PRIMER SERIES ON AI & COPYRIGHT

PRIMER 2

Human Authorship in Copyright Law

The Dialogue®
INFORM ENGAGE IDEATE

MAY 2026

PRIMER 2

Human Authorship in Copyright Law

Authored by Akshat Agrawal

The Dialogue is a public policy think tank with a vision to drive a progressive narrative in India's policy discourse. Founded in 2017, we believe in facilitating well-researched policy debates at various levels to help develop a more informed citizenry, on areas around technology and development issues. The Dialogue has been ranked as the world's Top 10 think tanks to watch out for, by the Think Tank and Civil Societies Programme (TTCSP), University of Pennsylvania in their 2020 and 2021 rankings.

About This Series

This primer is the first in a multi-part series examining the intersection of artificial intelligence and copyright law, with specific attention to the Indian legislative and judicial context. The series is intended for policymakers, legal practitioners, think-tank researchers, and industry stakeholders engaged in/contributing to shaping India's regulatory response to generative AI. Each primer is self-contained but designed to build upon the others.

For more information

visit thedialogue.org.in

Suggested Citation

Agrawal, A. (May, 2026). Primer 2: Human Authorship in Copyright Law. The Dialogue.®

Catalogue No.

TD/ET/PR/0526/06

Publication Date

May 14, 2026

Contents

Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
CDPA	Copyright, Designs and Patents Act
CCH	*CCH Canadian Ltd. v. Law Society of Upper Canada*
CIPPIC	Canadian Internet Policy and Public Interest Clinic
DABUS	Device for the Autonomous Bootstrapping of Unified Sentience
DALL·E	Digital Artwork and Language Learning Model
ECR	European Court Reports
EU	European Union
IP	Intellectual Property
KCC	Korea Copyright Commission
LLM	Large Language Model
PETA	People for the Ethical Treatment of Animals
PRC	People's Republic of China
TRIPS	Trade-Related Aspects of Intellectual Property Rights
U.K.	United Kingdom
U.S.	United States
WIPO	World Intellectual Property Organization

Executive Summary

Generative artificial intelligence has compelled copyright systems worldwide to confront a question that copyright doctrine has, until recently, taken for granted: whether the entity at the origin of a creative work must be a human being. Primer 1 of this series mapped the AI content pipeline and proposed a three-category taxonomy of AI outputs. Before the question of who, if anyone, owns those outputs can be resolved, two anterior questions must be answered.

The first is whether copyright should attach to AI-generated content at all, that is, whether the law should create a regime of exclusion in respect of works whose production does not require the kind of market-based recompense that copyright was instituted to provide. Property rights in expression are not free goods: their cost is the curtailment of speech, communication, and access. They are justified, where they are justified, by the need to prevent the disincentivisation of human creators in market societies. Whether that justification extends to AI outputs is the threshold question.

The second is, if some form of protection is to be conferred, why courts and legislatures have historically insisted on human authorship as a precondition for copyright, and what the doctrinal, theoretical, and practical content of that requirement is. This primer addresses the second question while keeping the first squarely in view.

The primer surveys the doctrinal foundations of the human authorship requirement in the Anglo-American, Continental European, and Indian traditions; the consistent refusal of courts to recognize non-humans, animals, mechanical processes, automated photographic apparatus, supernatural intelligences, and pre-generative computer programs, as authors; and the current global divergence on AI authorship across the United States, the United Kingdom, the European Union, China, Australia, Canada, Japan and South Korea, the Czech Republic and other civil-law and Commonwealth jurisdictions, and India.

Its central thesis is that the question of who should own AI outputs cannot responsibly be answered without first interrogating whether ownership ought to be recognised at all, and that any such answer must engage with, and not evade, the doctrinal reasons why copyright has, until now, been understood as a regime for human creative labour.

1. Introduction | The Threshold Question

Public discussion of the copyright treatment of AI-generated content has tended to skip a step. The most visible debate, who should own the output of a generative model: the developer, the operator, or the end user, assumes that some entity should own it.

Copyright is not a default state of affairs. It is a legislatively constructed regime of exclusion, justified historically and doctrinally by reference to a specific problem: the difficulty of sustaining human creative labour in market societies that deal in non-rivalrous, easily reproducible intangibles. Whether that problem arises in respect of AI-generated content is a question that has not been seriously addressed in most of the contemporary debate.

This primer therefore proceeds in two analytical stages. The first stage, addressed in Section 2, asks whether property rights in AI-generated content are justified at all, and what the implications of recognising or refusing such rights would be for the public domain, for human creators, and for the integrity of the copyright system. The second stage, addressed in Sections 3 to 7, examines how copyright doctrine has historically resolved the prior question of who counts as an author. Together, the two stages provide the framework necessary to evaluate the more granular questions of attribution and ownership that subsequent primers will address.

2. The Anterior Question | Why Property Rights at All?

2.1 The Disincentive-Avoidance Function of Copyright

Copyright is often described, including in judicial opinions, as an 'incentive' to create. That formulation, while widely deployed, is doctrinally and historically imprecise. Most creators do not produce because they expect a copyright. Engineers do not invent for patents, novelists do not write for royalties in the first instance, musicians do not compose anticipating mechanical reproduction rights. The empirical evidence on creative motivation, in industries from poetry and literary fiction to academic research and open-source software, suggests that intrinsic motivation, peer recognition, and vocational identity drive creative labour at least as powerfully as expected economic returns.[1]

The more accurate, and more historically grounded, account of copyright's function is one of disincentive avoidance[2]. Human creators in market societies face a particular structural problem: the resources they produce are intangible, infinitely replicable, and non-rivalrous in consumption. Without some legal mechanism enabling them to recover the cost of their creative labour, they will be unable to sustain that labour as a vocation, and will, in time, be displaced into other forms of work. Copyright is the legal mechanism – historically, the principal one – designed to prevent that *particular* mischief. It is calibrated to enable survival, not to maximise prosperity.

This distinction is important. A regime calibrated to prevent disincentivisation is conceptually narrow, it expands only to the point necessary to ensure that the creative vocation remains economically viable. A regime calibrated to provide positive incentives is conceptually open-ended – there is no natural limit on the size of the inducement that can be justified by reference to its incentive effect. The drift in copyright law over recent decades, towards longer terms, broader subject-matter coverage, and stronger remedies, owes much to the ascendancy of the incentive framing over the disincentive-avoidance framing.

[1] On the disincentive-avoidance framing as opposed to the positive-incentive framing, see generally Mark A Lemley, 'Faith-Based Intellectual Property' (2015) 62 UCLA L. Rev. 1328; James Boyle, The Public Domain: Enclosing the Commons of the Mind (Yale, 2008).

