FOREIGN TRAINED LAWYERS ASSOCIATION
Policy Paper No. II · April 2026

FTLA POLICY PAPER SERIES
NEW YORK · WASHINGTON · GLOBAL
INTELLECTUAL PROPERTY · ARTIFICIAL INTELLIGENCE · INNOVATION POLICY

Authorship, Originality, and the Machine-Generated Work

U.S. Copyright Doctrine, Patent Disclosure, and the Governance of Generative Artificial Intelligence

AN FTLA POLICY FRAMEWORK

Mariami Zhorzholiani
Founder, Foreign Trained Lawyers Association

An institutional analysis of the displacement of human authorship and inventorship in U.S. intellectual property law.

FOREIGN TRAINED LAWYERS ASSOCIATION
NEW YORK, NEW YORK

M M X X V I
II

ABSTRACT

The rapid deployment of generative artificial intelligence systems has produced the most significant stress test of U.S. intellectual property doctrine since the Copyright Act of 1976 and the Patent Act of 1952. Foundational doctrines that have organized the field for more than a century, the human authorship requirement that the Copyright Office reaffirmed in March 2023 and again in its Copyright and Artificial Intelligence Reports of January and May 2025, the Federal Circuit's holding in Thaler v. Vidal that an inventor must be a natural person, and the trade secret framework codified in the Defend Trade Secrets Act of 2016, were not designed for systems that produce expressive and inventive output at industrial scale without identifiable human contribution at each step. The result is a doctrinal landscape in which the most economically consequential outputs of contemporary technology, in many cases, fall outside the protective architecture that U.S. intellectual property law has historically extended to commercially significant creative and inventive work.

This paper argues that the United States should resist two equally unsatisfactory responses to this development: extending authorship and inventorship to non-human systems, which would untether the constitutional and statutory frameworks from their human-centered foundations; and leaving generative output entirely in the public domain, which would create disclosure-suppression incentives that undermine both the patent bargain and the integrity of training data markets. The paper develops a third pathway: a calibrated framework distinguishing protectable human contribution from unprotectable machine output, anchored in the existing statutory architecture, supplemented by a sui generis disclosure regime for substantial AI involvement, and coordinated with parallel developments in the European Union's AI Act and the United Kingdom's text and data mining consultation. The framework draws on the cross-border analytical method developed in FTLA Policy Paper No. 1 and applies it to the new transnational architecture of AI governance.

TABLE OF CONTENTS

GLOSSARY OF ABBREVIATIONS

The following abbreviations are used throughout this paper. Citations to source materials appear in the References.

Term	Full term
AI	Artificial Intelligence
AIA	Leahy-Smith America Invents Act of 2011
CDSM	Directive on Copyright in the Digital Single Market (EU 2019/790)
CMI	Copyright Management Information (17 U.S.C. § 1202)
DMCA	Digital Millennium Copyright Act of 1998
DTSA	Defend Trade Secrets Act of 2016
EPO	European Patent Office
EUIPO	European Union Intellectual Property Office
FTLA	Foreign Trained Lawyers Association
GPAI	General-Purpose Artificial Intelligence (under EU AI Act)
IPO	U.S. Intellectual Property Owners Association
LLM	Large Language Model
NIST	National Institute of Standards and Technology
PTAB	Patent Trial and Appeal Board
TDM	Text and Data Mining
UKIPO	United Kingdom Intellectual Property Office
USCO	United States Copyright Office
USPTO	United States Patent and Trademark Office
UTSA	Uniform Trade Secrets Act
WIPO	World Intellectual Property Organization

I. Introduction: Intellectual Property in the Generative Age

In December 2025, the United States Copyright Office issued the third installment of its Copyright and Artificial Intelligence Report, addressing the question of generative AI training and concluding a multi-year administrative inquiry that began with a Federal Register notice in August 2023. In the same period, the United States Patent and Trademark Office published its Inventorship Guidance for AI-Assisted Inventions (February 2024), the Federal Circuit reaffirmed its holding in Thaler v. Vidal that an inventor under the Patent Act must be a natural person, and the District Court for the District of Columbia entered judgment in Thaler v. Perlmutter, sustaining the Copyright Office's refusal to register a work generated entirely by an autonomous system. Across the same period the European Union's Artificial Intelligence Act entered into force, the United Kingdom Intellectual Property Office published a consultation on text and data mining and copyright reform, and the World Intellectual Property Organization convened the third session of its Conversation on Intellectual Property and Artificial Intelligence.

These developments are not coincidental. They reflect a sustained and largely synchronized effort by the major intellectual property jurisdictions to define how legal systems built around human authorship and human inventorship should respond to systems that produce, at industrial scale, output that closely resembles the protectable subject matter of those legal systems. The effort is not yet complete in any jurisdiction. The choices made over the next several years will shape the structure of innovation incentives, the allocation of returns to creative and inventive labor, and the cross-border architecture of intellectual property protection for a generation.

This paper, the second in the Foreign Trained Lawyers Association policy paper series, takes the position that U.S. intellectual property law should respond to generative AI through calibrated doctrinal refinement rather than wholesale recharacterization. The paper develops that position in three steps. It first identifies the specific doctrinal pressure points that generative AI imposes on the existing U.S. intellectual property architecture (Sections II and III). It then analyzes those pressure points in their economic and competitive context (Section IV) and situates the U.S. response against parallel developments in the European Union and the United Kingdom (Section V). It concludes with a four-part policy framework directed at codifying the human-contribution standard, creating a disclosure regime for substantial AI involvement, authorizing a narrow training-data fair use safe harbor, and strengthening trade secret remedies for training data and model architecture (Sections VI and VII).

A. The economic stakes

The intellectual property system is not an abstract regulatory architecture. It is the legal infrastructure through which the United States allocates returns to creative and inventive labor, and through which the United States establishes the conditions under which technology firms, content producers, pharmaceutical developers, software developers, and an extensive long tail of small and individual creators can recover their investment and continue to invest. The U.S. Patent and Trademark Office reported approximately 350,000 utility patent grants in fiscal year 2024. The U.S. Copyright Office registered approximately 480,000 works in calendar year 2024. The aggregate market capitalization of S&P 500 companies whose principal assets are intellectual property, software platforms, pharmaceutical patents, content libraries,

semiconductor designs, has grown to constitute a majority of total S&P 500 market capitalization. The intellectual property system is, in this sense, the legal infrastructure of the largest segment of the U.S. economy by market value.

