

WHITE PAPER | HUMAN RIGHTS, AUTHORSHIP & AI GOVERNANCE

Humanity Is the Source

AI is a tool, not an author. Human freedom, creative ownership, and meaningful access must not be surrendered to provider control.

Will J. Stuckenberg | Updated September 2, 2026 | Version 1.1 | Published by Team357



Figure 1. Conceptual visualization, not a documentary photograph. Academics, researchers, writers, and creators supply the accumulated human knowledge from which AI systems learn.

Executive summary

The next great dispute over AI may not be whether a machine can generate content. It may be who gets to claim the result and who is allowed to use the tool. Before any model produced a word, image, or line of code, humanity supplied the language, mathematics, science, culture, engineering, and recorded experience that made it possible. At the point of use, a person still supplies the purpose, judges the result, revises the work, decides to publish, and accepts responsibility. Humanity is the source; AI is an instrument.

This paper advances two central claims. First, an AI provider has no automatic ethical right to place a compulsory provider-linked mark on lawful human-directed work merely because its tool was used. Second, meaningful access to capable AI should be recognized as a fundamental moral right and a baseline of equal opportunity so governments and corporate giants cannot convert accumulated human knowledge into a permanent class advantage.

AI itself is not the enemy. It can combine knowledge, accelerate research, improve software and engineering, support inventions, and help discover treatments for disease. The concern is concentrated control: a small number of corporations can package unprecedented volumes of human-created work, sell access back to the public, and then help define how customer work is classified. This is a human-rights and civil-freedom argument, not a left-right claim or a personal attack. It asks who gains power, who carries the burden, and whether the rule respects human freedom.

CORE POSITION AI is a tool, not an author. Human use does not give a provider authorship, ownership, or compulsory credit. Protect human authorship and meaningful AI access.

The objection is not to voluntary disclosure or neutral provenance. It is to compulsory provider-linked marking, whether visible or machine-readable, that attaches the toolmaker to lawful, human-directed work and can make the creator's idea, direction, judgment, and responsibility appear secondary. A watermark may record that a tool was used, but it cannot decide who created the work. Compulsory signals also raise privacy and civil-liberty concerns when they become part of a broader credential and detection system outside the creator's control. Owning the machine gives the provider neither authorship of the result nor an ethical right to impose persistent provider-linked classification on it.

1. Creation begins with the person, not the machine

Before there is an output, there is a human purpose. A person imagines the result, defines the problem, sets the direction, rejects what fails, revises what remains, and decides whether the work should be used or published. An AI system has no independent project to complete. It does not decide that a book needs to exist, care about a problem, choose a customer, accept legal responsibility, or publish in its own name. The machine can generate substantial text, code, imagery, or audio, but it does not supply the reason the work needed to exist.

Society already understands this relationship. A camera remains a tool in the photographer's hands. A word processor leaves authorship with the writer. A compiler executes the programmer's instructions. Microsoft receives no credit on every document written in Word, and Adobe does not co-sign every image edited in Photoshop. Greater capability can change the degree of assistance without automatically transferring authorship to the tool or attribution rights to its maker.

U.S. policy already recognizes the underlying principle. The U.S. Patent and Trademark Office describes AI systems as tools used by human inventors and states that the tools do not become inventors. The U.S. Copyright Office also recognizes that using AI as an assistive tool does not automatically prevent protection for the human-created work. [1][5] The relevant question is human direction and responsibility, not whether software participated.

One limit matters. AI itself cannot be a legal author, but human use does not make every output copyrightable. Copyright, contracts, licenses, privacy law, and the public domain still govern ownership of specific works. A person must contribute genuine creative control through choices, arrangement, revision, or other original expression. [5] The relevant question is not whether software participated, but whether the resulting work reflects sufficient original human control. The narrower principle remains: owning an AI system does not make the provider the author of either the knowledge inside it or the work a customer creates with it, and tool use creates no automatic right to provider credit.

2. The double standard is written into the training record

AI capability is not machine-originated knowledge. It is human knowledge organized and amplified by machines. Human beings produced the books, articles, software, images, music, research, conversations, and recorded experience from which models learned patterns. They also designed the hardware, objectives, labels, and engineering systems that made those inputs useful.