[2] Goold, Patrick R. and Simon, David A. (2024) "On Copyright Utilitarianism," Indiana Law Journal : Vol. 99: Iss. 3, Article 2. <https://www.repository.law.indiana.edu/ilj/vol99/iss3/2>

For the purposes of this primer, the disincentive-avoidance framing is doctrinally and policy-relevantly correct, and the question whether AI-generated content warrants copyright protection should be answered against that framing.

2.2 Non-Rivalrous Intangibles and Market Survival

The structural problem that copyright addresses arises only because human creators must survive in a market society. The non-rivalrousness of intangibles is a problem for human creators because, absent a legal mechanism to internalise the value of their work, they cannot sustain themselves. Intangibles are not rivalrous goods like food or fuel; one person's enjoyment of a poem does not deprive another. In a market that prices goods by reference to scarcity, intangibles will, absent legal intervention, command no price.

The disincentive-avoidance function of copyright is a response to this market-failure framing. It creates artificial scarcity in respect of intangibles so that human creators can recover the costs of their creative labour and continue to produce.

The relevant question for the AI debate is whether AI systems face the same problem. As of now, they plainly do not. AI systems are not market participants. They do not require recompense for their continued operation; their operation is funded by the entities that develop, train, and deploy them. They do not have a vocation that they will be displaced from in the absence of revenue. They have no subjective interest in the survival of any particular form of life, creative or otherwise. The structural justification for copyright simply does not apply to them.

2.3 The Inapplicability of the Justification to AI Systems

It is sometimes argued that, even if AI systems do not require copyright for their own survival, the developers of those systems require copyright in the outputs to recover the substantial sunk costs of training and deploying generative models. This argument is structurally different from the classical disincentive-avoidance justification, and it is worth being explicit about that difference.

The classical justification is concerned with the conditions under which a particular kind of human labour, creative labour, can be sustained as a vocation. Investment in AI infrastructure is not creative labour in this sense. It is capital investment in tooling, calibrated to a wholly different economic logic. Copyright is not, on its classical justification, a mechanism for guaranteeing returns on capital, as was the regulatory regime prevalent during the time of the Stationers Company,

which was explicitly rejected in the late 1600s. It is a mechanism for ensuring that human creators can continue to create.

If the policy goal is to incentivise investment in AI infrastructure, copyright is, at best, a poor tool. Other regimes, patent (where eligibility criteria are met), trade secret law, contractual mechanisms, sui generis investment-protection rights of the kind found in the European Union's Database Directive,[3] exist precisely to address that goal. None of them depends on the recognition of authorship in machine output.

There is a second consideration. The conferral of copyright on AI-generated content does not flow downstream to AI systems (which do not benefit from it), or even reliably to AI developers (who, in the absence of statutory clarity, may not be the persons who 'cause the work to be created' within the meaning of Section 2(d)(vi)[4]). It flows, on most analyses, to the operators and end users who deploy AI systems in market settings. Whether those persons require copyright protection in order to be willing to deploy AI systems is, on the available evidence, doubtful: AI deployment is overwhelmingly driven by demand-side factors and infrastructural cost economics, not by expected copyright revenue from outputs. Indeed, AI deployment has scaled massively in jurisdictions that have not yet recognised copyright in pure AI outputs.

2.4 The Public Domain as Default

If the disincentive-avoidance justification does not extend to AI-generated content, the doctrinal default is the public domain. Copyright is a deliberate carve-out from a default rule that intangibles are freely available. The carve-out exists because of a particular policy problem. Where that problem does not arise, the default rule should govern.[5]

The public domain is not a legal failure. It is the substrate on which all creative practice, including the practice of human creators who later qualify for copyright protection, depends. The Berne Convention's exclusion of 'news of the day' and 'miscellaneous facts' from copyright is one acknowledgment of this; the idea-expression dichotomy codified in TRIPS Article 9(2) is another. Treating purely AI-generated content as presumptively in the public domain would simply extend an existing structural feature of copyright law to a new technological context.

[3] Compare Council Directive 96/9/EC of 11 March 1996 on the legal protection of databases (sui generis right) <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:31996L0009> last accessed on 29th April, 2026.
[4] Copyright Act, 1957, s. 2(d)(vi).
[5] On the public domain as default, see Pamela Samuelson, 'Mapping the Digital Public Domain: Threats and Opportunities' (2003) 66 Law & Contemp. Probs. 147; James Boyle, supra n. 1.

Several jurisdictions have moved, explicitly or implicitly, in this direction. The United States Copyright Office's Compendium of Practices specifies that the Office will refuse to register works produced by mechanical processes 'with no creative input or intervention from a human author'[6]. That position is consistent with treating purely AI-generated content as public-domain material. The Czech Republic's Municipal Court of Prague reached the same conclusion in respect of a DALL-E generated image in 2023[7].

2.5 The Cost of Exclusion in Speech and Communication

Property rights in expression are not costless. Their conferral creates exclusion: the rightsholder can prevent others from making, distributing, performing, or adapting the protected work. In the context of human-created works, that cost is justified by the disincentive-avoidance function. In the context of AI-generated content, the cost is the same, but the justification is absent.

The cost is significant. Generative AI systems can produce content at scale that vastly exceeds human creative capacity. If the outputs of those systems are routinely subject to copyright, the result is a significant enlargement of the universe of protected expression, protected without reference to any disincentive-avoidance rationale, and with substantial downstream effects on speech, education, research, journalism, and creative practice.

There is a competitive dimension as well. Human creators who do not use AI systems will, ex hypothesi, be competing with AI-assisted creators in the same markets. If AI-generated content attracts copyright protection on equivalent terms with human-created content, the structural advantage of AI-assisted creators, their ability to produce vastly more content for a given input cost, drawing on the corpus of works on which the model has been trained (and comparably much larger than that a human can ingest in a whole lifetime, consciously or subconsciously) will be reinforced rather than offset by copyright doctrine. Whether that is a desirable policy outcome is, at the very least, an open question.

The threshold question, then, is whether the cost of exclusion is justified in respect of AI-generated content. Only if the answer is affirmative does the question of who should hold the copyright become tractable. Sections 3 to 7 of this primer proceed on the assumption, for the sake of doctrinal exposition, that some answer to that secondary question may eventually be required, while keeping the prior question squarely in view.

---

[6] Compendium of U.S. Copyright Office Practices (3rd ed., 2021), § 313.2 <https://www.copyright.gov/comp3/> last accessed on 29th April, 2026.

[7] Municipal Court of Prague (Městský soud v Praze), judgment of 11 October 2023 (the so-called 'DALL-E case'); English-language summary in IPKat and at <https://www.taubel.cz/> last accessed on 29th April, 2026.

