Chapter 1 / Capítulo 1

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 5: Intelligent Visual Arts and Audiovisual Engineering: Generative AI, Computational Creativity and Superintelligence
ISBN (this volume): 978-9915-9966-5-3
DOI: 10.62486/978-9915-9966-3-9.Ch01
Pages: 1-11

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Authorship, copyright and creativity in machine-generated visual art

Autoría, derechos de autor y creatividad en el arte visual generado por máquinas

Carlos Alberto Gómez Cano[1]
Verenice Sánchez Castillo[2]

1 Corporación Universitaria de Asturias. Bogotá, Colombia.
2 Universidad de la Amazonia. Florencia, Colombia

ABSTRACT

The question of who authors a work produced by a machine, and whether copyright protects it, has moved from a theoretical curiosity to a central problem of intellectual property in less than a decade. This narrative review traces how legal and philosophical treatments of authorship, copyright and creativity in AI-generated art evolved from 2018 to 2026, a window that can be divided into four periods by technical and institutional inflection points. The GAN era and the first auction sales opened the question, the diffusion models of 2021 and 2022 made it a mass concern, the institutional decisions of 2023 established a human-authorship requirement, and the period from 2024 to 2026 consolidated that requirement alongside new regulation and training-data litigation. The review synthesizes the legal and philosophical literature across these periods and finds that the law has converged on a stable human-authorship standard while the philosophy has moved toward distributed and collaborative accounts of creativity, leaving a durable gap between the two that the field has not yet closed. The trajectory is interpreted as a shift in what the debate is about, from whether machines can be authors to how much human involvement suffices and who bears responsibility for what a model was trained on.

Keywords: AI-generated art; copyright; authorship; creativity; generative artificial intelligence; intellectual property

RESUMEN

La cuestión de quién es autor de una obra producida por una máquina, y si el derecho de autor la protege, ha pasado de ser una curiosidad teórica a un problema central de la propiedad intelectual en menos de una década. Esta revisión narrativa traza cómo evolucionaron los tratamientos jurídicos y filosóficos de la autoría, el derecho de autor y la creatividad en el arte generado por IA entre 2018 y 2026, una ventana que puede dividirse en cuatro períodos por puntos de inflexión técnicos e institucionales. La era de las GAN y las primeras subastas abrieron la cuestión, los modelos de difusión de 2021 y 2022 la convirtieron en una preocupación masiva, las decisiones institucionales de 2023 establecieron un requisito de autoría humana, y el período de 2024 a 2026 consolidó ese requisito junto a nueva regulación y litigios sobre datos de entrenamiento. La revisión sintetiza la bibliografía jurídica y filosófica de esos períodos y concluye que el derecho ha convergido en un estándar estable de autoría humana mientras que la filosofía se ha desplazado hacia relatos distribuidos y colaborativos de la creatividad, dejando una brecha duradera entre ambos que el campo aún no ha cerrado. La trayectoria se interpreta como un cambio en aquello de lo que trata el debate, de si las máquinas pueden ser autoras a cuánta intervención humana basta y quién responde por aquello con lo que se entrenó un modelo.

Palabras clave: arte generado por IA; derechos de autor; autoría; creatividad; inteligencia artificial generativa; propiedad intelectual

INTRODUCTION

For most of copyright's history, the attribution of authorship to a machine was a hypothetical that the law could comfortably set aside. The premise that a work must originate from a human being was embedded so deeply in the statutory and doctrinal structure of copyright that it rarely needed to be defended, and the occasional computer-generated output was treated as the product of the programmer or user who controlled the tool.[1,2] That premise became a live problem when generative systems began to produce images that were, to a viewer, indistinguishable from the work of a human artist, and the problem arrived with unusual speed. The period from 2018 to 2026 is the span in which the question moved from the margins of legal scholarship to the center of intellectual-property policy, and it is the span this review reconstructs.

The trigger was not a single event but a succession of them, each of which redefined what the debate was about and each of which arrived before the previous one had been fully digested. In 2018 an algorithmic portrait sold at auction for a price that forced the art world and the legal world to confront the possibility that a machine, or the people who built and prompted it, might be credited as the author of a work with commercial value.[3,4] Between 2021 and 2022 the release of text-to-image diffusion models made the production of images a matter of typing a sentence, which transformed the question from one about a small class of experimental works to one about a mass medium.[5,6] In 2023 courts and registration offices began to rule, and their rulings established a human-authorship requirement that the literature then spent the following years consolidating, extending and contesting.[7,8] By 2026 the debate had acquired a second center of gravity, the copyright status of the training data on which generative models are built, and a regulatory frame that did not exist when the window opened.(9,10)

