SCHOLARLY PORTABILITY IN THE AGE OF AI

Joseph W. Yockey*

ABSTRACT

For decades, the "traditional academic works" exception harmonized university intellectual property policies with academic freedom, ensuring faculty retain ownership of their scholarly expression. However, the rise of generative artificial intelligence (AI) fundamentally unsettles this balance. Modern research now generates a novel class of artifacts—including prompt libraries, conversation logs, and fine-tuned large language models (LLMs)—that defy classification under the legacy binary of "personal scholarship" versus "institutional resource." This Article addresses a critical, emerging governance gap: when a faculty member moves to a new institution, do these AI-era artifacts travel with them, or do they remain the property of the university?

The Article demonstrates that existing doctrinal frameworks—including copyright, trade secret, contract law, and data governance—fail to provide stable default rules for AI-mediated scholarship. To remedy this failure, the Article proposes a normative framework of "scholarly portability" based on functional decoupling. It argues that universities should distinguish between three categories of artifacts and then assign tailored rights to each. First, scholarly expression and personal know-how (e.g., prompts and interaction histories) should remain faculty-owned and fully portable. Second, AI operational artifacts (e.g., fine-tuned models utilizing significant institutional compute) should be portable subject to a retained institutional "shop right"—by analogy to patent shop-right principles—structured as a non-exclusive, royalty-free license that preserves institutional continuity while permitting faculty mobility. Third, regulated data (e.g., FERPA- or HIPAA-protected inputs) must remain structurally non-portable, subject to rights of reconstruction via redactions,

---

* F. Arnold Daum Chair in Corporate Law, University of Iowa College of Law. For outstanding advice, conversation, and support, the author would like to thank Diane Lourdes Dick. He is also grateful to Sydnee Flowers for her excellent research assistance.

secure access, or synthetic data. Ultimately, this framework aligns institutional stewardship with the realities of AI-mediated inquiry, preventing the restraint of scholarly knowledge while respecting non-negotiable compliance obligations.

INTRODUCTION

Universities have long balanced two commitments that occasionally sit in tension: institutional stewardship of research resources and faculty autonomy over the creation and dissemination of knowledge. For decades, that balance has been reflected in the "traditional academic works" exception common to university intellectual property policies.[1] Under this exception, faculty typically own the outputs of their scholarly expression even when those works are created with university support.[2]

Today, generative artificial intelligence (AI) unsettles this dynamic. The research process now often produces a new class of artifacts, including prompt libraries, conversation logs, system instructions, custom agents, and fine-tuned generative pre-trained transformer (GPT) models.[3] Many of these artifacts are created on institutionally financed platforms, use regulated data, or rely on models controlled or governed by third parties. This reality raises a critical question: when a faculty member leaves for another institution, can they take the AI-era artifacts they developed with them? To address this issue, this Article argues that universities should recognize a presumptive right of scholarly portability for AI-assisted research while also implementing clear guardrails for data governance, compliance, and institutional continuity.[4]

The problem at the core of this discussion is not simply one of ownership; it is one of categorization. Traditional policy categories like

1. *See* KEVIN L. SMITH, OWNING AND USING SCHOLARSHIP: AN IP HANDBOOK FOR TEACHERS AND RESEARCHERS 71–72 (2014); Donald R. Wagner et al., *Distance Education and Intellectual Property*, 85 ACADEME, May–June 1999, at 41, 43–44 (1999).

2. *See* Wagner et al., *supra* note 1, at 43.

3. For a description and discussion of GPTs, see Ivan Belcic & Cole Stryker, *What Is GPT (Generative Pretrained Transformer)?*, IBM, https://www.ibm.com/think/topics/gpt [https://perma.cc/6KML-4UR9] (last visited Jan. 30, 2026).

4. Throughout the Article, the term "portability" should be understood as capturing the bundle of separable rights available to faculty upon their departure from one institution to join another, including: a copy/export right (ability to extract the artifact); a continued-use right (ability to keep using it after departure); a reconstruction right (right to rebuild functional equivalents); or a verification right (ability to preserve provenance/reproducibility).

"expression," "data," "software," "portfolio," and "institutional resources" map imperfectly onto AI-assisted or augmented research. For example, a chat log is simultaneously a record of scholarly deliberation and an input to future model behavior, while a fine-tuned GPT model can embody years of a scholar's methodological judgment alongside datasets or staff time supplied by the university. If we were to treat all artifacts along these lines as either wholly private or wholly institutional, doing so could chill research, invite disputes at departure, and undermine reproducibility. A more nuanced framework is needed, one that preserves academic freedom and mobility, protects students and research subjects, and gives universities defensible compliance practices.

Doctrinally, these categorization challenges are intensified because at least five bodies of law and policy intersect here: copyright, trade secret, contract, data governance and privacy, and academic freedom. Copyright law remains centered on human authorship and the protectability of expression, not ideas or processes.[5] The human-authored portions of AI-assisted manuscripts fit comfortably within the traditional works exception—but prompts, logs, and agent workflows raise questions about whether they are protectable expression at all and, if so, by whom. Trade secret principles suggest that well-developed prompt libraries and the like may constitute protectable know-how if reasonable secrecy measures are maintained, yet enterprise AI tools often retain or log user inputs in ways that complicate secrecy.[6] Contracts play an outsized role in this context: enterprise terms of service for AI providers govern retention, export, and training on user content; university IP policies define "significant use of institutional resources"; and sponsored research agreements may allocate rights in models or datasets.[7] Data governance and privacy constraints may

5. *See* U.S. COPYRIGHT OFF., COMPENDIUM OF U.S. COPYRIGHT OFFICE PRACTICES §§ 302, 306, 313.3(A) (3d ed. 2021).

6. *See* Kathleen Johnson & Jonathan Madara, *Proprietary, Confidential Info, Trade Secrets, Know-How—Differences for Business Success*, FOX ROTHSCHILD (Mar. 21, 2023), https://www.foxrothschild.com/publications/proprietary-confidential-info-trade-secrets-know-how-differences-for-business-success [https://perma.cc/CKA7-NMSZ] ("Know-how generally refers to technical skills, including information, knowledge, techniques, and experiences, that are difficult to reduce to tangible form; it is often associated with trade secrets, but can include confidential or proprietary information."). *See generally* MELVIN F. JAGER & BRAD LANE, TRADE SECRETS LAW § 3:34 (2025).

7. *See, e.g.*, *Terms of Service – Enterprise*, XAI, (June 27, 2025), https://x.ai/legal/terms-of-service-enterprise [https://perma.cc/5CXQ-7BB2]; UNIV. OF IOWA, UNIVERSITY OF IOWA INTELLECTUAL PROPERTY POLICY § 30.3(b)(2)(b) (2019), https://opsmanual.uiowa.edu/administrative-

arise from institutional review board (IRB) protocols, data-management plans, and privacy focused statutes like the Family Educational Rights and Privacy Act (FERPA) or Health Insurance Portability and Accountability Act (HIPAA).[8] Finally, academic freedom norms and policies typically favor scholarly autonomy and portability, especially for materials essential to the continuation of a research agenda.[9]

Drawing on insights in these fields, this Article advances a core claim: universities should decouple three categories—scholarly expression, AI operational artifacts, and regulated data—and then design portability rules appropriate for each one. Faculty should presumptively own and be able to port their expression. This includes prompts and chat histories, since they are the functional equivalents of research notes or lab notebooks. Operational artifacts that embody material institutional investment, including bespoke fine-tuned models trained on institutional compute or staff-built agents, should be portable pursuant to a standardized transfer mechanism that preserves a shop right for the originating university. Finally, regulated data, such as student records, human-subjects data, and confidential sponsor data, should remain governed by existing IRB conditions and retention rules. These conditions and rules should include redaction and substitution procedures that allow departing scholars to continue their research without exporting protected content.

This proposed framework yields benefits across the board. For faculty, it secures a clear right to continue their research programs without reconstructing years of prompt engineering or model tuning from memory, thereby supporting scholarly mobility and collaboration. For universities, it reduces litigation and grievance risk by articulating predictable boundaries, and it strengthens compliance by segregating regulated data from portable know-how. For journals, funders, and the broader research ecosystem, the

financial-and-facilities-policies/university-iowa-intellectual-property-policy [https://perma.cc/FV23-NFNT]; *Data Sharing and Use,* UNIV. OF IOWA, https://dsp.research.uiowa.edu/types-contracts-and-model-contracts/data-sharing-and-use [https://perma.cc/X37P-GMUL] (last visited Feb. 2, 2026).

8. *See* U.S. DEP'T OF HEALTH AND HUM. SERVS., RESEARCH (Dec. 18, 2017), https://www.hhs.gov/hipaa/for-professionals/special-topics/research/index.html [https://perma.cc/D4J8-9E3V]; Longwood Rsch. Data Mgmt., *Data Management and Sharing Plan*, HARVARD MED. SCH., https://datamanagement.hms.harvard.edu/plan-design/data-management-and-sharing-plan [https://perma.cc/B2VY-G9ZH] (last visited Jan. 28, 2026).

9. *See* Donna R. Euben, *The Current Legal Landscape*, AM. ASS'N OF UNIV. PROFESSORS (May 2002), https://www.aaup.org/academic-freedom-professors-and-institutions [https://perma.cc/78TA-VW83]; Wagner et al., *supra* note 1.

framework should help establish practical disclosure and archiving standards that improve transparency without mandating public release of sensitive data.

The remainder of the Article proceeds as follows. Part I defines the scope of the challenge, detailing how the specific technical characteristics of AI artifacts strain legacy university policies and create ambiguity regarding ownership and control. Part II analyzes the complex legal landscape, examining the intersection of copyright, trade secret, contract law, data governance and privacy, and academic freedom to demonstrate why current doctrines fail to provide a clear, default rule for these new assets. Part III articulates the proposed scholarly portability framework in detail, arguing for a decoupled approach that distinguishes between expression, operational artifacts, and regulated data. This is followed by the conclusion. Ultimately, the Article's aim is to provide both the justification and the policy tools necessary to implement an institutional framework that is both faithful to academic freedom and responsive to the realities of AI-mediated research.

I. CATEGORIZATION CHALLENGES IN AI-MEDIATED RESEARCH AND SCHOLARSHIP

Understanding the pressures AI exerts on university policy first requires a brief account of the legacy categories that currently organize claims to scholarly work. From there, this Part will more fully explain the unique categorization issues that arise from AI-mediated research activities.

A. Legacy Categories

Universities have long relied on policy categories to allocate rights and responsibilities for scholarly work. Traditional intellectual property and research policies distinguish several broad types of research output, each with different default ownership and use regimes.

The first major legacy category is "scholarly expression" or "scholarly work." This category encompasses traditional works of authorship created by faculty or students in the course of academic work.[10] It includes journal

10. AM. ASS'N OF UNIV. PROFESSORS, INTELLECTUAL PROPERTY ISSUES FOR FACULTY AND FACULTY UNIONS 12, 25 (2020).

articles, books, lecture notes, artwork, musical compositions, and other writings. Under long-standing academic custom, often reflected in university policy, these scholarly and artistic works are owned by their creators rather than the institution.[11] University IP policies typically carve out such works from work-for-hire claims, recognizing them as the faculty member's intellectual property, with the university (at most) taking a nonexclusive license for internal use.[12] This custom safeguards academic freedom and the portability of a researcher's portfolio across institutions.