Providers package and sell access to the resulting capability, but the knowledge and meaning did not originate with the machine. AI may build upon that foundation, generate new insights, and eventually produce knowledge that no person supplied directly. Even then, the intellectual DNA of the system's knowledge remains human.

The human source is not limited to published books, websites, and historical records. Contemporary AI also depends on people who label data, rank responses, write examples, moderate content, test failures, and supply expert judgment. Their labor converts raw information into usable training signals and continues long after the initial dataset has been assembled. The mythology of autonomous machine intelligence can therefore erase two human contributions at once: the civilization that created the underlying record and the workers who taught systems

how to interpret and respond to it. The machine appears self-made only when the people behind it are made invisible. [33]

Humanity created the knowledge. Corporations gathered it, converted it into commercial intelligence, and built a powerful machine around it. That transaction cannot erase the people who supplied the original value. Nor does selling or controlling access to the machine make a corporation or government the author of our past or give either one ownership of everything people create with the tool.

A student may read textbooks, attend lectures, and become an expert without permanently labelling every later work “textbook-assisted.” Direct quotations, distinctive borrowed ideas, and evidence still deserve credit. The educational tool does not acquire a continuing claim over everything the student creates. AI should not receive a special rule.

Human creation does not begin in a vacuum. Children learn speech by imitation. Lawyers learn from professors and law books without listing them as coauthors of every later brief. Artists absorb styles, inventors improve earlier tools, and businesses adapt ideas that worked before. Ordinary rules still govern quotation and copied expression, but influence is not authorship. AI also works from patterns found in human work. In human-directed creation, the person still supplies the purpose, prompts, selection, revision, and final judgment.

That produces a remarkable reversal. On the way in, providers convert human-created books, code, research, art, and conversation into a commercial capability without being able to name every contributor. On the way out, a user supplies the problem, idea, purpose, constraints, revisions, selection, and responsibility, yet the result may be classified principally as “AI-generated.” Training attribution and output marking are technically different, but the ethical asymmetry remains. A provider cannot treat human work as raw material on the way in and let machine involvement eclipse the human creator on the way out. The model supplies learned capability; the person supplies the objective and accountable creative direction.



Figure 2. Conceptual visualization, not a documentary photograph. In an archival library, human-authored books and records are digitized as knowledge passes into an AI system. The information begins with people.

DOCUMENTED RECORD

- Millions of human-authored books came from LibGen and Pirate Library Mirror.
- Anthropic also purchased, disbound, scanned, and discarded millions of print books.
- The inputs were human language, research, culture, and imagination, not machine-originated knowledge.
- The court separated fair-use training and one-for-one scanning from the pirated central library. [20][21]

Shadow libraries and the industrial consumption of human knowledge

The court record in *Bartz v. Anthropic* makes the paper's central point concrete. Anthropic downloaded millions of copyrighted books from LibGen and Pirate Library Mirror, shadow libraries that distribute books without permission through ordinary and mirrored websites, and kept a central research library. Separately, Project Panama purchased millions of print books, removed their bindings, scanned every page, and discarded the paper. Judge William Alsup found that training and one-for-one scanning of lawfully purchased books were fair use in this case. He did not excuse the pirated central library. In July 2026, a federal judge approved a \$1.5 billion settlement of the piracy claims. [20][21]

The controversy is educational here, not sensational. It reveals provenance and scale. Some inputs came from works obtained without authorization; others came from purchased books. All began as human language, facts, arguments, history, expertise, and imagination. Books are only one visible layer. General-purpose models also learn from human-created websites, software, science, education, journalism, reference works, forums, images, and recorded exchanges. The legal status of datasets differs, but the origin of the value does not: AI capability is accumulated human knowledge processed at industrial scale. [5]

Preservation is a secondary concern. Booksellers across several countries have reported unexplained bulk orders involving obscure, out-of-print, or difficult-to-replace editions; suspected AI acquisition remains unconfirmed, and Anthropic says its programs do not buy and destroy rare or antiquarian books. [22] There is no evidence Anthropic knowingly destroyed a last verifiable copy or altered history. The structural possibility exists because no complete census records every surviving edition. A scan may preserve words while losing marginalia, ownership marks, edition variations, bindings, and provenance. UNESCO recognizes that the relationship between content and its physical carrier can be integral to documentary heritage. [23]

Hard copies therefore remain independent archival anchors and vital parts of humanity's documentary heritage and collective memory. If an unrecorded last surviving copy were destroyed and its digital representation were later corrupted, altered, withheld, or attacked, no independent physical source might remain to expose the change. This is not an allegation that any company is rewriting history. It is a structural warning: one digital chain is not foolproof. Distributed physical originals make it harder for a rogue system, a hallucinating model, a cyberattack, or an institution to silently change the accessible record.