The problem is at once legal and philosophical, and the two registers do not move in lockstep, a fact that is itself part of what the review has to explain rather than merely describe. The legal question is comparatively settled in its broad outline, namely that a work without sufficient human authorship is not protected, but it is unsettled at the boundary, where a human selects, arranges, edits or otherwise shapes machine output.[11,12] The philosophical question is more open, because it concerns what creativity is and whether the conditions that make human creativity valuable, intention, autonomy, a life to which the work is answerable, can be satisfied by a system, or by a system and its human users together.[13,14] A review that treated only the law would miss the philosophical movement beneath it, and a review that treated only the philosophy would miss the institutional facts that have forced the question, which is why both are synthesized here.

Surveys of the field have multiplied, but they tend to freeze the debate at a single moment, a particular case, a particular jurisdiction or a particular theory of authorship, rather than to follow its movement across time, and a reader of several such surveys is left with a series of stills rather than a film of the debate.[10,15] The chronological structure of this review is a deliberate choice, because the debate is best understood as a trajectory in which each period's questions set the terms of the next, and in which the meaning of a stable doctrine, such as the human-authorship requirement, changed as the technology around it changed.[2,16] What began as a question about a portrait sold at auction became a question about the distribution of authorship across a network of trainers, developers, platform operators and users, and that transformation, from a single object to a distributed system, is the story this review tells.

The objective is to trace the evolution of legal and philosophical treatments of authorship, copyright and creativity in AI-generated art from 2018 to 2026. Four specific objectives organize the account and correspond to the sections that follow: to periodize the debate into identifiable stages marked by technical and institutional inflection points; to synthesize how copyright law has addressed machine authorship, human involvement and training data across the period; to synthesize how philosophical accounts of creativity have been revised in response to generative systems; and to identify what changed, what persisted, and which open questions remain at the end of the window. Each period is developed in turn, and the trajectory is interpreted in the discussion.

METHODS

The review followed the SANRA guidance for narrative reviews, which requires a transparent account of sources, search terms, selection and organization without the formal screening counts of a systematic review. The object of the synthesis is the scholarly literature on authorship, copyright and creativity in AI-generated art, together with the adjacent literature on the philosophy of

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

machine creativity and on the regulation of generative systems that supplies context for the legal claims. The review is chronological by design, and its method is to divide the window into periods whose boundaries are marked by identifiable technical and institutional events, and to hold a fixed set of questions against each period in turn.

Three bibliographic sources were searched: OpenAlex, Crossref and Semantic Scholar. The queries combined terms for artificial intelligence and generative art with terms for copyright, authorship, creativity and originality, and included AI-generated art copyright authorship, artificial intelligence creativity copyright law, generative AI copyright training data litigation, machine authorship copyright office decision, and AI art philosophical creativity autonomy. The search window was 2018 through 2026 for the primary studies, which is the period in which the technical systems and the legal decisions the review covers were produced; earlier works on originality and on the philosophy of creativity were admitted only as background where a primary study explicitly builds on them. Records were deduplicated across sources by DOI and by normalized title, and the query log is retained in the working directory.

Selection followed three criteria applied to title and abstract. A study was retained if it addressed authorship, copyright or creativity in relation to AI-generated art from a legal, philosophical or empirical standpoint; if it was published within the search window; and if it was written in English or Spanish. Works on adjacent topics, such as AI regulation, machine creativity in domains other than art, or the governance of training data, were admitted where they contributed a finding that bears directly on the authorship question. The retained material was divided into a core of primary studies that populate the four periods and a supporting set of surveys and adjacent studies that supply context in the introduction and discussion.

The periodization is the central methodological decision, and it was made against explicit criteria. A boundary was placed where a technical release or an institutional decision changed the question the field was asking, rather than at arbitrary year intervals, because a periodization that simply divided the decade into equal segments would cut across the very developments the review aims to explain, separating a model release from the litigation it triggered or a ruling from the scholarship that followed it. On those criteria the window divides into four periods. The years 2018 to 2020 are the GAN era, bounded at the front by the first high-profile sales of algorithmic art and defined by the first wave of legal and philosophical scholarship on machine authorship. The years 2021 to 2022 are the diffusion period, defined by the release of text-to-image models that made generative art a mass medium. The year 2023 is an institutional turning point, defined by the first authoritative rulings on the human-authorship requirement. The years 2024 to 2026 are a period of consolidation and regulation, defined by the codification of the human-involvement standard and by the rise of training-data litigation and legislation.

