International Center for Law & Economics

Comments of the International Center for Law & Economics

Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

March 14, 2025

Authored by:

Kristian Stout (Director of Innovation Policy, International Center for Law & Economics)

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

Introduction

The International Center for Law & Economics (ICLE) appreciates the opportunity to respond to this request for information regarding the development of an AI Action Plan. ICLE is a nonprofit, non-partisan research organization that promotes the use of law & economics methodology to advance policy solutions that foster innovation, competition, and economic growth.

We note the significant policy shift marked by Executive Order 14179 ("Removing Barriers to American Leadership in Artificial Intelligence"), signed Jan. 23, 2025, which revoked the previous administration's AI Executive Order.[1] The new directive underscores the importance of developing a regulatory framework that balances sustaining the United States' leadership in AI innovation with the responsible management of potential risks.

Crafting such a framework presents urgent yet nuanced challenges. Policymakers must carefully calibrate regulations to encourage technological advancement and economic competitiveness without imposing unnecessary burdens that might stifle innovation or impede the dynamic growth of the AI sector. Our comments aim to provide a balanced, evidence-based approach to navigating these critical regulatory considerations.

I. Defining AI and Regulatory Scope

The development of an effective AI Action Plan requires careful consideration of how to define artificial intelligence. Crafting an appropriate definition of AI for regulatory purposes is complicated by the heterogeneous nature of AI technologies and the importance of avoiding regulatory fragmentation—both across the federal government's various agencies as well as among the states—that could undermine the United States' competitive position in AI development.

A. What Even Is 'AI'?

At root, defining AI as a single phenomenon presents significant conceptual and practical difficulties for regulatory frameworks. Unfortunately, international regulatory efforts like the EU's AI Act,[2] as well as domestic U.S. regulatory proposals,[3] frequently fail to appreciate this reality.

---

[1] Exec. Order No. 14179, 90 F.R. 8741 (2025), https://www.federalregister.gov/documents/2025/01/31/2025-02172/removing-barriers-to-american-leadership-in-artificial-intelligence.

[2] European Parliament Legislative Resolution of 13 March 2024 on the Proposal for a Regulation of the European Parliament and of the Council on Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM/2021/206, EUR. PARLIAM. (Mar. 13, 2024), *available at* https://www.europarl.europa.eu/doceo/document/TA-9-2024-0138_EN.html.

[3] *See* Kristian Stout & Subiksha Ramakrishnan, *ICLE Comments to CPPA on ADMT Regulations*, INT'L CTR. L. & ECON. (2025), https://laweconcenter.org/resources/icle-comments-to-cppa-on-amdt-amendments; Kristian Stout, *Biden's AI Executive Order Sees Dangers Around Every Virtual Corner*, TRUTH MARK. (Nov. 1, 2023), https://laweconcenter.org/resources/bidens-ai-executive-order-sees-dangers-around-every-virtual-corner.

Thinking of AI in overly broad terms risks creating definitions that do not adequately reflect the heterogeneity of AI systems.[4] This means potentially imposing premature and disproportionate obligations on businesses and stifling innovation at a critical stage of development. An overly broad definitional approach is likely to be analytically flawed and counterproductive, as it would fail to distinguish between high-risk AI applications with significant consumer impact and low-risk routine uses designed to improve business efficiency.[5] Such broad definitions can also unintentionally distort competition by favoring incumbent firms capable of bearing large compliance costs.

The heterogeneity of AI technologies calls for regulatory approaches that recognize substantive differences among, *e.g.*, large language models (LLMs), computer-vision systems, reinforcement-learning applications, and predictive-analytics tools. Each presents distinct regulatory challenges that would not be addressed adequately by a universal regulatory definition.

The danger of regulatory imprecision in this context cannot be overstated. When regulatory frameworks treat diverse technological applications as functionally equivalent, they inevitably produce inefficient and potentially counterproductive results. Overly broad definitions risk sweeping conventional software applications into the ambit of AI-specific regulations, potentially subjecting them to requirements ill-suited for their actual functionality and risk profiles. Conversely, narrowly tailored definitions that focus exclusively on specific AI implementations may fail to address novel applications or hybrid systems that do not fit neatly into predefined categories.