Generative artificial intelligence operates as both an input to and an output from this system. As an input, generative AI systems are trained on corpora that overwhelmingly comprise existing copyrighted works, creating an unresolved question about the lawfulness of the training process itself, a question that has generated more than thirty pending federal lawsuits against major AI developers as of early 2026. As an output, generative AI systems produce expressive and inventive material that, on its face, resembles the protectable subject matter of the copyright and patent systems, but that, under existing doctrine, in many cases falls outside the protective architecture entirely. The combined effect is a structural mismatch: generative AI consumes the existing intellectual property system as an input on terms that have not been judicially or legislatively settled, and it produces output that the same system has not been adapted to receive.

The economic consequences of leaving this mismatch unresolved are substantial. On the input side, the unresolved fair use question creates litigation exposure that, by some industry estimates, may exceed $100 billion in aggregate damages claims by the end of the current litigation wave. On the output side, the absence of clear protective architecture for AI-assisted creative and inventive work creates incentives for developers to either suppress disclosure of AI involvement (to preserve the appearance of human authorship and inventorship) or to retreat to trade secret protection that, while available, is structurally less suitable than the copyright and patent systems for the diffusion of technical knowledge that the U.S. innovation system depends upon.

B. The doctrinal stakes

Beyond the immediate economic consequences, the doctrinal stakes of the U.S. response to generative AI are significant in their own right. The human authorship requirement is not a peripheral technicality of U.S. copyright doctrine. It is, in the Copyright Office's view as expressed in the Compendium of U.S. Copyright Office Practices and reaffirmed across the 2023, 2025, and December 2025 Copyright and Artificial Intelligence Reports, a constitutional and statutory limit on what the copyright system can protect. The Constitution's Intellectual Property Clause authorizes Congress to secure exclusive rights to authors and inventors; the Copyright Act of 1976 implements that authority with respect to original works of authorship; the Patent Act, particularly as amended by the Leahy-Smith America Invents Act of 2011, implements that authority with respect to inventions and discoveries by inventors. Each of these statutory schemes is built on the structural assumption that the protected subject matter is the product of human creative or inventive activity.

That structural assumption is not a relic. It reflects a coherent theory of intellectual property as a system for incentivizing and rewarding the human contribution that the constitutional grant identifies as the object of protection. To extend authorship to non-human systems, or to inventorship as Thaler urged, would not merely expand the universe of protectable subject matter; it would alter the conceptual foundation of the protective architecture. The choice is not, as it is sometimes framed, between an inclusive system that

protects AI output and an exclusionary system that does not. The choice is between a system that retains its human-centered theoretical foundation and accommodates AI-assisted output through doctrinal calibration, and a system that abandons that foundation in pursuit of a wider scope of protection. This paper argues for the first.

II. The Problem Statement

The introduction of generative AI into the existing U.S. intellectual property architecture produces three distinct doctrinal pressure points. Each is a problem of fit between a system designed for human creative and inventive activity and outputs that involve substantial machine contribution. Each has generated, in the past three years, administrative guidance, litigation, and competing scholarly proposals. None has been resolved with the clarity that the underlying economic activity requires.

A. Three doctrinal pressure points

1. The authorship pressure point in copyright

Section 102(a) of the Copyright Act extends protection to original works of authorship fixed in any tangible medium of expression. The statute does not, in terms, define authorship. The Copyright Office, the federal courts, and the Supreme Court have, over more than a century, developed a body of doctrine treating authorship as the activity of a human originator who conceives the protectable expression. Burrow-Giles Lithographic Co. v. Sarony (1884) treated the author as the person to whom anything owes its origin. Feist Publications, Inc. v. Rural Telephone Service Co. (1991) located originality in independent creation accompanied by a modicum of creativity, with the modicum supplied by human judgment. The Compendium, since at least its third edition in 2014, has formalized this doctrinal trajectory by stating that the Copyright Office will refuse to register a claim if it determines that a human being did not create the work.

Generative AI introduces a category of output that, in many cases, cannot be traced to a human originator in the sense the doctrine contemplates. A user who supplies a prompt has supplied something, but in the typical case the prompt's relationship to the output is too attenuated to satisfy traditional originality analysis: a 200-word prompt that produces a 5,000-word short story has not been independently created by the user in the relevant sense, because the specific expressive choices that the originality analysis turns on, sentence structure, lexical selection, narrative voice, characterization, were made by the system. The Copyright Office's January 2025 Report on Copyrightability addressed this question directly, concluding that copyright protection extends only to material that reflects sufficient human creative control, with prompt-based generation generally falling below that threshold. The doctrinal result is that a substantial volume of expressive output produced through generative AI is, on current administrative practice, ineligible for copyright registration.

2. The inventorship pressure point in patent

Section 100(f) of the Patent Act, as amended by the AIA, defines inventor as the individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention. Section 115 requires that each inventor execute an oath or declaration. The Federal Circuit, in Thaler v. Vidal (2022), held that the term individual in this provision unambiguously means a natural person, and the Supreme Court denied certiorari. The USPTO, in its February 2024 Inventorship Guidance for AI-Assisted Inventions, accepted that AI-assisted inventions are not categorically unpatentable, but required that at least

one natural person have made a significant contribution to the conception of the claimed invention to satisfy the inventorship requirement.

The pressure on the patent system is structurally similar to the pressure on the copyright system, but with two important differences. First, the patent system has, through the inventorship guidance, articulated a more operationally specific standard than the Copyright Office's human-creative-control formulation: significant contribution to conception, drawing on the Pannu factors that the Federal Circuit developed in joint inventorship cases. Second, the consequences of failing the inventorship test are different in kind. A copyright claim that fails the human authorship test results in unprotected expression that any third party may use; a patent claim that fails the inventorship test results in either an unfilled application (if the failure is identified early) or, if the failure is identified after issuance, an invalid patent. The asymmetry creates a different set of disclosure incentives, examined further in Section IV.

3. The training data pressure point

The third pressure point arises not from the output of generative AI but from its input. Major generative AI systems are trained on corpora that comprise, in significant part, copyrighted works obtained without licenses from rights holders. The lawfulness of this training activity, under the existing fair use framework, is the question at issue in pending litigation including Authors Guild v. OpenAI (S.D.N.Y.), New York Times Co. v. Microsoft Corp. (S.D.N.Y.), and the consolidated Andersen v. Stability AI Ltd. (N.D. Cal.) actions. The Copyright Office's December 2025 Report on Generative AI Training surveys the doctrinal landscape and identifies the central interpretive question as whether the four-factor fair use analysis under Section 107, particularly the first factor (purpose and character of the use, including whether the use is transformative) and the fourth factor (effect on the potential market for or value of the copyrighted work), is satisfied by the training process.