A synthesis table (Table 1) sets out the inflection points that define the periods, and a second table (Table 2) summarizes the legal and philosophical positions that characterize each period. A timeline figure (Figure 1) presents the four periods and their defining events.

Table 1. Inflection points that define the four periods.

Period	Defining inflection point	Consequence for the authorship debate
2018-2020	First high-profile auction sale of algorithmic art	Authorship question becomes a commercial and policy concern
2021-2022	Release of text-to-image diffusion models	Generative art becomes a mass medium; the human-involvement boundary widens
2023	Copyright Office and court rulings on machine authorship	Human-authorship requirement established
2024-2026	Codification and training-data litigation	Human-involvement threshold refined; training data becomes central

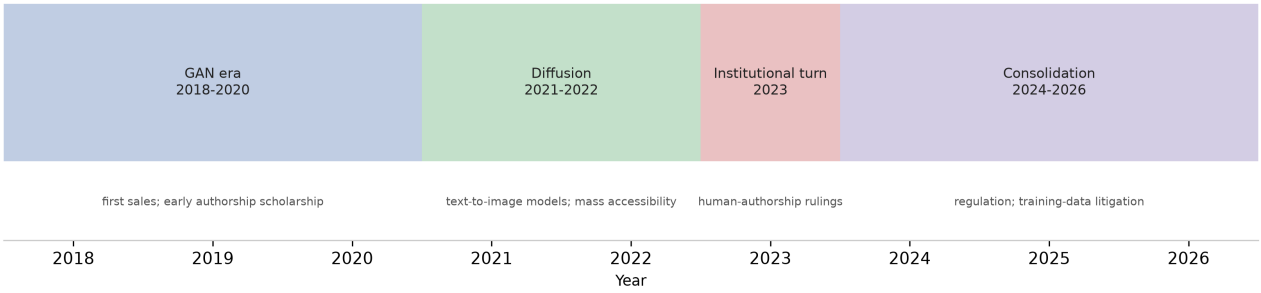
Table 2. Legal and philosophical positions across the four periods.

Period	Legal position	Philosophical position
2018-2020	Machines cannot be authors or legal persons	Authorship and creativity categories under pressure
2021-2022	Tool-author boundary; jurisdictional divergence	Machine agency and collaborative creativity
2023	Human-authorship requirement affirmed	Distributed co-authorship
2024-2026	Human-involvement threshold; training-data regulation	Distributed and posthuman authorship

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

Because the literature is heterogeneous in jurisdiction, method and theoretical orientation, the synthesis is narrative rather than quantitative, and no claim about the distribution of positions is made except where a cited study reports one.

Figure 1. The four periods of the authorship debate and their defining inflection points.



DEVELOPMENT

2018 TO 2020: THE GAN ERA AND THE FIRST AUTHORSHIP CONTROVERSIES

The window opens with a market event that gave the theoretical question a price. In 2018 an algorithmic portrait produced with a generative adversarial network sold at auction, and the sale, together with the surrounding attention, made concrete what had been abstract, that a machine-generated image could carry commercial value and therefore that the question of who owned it, and who authored it, could no longer be deferred.[3,4] The philosophical literature of this period took the sale as evidence that the category of authorship itself was under pressure, and it began to ask whether the concepts of authorship, originality and creativity inherited from romantic and post-romantic aesthetics could survive a system that produces works without a human hand in the traditional sense.(17)

The empirical work of the period complicated the assumption that the public would simply accept machine authorship. Studies of attribution showed that audiences and buyers perceived AI-generated art differently from human art, discounting it in some contexts and crediting less to the machine than the human collaborators, a finding that matters because the social recognition of authorship, not only the legal one, turns out to be contested.(4) The result is that the first period established the problem's shape rather than its solution: the technical capacity existed, the market had noticed it, and neither the law nor the philosophy had yet produced an account that everyone could accept.

Legal scholarship of the first period answered the framing question, whether a machine can be an author or a legal person, largely in the negative. The analyses that traced the doctrine across the common-law jurisdictions concluded that the human-authorship premise, grounded in the constitutional and statutory language of the major systems, left no room for a machine as an author in its own right, and that the more serious question was how to allocate the human contribution among the people who built, trained and directed the system.[1,18] This early work also named, without yet resolving, the second question that would come to dominate the debate, the copyright status of the data on which a model is trained, which the period framed as a question of infringement and exception that would grow in importance as the models grew in scale.