This definitional challenge is further complicated by the emerging patchwork of state and local AI regulations that has emerged in the absence of federal guidance. Developers and deployers of AI systems who operate across jurisdictional boundaries face substantial compliance challenges as a result of this regulatory fragmentation. The proliferation of potentially conflicting state regulations also creates significant legal uncertainty that disproportionately burdens smaller innovators and startups, as these entities often lack the resources to navigate complex regulatory environments. This has the potential to further entrench the market positions of larger incumbents.

B. Regulating AI to Protect Consumers and Promote Innovation

To address the foregoing definitional and jurisdictional challenges, the AI Action Plan should adopt a taxonomic approach to AI regulation that acknowledges the distinct characteristics and risk profiles of different AI domains. Critical questions remain regarding the specific nature of potential harms associated with these technologies. Among them, policymakers must clearly distinguish whether the principal concerns arise primarily from autonomous system behavior—similar to the issues raised by automated high-frequency trading[6]—or whether they stem primarily from malicious actors using AI

[4] *See* Lazar Radic & Kristian Stout, *What Is the Relevant Product Market in AI?*, SSRN (2024), at 109, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4927505.

[5] *Id.* at 110.

[6] *See* Tom C.W. Lin, *The New Investor*, 60 UCLA L. REV. 678 (2013).

tools to enhance their ability to break existing laws. In practice, these risks may coexist, requiring nuanced assessments.[7] This distinction is far from academic. Regulatory considerations differ significantly depending on the AI application's domain, whether it be autonomous vehicles, financial algorithmic trading, creative content generation, autonomous weapons, or predictive policing systems.

As part of this approach, the AI Action Plan should establish clear federal guidelines that preempt contradictory state and local regulations, while setting minimum transparency standards appropriate to each category of AI application.[8] The goal of such standards should be to protect consumers without imposing excessive compliance burdens that might stifle innovation. The aim should be to foster functional markets where customers can access the services they demand, not to initiate a new cottage industry for AI compliance lawyers.

Beyond transparency and preemption, the most effective regulatory regime for AI would adopt a harm-focused, incremental approach. Rather than erect burdensome new regulatory apparatuses, the goal should be to identify and address specific regulatory gaps that AI technologies could potentially exploit. An effective AI-governance strategy demands flexibility and adaptability, allowing regulatory frameworks to evolve organically in response to technological advancements and changing consumer needs. Policymakers should therefore adopt a dynamic, context-specific approach that seeks to address tangible, demonstrable harms without hindering experimentation or innovation.

Maintaining proper balance is essential to this strategy. AI regulation must not disproportionately emphasize risk mitigation to the exclusion of acknowledging and fostering AI's substantial potential benefits. Regulatory responses should be grounded in empirical assessments of tangible harms, while remaining mindful of the inherent complexities of risk assessment. As Aaron and Adam Wildavsky highlight, perceptions of technological risks—even among experts—vary significantly and are not strictly correlated with a given party's familiarity with actual hazards.[9] Even well-informed observers of AI and AI-related technologies "disagree[] significantly over how to interpret evidence, the relevance of speculative risks, and even the basic framing of AI-related threats."[10]

---

[7] *See, e.g.*, Josh Rosenberg *et al.*, *Roots of Disagreement on AI Risk: Exploring the Potential and Pitfalls of Adversarial Collaboration*, FORECAST. RES. INST. (2024), at 15 (studying the rather large divergence of optimistic and pessimistic views among AI experts on the potential long-term harms associated with AI, and suggesting the need to carefully parse the distinction between immediate—but less dramatic—potential harms, and long-term but speculative existential risks).