In contrast, "research data"—consisting of the factual results of experiments, observations, datasets, and records of research activity—is often deemed an institutional asset. Many universities assert ownership or stewardship over research data generated under their auspices. Columbia University's policy, for example, states that "[r]esearch [d]ata and other records of University Research belong to the University," subject to limited exceptions.[13] Faculty principal investigators (faculty members responsible for overseeing research projects) are viewed as custodians of data on the university's behalf.[14] The apparent rationale for this approach is that data, unlike a journal article, is not personal expression but a product of university resources and funding, often tied to sponsor agreements and compliance obligations the institution must fulfill. Consequently, when a researcher leaves, she may take copies of research data, but original datasets must usually remain with the institution or transfer only with formal approvals

11. *See id.* at 5; AM. ASS'N OF UNIV. PROFESSORS, STATEMENT ON INTELLECTUAL PROPERTY 2 (2014) [hereinafter IP STATEMENT], https://www.aaup.org/sites/default/files/aaupBulletin_StatementIntellectualPropJune5b_0.pdf [https://perma.cc/VR38-WHHA]; STANFORD UNIV., RESEARCH POLICY HANDBOOK §2.A. (2018), https://doresearch.stanford.edu/policies/research-policy-handbook/intellectual-property/copyright-policy [https://perma.cc/68SG-AE3H]; Weinstein v. Univ. of Ill., 811 F.2d 1091, 1094 (7th Cir. 1987).

> The University concedes in this court that a professor of mathematics who proves a new theorem in the course of his employment will own the copyright to his article containing that proof. This has been the academic tradition since copyright law began, see M. Nimmer, *Copyright* § 5.03[B][1][b] (1978 ed.), a tradition the University's policy purports to retain. The tradition covers scholarly articles and other intellectual property. When Saul Bellow, a professor at the University of Chicago, writes a novel, he may keep the royalties.

Id.

12. *See* IP STATEMENT, *supra* note 11.

13. *See Ownership of Data*, COLUM. UNIV., https://research.columbia.edu/content/ownership-data [https://perma.cc/3BKF-KUDV] (last visited Feb. 2, 2026).

14. *Id.*

and agreements.[15] This regime prioritizes institutional control and continuity of research records over individual portability.

Next, university policies have historically treated computer software and other patentable or commercializable innovations as a third distinct category. Software code straddles the line between copyrightable expression and functional invention.[16] Many institutions claim ownership of software developed with significant university resources or as part of research duties, as universities often do for patentable inventions.[17] Software may be treated as work-for-hire or subject to invention-assignment requirements, unless it is expressly deemed scholarly work.[18] Typically, if software is developed within the scope of employment or using substantial institutional support, the university will assert rights while providing for the option of granting the creator revenue-sharing or license rights per policy.[19] The way an institution defines "significant use" often dictates when the institution's claim is triggered.[20] Beyond customary support like salaries, offices, basic desktop computers, and library access, the use of extensive computing facilities, paid assistants, or special funding might render the resulting software an institutional work.[21] By contrast, small scripts or software analogs to academic papers may be treated as the author's property, akin to scholarly expression. For software, the line between institutional work and scholarly expression has often been blurry and

15. *Id.*

16. *See, e.g.*, Lotus Dev. Corp. v. Borland Int'l, Inc., 516 U.S. 233 (1996); Comput. Assoc. Int'l, Inc. v. Altai, Inc., 982 F.2d 693 (2d Cir. 1992). *See also* 17 U.S.C. § 102(b) (excluding "idea[s], procedure[s], process[es], system[s], [and] method[s] of operation" from copyright protection); 35 U.S.C. § 101 ("Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter . . . may obtain a patent therefor . . . .").

17. *See, e.g.*, STANFORD UNIV., *supra* note 11, at § 2.G; *Intellectual Property & Industry Research Alliance*, UNIV. OF CAL., BERKELEY, https://ipira.berkeley.edu/software-copyright-notice-and-disclaimer [https://perma.cc/R8MQ-FMSK] (last visited Jan. 29, 2026).

18. *See* IP STATEMENT, *supra* note 11, at 2; *see also Copyright Ownership and Management of Software*, IOWA ST. UNIV. POL'Y LIBR. (Sep. 5, 2013), https://www.policy.iastate.edu/softwarecopyright [https://perma.cc/D987-5RG8] (providing that non-patentable software is owned by its author(s) "unless . . . the work is a work for hire," and listing circumstances in which software "will be deemed a work for hire" and owned by the university).

19. *See, e.g.*, *Policy and Procedures for Intellectual Property*, WRIGHT ST. UNIV. (2001), https://www.wright.edu/research/technology-transfer/policy-and-procedures-for-intellectual-property [https://perma.cc/B6U4-T7SH] (describing the breakdown of inventor responsibility to assign rights and the university's reciprocal duty to distribute royalty income); STANFORD UNIV., *supra* note 11, at § 2.C.; UNIV. OF CAL., BERKELEY, *supra* note 17.

20. *See, e.g.*, STANFORD UNIV., *supra* note 11, at § 2.G.

21. *See, e.g.*, *id.*; IOWA ST. UNIV., *supra* note 18.

negotiated case-by-case. Plainly, software occupies an ambivalent space between academic publication and technical product, and university policies reflect this uncertainty.

Additionally, universities may identify certain works as categorically "institutional," either because they are commissioned by the institution or created by teams of university employees, or because they involve resources far beyond normal academic support.[22] These works might include an online course developed within university production studios, a database curated by a research center, or any project expressly designated as institutionally owned under a contract or grant. In such cases, the institution typically either owns the IP outright or reserves the right to assert ownership of the resulting copyright or patent.[23] Naturally, if a faculty member uses only normal, incidental resources, the work remains their own and would not fall within this category. Thus, the boundary between personal scholarly work and institutional work frequently turns on the degree of resource use and institutional direction. This category functions as a catch-all to protect the university's interest in what it considers enterprise-level or large-scale outputs.

Finally, falling outside of formal copyright and patent doctrines, academics also produce materials that form parts of their personal academic portfolios, including teaching materials. These portfolios often fall under the umbrella of scholarly or pedagogical works, presumed to belong to the individual by custom and tradition.[24] Absent extraordinary contractual agreements, universities generally do not claim ownership of lecture slides, assignments, or pedagogical strategies that faculty create, viewing these as part of the faculty member's professional toolkit.[25] Likewise, students retain rights to their coursework and creative projects, similarly forming a portfolio they carry with them.[26]

---

22. *See, e.g.*, STANFORD UNIV., *supra* note 11, at § 2.B.

23. *See, e.g.*, *id.*

24. *See id.* at §2.A; *Instructor-Created Course Materials*, GEO. UNIV., https://library.georgetown.edu/copyright/faculty-course-materials [https://perma.cc/7U46-YQ98] (last visited Feb. 2, 2026) ("Teaching materials prepared by instructors, such as syllabi, examination questions and answers, notes, PowerPoint and other presentation slides, and videos used in Georgetown University courses are the intellectual property of the creator of the material.").

25. *See* IP STATEMENT, *supra* note 11.

26. *Student Works*, UNIV. OF CAL., https://copyright.universityofcalifornia.edu/ownership/student-works.html [https://perma.cc/2BH3-28QW] (last visited Jan. 28, 2026); *Owning and Managing Academic Rights*, UNIV. OF MINN., https://www.lib.umn.edu/services/copyright/academic-

B. AI-Era Research Artifacts

The foregoing legacy categories—scholarly expression, research data, software, institutional resources, and personal portfolio materials—evolved in an era before generative AI and AI-driven research tools became mainstream. The categories assume a relatively clear taxonomy of scholarly products. In practice, scholars and universities have navigated this taxonomy through well-understood proxies: a journal article is "expression" (faculty-owned), a lab dataset is "data" (university-held), a piece of code is "software" (potentially university-owned), lecture notes are part of a personal portfolio (faculty-owned), and a patentable invention is an institutional resource outcome (assigned to the university under patent policy). While gray areas have always existed, the traditional framework has largely functioned well, seeing support through mutual understandings and the strong academic norms of openness and transparency in research and scholarship.

Modern AI-mediated research has introduced novel artifacts that do not fit neatly into the legacy categories. Four artifacts are especially illustrative: prompt libraries, AI conversation logs, fine-tuned AI models, and custom AI agents. Each artifact is a by-product or tool of scholarly activity in the age of large language models (LLMs), and each resists straightforward classification.

i. Prompt Libraries

Researchers and educators are amassing countless collections of prompts—carefully crafted input queries or instructions—to extract useful output from generative AI models.[27] In advanced scholarship, prompting has evolved from simple queries into a discipline of "prompt engineering," involving the design of and experimentation with complex instructions to control model behavior, reasoning, and formatting.[28] A robust prompt

ownership [https://perma.cc/N3W7-GN2N] (last visited Jan. 28, 2026).

27. *See, e.g.*, *Prompt Library*, UNIV. OF PA., https://gail.wharton.upenn.edu/prompt-library/ [https://perma.cc/9HJT-ANDJ] (last visited Feb. 2, 2026); *Teachers Reveal How AI Prompt Libraries Are Transforming Education*, YOUNGSTOWN ST. UNIV., https://online.ysu.edu/degrees/education/edd/educational-leadership/ai-prompt-libraries-in-education/ [https://perma.cc/YHN8-3GKN] (last visited Feb. 2, 2026) (finding that more than half of educators surveyed (52%) report using AI prompt libraries in their teaching workflow at least once a week).

28. SANDER SCHULHOFF ET AL., THE PROMPT REPORT: A SYSTEMATIC SURVEY OF PROMPT

library may contain hundreds of tested and optimized prompts that effectively "program" a model to perform specialized tasks like analyzing legal statutes for specific clauses or simulating historical geopolitical scenarios. Within this category, "system instructions" or "metaprompts" serve a foundational role. These are high-level directives that define an AI system or agent's persona, constraints, and operational logic.[29]

On one hand, prompt libraries look like textual expression: they are human-authored sequences of words, potentially subject to copyright as original text. Crafting an effective prompt can be a creative process, and a compilation of prompts may reflect significant scholarly expertise. In that sense, a prompt library resembles a set of research notes or teaching materials belonging to the individual scholar, which would traditionally be a component of their personal portfolio of know-how.

On the other hand, prompts serve a functional and iterative purpose that could be viewed as akin to software or data. Each prompt is a tool to perform a task with the AI, and a collection of prompts can resemble a dataset (i.e., a set of input–output examples), or even a form of source code instructing the machine. Universities might argue that a large prompt library developed in a lab, especially if used for research outcomes or built with student assistance, is a research resource or database that the university should own or at least have rights to use. The choice of classification has real implications: if treated as personal scholarly expression, a professor who compiles a prompt library could take it to a new institution or publish it freely. If treated as institutional data or software, though, the university might assert ownership or restrict dissemination—for instance, to protect a competitive advantage in grant applications or online course development.

Real-world practices already diverge. Some academics openly share prompt collections online as educational resources, while other organizations guard them closely. The ambiguity in academia can be seen in examples such as the Wharton School's "Generative AI Prompt Library,"

ENGINEERING TECHNIQUES 16–18 (2025), https://arxiv.org/pdf/2406.06608 [https://perma.cc/8FBF-H566].