2021 TO 2022: DIFFUSION AND MASS ACCESSIBILITY

The second period is defined by the release of text-to-image diffusion models, which changed the scale and the social meaning of the problem. Where the GAN era had produced images through a technical pipeline that remained the province of researchers, the diffusion models of 2021 and 2022 allowed a user to produce a photorealistic or stylized image from a short natural-language prompt, and the literature of the period describes this as the moment at which generative art became a mass medium rather than a research artifact.[5,6] The consequence for the authorship question was immediate: the range of plausible human contributions, a prompt, a selection among outputs, an edit, a curatorial arrangement, widened, and with it the range of

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

claims about who, among the many people in the chain, was the author.

The legal literature of the period turned to the boundary between the tool and the author. Analyses distinguished human-aided machine output, which the emerging consensus treated as unprotected, from machine-aided human output, which could be protected to the extent of the human contribution, and they began to test that boundary against the specific acts of prompting, selection and post-processing that the new tools made available.[2,19] Jurisdictional work examined how different systems, the United States with its constitutional human-authorship requirement, the United Kingdom with its specific provision for computer-generated works, and the European systems with their originality standards, would classify the same diffusion-produced image, and it found divergence rather than convergence.(20,21)

The philosophical literature of the period matured in a different direction, toward the conditions under which a machine could be an agent or an author at all. Accounts of machine agency tested the requirements of intention, autonomy and answerability against the capabilities of generative systems and concluded that the systems failed the strongest version of each requirement, while leaving open the weaker versions that a collaborative account might satisfy. (22) Alongside this, a line of work began to reframe creativity not as a property that either a human or a machine possesses but as something that emerges from their interaction, a position that would become more prominent in the following periods.(14,16)

2023: THE INSTITUTIONAL TURNING POINT

The year 2023 is the hinge of the review, because it is the year in which the question moved from the journals to the courts and the registration offices, and in which the law gave its first authoritative answers in a form that practitioners could rely on, a transition that the preceding periods had prepared but not yet delivered. The United States Copyright Office declined to register a purely machine-generated image, and the courts affirmed that refusal on the ground that the human-authorship requirement is rooted in the constitutional purpose of copyright, a ruling that the legal literature of the period read as establishing, in the American system at least, that there is no protection without a human author.[7,8] The same year produced the intermediate case, a comic whose images were machine-generated but whose text and selection were human, and the office's decision to protect the human elements while refusing the images gave the field its first worked example of the human-involvement threshold.(8,15)

The institutional decisions did not end the debate but relocated it. With pure machine authorship ruled out, the live questions became how much human involvement suffices, and what kind, and the literature of 2023 immediately began to map those questions onto the specific acts that the new tools made possible, prompting, selection, arrangement, modification and iterative refinement.[8,12] At the same time, the second center of gravity, training data, moved to the foreground, as litigation against model developers raised the question of whether the ingestion of copyrighted works for training is infringement or a permitted use, and as the first analyses of the European and American positions on that question appeared.(15,23)

The philosophical register of 2023 changed tone, from the question of whether a machine could be an author to the question of how humans and machines co-author. A widely cited manifesto for collaboration argued that the productive response to generative systems is neither to reject them nor to cede authorship to them but to reconceive creativity as a distributed practice in which human intention and machine capacity are jointly responsible, and the creativity literature of the period supplied empirical and conceptual support for that move.[16,24] The period thus left the law and the philosophy pointing in different directions, the law toward a strict human threshold and the philosophy toward a distributed account, a divergence that the following years would deepen rather than resolve.

2024 TO 2026: CONSOLIDATION AND REGULATION

The final period is one of consolidation, in which the human-authorship requirement was refined rather than reopened, and in which regulation and training-data litigation supplied the new frame. The registration office codified its position in a report that distinguished, with increasing precision, the human contribution that renders an AI-assisted work protectable from the machine output that does not, and the scholarly literature of the period worked through the report's implications for prompting, selection and modification.[11,12] The dominant question of the

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

period was no longer whether a human is required but how the requirement is operationalized, a shift that the review records as the clearest sign that the debate had matured.

The regulatory turn is the period's second defining feature. European and international instruments introduced transparency and disclosure obligations for generative systems, and the scholarly analyses of text-and-data-mining exceptions and of the legal status of training data became a central genre, asking whether the exceptions that enable research and data mining should be extended to commercial training and how a lawful training regime might be designed in practice.[9,10] Jurisdictional surveys compared the American, European and other approaches to the same questions and found that, while the human-authorship standard had converged, the training-data question had not, and that the two questions were increasingly bound together in the policy debate.(15,25)

The philosophical literature of the period consolidated the distributed turn. Accounts of machine creativity asked directly whether the output of a generative system can be creative in the sense that matters for authorship, and answered with a qualified account in which the creativity is real but is jointly attributable to the system, its trainers and its users rather than to any one of them alone.(13) Posthuman and sociotechnical framings extended this to a full account of authorship as distributed across a network, and the empirical and design literature supplied evidence that human and machine roles in creative work are already being renegotiated in practice.[26,27] The period ends, then, with a stable legal standard and a still-moving philosophical one, and with the training-data question open between them.