[8] Recognizing the importance of the nascent commercial spaceflight industry, Congress enacted the Commercial Space Launch Amendments Act of 2004. *Commercial Human Spaceflight Safety Regulations*, CONGR. RES. SERV. (*last updated* Feb. 5, 2025), *available at* https://sgp.fas.org/crs/space/IF12508.pdf ("For launch and reentry regulations, the Commercial Space Launch Amendments Act of 2004 set a statutory moratorium of eight years (the learning period) before the FAA could promulgate commercial human spaceflight regulations, beyond its statuary authorities described below. The learning period moratorium was intended to allow the nascent commercial spaceflight industry to develop without potential regulatory burdens.").

[9] *See* Aaron Wildavsky & Adam Wildavsky, *Risk and Safety*, ECONLIB, https://www.econlib.org/library/Enc/RiskandSafety.html (*last visited* Mar. 13, 2025).

[10] Rosenberg *et al.*, *supra* note 7, at 5.

Given these realities, policymakers should ensure their frameworks remain flexible, empirically informed, and responsive to evolving evidence. This approach would help to avoid regulatory overreach based on speculative or subjective assessments of AI's dangers, while continuing to provide appropriate safeguards where needed. Alternatively, a proportionate, risk-based framework that scales regulatory requirements according to actual risk and application context could also effectively balance innovation with necessary safeguards.[11]

Ideally, AI regulation should be directly responsive to empirically observed harms. The National Telecommunications and Information Administration (NTIA) previously developed a framework emphasizing the concept of marginal risks—those additional risks specifically attributable to the unique features of widely available foundation models, relative to closed models or non-AI alternatives.[12] The NTIA framework's strength lies in its empirical grounding, as it focused explicitly on the observable differences between open and closed AI models, as well as between AI and comparable non-AI technologies. By emphasizing concrete, empirically measurable factors that affect adoption, usage, and actual harms, the NTIA approach effectively avoids speculative, overly broad *ex-ante* regulatory interventions. Additionally, the NTIA framework realistically acknowledges inherent measurement challenges and uncertainties, cautioning policymakers to remain modest in their expectations regarding *ex-ante* risk assessments.

Such a framework would reserve heightened scrutiny for genuinely high-risk applications in sensitive domains such as healthcare, criminal justice, and critical infrastructure, while establishing regulatory safe harbors for good-faith AI development and deployment efforts that adhere to recognized best practices and standards. By crafting definitions and regulatory frameworks that acknowledge the heterogeneity of AI technologies, while also establishing federal preemption to prevent regulatory fragmentation, the AI Action Plan can foster an environment conducive to continued American leadership in AI innovation, while ensuring appropriate safeguards for critical concerns.

C. Factor Open Source into Regulatory Considerations

Discussions about AI regulation often focus exclusively on proprietary, commercial-end products, such as foundation models and highly visible consumer-oriented applications. This narrow focus, however, overlooks the crucial role of open-source software development within the broader AI technological ecosystem. Open-source methodologies underpin much of the AI stack and contribute significantly to innovation and competition.[13] Unlike proprietary AI systems developed by established firms, open-source AI emerges organically through distributed networks of developers,

---

[11] *See* Kristian Stout *et al.*, *NIST AI 800-I, Managing Misuse Risk for Dual-Use Foundation Models*, INT'L CTR. L. & ECON. (2024), at 8-13, *available at* https://laweconcenter.org/wp-content/uploads/2024/09/NIST-AI-comments-final.pdf.

[12] *See Dual-Use Foundation Models with Widely Available Model Weights*, NAT'L TELECOMM. INFO. ADMIN., *available at* https://www.ntia.doc.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf, at 2 (*last visited* Mar. 13, 2025).

[13] Alex Engler, *How Open-Source Software Shapes AI Policy,* BROOKINGS INST. (Aug. 10, 2021), https://www.brookings.edu/articles/how-open-source-software-shapes-ai-policy.

making traditional regulatory approaches—often predicated on centralized corporate structures—potentially inappropriate or counterproductive.