29. *See Safety System Messages*, MICROSOFT LEARN, https://learn.microsoft.com/en-us/azure/foundry/openai/concepts/system-message?tabs=top-techniques [https://perma.cc/Y9U6-WQTZ] (last visited Mar. 6, 2026) (explaining that a "system message" (also called a "metaprompt" or "system prompt") provides high-priority instructions that steer the model and establish consistent role/tone/boundaries).

shared as a community resource for educators.[30] One can easily imagine a different scenario where a faculty member's unique prompt set for a funded project is instead treated as part of the project's deliverables and is thus owed to the institution.

ii. AI Conversation Logs

Scholars increasingly engage in extended dialogues with chat-based AI models (e.g., ChatGPT, Gemini, Claude) when performing research, brainstorming ideas, debugging code, or conducting analyses in natural language.[31] These sessions produce conversation logs: time-stamped exchanges of prompts and the corresponding AI-generated responses.[32]

Conversation logs are complex artifacts. They record a research process, retain snippets of human expression, and embed substantial machine-generated text. Should these chats be considered research data, akin to observational notes or intermediate results? If so, many universities would claim at least a right to access and retain them as research records, and possibly even ownership of the compilation. Institutional data policies can easily be interpreted to cover "records of University research" in any form, which would include dialogues with an AI model conducted over the course of a research project.[33]

Alternatively, one might analogize these logs to personal notes or drafts that reflect exploratory or preliminary extensions of the researcher's mind. In many disciplines, the interaction history with an LLM is the locus of intellectual discovery: it records hypothesis testing, refinement of arguments, and the synthesis of literature. In the experimental sciences, the chat log can function as a "lab notebook," documenting the provenance of code, the parameters of simulations, or the statistical reasoning behind a conclusion. Historically, lab notebooks are the property of the laboratory

30. *See, e.g.*, UNIV. OF PA., *supra* note 27.

31. Mariana Lenharo, *ChatGPT Turns Two: How the AI Chatbot Has Changed Scientists' Lives*, NATURE (Dec. 2, 2024), https://www.nature.com/articles/d41586-024-03940-y [https://perma.cc/XC6T-3TUN].

32. *See generally* Aaron Chatterji et al., *How People Use ChatGPT* 1–4 (NBER Working Paper No. 34255, 2025), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5487080 [https://perma.cc/9SE3-TMLF]; *AI Chat Log Research,* MICROSOFT, https://www.microsoft.com/en-us/research/project/ai-chat-log-research/ [https://perma.cc/3Z7K-PN9T] (last visited Feb. 4, 2026).

33. *Ownership of Data*, COLUM. UNIV., https://research.columbia.edu/content/ownership-data [https://perma.cc/2EUX-SC5D] (last visited Feb. 4, 2026).

(and thus the university). Departing faculty can take copies, but originals remain with the institution to protect patent priority, ensure compliance, and support research integrity.[34]

AI complicates this binary. The conversation a researcher conducts with an AI chatbot is neither wholly private nor fully local; it typically resides on a third-party server, and the content may be mined by the provider.[35] The result is a dual character: what a researcher views as a private scholarly interaction or personal research record could simultaneously be treated as a contribution to an external model. University policies were not designed for this interplay. Classifying chat logs purely as institutional records subject to retention, e-discovery, and FOIA/state public-record law requests risks chilling academic freedom for scholars who use AI to explore sensitive, controversial, or tentative ideas. By the same token, treating them as wholly private property could deprive the university of the ability to investigate research misconduct (e.g., data fabrication via AI) or ensure reproducibility of published results.

These considerations force a choice between surveillance and opacity, neither of which serves the long-term interests of the academic enterprise. Some organizations have reacted by imposing strict boundaries, as illustrated by well-publicized corporate bans on employee use of ChatGPT after inadvertent disclosure of proprietary code.[36] While academia does not

---

34. *See, e.g.*, CORNELL UNIV., UNIVERSITY POLICY § 4.21 (Jan. 27, 2026), https://policy.cornell.edu/policy-library/research-data-retention [https://perma.cc/2EUX-SC5D] (data retention); *Records Management at UF*, UNIV. OF FLA., https://records.uflib.ufl.edu/record-retention/research-records-data/ [https://perma.cc/LQ43-KRY9] (last visited Feb. 4, 2026).

> Personnel who leave the University may be permitted to copy their laboratory notebooks and other research data and take the copies with them, and take samples of tangible property with them, although they are required to maintain the confidentiality of the data contained within the notebooks and the tangible property pursuant to this Intellectual Property Policy. The original notebooks and other research data will remain at the University.

Id.

35. A 2025 Stanford study found that six leading AI developers, including OpenAI and Anthropic, by default feed user inputs back into their models to improve training. *Study Exposes Privacy Risks of AI Chatbot Conversations*, STANFORD REP. (Oct. 15, 2025), https://news.stanford.edu/stories/2025/10/ai-chatbot-privacy-concerns-risks-research [https://perma.cc/V3GV-TQXA]. From the AI company's perspective, a conversation log is training data—part of an industrial data pipeline.

36. For example, Samsung famously banned employees from using ChatGPT at work after sensitive source code was inputted and effectively leaked to OpenAI's servers. *See* Kate Park, *Samsung Bans Use of Generative AI Tools Like ChatGPT After April Internal Data Leak*, TECHCRUNCH (May 2,

appear to have seen a similar widely publicized leak of confidential research via AI chat, the issue is the same: conversation logs can inadvertently shuttle proprietary information outside the institution, blurring the boundary between an individual's thought process and institutionally protected information.

iii. Fine-Tuned AI Models

A particularly complex new artifact is the fine-tuned model—an AI model (often a copy of a base LLM provided by developers like OpenAI or Anthropic) that a researcher trains further using specialized data to perform targeted tasks. Fine-tuning produces a set of model "weights," which consist of vast arrays of numerical parameters that encode both the original AI's "knowledge" and the new training data.[37]

In one view, model weights are a form of software: machine-readable files that, when loaded into the AI system, cause it to operate in a particular way, much like a compiled program. A model file is useless without the underlying AI architecture, typically software code, but together they operate as an executable system. Universities might be inclined to treat a fine-tuned model as they would any software developed in a lab: if using significant resources or grant funds, the institution could claim ownership, or at least rights to it, comparable to a software tool or database created in research.

In another view, the weights are derivative data, mathematically derived from training datasets and embodying statistical traces of those inputs. From an IP standpoint, model weights are difficult to fit into traditional definitions. They are not source code written by a human (the weights are generated algorithmically during training), and they are not expressive content intended for direct human interpretation (they are largely inscrutable numbers).[38] Commentators have questioned whether model

---

2023), https://techcrunch.com/2023/05/02/samsung-bans-use-of-generative-ai-tools-like-chatgpt-after-april-internal-data-leak/ [https://perma.cc/L4BX-KF3D].

37. *See, e.g.*, *Supervised fine-tuning*, OPENAI, https://platform.openai.com/docs/guides/supervised-fine-tuning [https://perma.cc/DSN7-FXER] (last visited Feb. 4, 2026) (stating that OpenAI uses the uploaded dataset "to update the model's weights" to produce outputs like the training data, including structured/format-specific outputs). *See, e.g.*, *Fine-tune Claude 3 Haiku in Amazon Bedrock*, CLAUDE (Nov. 1, 2024), https://claude.com/blog/fine-tune-claude-3-haiku [https://perma.cc/H42H-KUX8].

38. *See generally* Ho Tung Jeremy Chan & Eduardo Veas, *Importance Estimate of Features Via*

weights are protectable at all, expressing skepticism that anyone can be considered an "author" of the weights given the lack of direct human creative choices in the specific parameter values.[39] Each weight is a numerical parameter learned through optimization to implement the model's input–output function; to the extent particular values are dictated by functional performance requirements rather than expressive choice, they resemble unprotectable facts or methods under copyright law's merger principles.[40] For that reason, a fine-tuned model comprises arguably largely unprotectable data, or at best a compilation (if analogized to a database).[41] If the weights themselves carry little or no copyright, then that undercuts the notion of the model as a traditional "software" literary work.

Yet fine-tuned models are undeniably valuable—they represent knowledge and capability. Universities and companies therefore tend to protect them through secrecy or contract.[42] A university might treat a fine-tuned model as a trade secret or confidential research material, especially if it has commercial potential or contains sensitive training data. That treatment, however, conflicts with academic norms of openness. If a fine-tuned model is considered an ordinary scholarly result, one might expect it to be shared or published (as code or data often are). If it is treated as an institutional asset, by contrast, the university may prefer to patent some aspect (where patent-eligible). More commonly, the university may license

---

Analysis of Their Weight and Gradient Profile, 14 SCI. REPS. 23532 (2024), https://doi.org/10.1038/s41598-024-72640-4 [https://perma.cc/JKJ7-KRJA];

39. *See* Mark A. Lemley & Peter Henderson, *The Mirage of Artificial Intelligence Terms of Use Restrictions*, 100 IND. L.J. 1327, 1351–52 (2025); Nuno Sousa e Silva, *Are AI Models' Weights Protected Databases?*, KLUWER COPYRIGHT BLOG (Jan. 18, 2024), https://legalblogs.wolterskluwer.com/copyright-blog/are-ai-models-weights-protected-databases/ [https://perma.cc/5QUN-EAU3].

40. Sousa e Silva, *supra* note 39.

41. *See* U.S. COPYRIGHT OFF., COPYRIGHT AND ARTIFICIAL INTELLIGENCE, PART 2: COPYRIGHTABILITY 5 (2025) [hereinafter COPYRIGHT REPORT], https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf [https://perma.cc/NXR2-382X] (describing AI models as consisting of computer code and "numerical values (or 'weights')"); 17 U.S.C. § 101 (defining "compilation" to include assembled "data"); Feist Publ'ns, Inc. v. Rural Tel. Serv. Co., 499 U.S. 340, 347–49 (1991) (holding that facts are not copyrightable; compilation copyright is limited to original selection and arrangement); U.S. COPYRIGHT OFF., COPYRIGHT IN DERIVATIVE WORKS AND COMPILATIONS 2 (2020), https://www.copyright.gov/circs/circ14.pdf [https://perma.cc/2BC2-UDP3] (noting that compilation copyright extends only to selection, coordination, and arrangement, "not to the data itself").

42. *See* Camilla A. Hrdy, *Keeping ChatGPT a Trade Secret While Selling It Too*, 40 BERKELEY TECH. L.J. 75, 76 (2025).

the model (e.g., for industry use) or deploy it internally, limiting a researcher's ability to disseminate it.

External AI providers already highlight these tensions. OpenAI's policies for enterprise customers, for example, divide rights between the user and the company.[43] OpenAI allows a user to fine-tune its proprietary GPT models with the user's data and promises exclusive use of that fine-tuned model for that user; the training data the user provides remains the user's property, and outputs belong to the user as well.[44] Crucially, however, OpenAI retains the model itself.[45] Thus, the user "owns" the configuration but not the model in the traditional sense; one cannot take the fine-tuned model to a different platform or continue using it if OpenAI terminates access.[46]

Similar conflicts arise in academic settings. Suppose a professor fine-tunes an LLM on a university high-performance computing cluster using a proprietary base model under license. Does the resulting model belong to the professor, the university, or the company? Potentially all three stakeholders have claims. If the base model license forbids certain uses or distribution, the institution must enforce those limits even if the researcher prefers open release.[47]

From the faculty perspective, a fine-tuned model embodies the scholar's

---

43. *See Enterprise Privacy at OpenAI*, OPENAI (June 4, 2025) [hereinafter *Enterprise Privacy*], https://openai.com/enterprise-privacy/ [https://perma.cc/JV74-JG97] (stating that business users "own [their] inputs and outputs" and that "fine-tuned models are for [the customer's] use alone and never served to or shared with other customers."); OPENAI, SERVICES AGREEMENT §§ 2.2, 3.3, 9.1, 11.3 (2025) [hereinafter SERVICES AGREEMENT], https://openai.com/policies/services-agreement/ [https://perma.cc/DC7H-44PL] (granting only a limited right to access/use services; reserving IP rights in the services; prohibiting reverse engineering/extraction; and providing that rights cease immediately upon termination); *Model Optimization—Fine-Tuning*, OPENAI, https://platform.openai.com/docs/guides/model-optimization [https://perma.cc/3PTM-WUU9] (last visited Jan. 16, 2026) (describing creation of fine-tuned models "in the OpenAI platform" by uploading data to OpenAI and running a fine-tuning job).

