. The Copyright Office's report finds that building a training dataset from copyrighted works clearly implicates the reproduction right, and that model weights able to produce a nearly identical copy of a work may implicate it as well, according to Sidley Austin's review of the report. The Third Circuit applied the point in ROSS, rejecting the argument that copying headnotes as an intermediate step to train AI was excused under *Sega v. Accolade* and *Sony v. Connectix*, because ROSS could have built its training material from the unprotected judicial opinions themselves. In the court's words, "ease is not a justification for copying." That reasoning turns this

article's lawful-pathway logic around: where a lawful route to the same knowledge exists, a court may treat copying the protected version as unnecessary rather than inevitable.

The same opinion cuts against the article a second way. Under the fourth factor, the court treated the market for licensing headnotes as AI training data as a fast-developing derivative market and found that ROSS took Thomson Reuters's opportunity to enter it. If courts treat training licenses as a market the copyright owner may claim, the line this article draws between licensing harm and ordinary competition becomes harder to hold.

The opinion also meets the national security ground this article raises. ROSS invoked national security in its defense, and the court answered that "ROSS again presents no evidence connecting its legal-research platform to national security." It observed that the Justice Department, despite raising such concerns elsewhere, had not raised them in this case, and it held that while "some AI technology may implicate national security concerns, this does not give ROSS carte blanche to violate copyright law merely because it incorporates AI." On that reasoning, a national security claim needs evidence tying the specific system to the government's need, and the presence of AI alone does not supply it.

None of these points settles the question. The Third Circuit confined its holding to a legal-research service that competed directly with Westlaw and could not generate original expression, and the *Bartz* court, ruling in June 2025, found general-purpose training "spectacularly" transformative while setting ROSS apart as copying to build a directly competing tool. The ROSS decision is also not final. ROSS petitioned the full Third Circuit for rehearing in early November, with amici in support, and under Supreme Court Rule 13 the time to ask the Supreme Court to hear the case runs from the ruling on that petition. Unless the full court vacates it, the opinion stands as precedent within the Third Circuit, but a circuit decision binds only the courts of that circuit, and until the Supreme Court grants or denies review, or the time to ask runs out, the question remains open nationally. The facts of a single case also shape its reach, and the next case, built on a general-purpose model and a different record, may be presented in ways that lead courts to a different result.

The royalty proposal meets its own resistance. The Copyright Office's report warns that compulsory licenses could hamper flexible, market-based solutions and may conflict with the basic principle that using copyrighted work requires consent. It favors voluntary licensing and collective licensing with safeguards against anticompetitive conduct. From the other direction, a payment owed for every work ingested sits uneasily beside an argument that learning should not require permission. Proposing any payment may also point courts toward a training-license market, the very market the Third Circuit counted against ROSS under the fourth factor. Section 115 answers part of both objections, since payment without permission already exists in music, but Congress created that license by statute. Nothing comparable exists for AI training today.

Europe reads the technology differently. In *GEMA v. OpenAI*, decided November 11, 2025, the Munich Regional Court held that memorizing the lyrics of nine German songs inside a model's parameters was itself a reproduction under German copyright law. That holding challenges the line this article draws

between learning and copying, because it locates a copy inside the model rather than only in its outputs. The ruling applies German law and binds no U.S. court, but it shows how another major jurisdiction reads the same technology. The case concerned those nine songs, and the published case summary leaves open how far the reasoning reaches beyond them. The judgment remains on appeal before the Munich Higher Regional Court.

The criminal-law analogy carries its own limit. *Nix v. Williams* requires the prosecution to show that lawful means would have found the evidence, which is more than showing that they could have. That is why the Inevitable Knowledge Transmission Argument rests on the movement of ideas rather than on the inevitability of an identical model, and a court applying the doctrine's own standard would still ask whether the lawful human pathway would in fact have produced the same capability.

The Question the Courts Should Answer

The courts now face a choice with consequences beyond the technology industry. They can interpret computational learning primarily through the presence of copyrighted inputs, or they can examine the actual conduct, purpose, and market effects of the technology.

A reasonable application of existing copyright principles should recognize a strong fair use case for general-purpose model training that analyzes lawfully accessed materials to develop capabilities rather than distribute protected expression. Independent infringement arising from unauthorized acquisition, impermissible reproduction, or substantially similar outputs remains subject to legal accountability, and that principle is consistent with the treatment of human learning.

Authors deserve protection for their original expression, and publishers deserve legal remedies when their protected rights are violated. Society retains an interest in the free development and transmission of ideas, and artificial intelligence creates new technological questions about how those interests interact without erasing the centuries of human education that preceded it.

Knowledge has always moved from one mind to another, through books, classrooms, conversations, and institutions. Artificial intelligence introduces a faster computational pathway into that established system. The law should examine what the technology actually does. Copyright protects creative expression, but it should not become a perpetual license requirement for learning from the knowledge that expression conveys.

The method has changed. The human pursuit of knowledge continues.