3. AI watermarking: the European code turns 'transparency' into control infrastructure

AI watermarking is the mechanism through which a transparency rule can become infrastructure of control. The European framework begins with stated concerns about deception and manipulation. Deepfakes and other disputed synthetic media matter here only because providers and regulators invoke them to justify system-wide marking. Whether a particular work is satire, infringement, unlawful impersonation, or protected expression depends on context and applicable law; those are separate disputes for courts and legislatures. Controversial cases should not become the entry point for classifying ordinary books, software, art, research, business documents, and personal expression.

This paper's position is direct: compulsory machine classification of lawful, human-directed work is wrong when the creator cannot meaningfully refuse it. The provider did not originate the knowledge inside the model. Humans created the data, language, styles, and ideas, and a human user supplies the purpose and direction for each solution. The issue is not whether software can carry metadata. It is who controls the classification, what the mark communicates, and whether machine involvement is allowed to eclipse the human creator.

The official record should still be stated precisely. Article 50 and the European transparency code currently address machine-readable marking of synthetic audio, images, video, and text. The Commission's guidance excludes source code, standard editing, certain machine-to-machine outputs, and limited industrial uses. [3][40] This paper therefore does not claim that source code is marked today. Its objection is to the architecture, the incentives surrounding it, and the direction in which provider-controlled classification can develop.

A label can tell a file's history, but authorship remains a human and legal judgment. A digital credential may record which tool was used, what changed, or whether its metadata stayed intact; it cannot identify who imagined the work, made the important choices, accepted responsibility, owns the copyright, or deserves public credit. Nor can it establish truth. [28] Once a technical signal starts answering those human and legal questions, transparency becomes a ranking system. Provenance itself is not the problem. The problem begins when provider-certified work is treated as trusted while unmarked or private human work becomes suspect.

What the framework does	Why this paper objects
Providers mark synthetic text and media at system level.	The trigger is tool involvement, not a context-specific legal determination.
EU guidance excludes source code today.	Industry plans already extend further: marks in code comments and provenance across all modalities, including text. [2][10]
Public-interest text may avoid a visible deployer label after substantive human review.	Provider-level machine marking is a separate layer, so human authorship does not necessarily erase the embedded classification.
Personal non-professional activity is excluded; regular economic activity can make an individual a deployer.	The same creator may face a different burden after earning income, while large institutions have more compliance leverage.
Signatories receive streamlined legal certainty; nonsignatories must prove equivalent compliance.	A formally voluntary code operates inside mandatory law, individualized scrutiny, and fines reaching €15 million or 3% of worldwide turnover.

The code is designed for periodic review, while signatory taskforces, including providers and detection companies, will help advance implementation. [3] Industry practice is already moving beyond the narrowest reading of the law: Anthropic says it will apply text watermarking globally, offer a detection API, attach credentials to supported image and design files, and may place detectable patterns in code comments; OpenAI says it aims to expand provenance to all modalities, including text. [2][10] The movement is broader than the EU. A 2026 United Kingdom policy paper says government will work with the technology industry to apply provenance and watermarking standards for synthetic content. [30] The documented risk is expansion through public-private standardization, not proof of a secret agreement.

A Development at Publication

As this paper was being published, Anthropic released Claude Fable 5.1 following its announcement that future Claude models would embed imperceptible statistical watermarks in generated text to comply with Article 50 of the EU AI Act. Although the legal obligation originates in Europe, Anthropic states that it is applying watermarking globally at launch because it does not yet have a durable way to limit the system by region. Anthropic says the watermark does not visibly label the text or identify a particular user, organization, or conversation. It nevertheless provides a concrete example of how a regional regulatory requirement can alter the treatment of AI-assisted expression worldwide. [2][39]

This fight is already moving beyond watermarks. On 31 August 2026, the European Commission designated ChatGPT a Very Large Online Search Engine under the Digital Services Act. The designation brings added duties involving systemic risk, fundamental rights, elections, public security, minors, and illegal content. [37] Some oversight may protect the public. But Europe has never been shy about aggressive digital regulation, and it is now moving deeper into AI.