The training data question presents a doctrinal pressure of a different character than the authorship and inventorship questions. It does not require any modification of the human-centered foundations of U.S. intellectual property law. It requires, instead, the application of an existing doctrinal framework, fair use, to a new technological context. The challenge is that the existing framework was developed in cases involving identifiable individual uses (a parodist using portions of a song, a search engine indexing thumbnail images, a digitization project enabling full-text search of library holdings), whereas generative AI training involves the simultaneous use of millions of copyrighted works in a process whose outputs may, in some cases, substitute directly for the inputs. The doctrinal tools are familiar; the scale is new.

B. Why the existing answers are unstable

Each of the three pressure points has attracted a category of scholarly and policy proposals that, on closer analysis, prove unstable. This subsection identifies four such proposals and explains why they fail to provide a durable solution.

First, the proposal to extend authorship and inventorship to AI systems themselves, advanced by Stephen Thaler in the litigation that bears his name and by a small body of academic literature, fails on constitutional grounds. The Intellectual Property Clause authorizes Congress to secure rights to authors and inventors; the

Federal Circuit in Thaler v. Vidal and the District Court in Thaler v. Perlmutter have read these terms as requiring natural persons. Whether or not Congress could constitutionally extend authorship and inventorship to non-human entities is a question that the courts have, to date, declined to resolve in favor of such extension. As a matter of policy, the extension would also raise questions about the allocation of resulting rights, since the AI system itself cannot exercise rights, prosecute infringement, or receive damages. The system's owner, developer, or user would have to be designated as the rights holder, returning the analysis to the question of which human contribution merits protection.

Second, the proposal to leave AI-assisted output entirely outside the intellectual property system, treating it as public domain by default, fails on innovation policy grounds. The U.S. innovation system depends upon the disclosure-protection bargain that copyright and patent law establish: rights holders disclose their works in exchange for time-limited exclusive rights. If AI-assisted output is unprotected, the disclosure incentive collapses, and developers have a strong incentive to either suppress disclosure of AI involvement (preserving the appearance of human authorship) or retreat to trade secret protection (preserving the work's commercial value at the cost of its diffusion). Neither outcome is desirable from the perspective of the innovation system as a whole.

Third, the proposal to extend the work made for hire doctrine to AI output, treating the AI system's owner or operator as the statutory author, fails on doctrinal grounds. Section 101 of the Copyright Act defines a work made for hire as either a work prepared by an employee within the scope of employment or a specially commissioned work in one of nine enumerated categories. An AI system is neither an employee nor an independent contractor; it does not have the legal capacity to enter into the employment or commission relationship that the doctrine presupposes. The doctrine could be amended to encompass AI-generated work, but the amendment would constitute a significant expansion of the work made for hire concept beyond its historical scope, and would re-introduce the question of which human contribution justifies the resulting allocation of rights.

Fourth, the proposal to create a comprehensive sui generis regime for AI output, modeled on the European Union's database right under Directive 96/9/EC, fails on transition cost grounds. The creation of a new sui generis regime would require sustained legislative action, would generate years of interpretive uncertainty as the regime developed administrative and judicial gloss, and would create transition problems for existing AI developers operating under the assumption that current law governs their output. A more targeted intervention, retaining the existing copyright and patent architecture and refining its application to AI-assisted work, is, on the analysis offered here, the more durable approach.

III. The Legal Framework

Sound policy analysis presupposes accurate doctrinal description. This section sets out the existing U.S. intellectual property architecture as it bears on generative AI, drawing on the statutes, the leading administrative pronouncements, and the principal judicial authorities. The treatment is not encyclopedic; it is calibrated to the policy analysis that follows in Section IV.

A. Copyright: human authorship and the 1976 Act

The Copyright Act of 1976, as amended, extends protection to original works of authorship fixed in any tangible medium of expression, now known or later developed (17 U.S.C. § 102(a)). The statute enumerates eight non-exhaustive categories of protectable works: literary works; musical works; dramatic works; pantomimes and choreographic works; pictorial, graphic, and sculptural works; motion pictures and other audiovisual works; sound recordings; and architectural works. Protection is automatic upon fixation; registration with the Copyright Office is a precondition to suit for U.S. works under Section 411(a) but not to the existence of the right itself.

The human authorship requirement is not stated in those terms in the 1976 Act. It is derived from the constitutional grant of authority to Congress to secure exclusive rights to authors (Article I, Section 8, Clause 8), from the historical understanding of the term author as a human originator developed in cases including Burrow-Giles Lithographic Co. v. Sarony (1884), and from the Copyright Office's longstanding administrative practice. The Compendium of U.S. Copyright Office Practices, Third Edition (2014), states at Section 313.2 that the Office will not register works produced by a machine or mere mechanical process that operates randomly or automatically without any creative input or intervention from a human author.

This administrative position has been litigated and sustained. In Thaler v. Perlmutter (D.D.C. 2023), aff'd (D.C. Cir. 2025), the District Court rejected Stephen Thaler's challenge to the Copyright Office's refusal to register A Recent Entrance to Paradise, a work generated by Thaler's Creativity Machine. The Court held that human authorship is a bedrock requirement of copyright that has been consistently recognized by courts since the founding of the Republic. The D.C. Circuit affirmed in 2025, and the Supreme Court declined to grant certiorari. The administrative position has been further specified in the Copyright Office's three-part Copyright and Artificial Intelligence Report: Part 1 (Digital Replicas, July 2024) addressed unauthorized digital replicas; Part 2 (Copyrightability, January 2025) addressed the registration of works involving AI; and Part 3 (Generative AI Training, December 2025) addressed the training data question.

The Part 2 Report, on copyrightability, is of particular importance. The Report concludes that copyright protection is available for works that contain sufficient human-authored expression even when AI tools have been used in the creative process, but that purely AI-generated content, including content produced from text prompts, is not eligible for protection in the absence of substantial human creative control. The Report identifies several factors relevant to determining sufficient human creative control: the degree of human authorship in selecting and arranging AI-generated elements, the iterative refinement of AI output through human editing and modification, and the contribution of human-authored components to the final

work. Each factor is fact-specific, and the Report acknowledges that line-drawing in close cases will be difficult.