3. The Human Authorship Requirement | Historical and Doctrinal Foundations

3.1 The Anglo-American Tradition

The human authorship requirement is not a recent doctrinal innovation. It is implicit in the Statute of Anne, 1710, the first modern copyright statute, which conferred protection on the 'Author or Authors' of 'Books' in language that contemplated only human creators[8]. The Constitution of the United States, in its grant of power to Congress to 'promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries'[9], uses the word 'Authors' in a sense that the U.S. Supreme Court has read as denoting a human creative agent.

The leading early decision is Burrow-Giles Lithographic Co. v. Sarony[10]. The case concerned a photograph of Oscar Wilde, and the question was whether photographs produced by what was then a novel mechanical apparatus, were 'writings' of an 'Author' within the meaning of the constitutional grant. The U.S. Supreme Court answered in the affirmative, but on the basis that the photograph at issue reflected 'original mental conception', the photographer's choice of pose, costume, draping, lighting, and composition. The Court was explicit that the camera was the photographer's tool: the human exercise of creative judgment was what made the work copyrightable. By implication, if the photograph had been produced without that creative human input, by a self-tripping mechanical process, it would not have qualified.

That implication was developed more explicitly in subsequent cases. In Bleistein v. Donaldson Lithographing Co.[11], Justice Holmes wrote that copyright protection extended to a work because 'personality always contains something unique. It expresses its singularity even in handwriting.' The personality at issue is necessarily that of a human author.

In Feist Publications, Inc. v. Rural Telephone Service Co.[12], the U.S. Supreme Court rejected the 'sweat of the brow' doctrine and held that copyright requires 'originality,' which it defined as 'independent creation plus a modicum of creativity.' The Court located the source of that originality unambiguously in the human author. The Court's repeated invocation of the author as a creative agent making aesthetic choices is consistent only with a human authorship requirement.

---

[8] Statute of Anne, 1710 (UK).
[9] Constitution of the United States, art. I, § 8, cl. 8.
[10] Burrow-Giles Lithographic Co. v. Sarony, 111 U.S. 53 (1884) <https://supreme.justia.com/cases/federal/us/111/53/> last accessed on 29th April, 2026.
[11] Bleistein v. Donaldson Lithographing Co., 188 U.S. 239 (1903) <https://supreme.justia.com/cases/federal/us/188/239/> last accessed on 29th April, 2026.
[12] Feist Publications, Inc. v. Rural Telephone Service Co., 499 U.S. 340 (1991) <https://supreme.justia.com/cases/federal/us/499/340/> last accessed on 29th April, 2026.

The current U.S. doctrinal position is captured in the Compendium of U.S. Copyright Office Practices, Section 313.2, which provides that 'to qualify as a work of "authorship" a work must be created by a human being' and that the Office will not register works produced by mechanical processes or random selection without any creative input or intervention from a human author. The Office cites Burrow-Giles, the Trade-Mark Cases, and subsequent jurisprudence in support of this position. The Office's January 2025 Report, Copyright and Artificial Intelligence, Part 2: Copyrightability, reaffirms that 'human authorship is a bedrock requirement of copyright.'[13]

3.2 The Continental Tradition: Personality and "Author's Own Intellectual Creation"

The Continental European tradition is, if anything, more explicit in tying copyright to human authorship. The droit d'auteur tradition, particularly in France and Germany, conceives of the author as a person whose creative personality is expressed in the work. Moral rights, paternity, integrity, the right to withdraw, are predicated on the existence of a personality that can be honoured or affronted. They cannot, conceptually, attach to a machine[14].

The doctrinal expression of this tradition in EU law is the 'author's own intellectual creation' standard, articulated by the Court of Justice of the European Union in Infopaq International A/S v. Danske Dagblades Forening[15] and developed in subsequent cases including Painer v. Standard VerlagsGmbH[16], Football Dataco Ltd v. Yahoo! UK Ltd[17], and Levola Hengelo BV v. Smilde Foods BV[18].

In Painer, the Court of Justice held that the photographer's 'free and creative choices', in pose, framing, lighting, and post-production, were what gave a portrait photograph its character as the author's own intellectual creation. The Court's analysis was framed in terms that presuppose a human author exercising creative judgment.

In *Levola Hengelo*, the Court held that the taste of a cheese could not be a copyrighted work because it could not be 'expressed in a manner that makes it identifiable with sufficient precision and objectivity.' That holding has been read by some commentators as carrying a further implication: that copyright presupposes

[13] United States Copyright Office, Copyright and Artificial Intelligence, Part 2: Copyrightability (Jan. 2025) <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf> last accessed on 29th April, 2026.

[14] On moral rights and human authorship, see Mira T Sundara Rajan, Moral Rights: Principles, Practice and New Technology (OUP, 2011).

[15] Infopaq International A/S v. Danske Dagblades Forening, Case C-5/08 [2009] ECR I-6569 <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62008CJ0005> last accessed on 29th April, 2026.

[16] Eva-Maria Painer v. Standard VerlagsGmbH and Others, Case C-145/10 [2011] ECR I-12533 <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62010CJ0145> last accessed on 29th April, 2026.

[17] Football Dataco Ltd and Others v. Yahoo! UK Ltd and Others, Case C-604/10 [2012] <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62010CJ0604> last accessed on 29th April, 2026.

[18] Levola Hengelo BV v. Smilde Foods BV, Case C-310/17 [2018] <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62017CJ0310> last accessed on 29th April, 2026.

an author who can identify and acknowledge a work as theirs. A machine cannot, in the relevant sense, do so.

3.3 The Berne Convention and the Implicit Human Author

The Berne Convention for the Protection of Literary and Artistic Works does not, in terms, define 'author.' Article 2(1), however, defines the protected subject-matter as 'every production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression.' The reference to a 'production' by an 'author,' and the regime of moral rights in Article 6bis, have been understood throughout the Convention's history as presupposing a human creator.

This understanding is carried forward into the TRIPS Agreement, Article 9(1) of which incorporates Articles 1 to 21 of Berne. The implication for member states, including India, is that their copyright systems must be calibrated to a regime of human authorship if they are to remain consistent with their international obligations.

The WIPO Copyright Treaty, in Article 1(4), reaffirms members' obligations under Berne[19]. The position of WIPO, expressed in successive policy documents and in the proceedings of the WIPO Conversation on Intellectual Property and Artificial Intelligence[20], is that the question of AI authorship is one of national policy, but that any departure from the human authorship norm requires affirmative justification.

3.4 The Indian Statutory and Common-Law Position

The Copyright Act, 1957 does not, in its principal definitions, expressly require that an author be a human being. Section 2(d), which defines 'author,' lists categories of works and their respective authors: the person who creates a literary, dramatic, or musical work; the artist for an artistic work; the producer for sound recordings and cinematograph films. Section 2(d)(vi), introduced by the Copyright (Amendment) Act, 1994, defines the author of a 'computer-generated' literary, dramatic, musical, or artistic work as 'the person who causes the work to be created.' None of these provisions, on its face, mandates human authorship.