Taken together, the four periods describe a single movement with a clear direction. The debate began as a metaphysical question about whether a machine could be an author, and it ended as a practical question about contribution and liability, a movement in which each period's resolution passed the unsettled remainder to the next. What changed was not the underlying doctrine, which always presumed a human author, but the salience of the doctrine and the sophistication of the attempts to operationalize it, and what persisted was the difficulty of drawing the line between the human and the machine in a technology that was deliberately built to blur it. The movement is not complete, since the training-data question opened late and remains open, but its direction is legible, and the discussion that follows interprets it.

DISCUSSION

The trajectory across the four periods resolves into a single, steady observation: the debate changed what it was about, from whether a machine can be an author, which the law answered in the negative in 2023, to how much human involvement suffices, which remains open, and to who bears responsibility for the data on which the machine was trained, which opened later and remains the least settled of the three.[7,9] This reading is consistent with the way the surveys of the field describe the movement, as a shift from a metaphysical question about authorship to a practical question about contribution and liability, and the review's contribution is to date that shift and to show that it was driven by institutional decisions rather than by changes in the technology alone.(10,15)

The law's convergence on a human-authorship standard is the most stable finding of the review, and it deserves to be stated precisely. The United States system anchored the requirement in its constitutional purpose, and the major international analyses found the same premise operating, in different doctrinal forms, across the systems they surveyed, so that a purely machine-generated image is, by the end of the window, unprotected nearly everywhere the question has been decided.(1,8,25) The convergence is nevertheless thinner than it appears, because the systems diverge on the intermediate cases, and the divergence is where the practical action is. A human who prompts, selects and edits a diffusion output is, on the American approach, the author of the human contribution but not of the machine output, while other systems have provisions, such as the United Kingdom's computer-generated-works clause, that reach a different result for the same facts.[19,20] The review therefore treats the human-authorship standard as settled in name and contested in application, which is the honest reading of a literature that agrees on the principle and disagrees on its boundary.

The philosophical movement is the complement and, in part, the correction of the legal one. Where the law drew a line between human and machine contribution, the philosophy of the later periods argued that the line misdescribes how generative systems actually work, in which the human intention, the model's parameters and the training data are all jointly responsible for the output, and no single point in that chain is the author in the traditional sense.[13,24] The

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

empirical work on attribution gave this philosophical claim a social basis, by showing that audiences already perceive AI-made art as distributed rather than as the work of a single author, and the design literature showed that practitioners are building tools on the same assumption.[4,27] The divergence between a law that insists on a single human author and a philosophy that distributes authorship across a network is the deepest unresolved tension the review identifies, and it is unlikely to be resolved by further legal doctrine alone.

The training-data question is where the review's chronological structure is most illuminating, because it shows a problem that was present from the first period but dormant, and that was activated not by any new technical fact but by litigation and regulation. The earliest legal work named the ingestion of copyrighted works for training as a question of infringement and exception, but it remained secondary while the authorship question was unresolved.[1,18] Once the human-authorship standard was fixed, the training-data question moved to the center, and the literature of the final period treats it as the field's new frontier, asking whether the text-and-data-mining exceptions that exist for research can be extended to commercial training and how a lawful training regime might be designed.[9,23] The shift is a case in which the resolution of one question revealed the next, and the review records it as the clearest example of the trajectory's logic.

The trajectory also has a dimension that the legal literature tends to understate but that the philosophical and empirical work brings forward, the social meaning of authorship itself. The question of who authored a work is not only about who owns it but about who is credited, who is accountable, and how the public understands the work's provenance, and the empirical studies of attribution show that these social meanings do not reduce to the legal rule.[4,17] This matters for the policy reading of the trajectory, because a law that assigns authorship to a human user while the public perceives the machine as the author has not, in the relevant sense, settled the question at all. The review therefore treats the social and the legal questions as distinct and argues that the field's future work must hold them together.