B. Patent: invention, conception, and Thaler v. Vidal

The Patent Act, as comprehensively amended by the Leahy-Smith America Invents Act of 2011, provides patent protection for any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof (35 U.S.C. § 101), provided the claimed subject matter satisfies the further requirements of utility, novelty (Section 102), non-obviousness (Section 103), and adequate written description and enablement (Section 112). Section 100(f), added by the AIA, defines the term inventor as the individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention. Section 115 requires each inventor named in an application to execute an oath or declaration.

The Federal Circuit in Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022), addressed Stephen Thaler's challenge to the USPTO's refusal to issue patents naming his DABUS system as the inventor. The Court held that the term individual in Section 100(f) unambiguously means a natural person, relying on the dictionary meaning of the term, the Patent Act's repeated use of personal pronouns in connection with inventors, and the structural inference from Section 115's oath requirement that an inventor must be capable of executing an oath. The Supreme Court denied certiorari in 2023.

The USPTO's February 2024 Inventorship Guidance for AI-Assisted Inventions accepts the Thaler holding and addresses the inventorship analysis when AI has been used in the inventive process. The Guidance distinguishes inventions made by AI from inventions made with the assistance of AI. The former are not patentable; the latter may be patentable if at least one named inventor, who must be a natural person, has made a significant contribution to the conception of the claimed invention. The Guidance applies the Pannu factors developed in the Federal Circuit's joint inventorship case law: each named inventor must have contributed in some significant manner to the conception of the invention, made a contribution to the claimed invention that is not insignificant in quality when measured against the dimension of the full invention, and done more than merely explain to the real inventors well-known concepts and the current state of the art.

The Pannu framework applied to AI-assisted invention generates a fact-specific inquiry into the human contribution. The Guidance offers illustrative examples: a researcher who recognized a problem, designed an AI system to address it, and identified the patentable aspects of the system's output may satisfy the significant contribution standard; a researcher who merely instructed an AI system to generate an invention without further substantive engagement likely does not. As with the copyright human-authorship analysis, the patent inventorship analysis turns on identifying the locus of substantive human contribution, with the application of the standard varying with the facts of each case.

C. Trade secrets: the DTSA architecture

The Defend Trade Secrets Act of 2016 (Pub. L. No. 114-153) created a federal cause of action for trade secret misappropriation, supplementing the state-law trade secret regimes that derive from the Uniform Trade Secrets Act adopted in 48 states. Under the DTSA and parallel state law, a trade secret is information that derives independent economic value from not being generally known to, and not being readily ascertainable through proper means by, another person who can obtain economic value from its disclosure or use, and that is the subject of reasonable measures to maintain its secrecy.

The trade secret framework operates as a residual category for valuable commercial information that is not protected, or not most efficiently protected, by patent or copyright. AI training data, AI model architectures, AI training methodologies, and AI-generated output that has not been disclosed to the public are, in the typical case, eligible for trade secret protection if the requisite secrecy and value conditions are satisfied. This availability has substantial implications for the policy analysis offered in Section IV: where copyright and patent protection are unavailable for AI-related subject matter, the protective default is not the public domain but trade secret protection. The choice is therefore not between protection and non-protection, but among different protective regimes with different disclosure characteristics.

D. Adjacent doctrines: right of publicity, unfair competition, contract

Several adjacent legal frameworks bear on the AI question and warrant brief mention. The right of publicity, a state-law doctrine that varies in its scope across jurisdictions, governs the unauthorized commercial use of a person's name, likeness, voice, or other indicia of identity. The use of AI to generate digital replicas of identifiable persons, the subject of the Copyright Office's July 2024 Report on Digital Replicas, raises right of publicity questions that have generated state legislative responses including the Tennessee ELVIS Act of 2024 and proposed federal legislation including the NO FAKES Act.

Unfair competition doctrine, including federal claims under Section 43(a) of the Lanham Act and state-law claims for misappropriation and false designation of origin, may apply to certain AI-related conduct that does not fit within the copyright or patent frameworks. Contract law, particularly the use of terms of service and end-user license agreements by AI developers, governs much of the practical allocation of rights between AI developers, users, and the providers of training data. The contract-law dimension of AI governance is, by some measures, the most operationally significant, since the substantive rules of intellectual property law apply against a background of contractual allocations that determine who may bring claims and on what terms.

IV. Analysis

Sections II and III have established the doctrinal pressure points and the legal framework within which they arise. This section analyzes the resulting policy questions, drawing on the framework developed in the prior sections and on the comparative perspective developed in Section V. Four lines of analysis are pursued: the originality fault line in copyright; the fair use question in training data; the innovation incentive question in patent; and the trade-secret displacement effect that the absence of clear copyright and patent protection produces.

A. The originality fault line

The Copyright Office's January 2025 Report on Copyrightability articulates a doctrinal standard, sufficient human creative control, that is, in administrative practice, applied along a spectrum. At one end of the spectrum, a work created by a human using an AI tool to perform discrete subordinate tasks (a writer using an AI system to generate a list of synonyms, an illustrator using an AI system to remove background elements from a photograph) plainly retains human authorship in the substantive sense, and the resulting work is registrable. At the other end of the spectrum, a work generated by an AI system from a text prompt, with no further human involvement, plainly lacks the human authorship that the doctrine requires, and is not registrable.

The hard cases are in the middle. Consider the artist who, in producing a single illustration, employs an AI system to generate twenty initial drafts; selects one as the basis for further work; uses additional AI tools to modify specific regions; manually paints over portions of the modified image; instructs the AI to refine the painted portions; and concludes by manually selecting the final color palette and image dimensions. Is this a registrable work? The Copyright Office's framework indicates that the analysis depends on whether the human creative contribution, the selection, the manual painting, the iterative refinement, the final selections, is sufficient. There is no quantitative threshold; there is a qualitative inquiry that the registration examiner must conduct. Consistency across examiners, and predictability for applicants, will be difficult to achieve under this framework.

Two consequences follow. First, the originality fault line creates a registration-uncertainty zone in which significant volumes of commercially produced work fall. Industry estimates suggest that a substantial share of contemporary commercial illustration, photography, music, and short-form text production now involves some level of AI assistance; the registration eligibility of these works depends on case-specific examiner judgments that the applicant cannot predict in advance. Second, the uncertainty creates incentives for misrepresentation. An applicant whose work is, in fact, primarily AI-generated, has an incentive to characterize it as primarily human-authored, since the registration consequences of accurate characterization may include refusal. The Copyright Office has acknowledged this concern in the Part 2 Report and has indicated its intention to develop disclosure obligations for AI involvement, but the operational architecture of those disclosure obligations remains under development.

