REAL SAFETY AI FOUNDATION · WORKING DRAFT, APRIL 2026

AI COPYRIGHT · VICARIOUS LIABILITY · CONTRIBUTORY INFRINGEMENT · INDUCEMENT · GROKSTER · DMCA SAFE HARBOR · SONY DOCTRINE · MEMORIZATION

The Infringing Machine

Commercial Exploitation and AI Copyright Liability

Travis Gilly
Real Safety AI Foundation

Gilly, T. (2026). The Infringing Machine: Commercial Exploitation and AI Copyright Liability. Real Safety AI Foundation, Working Draft, April 2026.

THE CORE CLAIM

AI companies know their models memorize copyrighted content. They sell access to those models. That is the liability case.

EMPIRICAL PROOF

Stanford researchers found Claude reproduces 95.8% of *Harry Potter and the Sorcerer's Stone*. Gemini reproduces 76.8%. These are verbatim extractions, not paraphrases.

WHY TRAINING DATA IS A DISTRACTION

The method of memorization is a question of factual interest. The commercial sale of a system known to reproduce that content is a question of legal consequence.

SOURCES Ahmed, Cooper, Koyejo & Liang, Extracting Books from Production Language Models, arXiv:2601.02671 (2026).

PART I

The Liability Framework

Vicarious and Contributory Infringement

VICARIOUS LIABILITY

"When the right and ability to supervise coalesce with an obvious and direct financial interest in the exploitation of copyrighted materials, liability attaches even absent actual knowledge."

Shapiro, Bernstein & Co. v. H.L. Green Co., 316 F.2d 304, 307 (2d Cir. 1963)

PRONG 1: SUPERVISION

Companies implement real-time content filtering for violence, hate speech, and malware. They have the same technical capacity for copyright. That they have not exercised it is a business decision, not a technical limitation.

PRONG 2: FINANCIAL INTEREST

$20/month subscriptions, per-token API pricing, enterprise contracts in the millions. The infringing capability is not incidental to the commercial product. It is inseparable from it.

CONTRIBUTORY INFRINGEMENT

KNOWLEDGE

Companies possess actual knowledge. Peer-reviewed studies document memorization. Plaintiffs have filed lawsuits with verbatim reproduction exhibits. Most critically: companies have invested resources in attempting to remediate it. A party cannot work to eliminate a phenomenon it does not know exists. The remediation framing is itself an admission of knowledge.

MATERIAL CONTRIBUTION

Companies train the models. Host them. Provide API endpoints and user interfaces. Process payments. Deliver the infringing output to the user's screen. Without active and ongoing participation at every stage, the infringement cannot occur.

Gershwin Publ'g Corp. v. Columbia Artists Mgmt., Inc., 443 F.2d 1159, 1162 (2d Cir. 1971)

INDUCEMENT LIABILITY · GROKSTER

A party who distributes a product with the purpose of promoting its infringing use, as shown by affirmative acts designed to foster infringement, is liable for the resulting infringement by third parties.

MGM Studios, Inc. v. Grokster, Ltd., 545 U.S. 913, 936-37 (2005)

GROKSTER INDICIA

Targeting known infringers

GROKSTER INDICIA

Failure to develop filtering tools despite capacity

GROKSTER INDICIA

Business model dependent on infringing capability

PART II

International Precedent

GEMA v. OpenAI, Munich I Regional Court (2025)

GEMA V. OPENAI

Munich I Regional Court, November 2025: memorization of protected song lyrics within ChatGPT's model parameters constitutes copyright infringement under German and EU law. Liability assigned to OpenAI as commercial operator. The text-and-data-mining exception under the EU DSM Directive does not extend to memorization or reproduction of entire works.

GEMA v. OpenAI, Munich I Regional Court, Case No. 42 O 14139/24 (Nov. 11, 2025)

The court's reasoning maps onto U.S. vicarious liability doctrine with notable precision: supervision capacity established by architectural decisions; knowledge established by dataset selection; financial interest established by commercial deployment.

EU LAW APPLIED

German Copyright Act (Urheberrechtsgesetz) §§ 15, 16, 19a

EU InfoSoc Directive 2001/29/EC, Arts. 2 and 3

EU DSM Directive, Art. 4 (text-and-data-mining exception)

PART III

The Sony Defense Fails

On two independent grounds

WHY SONY CORP. V. UNIVERSAL CITY STUDIOS, 464 U.S. 417 (1984) DOES NOT APPLY

GROUND 1: PRE-LOADED CONTENT

A VCR ships as an empty device. An AI model does not ship empty. It contains copyrighted content already compressed into its weight parameters before the user sends a single prompt.

Sony protects manufacturers of neutral tools. AI models are not neutral tools.

GROUND 2: SERVICES VS. PRODUCTS

Sony addressed products sold at arms-length. Every AI prompt flows through the company's servers. Every output is generated on company infrastructure. Real-time control is exercised continuously.

A company cannot simultaneously claim Sony protection and market sophisticated content moderation capabilities.