References

Authors Alliance. (2024, September 5). *Hachette v. Internet Archive update: Second Circuit Court of Appeals rules against Internet Archive.* https://www.authorsalliance.org/2024/09/05/hachette-v-internet-archive-update-second-circuit-court-of-appeals-rules-against-internet-archive/

Authors Guild. (2026, July 21). *Court grants final approval of $1.5 billion Anthropic copyright settlement.* https://authorsguild.org/news/court-grants-final-approval-anthropic-copyright-settlement/

Authors Guild v. Google, Inc., 804 F.3d 202 (2d Cir. 2015). https://www.copyright.gov/fair-use/summaries/authorsguild-google-2dcir2015.pdf

Bartz v. Anthropic PBC, 787 F. Supp. 3d 1007 (N.D. Cal. 2025). https://www.govinfo.gov/content/pkg/USCOURTS-cand-3_24-cv-05417/pdf/USCOURTS-cand-3_24-cv-05417-0.pdf

ChatGPT Is Eating the World. (2026, November 12). *Amici briefs support ROSS Intelligence's bid for en banc review by Third Circuit.* https://chatgptiseatingtheworld.com/2026/11/12/amici-briefs-support-ross-intelligences-bid-for-en-banc-review-by-third-circuit/

Congressional Research Service. (2020, March 17). *Music licensing: The ASCAP and BMI consent decrees* (IF11463). https://www.everycrsreport.com/reports/IF11463.html

Copyright Act of 1976, 17 U.S.C. § 102. https://uscode.house.gov/us/usc/119-136/t17/s102/

Copyright Act of 1976, 17 U.S.C. § 107. https://uscode.house.gov/us/usc/119-136/t17/s107/

Copyright Act of 1976, 17 U.S.C. § 109. https://www.law.cornell.edu/uscode/text/17/109

Copyright Act of 1976, 17 U.S.C. § 115. https://uscode.house.gov/us/usc/119-136/t17/s115/

Hendrix, J. (2026, November 11). *Friends of the court flock to the OpenAI copyright fight.* Tech Policy Press. https://www.techpolicy.press/friends-of-the-court-flock-to-the-openai-copyright-fight/

Judiciary and Judicial Procedure, 28 U.S.C. § 1498. https://www.law.cornell.edu/uscode/text/28/1498

Lowenstein Sandler. (2026a, July 6). *SCOTUS ruling re-opens the door on the Copyright Office's AI training report.* https://www.lowenstein.com/news-insights/publications/client-alerts/scotus-ruling-re-opens-the-door-on-the-copyright-office-s-ai-training-report-ai

Lowenstein Sandler. (2026b, September 3). *U.S. government backs fair use for AI training in OpenAI copyright litigation.* https://www.lowenstein.com/news-insights/publications/client-alerts/us-government-backs-fair-use-for-ai-training-in-openai-copyright-litigation-intellectual-property

MediaCopilot. (2026, September 4). *Justice Department backs OpenAI in copyright fight with The New York Times.* https://mediacopilot.ai/?p=10916

MLex. (2026, November 3). *ROSS Intelligence asked US appeals court to revisit AI case over Westlaw.* https://www.mlex.com/mlex/articles/2545944/ross-intelligence-asked-us-appeals-court-to-revisit-ai-case-over-westlaw

MME. (2026, February 5). *The show is over (for the time being): OpenAI needs a licence for song lyrics.* https://www.mme.ch/en/magazine/articles/the-show-is-over-for-the-time-being-openai-needs-a-licence-for-song-lyrics

Nix v. Williams, 467 U.S. 431 (1984). https://www.law.cornell.edu/supremecourt/text/467/431

Open Markets Institute. (2025, January). *AI and market concentration* [Expert brief]. https://openmarketsinstitute.org/publications/expert-brief-ai-and-market-concentration-courtney-radsch-max-vonthun

Puglisi, B. C. (2026). *The minds that bend the machine: How to learn about AI, its bias, its power, its cost to humanity, and how I use it.* Digital Ethos. https://basilpuglisi.com/minds-that-bend-the-machine/

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM, 64*(12), 86–92. https://www.microsoft.com/en-us/research/publication/datasheets-for-datasets/

Khan, L. M. (2017). Amazon's antitrust paradox. *Yale Law Journal, 126*, 710–805. https://yalelawjournal.org/article/amazons-antitrust-paradox

Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control.* Viking.

Sidley Austin. (2025, May 28). *Generative AI meets copyright scrutiny: Highlights from the Copyright Office's Part III report.* Data Matters. https://datamatters.sidley.com/2025/05/28/generative-ai-meets-copyright-scrutiny-highlights-from-the-copyright-offices-part-iii-report/

Songtrust. (n.d.). *What kind of royalties will I earn from US radio?* https://help.songtrust.com/knowledge/what-kind-of-royalties-will-i-earn-from-us-radio

Supreme Court of the United States. (n.d.). *Rule 13: Review on certiorari, time for petitioning.* Legal Information Institute. https://www.law.cornell.edu/rules/supct/rule_13

The White House. (2026, March 20). *National policy framework for artificial intelligence: Legislative recommendations.* https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf

The White House. (2026, June 5). *Artificial intelligence in the national security enterprise* (National Security Presidential Memorandum 11). Daily Compilation of Presidential Documents, DCPD-202600387. https://www.govinfo.gov/content/pkg/DCPD-202600387/html/DCPD-202600387.htm

Thomson Reuters Enterprise Centre GmbH v. ROSS Intelligence Inc., No. 25-2153 (3d Cir. Sept. 29, 2026). https://www2.ca3.uscourts.gov/opinarch/252153p.pdf

U.S. Copyright Office. (n.d.). *Music licensing modernization.* https://www.copyright.gov/music-modernization/115

U.S. Copyright Office. (2025, May). *Copyright and artificial intelligence, Part 3: Generative AI training* (Pre-publication version). https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf

World Intellectual Property Organization. (2026). *GEMA v OpenAI, Case No. 42 O 14139/24* [Informal case summary]. https://www.wipo.int/wipolex/en/text/600712

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

Basil C. Puglisi, MPA
A Human-AI Collaboration

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.

Suggested citation: Puglisi, B. C. (2026). AI Copyright Debate and the Inevitable Knowledge Transmission Argument for LLMs and Future AI. basilpuglisi.com.