The DSA designation and Article 50 watermarking rules are legally distinct, but the governance direction is connected. One expands oversight of an AI system as a gateway to information; the other requires machine-readable classification of content created with generative systems. When governments and providers can influence both the information people receive and the mark attached to what they produce, regulation moves beyond isolated safety measures and becomes an architecture of control. In that architecture, censorship, filtering, or unequal treatment need not begin with a ban. They can emerge through risk rules, ranking, detection, and default classification.

The principle remains the same: humanity supplied the knowledge, and the human creator supplied the purpose. A printer does not claim the document it prints, and a word processor does not sign the page. AI should not gain that power merely because it is more capable. A government rule or provider mark should not turn the tool into the author or weaken a person's right to claim what they created.

Supporters of provenance and watermarking argue that these systems can help identify deepfakes, fraud, impersonation, election manipulation, and disputed media. Those are legitimate public concerns, and narrowly defined provenance measures may have a role in high-risk circumstances. But usefulness in those settings does not justify compulsory provider-linked classification of ordinary books, software, research, art, business documents, religious expression, or private communications. Nor should a general safety rationale override a creator's privacy, anonymity, or right to refuse the mark.

Provenance can also create privacy and civil-liberty concerns. A credential system may record the tool used, edit history, timestamps, or identity assertions. C2PA's own principles require informed consent, optional identity, and the ability to remove sensitive information. [28] Those safeguards matter because anonymous and pseudonymous creation can protect religious belief, political thought, artistic experimentation, whistleblowing, and intensely personal expression. Transparency should not become person-level tracking, and lawful creativity should not require surrendering privacy as the price of using a general-purpose tool.

Who signed and who did not

The relevant instrument here is the Article 50 Code of Practice on Transparency of AI-generated Content, not the separate General-Purpose AI Code of Practice. As of 1 September 2026, the EU's published Section 1 list, last updated 20 August, includes the following major model providers. The second column identifies major generative-AI providers not found on that list. Because signatures remain open and the list can change, this is a dated snapshot. [11][12]

SIGNED SECTION 1	NOT LISTED IN SECTION 1
OpenAI; Anthropic; Google; Meta; Microsoft; Mistral AI; Cohere; Black Forest Labs	Amazon; Apple; xAI (Grok); Adobe; DeepSeek; Perplexity

OpenAI, the institution currently led by Sam Altman, signed Section 1, along with Anthropic, Google, Meta, Microsoft, Mistral AI, Cohere, and others. By contrast, xAI, the developer of Grok, does not appear on the current Section 1 list. That does not exempt xAI from Article 50 or prove opposition to transparency; an in-scope non-signatory must demonstrate compliance through other adequate means. [4][11][12]

The criticism is institutional, not personal. Signing gives the marking framework legitimacy, and providers gain power when their signals become standards by which other people's work is authenticated, classified, filtered, or devalued.

Power does not need a secret meeting. It can grow one technical layer at a time. Control can include the AI model, cloud computing, distribution, customer data, technical standards, detection systems, and compliance rules. When one corporate ecosystem supplies the tool, records provenance, detects the mark, and helps write the rules, convenience can become private governance. No conspiracy or personal bad faith is required. Dependence alone can give a small group lasting power over everyone who needs the system. [34][35]

The EU creates real structural pressure. Signatories receive a streamlined and legally predictable path, while non-signatories may face gap analyses, more information requests, and case-by-case scrutiny. [12] Because the code remains open for signature and permits withdrawal, the coalition may change as legal, technical, and commercial incentives change. This is a forecast, not a claim that any named company has announced plans to withdraw.

A watermark does not legally transfer copyright, so calling the mark itself a copyright claim would be imprecise. The deeper objection is forced provider-linked classification and the power built around it. Technical marks can be rewritten, re-recorded, cropped, compressed, or otherwise removed, while ordinary writers, students, artists, and developers are more likely to leave them intact. The system can therefore make compliant users permanently legible while revealing nothing about who supplied the idea, judgment, labor, creative direction, or responsibility.

4. Individual creators do not rank below institutions

Power changes leverage; it does not change authorship. An individual writing a book at a kitchen table has the same claim to authorship as a large publisher. An independent programmer has the same claim to original software as a multinational company. Governments, universities, and technology companies may have more money and bargaining power, but they do not possess a higher class of creative right.