The empirical literature on how creativity is perceived adds a second, less visible movement to the trajectory, and it runs partly against the law's confidence. Studies of human evaluators find that people act as gatekeepers of creativity and are systematically biased against work they believe to be machine-made, discounting it even when the output is otherwise indistinguishable from human work, which suggests that the social recognition of authorship is itself a hurdle that no legal rule by itself removes.[4,28] Interviews and surveys of computer scientists and new-media artists show the practitioners themselves disagree about whether generative systems redefine creativity or merely interpolate within the space of what they were trained on, a disagreement that the philosophical literature of the same period encodes rather than resolves.[29,30] The philosophical tradition has its own long answer to this, distinguishing analytic from dialectical and crealectic accounts of creativity, and it reaches toward a position in which machine creativity is real but different in kind from human creativity, a distinction that the later, distributed accounts inherit.[31,32] The line of work that reads generative systems as instruments of emancipation rather than of replacement for digital artists completes this register, and together these studies show that the question of what creativity is remained open beneath the law's settling of who the author is.(33)

The legal literature of the window is broader than the leading American and European cases, and its breadth is itself a finding, because the same doctrinal pressure appears across very different systems. Studies of originality criteria show that the rise of generative systems has forced every system to re-examine what originality means and where the human contribution must be located, and the problem-and-solution literature across jurisdictions converges on the same set of unresolved points even where the doctrinal starting points differ.[34,35] Comparative work on the influence of artificial intelligence on copyright law reaches the same result from the other direction, that the technology changed the law's question rather than its answer, and analyses of the specific challenges and prospects of AI development in particular legal systems repeat the point in a narrower key.[36,37] The law-and-economics literature adds a frame that the doctrinal work often lacks, treating the authorship and training-data questions as problems of incentives and market design, and it concludes that the core challenge is not conceptual but economic, how to preserve the incentives for human creation while allowing the machine production the market clearly wants.(38) Against this, the institutional response has been cautious, a position captured in the observation that the field must rise to meet generative AI rather than legislate it away, and in the professional bodies' own measured responses to the technology.(39,40)

The training-data question, already identified as the new frontier, becomes clearer when the governance literature is read alongside the doctrinal work. The analyses that treat intellectual property and generative AI together argue that the ingestion of protected works for training is the

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

point at which the technology's promise and the law's limits collide, and the case-study work on large language and image models shows that the collision is general rather than specific to any one system.[41,42] The technical and governance literature complements this by showing that the training-data problem is not only legal but technical and ethical, since the provenance, consent and disclosure questions cannot be separated from the question of lawful use.(43) Jurisdictional detail matters here, because text-and-data-mining exceptions differ sharply across systems, and the analyses of the Korean, European and American positions show that a model's training can be lawful in one place and not in another, a divergence that the international-investment framing makes economically consequential.[44,45] Personality rights add a further layer, since the output of a generative model can implicate the likeness and identity of real people in ways that copyright doctrine was never designed to address, and the comparative work on personality rights therefore treats the training-data and output questions as a single problem of rights rather than as two.(46)

The domain-specific literature shows that the authorship problem is not confined to the fine-art cases that dominate the headlines but reproduces itself wherever a creative practice meets a generative tool. The analysis of tattooing is the sharpest example, because it forces the question of who authored a design that a client described, an artist rendered and a machine may have generated, and it shows that the human-authorship standard does not map cleanly onto practices that are already collaborative.(47) The study of automated copyright enforcement extends the concern to the output side, asking whether the systems that police infringement can themselves be trusted to make the authorship judgments that the new technology demands, and the fairness and privacy literature warns that the same generative capacities raise questions of bias and data protection that the authorship debate, narrowly construed, leaves out.[48,49] These adjacent literatures do not change the trajectory the review reconstructs, but they confirm that the trajectory is a general property of the encounter between generative systems and creative institutions, not a peculiarity of the art market.

The agenda that follows from the review is specific. First, the operationalization of the human-involvement threshold, what kind and how much human contribution suffices, is the most practical open problem, because it is where the settled principle meets the unsettled facts, and it is best pursued through comparative work on the intermediate cases across jurisdictions.[11,12] Second, the training-data question needs a lawful regime that the field has not yet produced, and the design analyses of text-and-data-mining exceptions and of collective licensing are the most concrete proposals on the table.[9,15] Third, the gap between the legal and the philosophical accounts, a single author against a distributed network, deserves work that the review can only flag, since it bears on whether the next generation of law will describe the reality of machine creativity or continue to legislate against it.[13,26] These three items are the routes by which the trajectory this review reconstructs would continue, and they are offered as concrete directions rather than as a generic call for more study.