44. *See id.*

45. SERVICES AGREEMENT, *supra* note 43; *Model Optimization—Fine-Tuning*, OPENAI, https://platform.openai.com/docs/guides/model-optimization [https://perma.cc/3PTM-WUU9] (last visited Jan. 23, 2026) (describing creation of fine-tuned models "in the OpenAI platform" by uploading data to OpenAI and running a fine-tuning job).

46. SERVICES AGREEMENT, *supra* note 43.

47. *See, e.g.*, Matthew Turk, *How Stanford Researchers Attempted to Make a New ChatGPT with less than $600*, STANFORD DAILY (Apr. 2, 2023), https://stanforddaily.com/2023/04/02/how-stanford-researchers-attempted-to-make-a-new-chatgpt-with-less-than-600/ [https://perma.cc/Q6TH-3LN5] (discussing a case illustrating the research and use challenges that come when faculty rely on proprietary third-party models or data to train or develop university-sponsored AI tools).

judgment in curating data, selecting parameters, and evaluating outputs. From the institution's perspective, a fine-tuned model is a valuable research asset built with significant computing resources and external licenses. From the provider's perspective, a fine-tuned model is an extension of proprietary infrastructure. This ambiguity can complicate the portability of expertise. If a faculty member leaves, can they bring a fine-tuned model to their new lab? Under traditional rules, a university might treat the model like specialized software or a unique biological reagent, asserting ownership and preventing unapproved transfer. Conversely, if the university disclaims any interest, a valuable AI model may walk out the door without consideration of student contributions or institutional investment.

iv. Custom AI Agents

Beyond standalone models, researchers are assembling custom agents and workflows that leverage AI.[48] A professor might script an "agent" that uses an LLM in combination with other software tools to conduct literature reviews, summarize case law, or run data analyses.[49] These systems can involve significant coding and bespoke workflow design unique to the researcher's lab. Recent work at Stanford, for example, deployed a "virtual lab" of specialized AI agents—including an AI "principal investigator"—that independently designed a novel nanobody-based vaccine candidate for SARS-CoV-2.[50]

In the agentic AI context, traditional policy would likely label the underlying code as software, potentially university-owned if developed with research funds. Yet many components of an agent—its workflow design, embedded prompts, and encoded habits or preferences—are closer to a researcher's methodology or tacit know-how. A lab-built agent might even incorporate the researcher's own writing or data (for instance, by fine-

48. *See generally* Joseph W. Yockey, *Algorithmic Tenure*, GEO. J.L. & PUB. POL'Y (forthcoming 2026); J. Byun & A. Stuhlmüller, *Artificial Intelligence in Science*, OECD (Jun. 26, 2023), https://www.oecd.org/en/publications/artificial-intelligence-in-science_a8d820bd-en/full-report/elicit-language-models-as-research-tools_fec8a6ab.html [https://perma.cc/8947-JU97].

49. Hanae Armitage, *Researchers Create 'Virtual Scientists' to Solve Complex Biological Problems*, STANFORD MED. (July 29, 2025), https://med.stanford.edu/news/all-news/2025/07/virtual-scientist.html [https://perma.cc/E3KW-RNYL]; Maojun et al., *A Survey on Large Language Model-based Agents for Statistics and Data Science*, AM. STATISTICIAN (Oct. 16, 2025), https://doi.org/10.1080/00031305.2025.2561140 [https://perma.cc/6ZME-2J3C].

50. Armitage, *supra* note 49.

tuning on prior publications to answer questions in that style), further blurring the line between institutional data and personal authorship.

If a custom agent is treated simply as code, standard IP rules will presumably apply: multiple co-creators may share authorship, and if created with institutional resources, the university may claim ownership. But that framing misses the extent to which agents function as amplifications of the researcher's own capacity. As the Stanford example illustrates, agentic AI is becoming part of how scholars think and work. In this sense, a custom agent resembles a scholarly tool that one would want to take when moving institutions—like a personal library or a piece of laboratory equipment that a new employee negotiates to obtain.

Legacy categories offer no consensus as to whether a multi-component AI agent is best understood as software (to be controlled and licensed by the university) or as an extension of the scholar's mind (to remain under their control). This uncertainty can deter researchers from investing in such agents unless they know they can continue to use them, and it raises data governance questions when agents connect to external services. Institutions may assert oversight for security and compliance reasons, again positioning the agent as an institutional resource rather than a purely personal creation.

C. Blurring Lines and Misclassification Risks

The foregoing examples demonstrate that AI-mediated research artifacts blend elements of personal scholarly know-how and institutional infrastructure in unprecedented ways. Legacy categories map imperfectly onto these hybrids. The result is growing tension and uncertainty in three key areas: ownership, governance, and portability.

i. Ownership Ambiguities

In the AI era, the question of ownership of artifacts and tools is increasingly contested. A prompt library could be characterized either as a faculty member's copyrighted material or as a university-commissioned research tool. A fine-tuned model may be nominally unprotected by copyright yet valuable enough to attract claims from the researcher (who supplied intellectual input and data), the university (which provided resources and the research context), and a third-party AI platform (which provided and controls the base model).

Without clear policy, ownership assertions are ad hoc and inconsistent. When an academic leaves one institution for another, their former university might claim that fine-tuned models or prompt libraries developed there are institutional IP, prohibiting the scholar from taking or using them elsewhere without a license. Conversely, if a university denies any claim, treating everything as the scholar's, it may inadvertently allow valuable research assets built with student labor and external funding to vanish or to be appropriated by private actors without oversight.

The available doctrinal tools are ill-suited to these disputes. Copyright, patent, and trade-secret law each have limitations in this context. Copyright does not readily recognize AI-generated portions as owned by anyone, and patent doctrine currently requires human inventorship.[51] Contracts and university policies thus become the primary determinants of ownership, which in turn puts greater pressure on universities to get their categorizations right in policy language. If they do not, litigation will fill the vacuum.

The problem is compounded by the structure of legacy IP policies. The "traditional academic works" exception—often framed as covering "books, articles, and lecture notes" that authors create independently—functions as a doctrinal shield of faculty ownership.[52] Its policy rationale is grounded in academic freedom and the absence of direct institutional direction in creating such works. Yet the exception is text-centric; it protects "expression." AI artifacts such as prompts and model weights sit on the boundary of expression and utility. If a university determines that a prompt library is a "functional tool" rather than a "literary work," the exception evaporates. Moreover, if the Copyright Office refuses to register AI-assisted works or prompts, there may be no copyright for the faculty member to own, rendering the exception moot and exposing the artifact to reclassification as

---

51. U.S. PATENT & TRADEMARKS OFF., REVISED INVENTORSHIP GUIDANCE FOR AI-ASSISTED INVENTIONS (Nov. 26, 2025), https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions [https://perma.cc/Z2SV-3H7W].

> The USPTO presumes that the inventors named on the application data sheet or oath/declaration are the actual, human-being inventors. AI systems, including generative AI and other computational models, are tools used by human inventors. Like any tool, while AI systems may assist inventors, such tools do not qualify for or elevate such assistance to inventor status.

Id.

52. *See* Wagner et al., *supra* note 1.

"data" or "know-how," where institutional ownership is often the default.

ii. Governance and Control

Classification also shapes the governance questions of who decides how an AI artifact is used, shared, or monitored. If a conversation log is deemed purely personal, a faculty member may feel free to delete it or keep it confidential. However, if it is treated as research data, the institution may insist on archiving it for compliance, subjecting it to retention rules, e-discovery, or Freedom of Information Act (FOIA) requests.

AI providers themselves are already exercising governance by absorbing user inputs for model training, sometimes without users fully understanding their rights regarding the data.[53] This external reality pressures universities to issue internal guidelines. Institutions such as Columbia University now warn faculty and staff not to input confidential information into AI tools lest the university lose IP rights or violate privacy laws.[54] Guidance along those lines implicitly frames AI interactions as implicating institutional resources and data (e.g., treating a prompt to ChatGPT that contains proprietary information as a potential transfer of institutional IP).[55] Thus, data security, privacy, research-misconduct, and open-science regimes converge on these AI artifacts.

Ambiguous classification makes it difficult to know which rules apply. Are AI outputs subject to plagiarism and citation rules as authored expression, or are they unowned data that can be used freely? Are prompt libraries subject to export-control or data-sharing agreements if they contain regulated information, or are they personal notes? If a fine-tuned model contains biased training data, who bears responsibility for auditing and remediation—the individual researcher (as "author") or the institution (as owner and provider)? Classification ambiguity can yield either governance

53. *See Study Exposes Privacy Risks of AI Chatbot Conversations*, STANFORD REP. (Oct. 15, 2025), https://news.stanford.edu/stories/2025/10/ai-chatbot-privacy-concerns-risks-research [https://perma.cc/QJ5T-JXVW].

54. *Generative AI Policy,* COLUM. UNIV. OFF. OF THE PROVOST, https://provost.columbia.edu/content/office-senior-vice-provost/ai-policy [https://perma.cc/S2GM-DJCF] (last visited Jan. 30, 2026) ("[I]f Generative AI is given access to confidential information or trade secrets, the University may lose its intellectual property (IP) rights to that information and the information may be disclosed to unauthorized third parties through their independent use of the Generative AI technology.").

55. *See id.*

gaps (where no one properly oversees the artifact) or governance overreach (where multiple parties assert control, creating conflict).

iii. Portability and Academic Freedom

Of immediate concern for the purposes of this Article, the mishmash of personal and institutional elements in AI research artifacts also threatens scholarly portability—the ability of scholars to carry their ideas, methods, and tools with them and disseminate them freely. The academic enterprise relies upon circulation of knowledge; when categories like "scholarly works" are left primarily to faculty ownership, it ensures researchers can publish and spread ideas without institutional bottlenecks. If key components of research now reside in prompt libraries, model weights, or interactive scripts, an unclear regime can effectively tie the scholar's hands. A researcher who has spent years developing a custom AI agent to analyze historical texts may find that, upon moving to a new university, they are not allowed to take the agent's code or the fine-tuned model it relies on because the former institution deems them institutional products or because platform-level contracts forbid transfer.

Legacy policies exacerbate the problem by hinging ownership on "significant," "substantial," or "more than incidental" categorizations of a scholar's use of university resources. As noted earlier, incidental use is typically defined as relying on an office, library, and standard desktop computer; whereas significant use involves specialized facilities, dedicated technical staff, or extraordinary funding.[56] In the pre-AI era, a humanist writing a book used incidental resources, whereas a physicist running simulations used significant resources. AI collapses this distinction. Fine-tuning a model for humanities research may require high-performance computing resources that far exceed the "desktop" standard. Strict application of "significant use" clauses can inadvertently capture the scholarship of AI-using faculty in fields where they historically retained ownership. The digital humanist who relies on advanced AI capabilities to analyze texts may lose ownership of core research tools simply because her method requires compute, while the traditional humanist retains ownership of a manuscript produced on a laptop. This dynamic creates a potential "penalty" for innovation.

56. *See infra* text accompanying notes 19–26.