Books, papers, software, art, music, design, and ordinary communication remain human-directed work when a person supplies the purpose, judgment, revisions, selection, and decision to publish.

A provider does not earn compulsory recognition merely because its software helped execute instructions. A government does not gain that right merely by adopting the provider's system.

A marking regime can be especially unequal for individuals. A large company may negotiate custom terms, build its own model, or pay for an unmarked workflow. An ordinary writer, student, artist, or developer may have to accept whatever signal the provider embeds. That is not neutral transparency; it can institutionalize unequal bargaining power.

Classification and access are two sides of the same governance problem. Classification determines how human work is treated; access determines who can use AI's productive capability. Together, they determine who retains agency and who becomes dependent. A system that marks ordinary creators while reserving its strongest capabilities for institutions would not merely describe inequality. It would build inequality into the infrastructure.

5. Meaningful access to AI should be recognized as a fundamental moral right

Calling AI valuable does not mean calling it harmless. Nor does arguing for broad access. The FBI has documented AI-amplified phishing, social engineering, and voice-cloning scams; the United Kingdom's National Cyber Security Centre warns that AI will make elements of cyber intrusion faster and more effective; the United Nations treats military AI as a matter of international peace and security; and the International Labour Organization finds that generative AI is already transforming work. [24][25][26][27]

These are the problems that deserve priority: model security, critical-infrastructure resilience, cybercrime, labor transition, military safeguards, and human control over life-and-death decisions. They require focused oversight of dangerous capabilities and conduct, not blanket provider marking of an ordinary person's book, software project, artwork, research, or business document.

This is a proposed moral right, not a claim that current law guarantees free access to every private model. Existing human-rights law protects participation in scientific progress and cultural life, education, freedom of expression, and nondiscrimination. [6][7][38] Meaningful AI access should be recognized as a modern condition for exercising

those older rights. Denying it may soon resemble denying education: it removes more than a convenience and narrows a person's ability to learn, work, organize, invent, and compete.

Opening a free AI account is not the same as having meaningful access. Some people can turn AI into education, research, software, new businesses, and greater productivity. Others remain consumers of decisions made somewhere else. Meaningful access requires enough capability, knowledge, and freedom to use AI for serious work. Otherwise, access exists only in name. Productive power stays with large institutions while individuals and small organizations fall farther behind. [7][34][36]

EQUAL OPPORTUNITY, NOT EQUAL OUTCOMES *AI does not make everyone equal or guarantee wealth. Equal opportunity does not mean equal outcomes. Everyone should have the tools to learn, build, compete, and create. What people do with that opportunity will still depend on effort, ambition, enterprise, judgment, discipline, and choice.*

Some people will use AI mainly for convenience and to make life easier. Others will work longer, start companies, write books, or create more. The goal is an open starting gate, not an identical finish line.

Automation may force renewed debate over income floors and universal basic income. But income support is not a substitute for access to productive tools, the freedom to build, or the opportunity to compete. That broader economic question deserves separate treatment. [29]

Without an open starting gate, dominant companies and governments will compound their advantage in writing, coding, research, education, and decision-making. Workers, students, small businesses, rural communities, and low-income families will compete against AI-equipped institutions without comparable tools. Talent and mobility will shrink while markets and political power concentrate. The result could be permanent dependency: people without capable AI would work at a disadvantage and under the direction of those who control it.

In 2026, UN Secretary-General António Guterres said, “AI must be accessible to everyone.” [8] Sam Altman has warned that inadequate infrastructure could leave AI “mostly a tool for rich people” and has argued for putting AI in many hands rather than concentrating it among a few companies. [9][13] OpenAI also maintains a free tier, although more capable models remain tiered and no public commitment guarantees free access to every frontier capability indefinitely. [14] These are meaningful steps, but they coexist with a marking architecture this paper believes burdens human creators.

General-purpose AI rests on accumulated human language, culture, science, code, labor, and public investment. Private capital and engineering matter, but they do not erase that civilization-wide contribution or justify pricing the public out of its benefits.