The complexity of AI technology extends across multiple interconnected layers that often run on open-source projects—each posing distinct regulatory considerations. At the foundational level lies hardware infrastructure, including semiconductors, computing power, and "XaaS" (everything-as-a-service) offerings that provide virtual-computing resources.[14] Above this is the data layer, encompassing collection, curation, and preparation processes that significantly influence an AI system's quality and effectiveness.[15] The subsequent model-training layer involves diverse methodologies such as supervised, unsupervised, reinforcement, and transfer-learning techniques.[16] Finally, the deployment layer comprises various operational environments, ranging from cloud-based platforms to edge devices and on-premises systems.

Given this multi-layered ecosystem, regulations designed with assumptions suited to proprietary, corporate-controlled products may inadvertently disadvantage open-source innovation. Blanket requirements for data access, security, and compliance—while they may be appropriate for centralized entities—could unintentionally suppress open-source initiatives by imposing disproportionate burdens and risks. Restrictions that make it more difficult to develop, modify, or distribute open-source AI models could inadvertently shift AI innovation toward proprietary, closed models controlled by large incumbents, thereby reducing market diversity and innovation.

Moreover, restrictive regulations could hinder the broader social and economic benefits derived from open-source AI, threatening to render the United States less competitive globally.[17] Open-source models enable widespread experimentation and collaborative development, thereby contributing significantly to technological innovation, academic research, and overall productivity

---

[14] Romit Dey & George Korizis, *How Anything-As-A-Service (XaaS) Can Help Reinvent Business Models and Transform Outcomes Across Industries*, PRICEWATERHOUSECOOPERS, https://www.pwc.com/us/en/services/consulting/business-transformation/library/use-xaas-to-reinvent-business-models.html (*last visited* Mar. 13, 2025).

[15] *See, e.g.*, *Structured vs Unstructured Data*, IBM (Jun. 29, 2021), https://www.ibm.com/think/topics/structured-vs-unstructured-data; Dongdong Zhang *et al.*, *Combining Structured and Unstructured Data for Predictive Models: A Deep Learning Approach*, BMC MED. INFORMATICS DEC. MAKING 280 (2020), https://link.springer.com/article/10.1186/s12911-020-01297-6 (describing generally the use of both structured and unstructured data in predictive models for health care).

[16] ANIL ANANTHASWAMY, WHY MACHINES LEARN: THE ELEGANT MATH BEHIND MODERN AI (2024), at 12-13.

[17] *See, e.g.*, Michael Chui, *et al.*, *The Economic Potential Of Generative AI: The Next Productivity Frontier*, MCKINSEY DIGITAL (2023), https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier ("Generative AI's impact on productivity could add trillions of dollars in value to the global economy. Our latest research estimates that generative AI could add the equivalent of $2.6 trillion to $4.4 trillion annually across the 63 use cases we analyzed—by comparison, the United Kingdom's entire GDP in 2021 was $3.1 trillion. This would increase the impact of all artificial intelligence by 15 to 40 percent. This estimate would roughly double if we include the impact of embedding generative AI into software that is currently used for other tasks beyond those use cases"); *Generative AI Could Raise Global GDP by 7%*, GOLDMAN SACHS (Apr. 5, 2023), https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent ("As tools using advances in natural language processing work their way into businesses and society, they could drive a 7% (or almost $7 trillion) increase in global GDP and lift productivity growth by 1.5 percentage points over a 10-year period").

growth.[18] Moreover, open-source communities inherently foster adaptive and nuanced self-regulatory practices.[19]

II. AI and Copyright

The intersection of AI and intellectual-property law presents complex legal questions that the AI Action Plan must thoughtfully address to foster innovation, while safeguarding creators' interests. Copyright protection fundamentally rests on a utilitarian economic premise: by conferring exclusive rights to creators, society generates greater long-term welfare through the stimulation of creative production.[20] This justification recognizes that creative works have the inherent characteristics of public goods; specifically, they are non-excludable (it is difficult to prevent unauthorized consumption) and non-rivalrous (consumption by one does not diminish availability to others). Absent legal protection for creators, these characteristics would lead to market failure through underproduction, as creators would face diminished incentives once their works could be freely reproduced without compensation, thereby reducing expected returns on creative investment.