The imperfect mapping of legacy categories onto AI-era artifacts demonstrates that the current framework is inadequate. It was built for a world in which the containers of knowledge—papers, datasets, code—were separable and relatively easy to classify. Now, AI has produced amalgams of text, data, and machine intelligence that defy those categorizations. The result could be over-regulation (treating everything as institutional property, which would discourage use of AI tools) or under-protection (treating everything as personal and informal, which allows important assets to slip away without oversight). Neither is tenable. Academia must therefore rethink and refine its classifications so that the aims behind legacy categories—balancing academic freedom, incentives for innovation, and institutional interests—are preserved in the AI context.

II. INTERSECTING DOCTRINAL FRAMEWORKS

As a first step toward addressing scholarly portability and related issues, it is important to recognize that AI-mediated research artifacts sit at the intersection of multiple legal doctrines rather than fitting neatly within any one of them. Copyright, trade secret, contract law, data governance and privacy, and academic freedom each speak, often only partially, to questions of ownership, control, and portability. This Part sketches the core principles of each framework and shows why, taken together, they leave universities without a clear default rule for AI-era scholarly artifacts. The doctrinal landscape instead pushes crucial classification choices back to institutional policy and bargaining, with significant consequences for academic freedom and scholarly mobility.

A. Copyright Law

U.S. copyright law begins with the premise that only humans can be authors.[57] The Constitution and Copyright Act contemplate human creators, and the United States Copyright Office has consistently refused to register works produced without human authorship, including recent AI-generated images.[58] Courts have followed suit, as in the much-discussed "monkey

57. *See* Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53, 57 (1884); *In re* Trade-Mark Cases, 100 U.S. 82, 94 (1879) (holding that copyright law only protects "the fruits of intellectual labor . . . founded in the creative powers of the mind.").

58. *See, e.g.*, Letter from Robert J. Kasunic to Van Lindberg regarding Registration of Zarya of

selfie" case, *Naruto v. Slater*, where a non-human photographer could not claim copyright.[59] Current guidance therefore treats AI-assisted works as protectable only to the extent a human actor shapes the expressive aspects of the work; purely machine-generated elements must be disclaimed in any registration.[60]

The familiar idea-expression dichotomy further narrows the scope of protection. Section 102(b) of the Copyright Act codifies the principle that copyright does not extend to ideas, methods, or processes, but rather only to the particular expression of those ideas.[61] That distinction maps awkwardly onto AI-era artifacts.

Prompts, system instructions, and agent workflows are written in natural language, but they function as operational commands—as methods for coaxing behavior from a model—rather than as traditional literary expression. Some prompts are elaborate enough to possess original wording that could be protected against verbatim copying.[62] But many others look more like uncopyrightable "methods of operation."[63]

Conversation logs likewise mix categories. A log contains human inquiries and commentary (potentially copyrightable if sufficiently original), non-copyrightable AI output, and often factual or functional material that falls outside copyright altogether. At most, a researcher might hold a narrow compilation right in the selection and arrangement of a lab-notes-like transcript, while the underlying conversational "method" remains free for anyone to use.

Applied to scholarly portability, the application of copyright therefore

---

the Dawn (U.S. Copyright Off. Feb. 21, 2023), https://www.copyright.gov/docs/zarya-of-the-dawn.pdf [https://perma.cc/NNJ9-E6HS].

59. *See generally* Naruto v. Slater, 888 F.3d 418 (9th Cir. 2018); *accord* Camila Domonoske, *Monkey Can't Own Copyright to his Selfie*, *Federal Judge Says*, NPR (Jan. 7, 2016), https://www.npr.org/sections/thetwo-way/2016/01/07/462245189/federal-judge-says-monkey-cant-own-copyright-to-his-selfie [https://perma.cc/9GE2-BZVT].

60. COPYRIGHT REPORT, *supra* note 41, at 8; U.S. COPYRIGHT OFF., COMPENDIUM OF U.S. COPYRIGHT OFFICE PRACTICES §§ 306, 313.2 (2021).

61. 17 U.S.C. § 102(b) ("In no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in such work.").

62. The U.S. Copyright Office has acknowledged that possibility. *See* COPYRIGHT REPORT, *supra* note 41, at 14–15; Mario Herger, *AI Prompts and Output Can Be Protected by Copyright Under Certain Circumstances*, TECHNOPHILOSOPH (Feb. 7, 2025), https://technophilosoph.com/en/2025/02/07/ai-prompts-and-output-can-be-protected-by-copyright-under-certain-circumstances/ [https://perma.cc/94DN-C3PA].

63. 17 U.S.C. § 102(b).

remains unclear in many situations involving AI. Human-authored portions of AI-assisted manuscripts, prompts, and notes—where the researcher exercises creative judgment in drafting, editing, or curating—fit comfortably within the traditional "scholarly works" exception that most universities use to leave authorship with faculty. The United States Copyright Office suggests they could receive legal copyright protection.[64] Pure AI output without significant human contribution, however, is unlikely to enjoy protection under current doctrine.[65] Leading commentators have noted that the human-authorship requirement is under conceptual strain in the AI era, with some predicting that truly autonomous systems may eventually force a reconsideration of authorship criteria.[66] For now, though, copyright supplies only a patchwork: it protects human expression at the margins of AI interaction when the extent of human contribution is significant "enough" (a level that remains undefined at present). But, it does not protect the model's output as such, and leaves key artifacts—prompts, workflows, and weights—either unprotected or in a state of uncertainty requiring a fact-specific, case-by-case analysis.[67] That patchwork neither resolves institutional-faculty disputes nor provides a stable or predictable rule for porting AI-era research tools across institutions.

B. Trade Secret Law

Where copyright protects expression, trade secret law protects valuable information that is kept secret. Under the Defend Trade Secrets Act and parallel state law, any technical or business information—formulas,

64. *See* COPYRIGHT REPORT, *supra* note 41, at 1 ("Neither the use of AI as an assistive tool nor the incorporation of AI-generated content into a larger copyrightable work affects the availability of copyright protection for the work as a whole.").

65. *See id.*

66. *See generally* Zachary L. Catanzaro, *Beyond Incentives: Copyright in the Age of Algorithmic Production*, 13 N.Y.U. J. INTELL. PROP. & ENT. L. 1 (2023).

67. *See* COPYRIGHT REPORT, *supra* note 41, at 2.

> For a work created using AI, like those created without it, a determination of copyrightability requires fact-specific consideration of the work and the circumstances of its creation. Where AI merely assists an author in the creative process, its use does not change the copyrightability of the output. At the other extreme, if content is entirely generated by AI, it cannot be protected by copyright. Between these boundaries, various forms and combinations of human contributions can be involved in producing AI outputs.

Id.

compilations of data, software, or prompts—can qualify for protection as a trade secret if it (a) derives independent economic value from not being generally known, and (b) is subject to reasonable measures to maintain secrecy.[68] In principle, many AI-era scholarly artifacts meet this definition. A carefully engineered prompt library that consistently produces superior outputs or a fine-tuned model with unique capabilities can plainly confer a competitive advantage to a lab or research center. If access is restricted and the material is treated as confidential, trade secret protection is seemingly available.

In practice, however, secrecy is difficult to maintain in the AI context. Public or consumer-facing AI services may log user inputs by default and reserve the right to use those prompts and outputs to improve the model.[69] From a trade secret perspective, entering a confidential prompt set or proprietary dataset into such a system without ironclad confidentiality terms is akin to disclosure to a third party. Commentators have warned that there should be no expectation of privacy when using public AI tools, and that failing to use reasonable measures—such as avoiding public web interfaces or insisting on enterprise terms that disable training—can destroy trade secret status.[70]

Universities therefore face a tension between the open-science norm and trade secret logic. In some settings—technology transfer, industry collaborations, or pre-publication work on potentially patentable AI tools—short-term secrecy is attractive. Labs may wish to treat prompt libraries, model weights, or training datasets as confidential know-how until a patent filing or publication. Doing so requires operational discipline: using secure local or on-premises models where possible, contracting for enterprise AI services that commit not to use inputs for training, labeling materials as

68. 18 U.S.C. § 1839(3); *see also, e.g.*, WASH. REV. CODE § 19.108.010(4) (2025).

69. *See* Nikki Goth Itoi, *Study Exposes Privacy Risks of AI Chatbot Conversations*, STANFORD REP. (Oct. 15, 2025), https://news.stanford.edu/stories/2025/10/ai-chatbot-privacy-concerns-risks-research [https://perma.cc/ZTC7-JK5Z]; *Privacy Policy*, OPENAI, (June 27, 2025), https://openai.com/policies/row-privacy-policy [https://perma.cc/V4AH-GAJA] (stating that OpenAI may use user-provided content for model training purposes but allows users to opt out).

70. Alexandra P. Moylan & Alisa L. Chestler, *OpenAI Updates Usage Policies: Key Considerations and Next Steps for Organizations Deploying AI*, BAKER DONELSON (Nov. 18, 2025), https://www.bakerdonelson.com/openai-updates-usage-policies-key-considerations-and-next-steps-for-organizations-deploying-ai [https://perma.cc/DP9R-2C3Z] ("[T]he leakage of sensitive confidential trade secrets or corporate data remains an ongoing risk for all companies. There should be no expectation of privacy when using public/open tools, and exposing confidential data is a significantly underappreciated organizational risk.").

confidential, and limiting access. Once that veil is lifted through publication, open sharing, or casual use of public AI interfaces, trade secret protection evaporates.

For purposes of scholarly portability, trade secret status can justify insisting that certain artifacts remain with the originating institution or transfer only under nondisclosure agreements. Thus, trade secrets uneasily coexist with academic norms that expect eventual disclosure of methods and tools. Accordingly, trade secret law does not eliminate the need for clear institutional policies about when AI-era artifacts are meant to be proprietary, when they are destined for open dissemination, and what conditions attach when a faculty member moves to a new institution.

C. Contract Law

Contract law does a lot of lifting in this space. Three sets of agreements are especially salient: (a) terms of service with AI providers, (b) university intellectual property and employment policies (including "significant use" provisions), and (c) sponsored research and collaboration agreements. Taken together, these agreements will often do more to determine ownership and portability than background IP doctrine.

First, platform terms of service govern the relationship between researchers and AI providers. Provider terms vary sharply by tier: consumer terms may permit training-use by default (subject to opt-out), whereas enterprise/API terms often prohibit training-use absent affirmative customer consent.[71] These contracts formalize two points developed above: they concede whatever copyright might exist in prompts and curated outputs to the user, but they also authorize the provider to reuse those artifacts internally. Importantly, the terms typically require users to promise that they have the necessary rights, licenses, or permissions to any content they

71. *See Enterprise Privacy*, *supra* note 43. *See also* SERVICES AGREEMENT, *supra* note 43, at §§ 2.2, 3.3, 9.1, 11.3; *Terms of Use*, OPENAI (Jan. 1, 2026), https://openai.com/policies/row-terms-of-use/; *Microsoft Copilot Terms of Use*, MICROSOFT, https://www.microsoft.com/en-us/microsoft-copilot/for-individuals/termsofuse [https://perma.cc/ZB7P-N3KS] (last visited Mar. 29, 2026). *See, e.g.*, *Copilot Terms of Service*, *Microsoft Service Agreement*, MICROSOFT, (July 1, 2025) https://www.microsoft.com/en-us/servicesagreement [https://perma.cc/GUF2-PN4Q]. *See also Expanded Legal Protections and Improvements to Our API,* ANTHROPIC (Dec. 19, 2023), https://www.anthropic.com/news/expanded-legal-protections-api-improvements [https://perma.cc/NVH9-TCZ8].

upload.[72] A faculty member who pastes confidential sponsor data or student work into such a system may therefore be breaching not only trade secret duties but also contractual promises to the sponsor or university.