CONCLUSIONS

The evolution of legal and philosophical treatments of authorship, copyright and creativity in AI-generated art from 2018 to 2026 divides into four periods, and the review traced them in sequence. The GAN era opened the question, the diffusion models made it a mass concern, the institutional decisions of 2023 fixed a human-authorship requirement, and the years from 2024 to 2026 consolidated that requirement while adding regulation and training-data litigation. Copyright law converged on the principle that a work needs a human author, but it diverges on the boundary where a human shapes machine output, and it has not yet resolved the status of training data. Philosophy moved in the opposite direction, toward distributed and collaborative accounts of creativity that the legal standard does not reflect. The review concludes that the debate changed what it was about, from whether a machine can be an author to how much human involvement suffices and who bears responsibility for what a model was trained on, and that the gap between a law that insists on a single author and a philosophy that distributes authorship remains the field's central unresolved tension. The practical consequence is that the law has given practitioners a stable answer to the first question, a rule that a work needs a human author, while leaving the questions that now matter, the threshold of human contribution and the status of training data, to the litigation and legislation that the final period of the window set in motion, and the trajectory reconstructed here is best read as the unfinished first chapter of a longer settlement rather than as a completed one.

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

REFERENCES

Chesterman S. ARTIFICIAL INTELLIGENCE AND THE LIMITS OF LEGAL PERSONALITY. ICLQ. 2020;69(4):819-844. doi: 10.1017/s0020589320000366
Xiao Y. Decoding Authorship: Is There Really no Place for an Algorithmic Author Under Copyright Law?. IIC. 2023;54(1):5-25. doi: 10.1007/s40319-022-01269-5
Deltorn JM, Macrez F. Authorship in the Age of Machine learning and Artificial Intelligence. SSRN Journal. 2018. doi: 10.2139/ssrn.3261329
Epstein Z, Levine S, Rand DG, Rahwan I. Who Gets Credit for AI-Generated Art?. iScience. 2020;23(9):101515. doi: 10.1016/j.isci.2020.101515
Cetinic E, She J. Understanding and Creating Art with AI: Review and Outlook. ACM Trans Multimedia Comput Commun Appl. 2022;18(2):1-22. doi: 10.1145/3475799
Lee HK. Rethinking creativity: creative industries, AI and everyday creativity. Media, Culture & Society. 2022;44(3):601-612. doi: 10.1177/01634437221077009
Murray MD. Generative AI Art: Copyright Infringement and Fair Use. STLR. 2023;26(2):259. doi: 10.25172/smustlr.26.2.4
Samuelson P. Generative AI meets copyright. Science. 2023;381(6654):158-161. doi: 10.1126/science.adi0656
Quintais JP. Generative AI, copyright and the AI Act. Computer Law & Security Review. 2025;56:106107. doi: 10.1016/j.clsr.2025.106107
Mazzi F. Authorship in artificial intelligence-generated works: Exploring originality in text prompts and artificial intelligence outputs through philosophical foundations of copyright and collage protection. World Intellectual Property. 2024;27(3):410-427. doi: 10.1111/jwip.12310
Samuelson P. How to Think about Remedies in the Generative AI Copyright Cases. Commun ACM. 2024;67(7):27-30. doi: 10.1145/3654699
Fritz J. Understanding authorship in Artificial Intelligence-assisted works. Journal of Intellectual Property Law and Practice. 2025;20(5):354-364. doi: 10.1093/jiplp/jpae119
Moruzzi C. Artificial Intelligence and Creativity. Philosophy Compass. 2025;20(3):e70030. doi: 10.1111/phc3.70030
Coeckelbergh M. The Work of Art in the Age of AI Image Generation. JHTR. 2023;1. doi: 10.59490/jhtr.2023.1.7025
Geiger C. Elaborating a Human Rights-Friendly Copyright Framework for Generative AI. IIC. 2024;55(7):1129-1165. doi: 10.1007/s40319-024-01481-5
Grilli L, Pedota M. Creativity and artificial intelligence: A multilevel perspective. Creat Innov Manage. 2024;33(2):234-247. doi: 10.1111/caim.12580
Sturm BLT, Iglesias M, Ben-Tal O, Miron M, Gómez E. Artificial Intelligence and Music: Open Questions of Copyright Law and Engineering Praxis. Arts. 2019;8(3):115. doi: 10.3390/arts8030115
Arriagada L. CG-Art: demystifying the anthropocentric bias of artistic creativity. Connection Science. 2020;32(4):398-405. doi: 10.1080/09540091.2020.1741514
Škiljić A. When Art Meets Technology or Vice Versa: Key Challenges at the Crossroads of AI-Generated Artworks and Copyright Law. IIC. 2021;52(10):1338-1369. doi: 10.1007/s40319-021-01119-w
Fedotova A, Romanov A, Kurtukova A, Shelupanov A. Authorship Attribution of Social Media and Literary Russian-Language Texts Using Machine Learning Methods and Feature Selection. Future Internet. 2021;14(1):4. doi: 10.3390/fi14010004
Tyagi K. Copyright, text & data mining and the innovation dimension of generative AI. Journal of Intellectual Property Law & Practice. 2024;19(7):557-570. doi: 10.1093/jiplp/jpae028
Mikalonytė ES, Kneer M. Can Artificial Intelligence Make Art?: Folk Intuitions as to whether AI-driven Robots Can Be Viewed as Artists and Produce Art. J Hum-Robot Interact. 2022;11(4):1-19. doi: 10.1145/3530875
Vesala J. Developing Artificial Intelligence-Based Content Creation: Are EU Copyright and Antitrust Law Fit for Purpose?. IIC. 2023;54(3):351-380. doi: 10.1007/s40319-023-01301-2
Vinchon F, Lubart T, Bartolotta S, Gironnay V, Botella M, Bourgeois-Bougrine S, et al. Artificial Intelligence & Creativity: A Manifesto for Collaboration. Journal of Creative Behavior. 2023;57(4):472-484. doi: 10.1002/jocb.597
Gaidartzi A, Stamatoudi I. Authorship and Ownership Issues Raised by AI-Generated Works: A Comparative Analysis. Laws. 2025;14(4):57. doi: 10.3390/laws14040057
Goodfellow P. The Distributed Authorship of Art in the Age of AI. Arts. 2024;13(5):149. doi: 10.3390/arts13050149
Hu X, Xing Y, Cai X, Zhao Y, Cook M, Borgo R, et al. Designing Interactions with Generative AI for Art and Creativity: A Systematic Review and Taxonomy. Proceedings of the 2025 ACM Designing Interactive Systems Conference. 2025:1126-1155. doi: 10.1145/3715336.3735843
Magni F, Park J, Chao MM. Humans as Creativity Gatekeepers: Are We Biased Against AI Creativity?. J Bus Psychol. 2024;39(3):643-656. doi: 10.1007/s10869-023-09910-x
Wingström R, Hautala J, Lundman R. Redefining Creativity in the Era of AI? Perspectives of Computer Scientists and New Media Artists. Creativity Research Journal. 2024;36(2):177-193. doi: 10.1080/10400419.2022.2107850
Hutson J, Lang M. Content creation or interpolation: AI generative digital art in the classroom. M. 2023;4(1):13. doi: 10.54517/m.v4i1.2158
Miranda LD. Artificial intelligence and philosophical creativity: From analytics to crealectics. Human Affairs. 2020;30(4):597-607. doi: 10.1515/humaff-2020-0053
Liu B. Arguments for the Rise of Artificial Intelligence Art: Does AI Art Have Creativity, Motivation, Self-awareness and Emotion?. Arte, individuo y sociedad (Internet). 2023;Avance en línea:1-11. doi: 10.5209/aris.83808