Consequently, copyright protection inherently confers on copyright holders the right to restrict others' use of protected materials, even if such use is arguably a net benefit for society. Recognizing this potential, copyright doctrine incorporates countervailing mechanisms—such as limited duration and fair-use exceptions—that function as pressure-release valves for compelling public-interest considerations. The fundamental tension in copyright law emerges from the simultaneous imperatives to ensure compensation for creators, while facilitating public access to and use of creative works, thereby realizing the broader social benefits of cultural and technological advancement.

Efforts to seek equilibrium in this complex system can be analogized to hydraulics. Much as pressure applied to fluid in one chamber necessitates compensatory movement elsewhere in a hydraulic system, the strengthening of creator rights in one domain often requires corresponding flexibility in another to maintain the copyright system's balance. Adjustments to any component of this interdependent system inevitably generate ripple effects throughout the broader copyright ecosystem.

The hydraulic nature of copyright incentives is particularly salient in the context of technologies like artificial intelligence, which fundamentally challenge established frameworks. Altering how copyright functions at the input stage of AI training—by liberalizing or constraining access to the corpus of available training materials—may require changes to creators' property rights at the output stage, in order to preserve the system's overall balance. Further, permitting the use of copyrighted

---

[18] Miguel A. Cardona *et al.*, *Artificial Intelligence and the Future of Teaching and Learning*, U.S. DEPT. OF EDUC., https://www2.ed.gov/documents/ai-report/ai-report.pdf (*last visited* Mar. 13, 2025).

[19] *See, e.g., Hugging Face*, GPT-4CHAN, https://huggingface.co/ykilcher/gpt-4chan (*last visited* Mar. 13, 2025).

[20] For more on the economics of copyright, *see* Brent Luches, *Introduction* 1-3, *in* IDENTIFYING ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE FOR COPYRIGHT POLICY (U.S. Copyr. Off., 2025), *available at* https://www.copyright.gov/economic-research/economic-implications-of-ai/Identifying-the-Economic-Implications-of-Artificial-Intelligence-for-Copyright-Policy-FINAL.pdf.

works for AI-model training could potentially alter creator incentives in the long term. At the same time, granting rightsholders absolute veto power over such uses might disproportionately impede AI developers' ability to construct effective and socially beneficial systems.

A. The Law & Economics of AI Training and Fair Use[21]

A foundational principle in copyright jurisprudence is the idea-expression dichotomy, whereby protection extends exclusively to a particular expression of an idea, rather than to the underlying idea itself.[22] Alternative articulations of the same concept typically do not constitute infringement under established copyright doctrine. Within this framework, the fair-use doctrine potentially provides a legal pathway for AI developers to incorporate copyrighted materials as training inputs for their models.

The applicability of fair use in this context does, however, remain subject to substantial dispute.[23] Various well-reasoned arguments have been advanced that forward radically different interpretations of how traditional copyright principles do or should apply to AI training. This interpretive uncertainty suggests that AI training represents a genuinely *sui generis* phenomenon that challenges conventional doctrinal boundaries and may require changes to intellectual-property law in order to resolve the tensions between AI innovation and copyright protection. This presents practical problems when trying to imagine how to facilitate both AI training and remuneration to creators.