The human authorship requirement in Indian copyright law is, however, a matter of judicial doctrine and statutory implication. In Tech Plus Media Private Ltd. v. Jyoti Janda[21], the Delhi High Court held that copyright in an underlying literary work cannot vest in a juristic person; only a natural person can be the author of such a work. The juristic person can hold the copyright, by operation of Section 17 (employment, contract of service, or assignment), but cannot be the author in the first instance.

---

[19] WIPO Copyright Treaty (1996), art. 1(4) <https://www.wipo.int/wipolex/en/text/295166> last accessed on 29th April, 2026.

[20] World Intellectual Property Organization, WIPO Conversation on Intellectual Property and Artificial Intelligence <https://www.wipo.int/about-ip/en/artificial_intelligence/conversation.html> last accessed on 29th April, 2026.

[21] Tech Plus Media Private Ltd. v. Jyoti Janda, CS(OS) No. 1098/2010 (Delhi High Court, decided 10 September 2014).

This holding is consistent with the Supreme Court's reasoning in Eastern Book Company v. D.B. Modak[22], in which the Court adopted the 'skill and judgment' standard from CCH Canadian and held that copyright requires the 'author's own intellectual creation.' The Court's analysis is framed throughout in terms of human creative labour.

Section 2(d)(vi) is often considered the principal point of textual divergence from the human authorship norm. By referring to 'the person who causes the work to be created,' the provision contemplates that there is, in some sense, a human author behind a computer-generated work. For present purposes, it is sufficient to note that Section 2(d)(vi) does not, on its terms, displace the human authorship requirement. It locates the author in a particular human role (the person who causes the work to be created), rather than dispensing with human authorship altogether.

[22] Eastern Book Company v. D.B. Modak (2008) 1 SCC 1 <https://indiankanoon.org/doc/1062099/> last accessed on 29th April, 2026.

4. Non-Human Creators in Copyright Jurisprudence

The question whether non-humans can be authors is not new. Before the advent of generative AI, courts in several jurisdictions had occasion to consider claims of authorship by animals, supernatural intelligences, natural processes, and pre-generative computer systems. The cases form a remarkably consistent pattern: across legal traditions and across decades, courts have refused to recognise non-human creators as authors for purposes of copyright protection. The pattern indicates that the human authorship requirement is not a contingent feature of any one jurisdiction's copyright statute, but a structural feature of the copyright concept itself.

4.1 Animals as Creators

The clearest illustration is Naruto v. Slater[23], the so-called ‘monkey selfie’ case. A crested macaque named Naruto, in Indonesia in 2011, took a series of self-portrait photographs using the camera of British wildlife photographer David Slater. People for the Ethical Treatment of Animals (PETA) sued Slater on Naruto's behalf, contending that the macaque was the author of the photographs and was entitled to copyright protection.

The U.S. Court of Appeals for the Ninth Circuit dismissed the suit. The Court held that, while Naruto might have Article III standing to sue, the Copyright Act did not authorise statutory standing for animals. The Court relied in part on a textual analysis of the Act's references to provisions presupposing human relations (the references to inheritance and survivorship within the Act), and concluded that the Act's protections were intended for human authors. The decision was foreshadowed by the U.S. Copyright Office's earlier statement, in the Compendium, that the Office ‘will not register works produced by nature, animals, or plants’.[24]

The reasoning in Naruto extends, in principle, beyond animals. If the doctrinal basis for refusing copyright is that the relevant creator is non-human, the same reasoning applies to other non-human creators, including, for present purposes, AI systems.

[23] Naruto v. Slater, 888 F.3d 418 (9th Cir. 2018) <https://cdn.ca9.uscourts.gov/datastore/opinions/2018/04/23/16-15469.pdf> last accessed on 29th April, 2026.
[24] Compendium of U.S. Copyright Office Practices (3rd ed., 2021), § 313.2 (the Office ‘will not register works produced by nature, animals, or plants’).

4.2 Spiritual or Supernatural Creators

The question whether spiritual or supernatural intelligences can be authors has been litigated more often than might be expected. In Urantia Foundation v. Maaherra, the Foundation sought copyright protection for The Urantia Book, a religious text alleged to have been produced through automatic writing by celestial beings transmitting through human intermediaries. The Ninth Circuit upheld the Foundation's copyright, but on the basis that the human intermediaries had exercised sufficient editorial and arrangement choices to qualify as authors. The Court located the human authorial input in the compilation, selection, and arrangement undertaken by the human transcribers; copyright, in effect, attached to the human contribution and not to the asserted celestial source.[25]

The English position is consistent. In Cummins v. Bond[26], an early twentieth-century decision involving spiritualist mediums claiming to have transcribed messages from the deceased, the Court of Chancery located authorship in the human medium rather than in the alleged spirit-author. The case is sometimes cited, with characteristic dryness, for the proposition that the law of copyright cannot conveniently reach beyond the present life.

4.3 Natural Processes and Found Objects

Natural processes, however aesthetically striking, do not generate copyright. In Kelley v. Chicago Park District[27], the U.S. Court of Appeals for the Seventh Circuit considered whether a wildflower garden, designed by a horticultural artist but maintained and grown by natural biological processes, was protectable as either a sculptural work or a horticultural work. The Court held that it was not. The garden was not ‘fixed’ in the relevant sense, and, more fundamentally, its current state was the product of natural growth and decay rather than human creative agency.

The implication for AI-generated content is quite relevant. If a wildflower garden, with its undeniable aesthetic qualities and substantial human design input at the planning stage, fails the authorship test because its present form is the product of non-human (in that case, biological) processes, then a fortiori an AI-generated image or text, produced through a stochastic generative process to which a human contributes only a prompt, should fail the same test. The biological processes at issue in Kelley are, if anything, more amenable to attribution to a human designer than the trained statistical distributions of a modern large language model.

[25] Urantia Foundation v. Maaherra, 114 F.3d 955 (9th Cir. 1997) <https://law.justia.com/cases/federal/appellate-courts/F3/114/955/532770/> last accessed on 29th April, 2026.
[26] Cummins v. Bond [1927] 1 Ch. 167 (Ch. D.).
[27] Kelley v. Chicago Park District, 635 F.3d 290 (7th Cir. 2011) <https://caselaw.findlaw.com/court/us-7th-circuit/1556063.html> last accessed on 29th April, 2026.