B. Training data and the fair use question

The fair use analysis under Section 107, as applied to generative AI training, has been the subject of an extensive recent literature and is the central question in the pending litigation. The Copyright Office's December 2025 Report on Generative AI Training analyzes the four statutory factors and reaches conclusions that, while not legally binding, articulate the administrative position that will be persuasive in the pending litigation.

On the first factor (purpose and character of the use), the Report distinguishes between training that produces an AI system whose outputs do not substitute for the training works (a research model used for non-commercial scientific purposes, for example) and training that produces an AI system whose outputs may directly substitute for the training works (a generative AI system trained on a particular author's body of work and used to generate works in that author's style). The first is, on the Report's analysis, more likely transformative; the second is less so. The fourth factor (market effect) tracks this analysis: the second category of training involves a more direct potential market harm than the first.

The Report acknowledges that the existing case law does not provide a clear answer to the training data question, and that the courts will, in the pending litigation, be developing the doctrinal framework as they apply it. The Report identifies licensing as the preferred long-term solution, with statutory licensing or collective licensing arrangements potentially providing a workable mechanism for compensating rights holders while permitting training to proceed. The Report does not, however, propose specific statutory action; it leaves that question to Congress and to the courts in the pending litigation.

From a policy perspective, the training data question raises a structural concern that the four-factor fair use analysis is not, by itself, well calibrated to. Generative AI training operates at a scale, and produces market effects of a kind, that the existing fair use case law has not had occasion to address. The largest training corpora include billions of works; the resulting AI systems may, depending on their architecture and use, produce output that competes directly with the training works in the same markets. A fair use analysis that operates work-by-work, examining the relationship between a particular use and a particular work, is poorly suited to assessing the aggregate effect of a use that simultaneously involves millions of works. A more structural analysis, addressing the cumulative market effect of training across the corpus as a whole, may be required, and may not be available within the four-factor framework as currently developed.

C. Innovation incentives in the absence of patent protection

The patent inventorship analysis developed in Section III creates a population of AI-assisted inventions that, under current law, may not be patentable because no natural person has made the requisite significant contribution to conception. The size of this population is uncertain; the USPTO does not maintain statistics on applications refused on inventorship grounds, and the inventorship analysis is sufficiently fact-specific that ex ante prediction is difficult. Industry surveys suggest, however, that a non-trivial share of contemporary research and development in pharmaceutical drug discovery, materials science, and software engineering involves AI systems whose contributions to specific inventive concepts may be difficult to characterize as either substantial human contribution or merely advisory.

Where patent protection is unavailable, the developer's choice is between disclosing the invention without patent protection (which makes the invention available to competitors at no cost) and protecting the invention through trade secret (which preserves the developer's competitive advantage but suppresses the diffusion of technical knowledge that the patent system is designed to facilitate). The structural concern is that, as the share of AI-assisted invention grows, the share of inventions protected by trade secret rather than patent will grow correspondingly, with adverse effects on the diffusion of technical knowledge through the U.S. innovation system.

The patent system is, in this sense, more vulnerable to AI-related distortion than the copyright system. Copyright operates automatically on fixation; AI-assisted expressive output that fails the human authorship test simply enters the public domain, where it is available for use by any party. Patent protection requires affirmative action through application, examination, and issuance; AI-assisted invention that fails the inventorship test is typically protected by trade secret instead, because the inventor has the option to choose between patent and trade secret protection at the time of disclosure. The asymmetry between copyright and patent in their default treatment of failed-test subject matter is a feature of the existing architecture that the policy analysis should attend to.

D. The trade-secret displacement effect

Combining the analyses in subsections A through C, a structural pattern emerges: as the share of creative and inventive activity that involves substantial AI contribution grows, the share that is protected by copyright or patent shrinks, and the share that is protected by trade secret grows correspondingly. This trade-secret displacement effect has consequences for the operation of the U.S. innovation system that warrant explicit policy attention.

First, trade secret protection is, by design, exclusive of disclosure. The patent bargain exchanges disclosure for time-limited exclusive rights, producing a cumulative public record of technical knowledge that subsequent inventors can build upon. Trade secret protection produces no such record. As the trade secret share of the protected universe grows, the disclosure share shrinks, and the cumulative knowledge available to subsequent inventors does not grow at the rate that a patent-dominant system would generate.

Second, trade secret protection is, by design, indefinite in duration. Patent protection expires after twenty years from the application filing date, after which the disclosed information enters the public domain. Trade secret protection persists as long as the requisite secrecy is maintained. As the trade secret share grows, the share of technical knowledge that will eventually enter the public domain shrinks, and the long-term knowledge stock available to the U.S. innovation system grows more slowly.

Third, trade secret protection generates litigation that is structurally different from patent litigation. Patent cases are litigated against a public record (the issued patent and the prior art); trade secret cases are litigated against a confidential record that the parties must reconstruct in litigation. The litigation cost differential is meaningful, and the distribution of the cost burden between parties may favor the trade secret holder over the alleged misappropriator in ways that distort the underlying merit of the dispute.

These three concerns, reduced disclosure, indefinite duration, and asymmetric litigation cost, are not new. They have been recognized in the trade secret literature for decades, and they are part of the standard policy analysis that has historically supported the patent system as the preferred protective regime for technical innovation. What is new is that generative AI accelerates the conditions under which trade secret displaces patent, by producing a larger volume of subject matter that fails the patent inventorship test. The policy framework set out in Section VI is responsive to this acceleration.

V. Comparative Perspective

The cross-border analytical method developed in FTLA Policy Paper No. 1 applies with full force to the AI and intellectual property question. The U.S. response to generative AI is unfolding in parallel with substantively important responses in the European Union and the United Kingdom, and the cross-border practitioner must read the U.S. framework against the foreign frameworks. This section sketches the principal points of comparison.

A. The European Union: the AI Act and the CDSM Directive

The European Union's Artificial Intelligence Act (Regulation (EU) 2024/1689), which entered into force on August 1, 2024, with phased applicability over the following 36 months, is the world's first comprehensive horizontal regulation of artificial intelligence. The AI Act addresses AI system risk classification, prohibited practices, transparency obligations, and conformity assessment for high-risk systems. While the AI Act is principally a product safety and fundamental rights regulation rather than an intellectual property regulation, it contains provisions of direct intellectual property significance, particularly its training data transparency obligations for general-purpose AI models under Article 53.