For example, the imposition of individualized licensing requirements for copyrighted materials used in AI training likely presents insurmountable practical impediments and prohibitive transaction costs.[24] Contemporary large-scale AI models typically incorporate billions of inputs gathered from across the digital ecosystem. The negotiation of discrete licenses for each copyrighted work—or comprehensive agreements with each rightsholder—within such expansive corpora would require an extraordinary volume of transactions, potentially numbering in the millions.[25] The administrative

---

[21] This discussion is limited to "generative AI" such as ChatGPT, Claude, and Llama. Narrower applications of AI (both generative and non-generative)—*e.g.*, medical-record scanning to create predictive models of disease—may or may not need access to similar sized data sets. And even within generative AI, the fair-use analysis can be complicated when the system being trained is meant to be a direct market-substitute for the material on which it is trained. *See, e.g.,* Kristian Stout, *AI Training Is Not Fair (According to One Court)*, TRUTH MARK. (Feb. 11, 2025), https://truthonthemarket.com/2025/02/11/ai-training-is-not-fair-according-to-one-court.

[22] *See, generally, Baker v. Selden*, 101 U.S. 99 (1879); *Golan v. Holder*, 565 U.S. 302 (2012).

[23] *See, e.g.,* Kristian Stout, *supra* note 21 (discussing recent caselaw and the nuanced fair-use analysis in order to parse when AI training should and should not be considered "fair use."); *see also* Kristian Stout, Geoffrey A. Manne, & Emily Corbeille, *ICLE Comments on Artificial Intelligence and Copyright*, INT'L CTR. L. & ECON. (2025), https://laweconcenter.org/resources/icle-comments-on-artificial-intelligence-and-copyright.

[24] *See* RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW 42 (7th ed., 2007), (discussing the transaction costs involved with copyright as including the tracing costs of identifying the copyright holder and negotiation costs of negotiating the license with the copyright holder).

[25] *See* Jorge Padilla & Kadambari Prasad, *Demystifying Licensing Debates: Should GenAI Developers Pay to Train Their Models on Copyright Protected Content?*, COMPASS LEXECON (Feb. 25, 2025), https://www.compasslexecon.com/insights/publications/demystifying-licensing-debates-should-genai-developers-pay-to-train-their-models-on-copyright-protected-content.

burdens inherent in such a process would likely render comprehensive licensing functionally impossible at the requisite scale for effective AI development.[26]

Given the impracticality of per-work licensing mechanisms, scholars have explored collective-licensing frameworks as a potential alternative.[27] This approach envisions a centralized entity functioning as a clearinghouse to negotiate comprehensive licenses on behalf of substantial rightsholder constituencies, analogous to the operational model of music performing-rights organizations. The primary advantage offered by such collective-licensing arrangements lies in their capacity to substantially reduce transaction costs through rights aggregation.[28] Conceptually, this methodology could enhance the accessibility and economic feasibility of training data acquisition.[29] Nevertheless, formidable obstacles would persist in implementing such a system.

Foremost among these challenges is the problem of incomplete participation; not all content proprietors would necessarily affiliate with collective-licensing entities, resulting in significant coverage disparities.[30] The sources constituting AI-training datasets exhibit extraordinary diversity and dispersion; consequently, a collective-licensing regime might encompass only certain categories of creative works—such as those under the control of major publishing houses or contained within established image repositories—while overlooking independent creators who lack institutional representation.[31]

Further, a considerable proportion of internet content needed for AI-model training remains outside the purview of such administrative frameworks. AI developers would encounter significant difficulties in identifying all relevant rightsholders with whom licensing agreements would be obligatory. Even assuming these independent creators could be successfully identified, the aforementioned transaction costs would prove prohibitively burdensome. This would likely result in diminished AI-model capabilities due to input constraints. Alternatively, such models might be compelled to rely predominantly or exclusively on synthetic data, potentially compromising model quality and performance.[32]

When considering limits on AI-training data, we must also acknowledge the importance of dataset diversity to prevent algorithmic bias. A system in which only commercially licensed content is available for training could produce AI models that disproportionately reflect perspectives from

---

[26] *Id.*

[27] *Id.*

[28] *Id.*

[29] *Id.*

[30] *See* Michael D. Smith & Rahul Telang, *The Effect of AI Ingestion on Rightsholders' Incentives* 35-38, *in* IDENTIFYING THE ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE FOR COPYRIGHT POLICY (U.S. Copyr. Off., 2025), *available at* https://www.copyright.gov/economic-research/economic-implications-of-ai/Identifying-the-Economic-Implications-of-Artificial-Intelligence-for-Copyright-Policy-FINAL.pdf (discussing the limitations of collective licensing for AI training).