4.4 Mechanical Reproduction and Automated Photography

The earliest engagements of the human authorship requirement, as noted in Section 3 above, arose in the context of mechanical reproduction technologies, first lithography, then photography, then mechanical music reproduction. In each case, the question was whether the mechanical process or the human operator was the author. In each case, the courts answered: the human operator, but only to the extent that the operator made substantive creative choices that the mechanical process executed. The implication for fully automated processes, early teletype-fed sports reports, early algorithmic music compositions, was that they did not generate copyright in the absence of identifiable human creative input. This understanding informed the early English authorities, including Walter v. Lane[28] (the verbatim shorthand reporter held to be the author of his transcript on the basis of skill and judgment) and University of London Press Ltd. v. University Tutorial Press Ltd.[29] (which articulated the foundational 'originated from the author' test).

4.5 Pre-Generative Computer-Generated Output

Two pre-generative-AI decisions illustrate the application of the human authorship requirement to computer-produced works. In Telstra Corporation Ltd v. Phone Directories Co. Pty Ltd[30], the Full Court of the Federal Court of Australia held that Telstra's White Pages and Yellow Pages directories, generated by a computer system from underlying data, did not attract copyright because no human author exercised 'independent intellectual effort' in their production. Special leave to appeal to the High Court of Australia was refused. The decision built on the High Court's earlier reasoning in IceTV Pty Ltd v. Nine Network Australia Pty Ltd[31]. In Acohs Pty Ltd v. Ucorp Pty Ltd[32], the Full Court reached an analogous conclusion in respect of HTML source code generated automatically by a software application.

In the United Kingdom, the position is shaped by Section 9(3) of the Copyright, Designs and Patents Act 1988[33], 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.' Section 178 defines 'computer-generated' as 'generated by computer in

[28] Walter v. Lane [1900] AC 539 (HL).
[29] University of London Press Ltd. v. University Tutorial Press Ltd. [1916] 2 Ch. 601.
[30] Telstra Corporation Ltd v. Phone Directories Co. Pty Ltd [2010] FCAFC 149 <https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/FCAFC/2010/149.html> last accessed on 29th April, 2026; special leave to appeal refused, [2011] HCATrans 248.
[31] IceTV Pty Ltd v. Nine Network Australia Pty Ltd (2009) 239 CLR 458; [2009] HCA 14 <https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/HCA/2009/14.html> last accessed on 29th April, 2026.
[32] Acohs Pty Ltd v. Ucorp Pty Ltd [2012] FCAFC 16 <https://www.austlii.edu.au/cgi-bin/viewdoc/au/cases/cth/FCAFC/2012/16.html> last accessed on 29th April, 2026.
[33] Copyright, Designs and Patents Act 1988 (UK), s. 9(3) and s. 178 <https://www.legislation.gov.uk/ukpga/1988/48/contents> last accessed on 29th April, 2026.

circumstances such that there is no human author of the work.' This is a deeming provision: it acknowledges the absence of a human author and creates a statutory fiction to allocate copyright to the human entity that orchestrated the production. The provision has been applied only sparingly, most notably in Nova Productions Ltd v. Mazooma Games Ltd[34], in which the English Court of Appeal located authorship of dynamically generated video-game frames in the game designer rather than in the player or the system. The provision has been criticised on the ground that it sits awkwardly with the Berne Convention's implicit human authorship requirement and with the EU 'author's own intellectual creation' standard. The U.K. Intellectual Property Office, in its 2022 consultation on AI and IP, declined to propose changes to section 9(3) while acknowledging ongoing concerns about the provision's continued viability[35].

[34] Nova Productions Ltd v. Mazooma Games Ltd [2007] EWCA Civ 219 <https://www.bailii.org/ew/cases/EWCA/Civ/2007/219.html> last accessed on 29th April, 2026.

[35] UK Intellectual Property Office, Artificial Intelligence and Intellectual Property: copyright and patents — Government response to consultation (June 2022) <https://www.gov.uk/government/consultations/artificial-intelligence-and-ip-copyright-and-patents> last accessed on 29th April, 2026.

5. Comparative Jurisdictional Analysis | The Treatment of AI Outputs

This Section maps the current position on AI-generated content in major copyright jurisdictions. The pattern that emerges is one of broad doctrinal convergence, most jurisdictions require a meaningful human creative contribution as the basis for copyright protection, coupled with significant variation in how that requirement is applied to specific fact patterns.

5.1 United States

The United States has produced the most developed body of recent jurisprudence and administrative practice on AI authorship. The principal sources are the following.

First, in Thaler v. Perlmutter, Judge Howell of the U.S. District Court for the District of Columbia held in August 2023 that copyright 'requires human authorship' and that an image autonomously generated by Dr. Stephen Thaler's 'Creativity Machine', without human creative input, was not eligible for copyright protection. The decision was affirmed by the U.S. Court of Appeals for the District of Columbia Circuit in March 2025, in an opinion that carefully reviewed the constitutional and statutory text and concluded that the human authorship requirement was a 'fundamental' element of U.S. copyright doctrine. [36]

Second, in the Zarya of the Dawn matter[37], the U.S. Copyright Office partially cancelled the registration of a graphic novel that combined human-authored text and AI-generated images (produced using Midjourney). The Office held that the human author was entitled to copyright protection in the text and in the selection and arrangement of the images, but not in the images themselves.

Third, in the Théâtre D'opéra Spatial matter[38], the U.S. Copyright Office Review Board affirmed the rejection of a registration for an AI-generated image (produced using Midjourney) that had won the 2022 Colorado State Fair fine art competition. The Board held that the prompting and iterative refinement undertaken by the applicant did not constitute sufficient human authorship.

[36] Thaler v. Perlmutter, 687 F. Supp. 3d 140 (D.D.C. 2023), aff'd 134 F.4th 1 (D.C. Cir. 2025) <https://media.cadc.uscourts.gov/opinions/docs/2025/03/23-5233.pdf> last accessed on 29th April, 2026.
[37] United States Copyright Office, Zarya of the Dawn, Registration No. VAu001480196 (cancellation letter dated 21 February 2023) <https://copyright.gov/docs/zarya-of-the-dawn.pdf> last accessed on 29th April, 2026.
[38] United States Copyright Office Review Board, Re: Second Request for Reconsideration for Refusal to Register 'Théâtre D'opéra Spatial' (Sept. 5, 2023) <https://www.copyright.gov/rulings-filings/review-board/docs/Theatre-Dopera-Spatial.pdf> last accessed on 29th April, 2026.

Fourth, the U.S. Copyright Office's three-part Report on Copyright and Artificial Intelligence[39] sets out the Office's settled position. The Office holds that human authorship is required; that the use of AI as a tool does not defeat copyright if the human contribution is sufficient; and that the prompting of generative AI systems, without more, will not generally meet the human authorship threshold.