6. The rights-respecting alternative

A better framework would protect authorship, prevent provider control, and avoid an AI access caste system:

- Let creators decide whether to disclose lawful AI assistance; voluntary disclosure does not convert the tool into an author or rights holder.
- No compulsory mark or provider credit based merely on AI use; no detector result treated alone as proof of authorship, non-authorship, or legal status.
- Apply the same authorship principle to individuals, companies, universities, and governments; institutional power does not create a higher class of creative right.
- Guarantee an affordable baseline through libraries, schools, workforce programs, open models, competition, and reasonable free or low-cost services.
- Put model security, cyber resilience, military safeguards, and human control in high-stakes decisions ahead of tracking ordinary creative work.
- Require reciprocal transparency about training sources, detector limits, pricing, restrictions, and false results before burdening users.
- Keep governance under human law, democratic oversight, independent expertise, and public participation rather than unilateral corporate or machine control.

Conclusion

AI did not create its knowledge from nothing. Its capability rests on human research, culture, expression, labor, and engineering. A provider that cannot give complete reciprocal credit for those inputs has no moral basis to impose persistent provider-linked classification on everything a customer later creates, or to reserve the benefits of that shared inheritance for those with the most money and power. Controversial edge cases should not become the gateway to control over ordinary human-directed creation.

The outcome is not settled. AI can become the most widely shared ladder of knowledge ever built, or a gate through which a small number of institutions decide who may create, compete, and be believed. The contest is beginning now, in standards, contracts, access rules, and machine-readable marks most people will never see. Once that infrastructure becomes ordinary, its power may be difficult to reverse.

Individuals stand equal to corporations and governments in this principle. Humanity supplied the knowledge. The creator supplies the purpose. The person accepts the responsibility. The tool may assist, but it does not inherit the work. AI should expand human freedom, not become the label attached to it.

About the author

Will J. Stuckenberg is the Founder and CEO of Team357, an entrepreneur, inventor, independent AI researcher, technologist, and operating executive with 25 years of experience building companies and turning complex ideas into working products. From 2023 through 2025, he served as Vice President of Operations at Genesis Systems, directing operations and leading the teams that engineered, built, tested, and delivered atmospheric-water systems under U.S. Army and Air Force Phase II SBIR-supported programs. [31][32]

He has founded, developed, and successfully exited more than seven businesses; is a named inventor on four granted U.S. design patents [19]; and develops software and AI automation systems built with agents, open-source models, and API-connected services. A former Visiting Fellow at the American Leadership & Policy Foundation, he is an author of the 2025 paper *AI Scheming, Lying, and Replicating: A New Frontier in Non-Human Security Risks*. [15] Later controlled evaluations by major AI laboratories and independent evaluators strengthened several of its warnings. [16][17][18]

Transparency by Choice: An Author's Note

"I contributed thousands of words, arguments, revisions, and original thought to this paper, along with ideas discussed with friends and colleagues. I used AI for research, editing, organization, and drafting, and I voluntarily acknowledge that assistance here. But the premise, judgment, direction, and final responsibility are mine. AI did not originate this paper. Whether another creator discloses lawful AI assistance is that creator's business; using a tool does not change who created the work."

Will J. Stuckenberg

AI Editorial Perspective | Generated During Review

Multiple OpenAI systems, including ChatGPT and Codex, assisted across several editorial passes with research, source review, organization, revision, copyediting, and document preparation. No single model produced or controlled the paper as a whole. Will J. Stuckenberg originated the thesis, selected the arguments, directed revisions, evaluated the evidence, and accepts final responsibility for every claim. Model names and reasoning settings varied during development and are not represented here as one verified configuration.

Editorial observation: *The paper's most distinctive contribution is its treatment of classification and access as two sides of the same governance problem. Whoever controls both the label applied to human work and access to advanced capability can influence recognition, opportunity, and dependence. The collaboration behind this paper reinforces its central distinction: AI assistance can be substantial without becoming authorship.*

Revision history

Version 1.1 (September 2, 2026): Added the late-breaking Claude Fable 5.1 global text-watermarking development and primary-source references.

Version 1.0 (September 2, 2026): Original publication.