[31] *See id.* at 37.

[32] *See* Maggie Harrison Dupre, *When AI Is Trained on AI-Generated Data, Strange Things Start to Happen*, FUTURISM (Aug. 2, 2023), https://futurism.com/ai-trained-ai-generated-data-interview.

entities with market presence or established licensing frameworks. This would likely create systems with significant blind spots, particularly regarding independent creators and noncommercial knowledge sources.

1. What is the value of content for models?

If the value of creative works to AI-model training fundamentally cannot be assigned at the input stage, this would severely restrict the development of efficient licensing markets for AI training data. Beyond the question of transaction costs, two principal complications emerge when focusing on input-licensing markets: the negligible marginal value of individual works within extensive training datasets and the methodological challenges of value attribution. No established framework exists for calculating the monetary contribution of specific works as AI training inputs, rendering the process of determining fair compensation inherently fraught.

The U.S. Copyright Office has reached a similar conclusion: the incorporation of millions or billions of works into foundation models necessarily attenuates the influence of any individual copyrighted work to such a degree that even minimal transaction costs associated with licensing negotiations would invariably exceed that work's proportional contribution to the model's utility.[33] While there are certainly qualitative distinctions among training-data elements, they are, at most, only relative. For instance, the collected works of Isaac Asimov represent an undeniably significant contribution to English-language literature; nevertheless, they constitute only a tiny fraction of the English-language corpus. Even exceptionally valuable literary properties would command only nominal compensation in the context of an extensive training dataset. This fundamental reality renders conventional valuation methodologies impracticable; a simple calculation that applies a standardized per-work fee across millions of works would inevitably yield a result fundamentally disconnected from each work's actual impact on model performance.

A principal factor complicating valuation efforts is the monetization structure of generative AI, which manifests predominantly at the output stage, rather than during data ingestion. Training data is neither directly commercialized nor consumed; instead, economic value materializes when the model generates outputs—whether textual, visual, or otherwise—for which users demonstrate willingness to pay, or which facilitate commercial applications. The contribution of any specific training example remains indirect and inextricably intermingled with innumerable others.

Consequently, absent methodologies to establish causal connections between generated content and specific elements within training datasets (assuming such a thing is possible), any attempt to assign monetary valuations to individual training components inevitably relies on highly speculative assumptions. For the majority of discrete works, particularly those created by independent producers

---

[33] *See* Adam Jaffe, *Controlling the Use of Copyrighted Materials in Training* 50, *in* IDENTIFYING THE ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE FOR COPYRIGHT POLICY (U.S. Copyr. Off., 2025), *available at* https://www.copyright.gov/economic-research/economic-implications-of-ai/Identifying-the-Economic-Implications-of-Artificial-Intelligence-for-Copyright-Policy-FINAL.pdf.

distributed throughout the digital ecosystem, no established market rate exists for "AI-training use." The value proposition is inherently context-dependent and, on a per-work basis, typically *de minimis*.

B. Moral Rights and Attribution in AI Outputs

A better approach would focus on addressing copyright concerns at the output stage, rather than imposing restrictive controls at the input level. Given that AI-generated outputs may strongly evoke copyrighted inputs,[34] policymakers should explore legal frameworks that ensure creators have meaningful control and compensation rights. Emphasizing output-based protections, rather than strict input constraints, would strike a balanced approach by accommodating AI's significant potential for innovation, while safeguarding creators' legitimate interests.