Second, university IP and employment policies mediate ownership between faculty and the institution. Returning to the discussion above, faculty generally own their scholarly writings unless there is a specific work-for-hire agreement. At the same time, policies often assert institutional rights in certain categories: commissioned works, patentable inventions, software or databases created with "significant use" of university resources, or outputs covered by sponsored research obligations. Policies at institutions like Stanford University illustrate the pattern: normal salary, office space, and library use do not trigger institutional ownership, but extensive use of laboratories, specialized computing infrastructure, or dedicated staff often does.[73] AI-era research complicates this issue. Fine-tuning a model or building a research agent may require high-performance computing clusters or staff support that look like "significant use" even in disciplines, such as the humanities, where faculty historically owned their work. If universities interpret "significant use" broadly, they will risk sweeping AI-mediated scholarship into institutional ownership simply because it depends on expensive compute. Such allocation would penalize methodological innovation.

Third, sponsored research agreements add another contractual layer. Under the Bayh-Dole Act, universities may elect to retain title to patentable inventions arising from federally funded research, but *Stanford v. Roche* famously held that those rights are not automatic and depend on valid assignment contracts with inventors.[74] Industry-funded projects raise additional issues. Sponsors often seek options or licenses to inventions, rights to confidential data, or some control over publication. Many universities, following principles endorsed by the American Association of University Professors (AAUP), insist on preserving faculty freedom to publish subject only to short delays for patenting or removal of sponsor

72. *See, e.g.*, *Terms of Use*, OPENAI (Jan. 1, 2026), https://openai.com/policies/row-terms-of-use/; *Microsoft Copilot Terms of Use*, MICROSOFT, https://www.microsoft.com/en-us/microsoft-copilot/for-individuals/termsofuse [https://perma.cc/ZB7P-N3KS] (last visited Mar. 29, 2026).

73. *See, e.g.*, STANFORD UNIV., *supra* note 11, at § 2.G.

74. Bd. of Trs. of the Leland Stanford Junior Univ. v. Roche Molecular Sys., 563 U.S. 776, 790 (2011). *See also* 37 C.F.R. § 401.14(f)(2) (2026).

trade secrets.[75] AI-era projects—such as building a custom model or dataset for an industry partner—will be governed by these negotiated terms. An agreement might, for example, require that a fine-tuned model be delivered to the sponsor as a confidential deliverable, or grant the sponsor broad rights in project data.

From the standpoint of scholarly portability, the upshot is sobering. A faculty member's ability to "take" a model, prompt library, or AI agent to a new institution may be constrained not by copyright doctrine but by a web of contractual commitments, including university policies, sponsor agreements, and AI-provider terms. When faculty move, portability often must be negotiated through inter-institutional agreements, material-transfer agreements, or explicit licenses.[76] Without such arrangements, a crucial AI-era research tool may either remain locked at the originating institution or walk out the door under a cloud of unresolved claims.

D. Data Governance and Privacy

Even when IP and contracts law would permit portability, data-governance and privacy rules may not. AI-mediated artifacts frequently contain human-subject, health, educational, or confidential third-party information. In those cases, IRB rules, HIPAA, FERPA, sponsor confidentiality clauses, and funding-agency data-sharing mandates introduce potential constraints on where data can be stored, how it can be processed by AI tools, and whether it can travel with a departing researcher.

Human-subjects protections and IRB oversight impose baseline duties

---

75. *See, e.g.*, *Guide for Industry*, NW. UNIV. OFF. FOR RSCH., https://sponsoredresearch.northwestern.edu/agreements/industry-guide.html [https://perma.cc/3RT9-QSLP] (last visited Mar. 6, 2026); AM. ASS'N UNIV. PROFESSORS, RECOMMENDED PRINCIPLES TO GUIDE ACADEMY-INDUSTRY RELATIONSHIPS (2014), https://www.aaup.org/reports-publications/aaup-policies-reports/topical-reports/recommended-principles-guide-academy [https://perma.cc/3U3X-MLRC].

76. *See, e.g.*, PRINCETON UNIV. OFF. OF THE DEAN OF THE FACULTY, CHECKLIST FOR DEPARTURE OF FACULTY AND RESEARCHERS 4 (2025), https://dof.princeton.edu/document/3161 [https://perma.cc/WQH2-8RXR] (stating that transferring "materials . . . to another institution may be subject to third party obligations or require the execution of a Material Transfer Agreement" and that transfer of data may require completion or revision of a data transfer/use agreement with a sponsor or collaborator); *IP Agreement Flowchart*, GEO. UNIV. OFF. OF TECH. COMMERCIALIZATION, https://otc.georgetown.edu/agreement-flowchart/ [https://perma.cc/NJ6X-LXQK] (last visited Mar. 6, 2026) (explaining that where IP is "jointly owned between the institutions," the parties "typically endeavor to enter into an inter-institutional agreement to jointly manage the intellectual property," and describing "License Agreements" for explicit use of technology).

of confidentiality, data minimization, and protocol adherence.[77] If a model is trained on interview transcripts or other identifiable human-subject data, researchers generally cannot upload those materials to public AI services or move them to a new institution without IRB approval, and often they will need a new data-use agreement between institutions. Consent forms may specify whether AI tools will be used and what level of data use is permitted. Some universities now explicitly warn that feeding research data into external AI tools is tantamount to sharing it with a third party.[78]

HIPAA and FERPA add sector-specific restrictions. HIPAA tightly regulates disclosure of protected health information (PHI).[79] Consumer AI services like ChatGPT are generally not HIPAA-compliant for covered-entity use because the vendor will not enter into a HIPAA-required business associate contract for that product, making entry of PHI an impermissible disclosure.[80] Similarly, FERPA prohibits the disclosure of education records—such as graded assignments, advising notes, or identifiable classroom interactions—to third parties without consent, and public AI tools generally will not qualify as "school officials" under FERPA's exceptions absent a FERPA-compliant contractor arrangement.[81] As a result, faculty cannot simply export student-containing conversation logs, AI-graded assignments, or analytics-driven datasets when they change institutions. Those records remain under the university's custodial control and must be handled according to its retention and destruction policies.[82]

Moreover, confidentiality obligations to sponsors and data providers can constrain AI usage and portability. Datasets or materials provided under

77. 45 C.F.R. § 46.111(a)(7) (2026).

78. *See, e.g.*, *Generative AI Policy*, COLUM. UNIV. OFF. OF THE PROVOST, https://provost.columbia.edu/content/office-senior-vice-provost/ai-policy [https://perma.cc/S2GM-DJCF] (last visited Mar. 6, 2026).

79. 45 C.F.R. § 164.502(a) (2024).

80. Moylan & Chestler, *supra* note 70; *see also* 45 C.F.R. § 164.504(e) (2024).

81. 20 U.S.C. § 1232g (2013); 34 CFR §§ 99.3, 99.30 (2004); *see also* UNIV. OF WISC.-MADISON, OFF. OF THE REGISTRAR, FERPA AND ARTIFICIAL INTELLIGENCE (AI), https://registrar.wisc.edu/ferpa-and-artificial-intelligence-ai [https://perma.cc/R5SU-LWQC] (last visited Mar. 6, 2026) ("Sharing sensitive, restricted, or otherwise protected institutional data, including that covered by FERPA, with generative AI tools is prohibited under UW–Madison policy unless the tool has undergone appropriate review.").

82. *See generally* U.S. DEP'T OF EDUC., PRIVACY TECH. ASSISTANCE CTR., BEST PRACTICES FOR DATA DESTRUCTION (2019) (explaining that FERPA generally leaves internal record disposition to institutions (subject to safeguarding requirements) and that schools often establish record-retention policies, including time frames for eventual destruction).

nondisclosure agreements—such as proprietary code, customer data, or restricted government datasets—may be used only for specified purposes and cannot be disclosed to additional parties, including AI platforms, without consent.[83] Agreements often require destruction or return of such data at project end.[84] A faculty member who leaves may have no unilateral right to take those datasets or models trained on them, and any attempts to do so could trigger breach-of-contract or trade secret claims.

Finally, agency data-management and sharing policies pull in the opposite direction by requiring some forms of portability and public access. The NIH Data Management and Sharing Policy, for example, expects investigators to make scientific data available after publication, consistent with privacy and security obligations.[85] NSF and other funders adopt similar expectations.[86] For AI-era artifacts, that may mean sharing de-identified training datasets or code through repositories or controlled-access mechanisms. These mandates underscore that decisions about whether an AI-era artifact "belongs" to a particular researcher or institution are nested within broader public obligations. Faculty cannot, for instance, simply convert a federally funded dataset into a proprietary asset of a new private venture if the grant requires public sharing.

E. Academic Freedom

Overlaying these positive-law regimes is the normative framework of academic freedom. While not itself a comprehensive legal doctrine, academic freedom—articulated most prominently in American Association

83. *See, e.g.*, CURATORS OF UNIV. OF MO., MASTER SPONSORED PROJECT AGREEMENT §§ 8.2, 8.5 (2022), https://docs.research.missouri.edu/ospa/MU_Master_Agreement_Template.pdf [https://perma.cc/2878-CLA5] (confidential information may be used only for the project; may not be disclosed outside the recipient's organization without prior written approval; and upon request must discontinue use and return or destroy).

84. *See id.* at § 8.5.

85. *See generally Final NIH Policy for Data Management and Sharing*, NAT'L INSTS. OF HEALTH (Jan. 25, 2023), https://grants.nih.gov/policy-and-compliance/policy-topics/sharing-policies/dms [https://perma.cc/NNU3-DG5G] (requiring data sharing no later than publication, subject to privacy limitations); EXEC. OFF. OF PRESIDENT BIDEN, OFF. OF SCI. AND TECH. POL'Y, ENSURING FREE, IMMEDIATE, AND EQUITABLE ACCESS TO FEDERALLY FUNDED RESEARCH (Aug. 25, 2022) [hereinafter FEDERALLY FUNDED RESEARCH], https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/08/08-2022-OSTP-Public-Access-Memo.pdf [https://perma.cc/KYY3-BKXH].

86. NAT'L SCI. FOUND., PROPOSAL & AWARD POLICIES & PROCEDURES GUIDE § XI.D.4 (2024), https://www.nsf.gov/policies/pappg [https://perma.cc/MQE3-3C79] (mandating dissemination of research data); Federally Funded Research, *supra* note 85.

of University Professors (AAUP) statements and incorporated into many university faculty handbooks—functions as both a constraint on and a lens for interpreting copyright, trade secret, contract law, and data-governance rules in the university setting.

The AAUP's 1940 *Statement on Academic Freedom and Tenure* provides that faculty are "entitled to full freedom in research and in the publication of the results, subject to the adequate performance of their other academic duties."[87] Subsequent AAUP reports explicitly link IP and academic freedom, warning that aggressive institutional claims to faculty creations will risk converting independent scholars into mere employees and undermining their control over how their work is used and disseminated.[88] Many universities echo these principles in policy language, committing to not sign research agreements that give sponsors veto power over publication, and structuring IP regimes so that faculty retain copyright in scholarly and artistic works, with the institution receiving at most a nonexclusive license for internal use.[89]

In ambiguous cases, academic freedom functions as a tiebreaker. When it is unclear whether a particular AI-era artifact is best characterized as a "traditional academic work" (presumptively faculty-owned) or an institutional resource, principles of scholarly autonomy counsel caution before reclassifying it as institutional property solely because it has commercial potential or relies on expensive infrastructure. Likewise, when interpreting "significant use" clauses or negotiating sponsored research terms, universities attentive to academic freedom ought to favor arrangements that allow faculty to publish results, reuse tools in future research, and continue their research programs when they change institutions, subject only to genuine compliance obligations. The AAUP's 2013 report on "Defending the Freedom to Innovate," for example,

87. AM. ASS'N OF UNIV. PROFESSORS, 1940 STATEMENT OF PRINCIPLES ON ACADEMIC FREEDOM AND TENURE WITH 1970 INTERPRETIVE COMMENTS, POLICY DOCUMENTS AND REPORTS 13 (11th ed. 2014).