Article 53 requires providers of general-purpose AI models to draw up and make publicly available a sufficiently detailed summary of the content used for training the model, prepared according to a template provided by the AI Office. The provision does not, in terms, regulate the lawfulness of training data use; it requires transparency about the training corpus. The transparency obligation interacts with the existing copyright framework under the Directive on Copyright in the Digital Single Market (Directive (EU) 2019/790), which establishes a text and data mining exception (Article 4) that allows TDM for any purpose, subject to an opt-out mechanism that rights holders may exercise. The combination of the CDSM Article 4 opt-out and the AI Act Article 53 disclosure produces a workable mechanism: rights holders may exercise an opt-out to prevent their works from being used for training; the disclosure obligation makes it possible for rights holders to ascertain whether their works have, in fact, been used.

The EU framework therefore differs from the U.S. framework in two important respects. First, the EU has, through the CDSM Directive, statutorily addressed the training data question through an opt-out mechanism, whereas the U.S. has left the question to fair use analysis under Section 107. Second, the EU has, through the AI Act, imposed a transparency obligation that the U.S. has not, to date, imposed at the federal level. Both differences create cross-border compliance complexity for AI developers operating in both markets, since the operational compliance posture in the EU may differ from the operational compliance posture in the U.S. with respect to the same training corpus and the same model.

B. The United Kingdom: text and data mining and the 2025 consultation

The United Kingdom, having departed from the European Union, has pursued a separate but parallel reform trajectory. The UK Intellectual Property Office issued a consultation in December 2024 on copyright and AI, addressing the text and data mining question and proposing a framework similar in structure to the EU's CDSM Article 4: a TDM exception with an opt-out mechanism. The consultation closed in February 2025,

and as of early 2026 the UK government has not yet announced the substantive direction of its response. The UK's approach has been characterized by some commentators as more developer-friendly than the EU's, in keeping with the UK's broader post-Brexit positioning as a hospitable jurisdiction for AI development.

The UK is also notable for its limited statutory recognition of computer-generated works under Section 9(3) of the Copyright, Designs and Patents Act 1988, which provides that, in the case of a literary, dramatic, musical or artistic work which is computer-generated, the author shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken. This provision, dating from the 1988 Act's original enactment, has been read by some commentators as providing a textual basis for protecting AI-generated work, although its application to contemporary generative AI has not been definitively addressed by the UK courts. The UK's approach therefore differs from the U.S. approach in a way that is directly relevant to the human authorship question: the UK has, since 1988, had a statutory framework that contemplates non-human computer-generated authorship, while the U.S. has not.

C. Convergence and divergence

The U.S., EU, and UK frameworks are converging on certain points and diverging on others. The convergence is most visible on the inventorship question: the European Patent Office's Legal Board of Appeal in J 8/20 (DABUS) and the UK Supreme Court in Thaler v. Comptroller-General (2023) both reached the same result as the U.S. Federal Circuit, holding that an inventor must be a natural person. On this point the major intellectual property jurisdictions are aligned. The divergence is most visible on the training data question: the EU has chosen statutory regulation through opt-out, the UK is consulting on a similar approach, and the U.S. has left the question to fair use litigation.

From a policy perspective, the divergence creates risks of regulatory arbitrage and operational complexity that the cross-border practitioner must manage. An AI developer training a model in the U.S. on a corpus that includes EU-origin works must, in the EU, comply with the CDSM Article 4 opt-out and the AI Act Article 53 disclosure; in the U.S., the same training activity is subject to fair use analysis with no statutory disclosure obligation. The same model deployed in both markets generates different compliance postures with respect to the same activity. The policy framework set out in Section VI is informed by the desirability of reducing this cross-border complexity, while preserving the substantive distinctness of the U.S. framework.

VI. Policy Recommendations

The diagnosis offered in Sections II through V supports a policy framework with four components. Each is calibrated to address a specific pressure point identified in the diagnosis. None requires the abandonment of the human-centered foundations of U.S. intellectual property law; each builds on the existing architecture. The framework is offered, in the manner of FTLA Policy Paper No. 1, as a structural proposal for stakeholder consultation rather than as a fully specified legislative or regulatory program.

A. Codify the human-contribution standard

The Copyright Office's January 2025 Report on Copyrightability and the USPTO's February 2024 Inventorship Guidance for AI-Assisted Inventions have, in administrative practice, converged on a similar substantive standard: protection is available where a human contribution is sufficient (in copyright) or significant (in patent), and is not available where the human contribution falls below that threshold. The two standards are administrative pronouncements, not statutory text. They are vulnerable to judicial reformulation, administrative reversal, and the kind of inter-agency divergence that complicates cross-border analysis.

FTLA proposes that Congress codify a unified human-contribution standard, applicable to both copyright and patent, that incorporates the existing administrative analysis. The proposed statutory language would specify that protection is available where a natural person has made a substantial creative or inventive contribution to the protected subject matter, with the determination of substantial contribution to be made on a case-specific basis using factors that the statute would enumerate. The factors would draw on the Copyright Office's sufficient creative control analysis and the USPTO's Pannu factors, harmonized in a single statutory provision. Codification would not change the substantive direction of the existing administrative analysis; it would stabilize that analysis against reversal and provide a textual anchor for the cross-border practitioner explaining the U.S. framework to foreign clients.

B. Create a disclosure regime for substantial AI involvement

The Copyright Office has indicated its intention to develop disclosure obligations for AI involvement in registered works. The current registration form does not require an applicant to disclose the use of AI in the creative process; the Copyright Office's administrative practice has been to ask follow-up questions when an application appears to involve AI, but the inquiry is initiated by the examiner rather than the applicant. The result is a disclosure regime that is, in operational practice, opt-in by examiner inquiry rather than opt-out by applicant disclosure, with predictable adverse selection consequences.

FTLA proposes that the Copyright Act be amended to require, as part of the registration process, affirmative disclosure of substantial AI involvement in the creation of the work. The disclosure would not, by itself, affect registrability; the human-contribution analysis would proceed on its merits with the disclosure as one input. The disclosure would, however, be a condition of registration, with material misrepresentation in the disclosure providing a basis for invalidation of the registration and (in egregious cases) referral for prosecution under 17 U.S.C. § 506(e). A parallel disclosure obligation would apply to patent applications

under the Patent Act, with material misrepresentation providing a basis for invalidation under the inequitable conduct doctrine.

The disclosure regime would address two of the policy concerns identified in Section IV. It would reduce the misrepresentation incentive that the current uncertain registration framework creates. It would generate, over time, a public record of the operational role of AI in commercially produced creative and inventive work, which would provide an empirical foundation for further policy refinement. And it would align the U.S. framework with the disclosure direction of the EU AI Act, reducing cross-border compliance complexity for developers operating in both markets.