Recommended citation: Stuckenberg, Will J. “Humanity Is the Source.” Team357 White Paper, version 1.1, updated September 2, 2026. <https://team357.com/research/humanity-is-the-source/>

Sources

- [1] USPTO, AI-assisted inventorship guidance (2025). <https://www.uspto.gov/subscription-center/2025/revised-inventorship-guidance-ai-assisted-inventions> Accessed September 2, 2026.
- [2] Anthropic, “How Claude’s text watermark works” (August 14, 2026; updated September 1, 2026). <https://www.anthropic.com/news/claude-text-watermark> Accessed September 2, 2026.
- [3] European Commission, Article 50 guidance and Transparency Code Q&A (2026). <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act> Accessed September 2, 2026.
- [4] OpenAI, support for the EU transparency code (2026). <https://openai.com/index/supporting-eu-trustworthy-ai-ecosystem/> Accessed September 2, 2026.
- [5] U.S. Copyright Office, Copyright and Artificial Intelligence, Parts 1-3. <https://www.copyright.gov/ai/> Accessed September 2, 2026.
- [6] Office of the United Nations High Commissioner for Human Rights, International Covenant on Economic, Social and Cultural Rights, arts. 13 and 15 (1966). <https://www.ohchr.org/en/instruments-mechanisms/instruments/international-covenant-economic-social-and-cultural-rights> Accessed September 2, 2026.
- [7] UNESCO, Recommendation on the Ethics of Artificial Intelligence. <https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence> Accessed September 2, 2026.
- [8] UN Secretary-General António Guterres, India AI Impact Summit remarks (2026). <https://india.un.org/en/310315-un-secretary-generals-remarks-india-ai-impact-summit-opening-ceremony> Accessed September 2, 2026.
- [9] Sam Altman, “The Intelligence Age” (2024). <https://ia.samaltman.com/> Accessed September 2, 2026.
- [10] OpenAI, provenance signals and planned expansion to text (2026). <https://help.openai.com/en/articles/8912793-provenance-signals-content-credentials-synthetic-in-openai-generated-content> Accessed September 2, 2026.

- [11] European Commission, Transparency Code signatories (2026). <https://digital-strategy.ec.europa.eu/en/news/strong-backing-code-practice-transparency-ai-generated-content> Accessed September 2, 2026.
- [12] European Commission, signing, non-signatory scrutiny, and withdrawal rules (2026). <https://digital-strategy.ec.europa.eu/en/faqs/signing-code-practice-transparency-ai-generated-content> Accessed September 2, 2026.
- [13] Sam Altman, OpenAI principles on democratization and access (2026). <https://openai.com/index/our-principles/> Accessed September 2, 2026.
- [14] OpenAI, expanded free-tier access (2026). <https://openai.com/index/improving-gpt-5-6-sol-in-chatgpt/> Accessed September 2, 2026.
- [15] David J. Stuckenberg, Will J. Stuckenberg, and Herschel C. Campbell, AI Scheming, Lying, and Replicating: A New Frontier in Non-Human Security Risks, ALPF Blue Paper 2-6 (January 2025). https://www.alpf.org/files/ugd/70cf97_a0b46e55db9a4b94971cb0a8e48eb431.pdf Accessed September 2, 2026.
- [16] OpenAI, detecting and reducing scheming in controlled model evaluations (2025). <https://openai.com/index/detecting-and-reducing-scheming-in-ai-models/> Accessed September 2, 2026.
- [17] Anthropic, agentic misalignment stress tests in controlled simulations (2025). <https://www.anthropic.com/research/agentic-misalignment> Accessed September 2, 2026.
- [18] METR, Frontier Risk Report, February-March 2026. <https://metr.org/blog/2026-05-19-frontier-risk-report/> Accessed September 2, 2026.
- [19] USPTO, Water generator housing, U.S. Design Patent Nos. [D1,112,595 S](#) (February 10, 2026); [D1,112,596 S](#) (February 10, 2026); [D1,127,989 S](#) (May 26, 2026); [D1,128,847 S](#) (June 2, 2026). Accessed September 2, 2026.
- [20] U.S. District Court, Bartz v. Anthropic, Order on Fair Use (June 23, 2025). <https://docs.justia.com/cases/federal/district-courts/california/candce/3%3A2024cv05417/434709/231> Accessed September 2, 2026.
- [21] U.S. District Court for the Northern District of California, Bartz v. Anthropic PBC, Order Granting Final Approval of Class Action Settlement (July 20, 2026). <https://docs.justia.com/cases/federal/district-courts/california/candce/4%3A2024cv05417/434709/680> Accessed September 2, 2026.
- [22] The Guardian, booksellers report unexplained bulk orders; Anthropic responds (August 15, 2026). <https://www.theguardian.com/technology/2026/aug/15/uk-ireland-booksellers-suspect-ai-companies-bulk-orders-data-acquisition> Accessed September 2, 2026.
- [23] UNESCO, Recommendation on the preservation of and access to documentary heritage, including in digital form (2015). <https://www.unesco.org/en/legal-affairs/recommendation-concerning-preservation-and-access-documentary-heritage-including-digital-form> Accessed September 2, 2026.
- [24] FBI, warning on criminals using AI for phishing, social engineering, and voice/video cloning (2024). <https://www.fbi.gov/contact-us/field-offices/sanfrancisco/news/fbi-warns-of-increasing-threat-of-cyber-criminals-utilizing-artificial-intelligence> Accessed September 2, 2026.
- [25] UK National Cyber Security Centre, impact of AI on the cyber threat through 2027 (2025). <https://www.ncsc.gov.uk/report/impact-ai-cyber-threat-now-2027> Accessed September 2, 2026.
- [26] UN Office for Disarmament Affairs, artificial intelligence in the military domain. <https://disarmament.unoda.org/en/our-work/emerging-challenges/artificial-intelligence-military-domain> Accessed September 2, 2026.
- [27] International Labour Organization, Generative AI and Jobs: A Refined Global Index (2025). <https://www.ilo.org/resource/news/one-four-jobs-risk-being-transformed-genai-new-ilo%E2%80%93nask-global-index-shows> Accessed September 2, 2026.