88. *Statement on Intellectual Property*, AM. ASS'N OF UNIV. PROFESSORS (2013), https://www.aaup.org/reports-publications/aaup-policies-reports/policy-statements/statement-intellectual-property [https://perma.cc/C3FW-H8FH].

89. *See, e.g.*, MASS. INST. OF TECH., POLICIES AND PROCEDURES § 13.1.4; COLUM. UNIV., FACULTY HANDBOOK § 4.3 (2025) (Obligations and Responsibilities) (providing that an outside party may not be given the power to censor, exercise an effective veto over, or unreasonably delay publication or dissemination of project results), https://facultyhandbook.columbia.edu/content/officers-instruction/obligations-and-responsibilities [https://perma.cc/JBR7-ND5H].

criticizes policies that claim sweeping institutional rights in online course content or research outputs and urges approaches that rely on licenses rather than automatic assignments.[90]

Academic freedom does not, of course, license violations of FERPA, HIPAA, IRB requirements, or contractual confidentiality. The AAUP itself recognizes that research must comply with legal and ethical obligations to subjects and data providers.[91] But academic freedom does constrain how those obligations are implemented. It counsels against using privacy or trade secret claims as pretexts to suppress inconvenient findings, it argues for publication-friendly sponsored research terms even when industry funders would prefer tighter control, and it supports faculty efforts to study, critique, and share AI systems, including those developed by their own institutions or sponsors.

In the AI context, embracing academic freedom has three primary implications. First, the principles of academic freedom suggest that universities should be wary of treating AI tools and artifacts primarily as proprietary software assets, particularly when they encode a faculty member's scholarly methods and are central to that scholar's research trajectory. Second, these principles support designing ownership and portability rules that enable faculty to continue using and refining their AI-based tools across institutional moves. Third, these principles underscore the importance of faculty governance in drafting and applying AI-related IP and data policies so that decisions about classification, shop-rights, and portability are not left solely to administrative or commercial actors.

* * * * *

Across the five frameworks listed above, several broad conclusions emerge. Copyright draws a narrow circle around human expression and leaves most AI-era artifacts either unprotected or cloaked in uncertainty and

90. *Defending the Freedom to Innovate: Faculty Intellectual Property Rights after Stanford v. Roche*, AM. ASS'N OF UNIV. PROFESSORS (June 2013), https://www.aaup.org/reports-publications/aaup-policies-reports/topical-reports/defending-freedom-innovate-faculty [https://perma.cc/T96U-J8TX].

91. *Statement on Professional Ethics*, AM. ASS'N OF UNIV. PROFESSORS (2009), https://www.aaup.org/reports-publications/aaup-policies-reports/policy-statements/statement-professional-ethics [https://perma.cc/SRL3-79UW] (stating that professors "respect the confidential nature of the relationship between professor and student" and that, as members of an academic institution, professors "observe the stated regulations of the institution," so long as those regulations do not contravene academic freedom).

unpredictability. Trade secret law offers protection only if secrecy is actively maintained, which sits uneasily with academic norms of openness and is difficult to reconcile with public AI services. Contract law—through provider terms, university policies, and sponsored research agreements—does most of the heavy lifting in allocating rights, but in a highly contingent, institution-specific way. Data-governance and privacy rules impose hard limits on the movement and use of artifacts containing regulated information. Finally, academic-freedom principles push back against limits on transparency or sharing and urge classifications that preserve scholarly autonomy.

Going forward, what these regimes do not supply is a coherent, shared portability process or default for AI-mediated research artifacts. The next Part turns to that task, arguing for a portability framework that decouples scholarly expression, operational AI artifacts, and regulated data. Additionally, the framework aligns institutional policy with both compliance obligations and the core commitments of academic freedom.

III. DESIGNING A SCHOLARLY PORTABILITY FRAMEWORK

The preceding Parts show that AI-mediated artifacts do not fit comfortably within legacy categories of "scholarly expression," "data," "software," or "institutional resources," and that existing doctrinal frameworks provide no stable default rule for allocating rights or responsibilities. Universities are thus left to improvise in the moment of conflict—for example, as a faculty member seeks to move a fine-tuned model or custom agent to a new institution—often with little guidance beyond indeterminate "significant use" clauses or ad hoc interpretations of contractual agreements.

This Part advances a different approach. The core proposal is a portability framework that decouples AI-era research outputs into three analytically distinct categories and then assigns a tailored portability rule to each. The aim is not to displace copyright, trade secret, contract law, or data-governance doctrine, but to give institutions and faculty a coherent way to classify and manage AI artifacts that is consistent with those doctrines and with academic-freedom norms.

A. Decoupling as an Organizing Principle

The central insight animating the proposed portability framework is that AI-mediated research artifacts are composites. A single project may generate human-authored manuscripts and prompts, machine-generated text, fine-tuned model weights, custom agents, and datasets that include protected information. Treating this entire bundle as a monolithic "work" for ownership purposes obscures the fact that different components raise very different legal and normative concerns.

Decoupling offers a way to respect those differences. Instead of forcing each artifact into the same box—and then deciding whether that box is "faculty-owned" or "institutional"—the proposed framework asks first which of the following three categories best describes the artifact at issue:

(1) *Category One, Scholarly expression and personal know-how*: human-authored materials and interaction histories that function as the modern equivalent of methods notes, lab notebooks, or teaching materials;

(2) *Category Two, AI operational artifacts*: models, agents, and workflows that embody significant institutional investment and that are used as durable research infrastructure; or

(3) *Category Three, Regulated data*: any information subject to external or internal governance regimes.

Once an artifact is placed in one of these categories, the relevant portability rule follows. Category One artifacts are presumptively portable under a strengthened "traditional academic works" norm. Category Two artifacts are portable subject to a standardized transfer mechanism that preserves a shop right—a non-exclusive, royalty-free license—for the originating university. Category Three artifacts are not themselves portable, but the framework requires institutions to support redaction, substitution, and reconstruction strategies that allow research programs to travel without exporting regulated data.

Decoupling has three advantages. First, it forces *ex ante* clarity: policies describe which artifacts fall where, reducing the scope for opportunistic reclassification when a particular model or agent becomes commercially

valuable. Second, it aligns institutional practice with the underlying doctrinal landscape. As Part II above demonstrated, copyright, trade secret, contract law, and data-governance rules already differentiate among expression, tools, and regulated information. The proposed categories formalize those distinctions in university policy. Third, decoupling gives content to academic-freedom commitments by specifying which aspects of an AI-augmented research program must remain under faculty control and which may justifiably be treated as shared infrastructure or tightly constrained data.

B. Category One, Scholarly Expression and Personal Know-How

The first portability category encompasses human-authored materials and interaction histories that primarily reflect a scholar's intellectual labor and are essential to the continuity of their research agenda. It includes:

- Manuscripts, articles, books, lecture notes, syllabi, and other traditional forms of scholarly expression;
- Prompt libraries, system instructions, and meta-prompts developed by the scholar for research or teaching;
- AI conversation logs and chat histories to the extent they document the scholar's questions, hypotheses, and reasoning; and
- Pedagogical materials, including AI-assisted exercises and simulations, that form part of the instructor's personal teaching toolkit.

These artifacts may incorporate AI-generated text or rely on interactions with a proprietary platform, but their core function is to capture the scholar's methods, ideas, and arguments. They are the digital heirs of lab notes and the file cabinet full of draft articles and lecture plans.

For Category One artifacts, the default rule should be presumptive faculty ownership and full portability, subject only to pre-existing contractual commitments or legal requirements that allocate specific rights (e.g., a publisher's license in a manuscript). This default would also be subject to compliance obligations associated with embedded regulated data.

University IP policies should accordingly be revised to make explicit that the "traditional academic works" exception covers not only manuscripts and lecture notes, but also prompt libraries, system instructions, and AI-era methods notes and chat histories, barring specific contractual exceptions.

The portability consequence is straightforward. When a faculty member moves institutions, she may take her prompt library, chat histories (subject to provider access and retention policies), and related methods notes, copy them, and continue to use them in future research and teaching. The originating institution retains no claim to control or veto that portability, beyond any non-exclusive rights it already holds to use those materials internally.

Several considerations support this rule. First, Category One artifacts are integral to academic freedom and scholarly identity. They encode the scholar's approach to framing questions, structuring arguments, and deploying AI as a research assistant. To deny portability would be tantamount to preventing a chemist from taking her notes or a historian from taking her annotated archives. Because modern scholarship increasingly depends on iterative interaction with AI, locking those interactions to a particular institution would chill the adoption of AI tools and distort the market for academic labor.

Second, Category One artifacts typically involve minimal institutional investment beyond salary and ordinary support. While AI tools may run on university-paid licenses, the marginal cost of a particular prompt library or set of chat histories is low. Heavy-handed institutional claims over such materials would therefore provide little additional protection of legitimate institutional interests while substantially burdening faculty mobility.

Third, treating Category One artifacts as portable is consistent with the positive-law landscape described in Part II. Copyright protects human-authored expression in prompts, notes, and the human contributions in chat logs; and existing university policies already recognize such expression as faculty-owned "traditional academic works." Trade-secret protection is unlikely to attach to widely used prompt templates, and provider terms of service generally assign ownership of inputs and outputs to the user while reserving broad licenses for the company. Thus, there is no strong doctrinal argument for institutional appropriation of these artifacts.

Finally, the Category One rule helps discipline the scope of "significant use" clauses. A faculty member's reliance on AI tools—even on university-

licensed platforms or high-performance computing clusters for text analysis—should not, without more, reclassify their prompts and methods notes as institutional property. The fact that a historian needed university-owned computers to analyze a corpus of texts should not strip her of ownership over the prompt strategies and interpretations that constitute her scholarly contribution.

C. Category Two, AI Operational Artifacts and Shared Institutional Shop Rights

The second portability category captures AI artifacts that function as durable research infrastructure and embody significant institutional investment. Examples include:

- Fine-tuned LLMs trained on large datasets using university-owned or grant-funded high-performance computing clusters;
- Custom research agents built with substantial assistance from university staff, data scientists, or research computing centers;
- AI-driven platforms maintained by a laboratory or center for use by multiple investigators; and
- Institutional "virtual labs" or domain-specific models deployed as shared resources across departments.

These artifacts differ from those in Category One in at least two respects. First, their value often depends heavily on institutional contributions, including compute time, licenses to proprietary base models, staff design and maintenance, and integration with other core institutional systems. Second, they are typically designed for use by multiple researchers or projects, not just the work of a single scholar. Thus, they look more like shared infrastructure or large-scale software projects than individual scholarly works. From the perspective of drafting potential university policy, one way to capture this distinction is to provide that an artifact falls into Category Two only if it could not have been created without the specific, non-commodity infrastructure provided by the university (e.g., high-performance computing clusters, proprietary datasets, or dedicated

engineering staff). Artifacts created through the use of commodity tools (standard laptops, standard email/cloud accounts) should remain in Category One.