C. Authorize a narrow training-data fair use safe harbor

The fair use analysis of generative AI training, currently being litigated in dozens of pending federal cases, will produce judicial answers over the next several years. Those answers will be valuable, but they will also be the product of fact-specific litigation in particular cases, with the resulting doctrinal framework calibrated to the particular litigation strategies of the parties to those cases. A more deliberate legislative response is available.

FTLA proposes that Congress authorize a narrow training-data fair use safe harbor with the following elements. First, the safe harbor would apply to training of AI models on copyrighted works obtained through lawful means (works available on the public internet, works licensed to the developer, works obtained through other lawful channels). Second, the safe harbor would condition fair use protection on the developer's compliance with a disclosure obligation (modeled on the EU AI Act Article 53) requiring publication of a sufficiently detailed summary of the training corpus. Third, the safe harbor would be subject to an opt-out mechanism, available at the work level or at the corpus level, that rights holders could exercise through technical or procedural means. Fourth, the safe harbor would not extend to training that is reasonably characterized as designed to produce output that substitutes for specific identified rights holders' works in those rights holders' markets, with the substitution analysis conducted under the existing Section 107 fourth factor framework.

The safe harbor would not displace the four-factor fair use analysis in cases falling outside its scope. It would provide a more predictable framework for the principal training activities at issue, while preserving the case-specific analysis for borderline cases. It would also align the U.S. framework with the EU and UK directions, reducing cross-border complexity and (potentially) supporting the development of an international consensus through the WIPO process.

D. Strengthen trade secret remedies for training data

The trade secret displacement effect identified in Section IV.D is a structural concern that the policy framework should address through means that do not require the abandonment of the trade secret framework. The trade secret system performs important functions, particularly in the protection of business-sensitive information that the patent and copyright systems are not designed to protect. The concern is not

that trade secret protection exists, but that its scope expands at the expense of patent and copyright as the share of AI-assisted invention and authorship grows.

FTLA proposes two refinements to the DTSA framework targeted at AI-related subject matter. First, the DTSA's reasonable measures requirement should be specified, in administrative guidance from the Department of Justice or in legislation, to address the distinctive features of AI training data: the volume of the data, the multiple parties typically involved in its assembly, and the operational difficulty of maintaining secrecy across the training pipeline. The specification would clarify the operational requirements that AI developers must satisfy to maintain trade secret protection over their training data and methodologies. Second, the DTSA's misappropriation remedies should be supplemented to provide for prospective injunctive relief in cases of inadvertent disclosure of AI training data through model output, with the relief calibrated to the particular technical features of the disclosure mechanism. The supplementation would address the operational reality that AI models may, in some circumstances, reveal training data through their output in ways that traditional trade secret doctrine has not adequately addressed.

VII. Implementation Considerations

The four policy recommendations set out in Section VI are calibrated to be implementable through existing legislative and administrative channels, without the creation of new federal agencies or comprehensive sui generis regimes. This section addresses the principal implementation considerations: legislative feasibility, administrative capacity, transition arrangements, and risks.

Legislative feasibility. The four recommendations require legislative action of varying degrees of magnitude. Recommendation A (codification of the human-contribution standard) is the most modest, in that it codifies existing administrative practice without altering substantive direction. Legislation of this character has, in past intellectual property amendments, been pursued through the omnibus appropriations vehicle or as part of broader intellectual property modernization legislation. Recommendation B (disclosure regime) is similarly modest, requiring an amendment to the Copyright Act and the Patent Act to add a disclosure obligation. Recommendation C (training-data fair use safe harbor) is more substantial, requiring a new statutory provision in the Copyright Act and (likely) implementing regulations from the Copyright Office and the USPTO. Recommendation D (trade secret remedies) requires an amendment to the DTSA. Each recommendation is, on its own, within the customary scope of intellectual property legislation. The combined package would require sustained Congressional attention but does not require a new statutory framework.

Administrative capacity. The recommendations rely on the existing administrative capacity of the Copyright Office and the USPTO. The disclosure regime under Recommendation B would require modifications to the registration and application forms, additional examiner training, and the development of standards for evaluating disclosed AI involvement. The Copyright Office and the USPTO have, in their respective administrative practices over the past three years, developed substantial expertise on these questions. The training-data safe harbor under Recommendation C would require the Copyright Office to develop the template for the corpus summary, modeled on the EU AI Office template that is being finalized in 2026. The Copyright Office's existing engagement with these questions through the Copyright and Artificial Intelligence Report process positions it to undertake this work.

Transition arrangements. The recommendations would, if enacted, apply prospectively to registrations and applications filed after their effective dates. The training-data safe harbor under Recommendation C raises a particular transition question: training conducted before the effective date would remain subject to fair use analysis under the pre-amendment framework, while training conducted after the effective date would be subject to the safe harbor. The transition arrangement should specify the treatment of models trained before the effective date and deployed thereafter, since the deployment activity may, on some characterizations, constitute a continuing use that the post-amendment framework would govern. FTLA recommends that the transition arrangement address this question explicitly, with the safe harbor's applicability tied to the training activity rather than the deployment activity.

Risks. The principal risks of the proposed framework are three. First, the codification of the human-contribution standard, while stabilizing the existing administrative analysis, may freeze that analysis in ways that make it less responsive to technological developments. The risk is mitigated by drafting the

statutory provision in terms general enough to accommodate developments while specific enough to provide guidance. Second, the disclosure regime creates risks of strategic non-disclosure or false disclosure, with the strength of the resulting enforcement architecture determining whether the regime achieves its informational objectives. The risk is mitigated by linking material misrepresentation to invalidation, providing a strong incentive for accurate disclosure. Third, the training-data safe harbor may, depending on its specification, prove either too narrow (failing to provide meaningful clearance for principal training activities) or too broad (foreclosing legitimate fair use challenges). The risk is mitigated by careful drafting of the safe harbor's scope, with reference to the developing case law and the EU and UK frameworks.

These risks are real but manageable. They do not, in FTLA's analysis, outweigh the benefits of the framework: a stabilized doctrinal architecture, reduced misrepresentation incentives, a more predictable training framework, and a measured response to the trade-secret displacement effect. The framework is offered as a structured starting point for the further policy development that the question demands.