5.2 United Kingdom

As noted in Section 4.5, the United Kingdom has a statutory deeming provision (CDPA 1988, s. 9(3)) that allocates authorship in computer-generated works to the person making the necessary arrangements. The provision is, in form, distinctive within the major copyright jurisdictions. In substance, however, the U.K. position is converging with that of other jurisdictions: the U.K. Intellectual Property Office has noted that section 9(3) will be applied with regard to whether the alleged 'arrangements' reflect the kind of human creative input that copyright is meant to protect, and recent academic commentary suggests that purely AI-generated content may not qualify as a 'work' within the meaning of the Act in any event.

The U.K. position remains the most accommodative of AI-generated content among the major Anglophone jurisdictions, but the practical effect of this accommodation is currently uncertain. Following the U.K.'s departure from the European Union, the application of the EU 'author's own intellectual creation' standard in U.K. law is itself a contested question.

5.3 European Union

In the European Union, the 'author's own intellectual creation' standard governs the question of copyright subsistence. The standard requires free and creative choices by a human author. AI-generated content that does not reflect such choices does not satisfy the standard.

The EU AI Act (Regulation (EU) 2024/1689)[40] does not, in terms, address the copyright status of AI-generated content, but its transparency provisions (Article 50) require disclosure of AI-generated content in certain contexts. The AI Act and the existing copyright acquis are designed to operate in parallel.

EU member states have begun to address the question through their national courts. The German Federal Court of Justice (Bundesgerichtshof), in its decision of 11 June 2024 on the DABUS patent application[41], addressed the question of AI inventorship in patent law, holding that an inventor must be a natural person

[39] United States Copyright Office, Copyright and Artificial Intelligence, Parts 1, 2 and 3 (2024–2025) <https://www.copyright.gov/ai/> last accessed on 29th April, 2026.
[40] Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (the EU AI Act) <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689> last accessed on 29th April, 2026.
[41] Bundesgerichtshof, Decision of 11 June 2024, X ZR 39/23 (DABUS) (German Federal Court of Justice).

but allowing the human applicant to be named as inventor. The decision is in the patent context, and its direct copyright relevance is therefore limited; but the underlying conceptual move, naming the human who prompted the AI as the relevant rightsholder, is one that may feature in copyright debates as well.

5.4 China

China has produced some of the most consequential recent jurisprudence on AI-generated content. The position is more permissive than that of the United States or the European Union, but it is not unconditional.

In Beijing Feilin Law Firm v. Beijing Baidu Wangxun Technology Co., Ltd., the Beijing Internet Court considered an analytical report partially generated by a software tool. The Court held that the report was not an original work to the extent the underlying analysis was the output of an automated process, but the human-authored elements (the introduction, conclusions, and arrangement) were protected.

In Shenzhen Tencent Computer System Co., Ltd. v. Shanghai Yingxun Technology Co., Ltd.[42], the Nanshan District People's Court of Shenzhen held that a financial article generated by Tencent's Dreamwriter AI system was a 'written work' protected by copyright, on the basis that the human team had made sufficient creative choices in selecting data sources, designing the article structure, and supervising the output.

In Li v. Liu[43], in a case concerning an AI-generated image produced using Stable Diffusion, the Beijing Internet Court held in November 2023 that the image attracted copyright protection. The Court reasoned that the plaintiff's selection of prompts, parameters, and iterations reflected creative choices sufficient to ground authorship. The decision is the most permissive judicial holding on AI image authorship in any major jurisdiction to date and has been the subject of significant international academic comment.

The Chinese position should be read in the context of the country's broader policy emphasis on AI development as a strategic priority. The cases reflect a pragmatic accommodation rather than a wholesale doctrinal shift; the human authorship requirement is preserved in form, but the threshold of human creative contribution is set lower than in U.S. or E.U. jurisprudence.

5.5 Australia

Australia has not adopted a U.K.-style deeming provision. The Australian position is that copyright requires identifiable human authors who exercise independent intellectual effort. There has been no Australian appellate decision on AI-generated content specifically, but the existing authorities suggest that purely AI-

[42] Shenzhen Tencent Computer System Co., Ltd. v. Shanghai Yingxun Technology Co., Ltd., (2019) Yue 0305 Min Chu No. 14010 (Nanshan District People's Court, Shenzhen).
[43] Li v. Liu, (2023) Jing 0491 Min Chu No. 11279 (Beijing Internet Court, judgment of 27 November 2023).

generated content would not attract copyright. The Australian position is, in substance, among the most restrictive of the major jurisdictions.

5.6 Canada

Canada's position is anchored in the Supreme Court's decision in CCH Canadian Ltd v. Law Society of Upper Canada[44], which held that copyright requires the exercise of 'skill and judgment' by an author. The 'skill and judgment' standard is, like its U.K. and Indian counterparts, framed in terms that presuppose a human author.

In a development of immediate relevance, the Canadian Internet Policy and Public Interest Clinic (CIPPIC) initiated proceedings in the Federal Court of Canada in July 2024 challenging the registration of an AI-generated work produced by Mr. Ankit Sahni's RAGHAV system (a parallel registration to that obtained from the Indian Copyright Office)[45]. The litigation is pending; its outcome will be the first authoritative judicial pronouncement in Canada on AI authorship.

5.7 Japan and South Korea

Japan's Copyright Act defines a 'work' as 'a production in which thoughts or sentiments are creatively expressed and which falls within the literary, scientific, artistic or musical domain'[46]. The reference to 'thoughts or sentiments' has been interpreted as implying a human author, since machines do not have thoughts or sentiments in the relevant sense. The Japanese Agency for Cultural Affairs, in its policy statements on AI and copyright, has indicated that pure AI-generated outputs are not copyrightable but that AI-assisted works with sufficient human creative input may be[47].

South Korea's Copyright Act is similar. The Korea Copyright Commission has taken the position that copyright protection requires a human author, and that AI-generated content without sufficient human creative input is not eligible for protection[48].

---

[44] CCH Canadian Ltd. v. Law Society of Upper Canada, 2004 SCC 13, [2004] 1 SCR 339 <https://decisions.scc-csc.ca/scc-csc/scc-csc/en/item/2125/index.do> last accessed on 29th April, 2026.

[45] Canadian Internet Policy and Public Interest Clinic (CIPPIC), Application for judicial review filed in the Federal Court of Canada (July 2024) challenging the registration of a work generated by the RAGHAV system; reported in CIPPIC press release (2024) <https://cippic.ca/> last accessed on 29th April, 2026.

[46] Copyright Act of Japan, Act No. 48 of 6 May 1970 (as amended), art. 2(1)(i) <https://www.japaneselawtranslation.go.jp/en/laws/view/4207> last accessed on 29th April, 2026.

[47] Agency for Cultural Affairs (Japan), Approach to Copyright in the Age of AI (May 2024).