For Category Two artifacts, the default rule is qualified portability. A departing faculty member would have a right to continue using and developing AI operational artifacts that they played a substantial role in creating, but that right would coexist with a non-exclusive, royalty-free license in favor of the originating institution—a university-level analog to the "shop right" concept in patent law.[92] Under this standardized mechanism, the faculty member may take a copy of the model or agent configuration to their new institution, subject to any non-negotiable constraints in base-model licenses or sponsored research agreements. The originating university would retain a perpetual right to use, maintain, and further develop the artifact for institutional purposes, including fulfilling grant obligations and supporting students or collaborating investigators. From there, any commercialization of the artifact by either party would trigger the revenue-sharing, conflict-of-interest, and technology-transfer policies that would ordinarily apply to inventions and software tools.

This structure could be implemented contractually through modest amendments to existing IP policies and technology-transfer agreements. Policies would need to define "AI operational artifacts" and specify that, absent contrary terms in a particular grant or license, creators and the institution hold mirror-image rights whereby each may use the artifact internally and in future research and seek external licenses subject to disclosure and revenue-sharing rules. Crucially, though, the portable shop-right framework must also account for the full commercial trajectory of AI operational artifacts after a scholar departs. If a faculty member transfers a fine-tuned model or custom agent to a new institution and subsequently seeks to commercialize it—whether through a startup or a third-party license—the originating university's foundational investment warrants recognition beyond a mere internal-use license. To address this, the framework contemplates that any external commercialization of a Category Two artifact will require an Inter-Institutional Agreement (IIA) between the

92. In United States patent law, a "shop right" is a common law equitable doctrine that grants an employer a limited right to use an invention created by an employee, even if the employee technically owns the patent. *See generally* United States v. Dubilier Condenser Corp., 289 U.S. 178 (1933); McElmurry v. Ark. Power & Light Co., 995 F.2d 1576 (Fed. Cir. 1993).

old and new universities. Modeled on standard IIAs for jointly owned patents, these agreements would establish equitable revenue-sharing and management protocols, ensuring that the originating institution shares in the financial upside of the asset it helped create without retaining the power to block the scholar's innovation or entrepreneurial activities at their new home.[93]

The Category Two rule is designed to balance several competing interests. First, it protects scholarly mobility. A faculty member whose research agenda centers on a particular fine-tuned model or custom agent should not be forced to abandon that tool when changing institutions. Doing so would waste public and private investments and dissuade scholars from building ambitious AI infrastructure in the first place. Additionally, the rule respects institutional investments. Because Category Two artifacts often depend on substantial institutional backing, it would be problematic for a single researcher to walk away with exclusive control. Universities must be able to honor grant commitments, support students and collaborators who remain, and maintain critical infrastructure. A shop-right-style license ensures that the institution's ability to fulfill these duties does not depend on the continued employment of a single individual.

Lastly, the rule fits comfortably within the doctrinal environment. Copyright offers, at best, a thin layer of protection for model weights, while trade-secret protection is available only so long as secrecy is maintained. In practice, contract law governs ownership and control of AI operational artifacts—base-model licenses, sponsored research agreements, and internal policies. By specifying a default of shared, non-exclusive rights, universities can avoid assignment conflicts while preserving flexibility to negotiate alternative arrangements where funders or providers demand them. The shop-right framework also provides a clear answer to questions about resource-based claims. Where institutional contributions cross the "significant use" threshold, the artifact will fall into Category Two, and the institution's license is guaranteed. Where similar levels of contribution are absent or incidental, the artifact remains in Category One. This boundary setting should reduce the temptation to redraw lines *ex post* to capture

93. *See Model IIA Project*, ASS'N OF UNIV. TECH. MANAGERS, https://autm.net/surveys-and-tools/agreements/model-iia-project [https://perma.cc/E4AP-MW9V] (last visited Mar. 6, 2026) (providing standard frameworks for designating a lead licensing institution to facilitate commercialization while ensuring equitable royalty distribution).

successful tools.

D. Category Three, Regulated Data and Compliance-Driven Limits on Portability

The third proposed category consists of regulated data: information whose collection, storage, and dissemination are governed by external legal regimes or binding institutional commitments. It includes:

- FERPA-protected education records;
- HIPAA-regulated protected health information (PHI);
- IRB-regulated human-subjects data, including interview transcripts and survey responses;
- Confidential or proprietary data provided by sponsors or collaborators under nondisclosure agreements; and
- Other sensitive datasets covered by agency-level data-sharing policies or national security controls.

As Part II discussed, these regulatory regimes impose rigid constraints on where data can reside and how it can be processed by AI tools.[94] They also often assign responsibility to the institution rather than the individual faculty member, reflecting the view that universities are stewards of research data on behalf of subjects, students, and society.[95]

For Category Three materials, the portability rule must be compliance-forward, meaning that regulated data is not itself portable by default. But institutions should facilitate portability through redaction, anonymization, synthetic data, and controlled-access mechanisms. In practice, this means a departing faculty member may not unilaterally remove or copy identifiable FERPA, HIPAA, or IRB-covered data from the originating institution without explicit authorization through IRB amendments, data-use agreements, or sponsor consent.

The institution must, however, act in good faith to enable the scholar to reconstruct or continue their research at the new institution, consistent with subject consent and legal obligations. Mechanisms for doing so might

94. *See infra* text accompanying notes 76–82.
95. *Id.*

include: providing de-identified or aggregated versions of datasets; generating synthetic datasets that preserve key statistical properties without identifying individuals; entering into inter-institutional data-use agreements that permit controlled access from the new institution; or housing the dataset in a neutral repository or secure enclave accessible to authorized collaborators.[96] IRBs and data-governance offices already manage many of these arrangements in other contexts, such as when external collaborators join a multi-site study or when a principal investigator retires. The proposed framework simply standardizes the expectation that portability disputes be resolved through these well-understood tools rather than by informal practices or unilateral copying of regulated data.

The Category Three rule reflects the reality that compliance obligations are non-negotiable. Treating regulated data as categorically non-portable by default prevents well-intentioned faculty from inadvertently triggering breaches when they move institutions or when they export AI conversation logs that embed student or subject information. At the same time, the rule rejects the view that compliance concerns justify locking down entire research programs. By obligating institutions to pursue redaction, anonymization, synthetic data, and inter-institutional agreements when feasible, the framework honors the idea that knowledge generated from publicly funded research should remain mobile and accessible.

Finally, clarifying that regulated data is a distinct category helps cleanly separate disputes about who may use an AI tool from disputes about who may access a dataset. A fine-tuned model trained on clinical notes, for example, may be portable as a Category Two artifact (subject to shop-right and licensing constraints) even if the underlying PHI remains locked within the originating institution's secure environment.

96. Admittedly, some of these mechanisms may prove more technically difficult and expensive to implement than others. Many university IRBs and IT departments might not have the budget or expertise to generate HIPAA-compliant synthetic data for a departing faculty member, for example. Acknowledging the risk of resource limitations, it ought to be sufficient to require that a university permit *access* to the data (via secure enclave or remote desktop) rather than necessarily paying to *transform* the data for export. This approach can be framed in university policy as a "right of continued access" rather than as a "right of export."

E. Advantages of the Proposed Framework

The principal advantage of the three-part portability framework proposed is that it simplifies classification decisions for universities and faculty. It does so by re-describing AI-era artifacts in functional terms, as well as by assigning clear, differentiated rules for each category.

Consider prompt libraries. Under the new framework, prompts and meta-prompts represent presumptively Category One artifacts: they are methods notes, not institutional infrastructure. Universities need not decide whether they are "software," "data," or "expression" for IP purposes. Instead, what matters is their function as personal research tools. Therefore, faculty may take them when they leave, notwithstanding the use of institutional AI licenses or compute.

Conversation logs follow a similar pattern. To the extent they record a scholar's interaction with AI as part of research or teaching, they fall under Category One as personal methods notes. To the extent they embed regulated data (for example, identifiable student work or human-subject interviews), those portions are subject to Category Three limits and must be redacted or handled through controlled-access mechanisms. This dual classification avoids both extremes: logs are not treated as wholly private, and therefore immune from research-integrity or compliance review, nor are they automatically swept into institutional records systems subject to broad access and disclosure.

In contrast, fine-tuned models and custom agents are paradigmatic Category Two artifacts when they rely on significant institutional resources or are designed for shared use. The framework instructs universities not to shoehorn them into "software" or "data" categories as a pretext for exclusive control. Instead, policies should recognize shared rights: the university's license (i.e., a shop right) protects institutional continuity while the faculty member's equal license preserves mobility. Where a model has been fine-tuned primarily with the faculty member's own time and resources—say, on a personal cloud account without substantial institutional support—it may instead remain in Category One, with full portability. The decisive factor is institutional investment and infrastructural function, not *ex post* assessments of commercial potential.

By specifying these categories and rules *ex ante*, universities can draft policies that avoid opportunistic reclassification when conflicts arise. A

custom agent that begins as a personal tool but grows into a shared institutional platform can be re-categorized from Category One to Category Two through a documented process that includes recognition of faculty and institutional contributions, rather than through unilateral administrative edict. Conversely, administrators cannot suddenly declare that all prompt libraries are institutional "software" once they realize their value. The policy presumption of Category One status stands unless clearly overridden by a negotiated agreement.

In sum, a decoupled scholarly portability framework offers a principled way to govern AI-mediated research artifacts. By distinguishing among scholarly expression, AI operational artifacts, and regulated data—and then assigning each a tailored portability rule grounded in academic freedom norms and existing law—the framework preserves faculty mobility, safeguards institutional stewardship, and respects non-negotiable compliance obligations.

CONCLUSION

The rise of generative and agentic AI has exposed a gap at the heart of university governance. Historically, institutions have mediated the tension between institutional stewardship and faculty autonomy through relatively stable categories of research-related outputs like "traditional academic works," "research data," "software," and "institutional resources." AI-era artifacts do not fit neatly in those boxes. Prompt libraries, conversation logs, fine-tuned models, and custom agents are at once expressions of scholarly judgment, components of research infrastructure, and potential repositories of regulated data. If universities treat these composites as wholly personal (belonging to faculty), they risk losing control over core assets and failing in their compliance obligations. If they treat them as wholly institutional, they risk chilling experimentation, undermining scholarly mobility, and eroding academic freedom.

This Article has argued that the way out of that dilemma is to reframe the problem as one of functional decoupling and portability, distinguishing among scholarly expression and personal know-how, AI operational artifacts, and regulated data, and then assigning a tailored portability rule to each category. This approach clarifies when a scholar's interest in portability should dominate, when institutional stewardship justifies a

retained license or shared control, and when compliance properly limits both sides.

Moreover, this framework aligns the incentives of faculty and institutions. Clear rules that guarantee portability of prompts, methods, and know-how encourage faculty to invest in AI-mediated research without fear that their work will be stranded at a single institution. Shared rights in operational artifacts give universities a defensible claim to core infrastructure while avoiding zero-sum fights when faculty move or when tools transition from personal experiments to campus-wide platforms. Robust, category-specific protections for regulated data ensure that neither side can invoke "innovation" or "ownership" to sidestep non-negotiable privacy and ethical commitments.

The Article's broader normative message is that universities cannot afford to be purely reactive in this space. AI is quickly becoming part of the ordinary fabric of scholarly work, not an exotic add-on. Institutional norms and policies that were built for a paper-and-mainframe era will either be reinterpreted hastily and under the shadow of conflict or deliberately redesigned to match the realities of AI-mediated inquiry. The scholarly portability framework offered here aims to enable the latter path. It provides governing boards, general counsels, faculty senates, and research offices with a vocabulary and set of defaults that can be translated into policy language, promotion and tenure criteria, and template agreements with providers and sponsors. Properly implemented, the framework can help ensure that AI becomes a catalyst for new forms of inquiry, collaboration, and teaching, rather than a pretext for hindering the very mobility of ideas that universities exist to promote.