PART IV

Defenses and Responses

Five defenses, five failures

FIVE DEFENSES AND WHY EACH FAILS

The Guardrail Defense	By implementing controls for other content categories, companies establish supervisory capacity. Inadequacy for copyright does not negate the capacity; it establishes the choice. *Shapiro, Bernstein & Co. v. H.L. Green Co., 316 F.2d 304 (2d Cir. 1963).*
The "Just a Tool" Defense	If a tool, Sony applies and fails. If not a tool, the model's operational autonomy does not eliminate vicarious liability for outputs the company profits from. *Sony Corp. v. Universal City Studios, 464 U.S. 417 (1984).*
Transformative Use	Transformativeness is evaluated output-by-output, not system-wide. Verbatim reproduction of Harry Potter fails all four factors of 17 U.S.C. § 107. Campbell v. Acuff-Rose Music, 510 U.S. 569 (1994).
Probabilistic System	Unpredictability of specific instances does not eliminate liability for foreseeable categories of harm. Companies know their models reproduce copyrighted content. The category is entirely foreseeable. *A&M Records v. Napster, 239 F.3d 1004 (9th Cir. 2001).*
DMCA Safe Harbor	§ 512(c) protects material placed "at the direction of a user." Memorized training data is not user-uploaded content; it is incorporated by the company's own architectural decisions. Viacom Int'l v. YouTube, 676 F.3d 19 (2d Cir. 2012).

REMEDIES

INJUNCTIVE RELIEF (17 U.S.C. § 502(A))

Courts could enjoin commercial deployment until infringement is cured, require licensing as a condition of operation, mandate output filtering for specific protected works, or order regular memorization auditing. The prospect of injunction creates powerful pre-litigation settlement pressure.

DAMAGES (17 U.S.C. § 504)

Actual damages include lost licensing fees. Defendant's profits include revenue attributable to infringing capabilities. Statutory damages up to $150,000 per work infringed in cases of willful infringement under § 504(c)(2).

STRUCTURAL REMEDIES

Training data audits, output monitoring systems, licensing mandates, periodic compliance reporting. Unlike a discrete infringement, AI models continue generating outputs indefinitely across millions of interactions.

The liability case is more straightforward than prevailing discourse suggests.

KNOWLEDGE

Empirical research, litigation exhibits, and their own remediation efforts confirm it.

FINANCIAL INTEREST

Subscription fees, API pricing, enterprise contracts valued in the millions annually.

SUPERVISORY CAPACITY

Demonstrated daily by the content moderation systems they operate for other harmful content categories.

MATERIAL CONTRIBUTION

Training, hosting, delivery. Without active participation at every stage, the infringement cannot occur.

SOURCES Shapiro, Bernstein & Co. v. H.L. Green Co. (1963); Gershwin Publ'g Corp. v. Columbia Artists Mgmt. (1971); MGM Studios v. Grokster (2005).

Ahmed Ahmed, A. Feder Cooper, Sanmi Koyejo & Percy Liang, *Extracting Books from Production Language Models*, arXiv:2601.02671 (Jan. 6, 2026).

Nicholas Carlini et al., *Extracting Training Data from Large Language Models*, 30th USENIX Security Symposium (2021).

A. Feder Cooper et al., *Extracting Memorized Pieces of (Copyrighted) Books from Open-Weight Language Models*, arXiv:2505.12546 (2025).

OpenAI, *OpenAI and Journalism* (Jan. 8, 2024), https://openai.com/index/openai-and-journalism/.

EU MATERIALS

German Copyright Act (Urheberrechtsgesetz), §§ 15, 16, 19a.

EU InfoSoc Directive 2001/29/EC, Arts. 2 and 3.

EU DSM Directive, Art. 4 (text-and-data-mining exception).

A&M Records, Inc. v. Napster, Inc., 239 F.3d 1004 (9th Cir. 2001)

Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569 (1994)

Cartoon Network LP, LLLP v. CSC Holdings, Inc., 536 F.3d 121 (2d Cir. 2008)

CoStar Grp., Inc. v. LoopNet, Inc., 373 F.3d 544 (4th Cir. 2004)

GEMA v. OpenAI, Munich I Regional Court, Case No. 42 O 14139/24 (Nov. 11, 2025)

Gershwin Publ'g Corp. v. Columbia Artists Mgmt., Inc., 443 F.2d 1159 (2d Cir. 1971)

Greenman v. Yuba Power Prods., Inc., 59 Cal.2d 57 (1963)

MGM Studios, Inc. v. Grokster, Ltd., 545 U.S. 913 (2005)

N.Y. Times Co. v. Microsoft Corp., No. 1:23-cv-11195 (S.D.N.Y. Dec. 27, 2023)

Perfect 10, Inc. v. Amazon.com, Inc., 508 F.3d 1146 (9th Cir. 2007)

Shapiro, Bernstein & Co. v. H.L. Green Co., 316 F.2d 304 (2d Cir. 1963)

Sony Corp. of America v. Universal City Studios, Inc., 464 U.S. 417 (1984)

Viacom Int'l, Inc. v. YouTube, Inc., 676 F.3d 19 (2d Cir. 2012)

UNITED STATES CODE

17 U.S.C. § 107 (Fair Use)

17 U.S.C. § 502(a) (Injunctive Relief)

17 U.S.C. § 504 (Damages)

17 U.S.C. § 504(c)(2) (Statutory Damages, Willful)

17 U.S.C. § 512 (DMCA Safe Harbors)

17 U.S.C. § 512(a) (Transitory Communications)

17 U.S.C. § 512(b) (System Caching)

17 U.S.C. § 512(c) (User-Direction Storage)

17 U.S.C. § 512(c)(1) (Conditions for Safe Harbor)

17 U.S.C. § 512(c)(1)(A)-(C) (Knowledge Conditions)

17 U.S.C. § 512(d) (Information Location Tools)

EU AND GERMAN LAW

German Copyright Act (Urheberrechtsgesetz), §§ 15, 16, 19a

EU InfoSoc Directive 2001/29/EC, Arts. 2 and 3

EU DSM Directive, Art. 4