VIII. Conclusion

This paper has argued that U.S. intellectual property law should respond to generative artificial intelligence through calibrated doctrinal refinement rather than wholesale recharacterization. The human authorship requirement in copyright and the natural-person inventorship requirement in patent are not technical obstacles to be removed in pursuit of broader protective scope; they are the constitutional and statutory foundations of the protective architecture, and their preservation is essential to the coherence of the system as a whole. The administrative authorities with primary responsibility for these questions, the Copyright Office and the USPTO, have, in their work since 2023, developed an analytical framework that is sound in principle and that should be codified, supplemented, and stabilized through the policy framework set out in Section VI.

Three points warrant emphasis as the paper closes.

First, the framing offered here is consistent with the established U.S. policy posture across administrations. The human-centered foundations of U.S. intellectual property law are bipartisan in their origin and bipartisan in their continued operation. The Copyright Office's three Reports on Copyright and Artificial Intelligence span both Democratic and Republican administrations. The USPTO's Inventorship Guidance was issued under one administration and remains in force under another. The Federal Circuit's holding in Thaler v. Vidal has been accepted across the executive branch and the courts. The framework offered in this paper builds on this consensus rather than challenging it.

Second, the framing offered here is grounded in the U.S. national interest in a functioning innovation system, not in the interests of any particular constituency. The case for the framework is not that AI developers deserve more or less than they currently receive, or that rights holders deserve more or less than they currently receive. The case is that the United States gets better outcomes, more reliable innovation incentives, more efficient cross-border legal architecture, more measured trade-secret displacement, when the protective framework is calibrated rather than left to evolve through litigation alone. The case is, in the precise sense, an argument about the operation of the U.S. innovation system in the same register as the SEC's mission of protecting investors, the Treasury Department's mission of safeguarding financial system integrity, and the Commerce Department's mission of promoting American economic competitiveness.

Third, the work proposed here is institutional and incremental, not regulatory in any expansive sense. The policy framework set out in Section VI does not call for new federal agencies, new sui generis protective regimes, or comprehensive regulation of artificial intelligence as a technology. It calls for modest amendments to the Copyright Act, the Patent Act, and the DTSA, building on the administrative work that the Copyright Office and the USPTO have already substantially completed. The marginal cost of implementing this framework is modest. The marginal benefit, by the analysis offered, is substantial: a U.S. intellectual property system that retains its human-centered foundations while accommodating the technological developments that have occupied the field over the past three years.

The Foreign Trained Lawyers Association will, in its applied scholarship and advocacy work over the coming years, continue to engage with the AI and intellectual property question. The Association's interest

in the question is grounded in the cross-border practice perspective developed in FTLA Policy Paper No. 1: the AI question is not, and will not be, resolved within any single jurisdiction, and the practitioners best positioned to advise on its evolving architecture are those who can read the U.S. framework against the EU, UK, and emerging WIPO frameworks. The Association looks forward to contributing to the policy development that the question requires.

References

This paper draws principally on U.S. government sources, including reports of the United States Copyright Office, guidance documents of the United States Patent and Trademark Office, statutes enacted by Congress, and decisions of the federal courts. Where comparative material is cited, the cited sources are official publications of the European Union, the United Kingdom, and the World Intellectual Property Organization. URLs and pinpoint citations are accurate as of the date of publication and may evolve thereafter.

1. U.S. Copyright Office, Copyright and Artificial Intelligence, Part 1: Digital Replicas (July 2024).

2. U.S. Copyright Office, Copyright and Artificial Intelligence, Part 2: Copyrightability (January 2025).

3. U.S. Copyright Office, Copyright and Artificial Intelligence, Part 3: Generative AI Training (December 2025).

4. U.S. Copyright Office, Compendium of U.S. Copyright Office Practices, Third Edition (2014, as updated).

5. U.S. Patent and Trademark Office, Inventorship Guidance for AI-Assisted Inventions, 89 Fed. Reg. 10,043 (Feb. 13, 2024).

6. Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022), cert. denied, 143 S. Ct. 1783 (2023).

7. Thaler v. Perlmutter, No. 22-cv-1564 (D.D.C. Aug. 18, 2023), aff'd, No. 23-5233 (D.C. Cir. 2025).

8. Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53 (1884).

9. Feist Publications, Inc. v. Rural Telephone Service Co., 499 U.S. 340 (1991).

10. Pannu v. Iolab Corp., 155 F.3d 1344 (Fed. Cir. 1998).

11. Copyright Act of 1976, Pub. L. No. 94-553, 90 Stat. 2541 (codified as amended at 17 U.S.C. §§ 101 et seq.).

12. Patent Act, 35 U.S.C. §§ 1 et seq., as amended by the Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011).

13. Defend Trade Secrets Act of 2016, Pub. L. No. 114-153, 130 Stat. 376.

14. Digital Millennium Copyright Act, Pub. L. No. 105-304, 112 Stat. 2860 (1998).

15. Authors Guild v. OpenAI, Inc., No. 23-cv-08292 (S.D.N.Y., consolidated 2024).

16. New York Times Co. v. Microsoft Corp., No. 23-cv-11195 (S.D.N.Y., filed Dec. 2023).

17. Andersen v. Stability AI Ltd., No. 23-cv-00201 (N.D. Cal., consolidated 2024).

18. Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), 2024 O.J. (L 168) 1.

19. Directive (EU) 2019/790 of the European Parliament and of the Council on Copyright in the Digital Single Market, 2019 O.J. (L 130) 92.

20. European Patent Office, Decision of the Legal Board of Appeal, J 8/20 (DABUS) (Dec. 21, 2021).

21. Thaler v. Comptroller-General of Patents, Designs and Trademarks, [2023] UKSC 49.

22. Copyright, Designs and Patents Act 1988 (U.K.), § 9(3).

23. United Kingdom Intellectual Property Office, Copyright and AI: Consultation (Dec. 2024).

24. World Intellectual Property Organization, Conversation on Intellectual Property and Artificial Intelligence (Sessions 1-3, 2020-2024).

25. Tennessee ELVIS Act, Tenn. Code Ann. § 47-25-1101 et seq. (2024).

26. Mariami Zhorzholiani, Foreign Trained Legal Expertise as Critical Infrastructure for U.S. Cross-Border Compliance, Capital Markets Integrity, and Inbound Foreign Direct Investment, FTLA Policy Paper No. 1 (April 2026).

27. Mariami Zhorzholiani, The Function of U.S. Corporate Governance in Investor Protection and Capital Markets Integrity, International Journal of Business Research and Management (2025).

Suggested citation: Mariami Zhorzholiani, Authorship, Originality, and the Machine-Generated Work: U.S. Copyright Doctrine, Patent Disclosure, and the Governance of Generative Artificial Intelligence, FTLA Policy Paper No. 2 (Apr. 2026).

— ❖ —