[48] Copyright Act of the Republic of Korea, art. 2 (definitions) and art. 4 (works); guidance issued by the Korea Copyright Commission, AI-Generated Content Copyright Registration Guidelines (Dec. 2023).

5.8 Czech Republic, Germany, and Other Jurisdictions

In the Czech Republic, the Municipal Court of Prague held in October 2023 (in the so-called 'DALL-E case') that an image generated by the DALL-E system was not a 'work' within the meaning of Czech copyright law because no human author had exercised the requisite creative choices. The decision is notable as one of the first European judicial holdings on the copyright status of generative AI content.

In Germany, the Federal Court of Justice has addressed AI inventorship in the patent context. There is no comparable copyright decision yet, but the Bundesgerichtshof's reasoning is consistent with the broader European position that originality requires a human author.

Other civil-law jurisdictions, including France, Italy, Spain, the Netherlands, and Belgium, have not yet produced authoritative judicial determinations on AI authorship. Their copyright systems are, however, anchored in the EU 'author's own intellectual creation' standard, and the analytical pressure on AI-generated content is, accordingly, similar in those jurisdictions.

Within the broader Commonwealth, several jurisdictions have adopted statutory deeming provisions modelled on section 9(3) of the U.K. CDPA 1988. These include New Zealand (Copyright Act 1994, s. 5(2)(a)), Ireland (Copyright and Related Rights Act 2000, s. 21(f)), Hong Kong (Copyright Ordinance Cap. 528, s. 11(3)), and South Africa (Copyright Act 98 of 1978, definition of 'computer-generated work'). The provisions in each case allocate authorship of computer-generated works to the person undertaking the arrangements necessary for the work's creation. As with the U.K. provision, however, the practical application of these provisions to outputs of generative AI systems is contested, and no authoritative judicial pronouncement has yet been made in any of these jurisdictions.

5.9 India

India's position is, as set out in Primer 1, defined by the interaction of (a) the Copyright Act, 1957, particularly Sections 2(d), 2(d)(vi), 13, and 17; (b) the judicial doctrine, anchored in Eastern Book Company v. D.B. Modak and Tech Plus Media v. Jyoti Janda, that copyright requires the 'author's own intellectual creation' and that the author of an underlying literary work must be a natural person; and (c) the Indian Copyright Office's handling of the RAGHAV/Suryast registration[49].

The RAGHAV episode is illustrative. The Office initially rejected an application listing the AI system RAGHAV as sole author; it subsequently accepted an amended application listing Mr. Ankit Sahni and RAGHAV as co-authors (Registration No. A-135120/2020); and on 25 November 2022 it issued a withdrawal notice questioning RAGHAV's status as an 'artist' within the meaning of Section 2(d). The withdrawal has been

[49] Indian Copyright Office, Registration No. A-135120/2020 (RAGHAV / 'Suryast'); Withdrawal Notice dated 25 November 2022; response dated 8 December 2022; status as of the date of this primer: contested.

contested by the registrant, and the matter remains unresolved at the time of writing. The episode illustrates the unsettled state of Indian administrative practice; it does not yet provide binding doctrinal guidance.

The interaction of Section 2(d)(vi)'s 'person who causes the work to be created' language with the broader human authorship requirement in the Act remains untested by any Indian court. The Indian position is, in form, more permissive than the U.S. or E.U. positions but more restrictive than the U.K. position. In substance, it remains the most uncertain among the major jurisdictions, and the existing administrative practice is itself in flux.

5.10 Comparative Overview

The following table summarises the position in each of the principal jurisdictions covered in this Section. The summary is high-level and is intended to facilitate comparison; it does not displace the more detailed treatment in the preceding sub-sections or the doctrinal qualifications discussed there.

Jurisdiction	Statutory / Doctrinal Anchor	Pure AI Output	AI-Assisted Output
United States	Constitutional grant; Compendium § 313.2; Thaler (D.D.C. 2023, aff'd D.C. Cir. 2025)	Not protectable	Protectable in respect of human contribution; mere prompting insufficient
United Kingdom	CDPA 1988, ss. 9(3), 178 (deeming provision)	In form, protectable; in substance, doubtful post-Infopaq	Protectable subject to originality
European Union	'Author's own intellectual creation' (Infopaq, Painer); EU AI Act 2024 (transparency only)	Not protectable	Protectable in respect of human creative choices
China	Copyright Law (PRC); Feilin (2018), Tencent (2019), Li v. Liu (2023)	Protectable in some cases (low threshold of human input)	Protectable
Australia	Copyright Act 1968; IceTV (2009); Telstra (2010); Acohs (2012)	Not protectable	Protectable subject to identifiable human authors and effort
Canada	Copyright Act; CCH Canadian (2004) ('skill and judgment')	Likely not protectable; CIPPIC RAGHAV challenge pending (2024)	Protectable subject to skill and judgment

Japan	Copyright Act, art. 2(1)(i) ('thoughts or sentiments')	Not protectable	Protectable subject to human creative input
South Korea	Copyright Act, art. 2; KCC guidelines	Not protectable	Protectable subject to human creative input
Czech Republic	Czech Copyright Act; Municipal Court of Prague 'DALL-E' case (2023)	Not protectable	Protectable subject to human creative input
India	Copyright Act 1957, ss. 2(d), 2(d)(vi); Eastern Book Co. (2008); Tech Plus Media (2014); RAGHAV (administrative, contested)	Unsettled; statutory text potentially permissive but doctrinally restrictive	Protectable subject to originality

The comparative pattern is one of broad convergence on the human authorship requirement, with significant variation in how the threshold of human creative input is set. The United Kingdom and China stand out as the principal outliers from the doctrinal majority, although for different reasons: the U.K. by virtue of its statutory deeming provision, China by virtue of judicial willingness to set the human-input threshold low.

6. Why Human Intellectual Creativity is the Anchor

The doctrinal preference for human authorship is not without theoretical underpinning. The three principal justifications for copyright, personality, labour, and utility, each presuppose a human creator, and none extends to autonomously generated machine output.

The personality account, associated with Hegel and Kant, conceives of intellectual property as an extension of the personhood of the creator. Machines do not have personhood; the output of a generative AI system is a probabilistic computation, not the projection of a self. The Continental moral-rights regime, and its Indian counterpart in section 57 of the Copyright Act, protects the rights of paternity and integrity, and these rest on the existence of an authorial personality whose attribution and reputation can be honoured or affronted. They cannot, conceptually, attach to a machine.[50]

The labour account, associated with Locke, justifies property as the labourer's entitlement to the fruits of effort. The human prompter may be said to engage in some labour, but the generation of the expression itself is performed by the machine. Even on the labour account, the human's residual contribution, selecting subject-matter, providing prompts, choosing among outputs, may suffice to protect the prompt or the selection, but not the output itself. The Lockean argument presupposes a human labourer; machines are not such subjects.