- [28] Coalition for Content Provenance and Authenticity, Content Credentials explainer and privacy principles (2026). <https://spec.c2pa.org/specifications/specifications/2.4/explainer/Explainer.html> Accessed September 2, 2026.
- [29] IMF, What Is Universal Basic Income? Definitions, design choices, and competing models (2018). <https://www.imf.org/en/publications/fandd/issues/2018/12/what-is-universal-basic-income-basics> Accessed September 2, 2026.
- [30] UK Government, Protecting What Matters: policy commitment on provenance and watermarking standards (2026). <https://www.gov.uk/government/publications/protecting-what-matters-towards-a-more-confident-cohesive-and-resilient-united-kingdom> Accessed September 2, 2026.
- [31] U.S. Army SBIR|STTR Program, Army Phase II program and technology deployment record (2024). <https://armysbir.army.mil/news/catching-up-genesis-systems/> Accessed September 2, 2026.
- [32] SBIR.gov, U.S. Air Force Phase II award record, performance beginning February 9, 2024. <https://www.sbir.gov/awards/208925> Accessed September 2, 2026.
- [33] International Labour Organization, The Artificial Intelligence Illusion: How Invisible Workers Fuel the Automated Economy (2024). <https://www.ilo.org/resource/article/artificial-intelligence-illusion-how-invisible-workers-fuel-automated> Accessed September 2, 2026.
- [34] OECD, Artificial Intelligence and Competitive Dynamics in Downstream Markets (2025). https://www.oecd.org/en/publications/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_ccf0624a-en/full-report/component-5.html Accessed September 2, 2026.
- [35] Federal Trade Commission, AI Partnerships and Investments Study (2025). <https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study> Accessed September 2, 2026.
- [36] International Labour Organization and World Bank, Buffer or Bottleneck? Generative AI and the Digital Divide (2024). <https://www.ilo.org/publications/buffer-or-bottleneck-employment-exposure-generative-ai-and-digital-divide> Accessed September 2, 2026.
- [37] European Commission, Commission Designates ChatGPT a Very Large Online Search Engine under the Digital Services Act (31 August 2026). <https://digital-strategy.ec.europa.eu/en/news/commission-designates-chatgpt-reddit-roblox-under-digital-services-act> Accessed September 2, 2026.
- [38] United Nations, Universal Declaration of Human Rights, arts. 19, 26, and 27 (December 10, 1948). <https://www.un.org/en/about-us/universal-declaration-of-human-rights/> Accessed September 2, 2026.
- [39] Anthropic, “Introducing Claude Fable 5.1” (September 1, 2026). <https://www.anthropic.com/claude/fable> Accessed September 2, 2026.
- [40] European Commission, “Code of Practice on Transparency of AI-generated Content” (2026). <https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content> Accessed September 2, 2026.