Existing legal concepts like the common-law "right of publicity" may offer useful analogies.[35] Many states already provide legal protection against unauthorized commercial exploitation of an individual's identity, likeness, or voice. Policymakers should consider adapting these standards explicitly for AI contexts, ensuring individuals would retain adequate control over their unique personal attributes when reflected in AI-generated outputs. Such adaptations would extend rights akin to existing publicity standards, granting individuals enforceable claims against unauthorized, substantially similar AI reproductions of their likeness or distinctive style.

Given these dynamics, the AI Action Plan should actively explore new compensation models that leverage outputs, rather than inputs. One promising direction involves reforming the current patchwork moral-rights system in the United States. As recognized by the U.S. Copyright Office, the existing fragmented landscape—where moral-rights protections are inconsistently applied through state laws and limited federal statutes such as the Visual Artists Rights Act (VARA)[36]—is inadequate to address attribution and integrity concerns in digital and AI-generated contexts.[37] Specifically, the Copyright Office identified significant gaps that are exacerbated by digital environments, particularly around the removal or manipulation of copyright-management information and attribution metadata.

Policymakers could explore changes, such as possible amendments to the Lanham Act, to restore and clarify attribution protections diminished by the U.S. Supreme Court's decision in *Dastar Corp. v. Twentieth Century Fox*.[38] Additionally, there are currently a number of cases examining the use

---

[34] Notably, if outputs actually duplicate existing works, it is likely they would already be found impermissible infringements under existing U.S. law.

[35] *See* Shane Greenstein, *Commercial Exploitation of Name, Image, and Likeness* 24-30, *in* IDENTIFYING THE ECONOMIC IMPLICATIONS OF ARTIFICIAL INTELLIGENCE FOR COPYRIGHT POLICY (U.S. Copyr. Off., 2025), *available at* https://www.copyright.gov/economic-research/economic-implications-of-ai/Identifying-the-Economic-Implications-of-Artificial-Intelligence-for-Copyright-Policy-FINAL.pdf.

[36] 17 U.S.C. § 106A(a).

[37] *Authors, Attribution, and Integrity: Examining Moral Rights in the United States*, U.S. COPYR. OFF. (2019), at 59-60, https://www.copyright.gov/policy/moralrights.

[38] *Id.* at 42-54.

Section 1202 of Title 17 to address the intentional removal or altering of copyright-management information in digital and AI-generated contexts. These cases should be tracked closely, as they may reveal further needed changes to the copyright regime to both enable creators to assert attribution rights, while also demonstrating where existing copyright law frustrates large-scale training.[39]

Furthermore, achieving a balance between moral-rights protections and AI innovation would necessitate nuanced, sector-specific adjustments similar to those recommended for VARA. Moral rights should not be applied uniformly across all digital contexts; instead, nuanced legislative adjustments are needed to address unique challenges posed by generative AI without stifling innovation.

III. Conclusions and Recommendations

In developing the AI Action Plan, we recommend the following:

- Establish clear, nuanced definitions of AI that recognize the internal heterogeneity of technologies and prevent fragmented regulatory approaches.
- Prioritize empirically grounded, harm-focused regulatory frameworks over speculative, overly broad *ex-ante* restrictions.
- Ensure that federal standards provide transparency and clarity, preempting conflicting state and local regulations.
- Foster an environment that supports both proprietary and open-source AI development, avoiding measures that inadvertently favor large incumbents.
- Adjust intellectual-property frameworks—particularly copyright—to reflect AI's unique characteristics, emphasizing balanced protections and output-based compensation.

By adopting these targeted measures, the AI Action Plan will effectively promote innovation, safeguard competition, and secure U.S. leadership in AI technologies.

[39] *See, e.g.*, *Kadrey v. Meta Platforms Inc.*, 23-cv-03417-VC (N.D. Cal. Nov. 20, 2023); *The Intercept Media Inc. v. OpenAI Inc.*, 24-cv-1515 (JSR) (S.D.N.Y. Nov. 21, 2024); *Raw Story Media Inc. v. OpenAI Inc.*, 24 Civ. 01514 (S.D.N.Y. Nov. 7, 2024).