The utilitarian or disincentive-avoidance account, developed at length in Section 2 above, justifies copyright as the means by which human creative labour is sustained as a vocation in market societies in which intangible expression is non-rivalrous and infinitely replicable. AI systems do not face that problem; they do not require copyright to continue operating. The mischief that does exist, that AI-assisted output will outcompete unaided human creators, argues against, not for, the conferral of copyright on AI outputs.

The common thread across the three theories is the centrality of the human creator: as a personhood whose self is projected, as a labourer whose effort is rewarded, or as a vocation whose continuation is to be enabled. None of the theories, in its classical formulation, contemplates a non-human creator. This is not an oversight; it reflects the structural problem that copyright was instituted to solve. Any extension of copyright to autonomous AI output must therefore be undertaken on a different theoretical basis than copyright's classical justifications support, a basis the burden of articulating which lies on those who advocate the extension.

[50] GWF Hegel, Elements of the Philosophy of Right (1820); Kant, Of the Wrongfulness of Unauthorised Publication of Books (1785); see generally Margaret Jane Radin, 'Property and Personhood' (1982) 34 Stan. L. Rev. 957.

7. Implications for the AI Outputs Debate

7.1 Recasting the Question

The implication of Sections 2 to 6 is that the most-discussed question, who should own AI outputs, is the wrong starting point. The right starting point is whether AI outputs should be subject to ownership at all. If the answer is no, or qualified no, the further question of attribution is moot in the relevant range of cases. If the answer is yes, the basis for the answer must be made explicit, and it cannot be one of copyright's classical justifications without significant doctrinal extension. The contemporary debate has tended to assume the answer to the threshold question and to focus on the secondary one. The result is a debate that proceeds at one remove from the analytical foundation that ought to anchor it.

7.2 The "Tool" vs. "Co-Author" Question

Primer 1 introduced the question whether AI systems are best characterised as tools used by human authors or as co-authors in their own right. None of copyright's principal justifications supports the recognition of an AI system as a co-author. AI systems are not personhoods; they are not labourers entitled to the fruits of their effort; their continued operation does not depend on legal protection. The ‘tool’ framing is the only one consistent with the doctrinal foundations of copyright.

The ‘tool’ framing has, however, its own difficulties. A tool is something a human uses; the human's creative input must be sufficient, in the relevant sense, to ground authorship. For most uses of generative AI, the human input, a prompt, perhaps iteratively refined, is qualitatively different from the input of a photographer using a camera or a painter using a brush. The photographer specifies the entire creative content of the photograph; the painter executes each stroke. The prompter specifies only the conceptual constraints, leaving the model to generate the expression. Whether this is sufficient to ground authorship is a question to which subsequent primers will return.

7.3 Moral Rights and Verifiability

Two practical problems flow from the human authorship requirement. The first is moral rights: the rights of paternity and integrity rest on the existence of an authorial personality whose attribution and reputation can be honoured or affronted, and machines do not have such personalities. If AI-generated content were to attract copyright on the same terms as human-created content, but moral rights could not attach to the AI

system, the result would be a doctrinal asymmetry, economic rights but no moral rights, that the Continental tradition, including India under section 57 of the Copyright Act, would find difficult to accept.

The second is verifiability. Most copyright systems do not require proof of authorship at the registration stage; the applicant's declaration is presumed accurate. With AI-generated content, that presumption becomes harder to sustain: there is no reliable technical means of distinguishing, after the fact, a human-created work from an AI-generated work or an AI-assisted work with substantial human input[51]. If the human authorship requirement is in practice unverifiable, registrants will have an incentive to misdeclare. The integrity of the copyright system depends on the alignment of declared authorship with actual authorship; the AI context strains that alignment. The U.S. Copyright Office's response has been to require disclosure of AI-generated material in registration applications; whether such disclosure can reliably be enforced, and what the consequences of non-disclosure should be, are questions to which subsequent primers will return.

[51] On the verifiability problem, see United States Copyright Office, Copyright and Artificial Intelligence, Part 2: Copyrightability (Jan. 2025), at pp. 30–41 <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf> last accessed on 29th April, 2026.

8. Conclusion

The human authorship requirement is not an arbitrary feature of copyright law. It is the doctrinal expression of a particular structural problem, the difficulty of sustaining human creative labour in market societies dealing in non-rivalrous intangibles, and of the theoretical justifications that copyright systems have adopted to address that problem. The requirement is implicit in the Berne Convention, explicit in the case law of the United States, the European Union, Australia, and Canada, and embedded in the Indian doctrinal tradition through Tech Plus Media and Eastern Book Company.

Generative artificial intelligence has compelled a re-examination of the requirement, but the re-examination has, in most jurisdictions, confirmed rather than displaced it. With the partial exception of the United Kingdom (whose section 9(3) is in any event of contested practical effect) and the partial exception of China (whose courts have been more permissive in specific cases), the human authorship requirement remains the doctrinal anchor of copyright systems worldwide.

This primer has argued that the question of AI authorship cannot responsibly be addressed without first confronting the prior question: whether copyright protection should extend to AI-generated content at all. The classical justifications do not extend, in their own terms, to autonomous machine production; the cost of exclusion in human speech and communication that copyright imposes is not justified, on those theories, by the conferral of property rights on autonomously generated AI content. The choice between conferring property rights on AI-generated content and treating such content as presumptively in the public domain is a policy choice that legislatures and, in their absence, courts will have to make. The purpose of the present primer has been to clarify the doctrinal landscape against which the choice will be made, and to insist that the choice is a real one, not a foregone conclusion. The next primer in this series will turn to the more granular question of how, if some form of protection is to be conferred, the copyright system might calibrate that protection, and to the structural alternatives, including sui generis investment-protection rights, contractual mechanisms, and a presumptive public domain default, that ought to be considered alongside any extension of conventional copyright to AI-generated content.

Author

Akshat Agrawal

Founder, AASA Chambers | Visiting Fellow, The Dialogue

Akshat Agrawal is the Founder of AASA Chambers and a PhD Candidate at the University of Cambridge (Faculty of Law), where he researches copyright reform under Professor Lionel Bently. He holds an LLM from UC Berkeley; his award-winning thesis on copyright (Google Global Award, 2023) was published in the Berkeley Technology Law Journal. He is also a Wikimedia–IIPSJ Fellow (2023–24).

His work spans intellectual property, technology law, and public policy. He has worked with the Additional Solicitor General of India, senior advocates at the Delhi High Court, and as a law researcher to a Delhi High Court judge. He was previously a Senior Associate at Saikrishnaa & Associates.

thedialogue.org.in

LinkedIn | The Dialogue

X | The Dialogue

Whatsapp | The Dialogue

Instagram | The Dialogue