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

Matias CF. Generative AI, Copyright and Emancipation: The Case of Digital Art. Law Tech Hum. 2024. doi: 10.5204/lthj.3567
KULINICH O, KONDYK D. Impact of Artificial Intelligence on the Criteria of Originality in Copyright Law. UNZ. 2023:43-56. doi: 10.37491/unz.93-94.4
Institute of law and public relations, Open International University of Human development “Ukraine, Kyiv, Ukraine, Valchuk D. ARTIFICIAL INTELLIGENCE COPYRIGHT: PROBLEMS AND SOLUTIONS. MSJ. 2024;3(3):7-15. doi: 10.36994/2786-9008-2024-3-1
Çebi E, Reisoğlu P, Goktas E. The Influence of Artificial Intelligence on Copyright Law. ISSLP. 2023;2(2):33-41. doi: 10.61838/kman.isslp.2.2.5
Waris M, Kim SI. A Study of Analyzing the Challenges and Prospects of Artificial Intelligence Development and its Impact on Copyright Law in the Digital Age. JDMCT. 2024;4(1):1-12. doi: 10.29056/jdmct.2024.06.01
Cooper Z, Martens B, Peukert C, Stocker V. The Law and Economics of Generative AI and Copyright: A Primer to Core Challenges for Our Digital Future. Review of Network Economics. 2025;24(3):161-176. doi: 10.1515/rne-2025-0039
Scott I. Rising to Meet the Challenge of Generative AI. J Legal Studies Education. 2024;41(1):29-37. doi: 10.1111/jlse.12141
Kretschmer M, Meletti B, Porangaba LH. Artificial intelligence and intellectual property: copyright and patents—a response by the CREATe Centre to the UK Intellectual Property Office’s open consultation. Journal of Intellectual Property Law & Practice. 2022;17(3):321-326. doi: 10.1093/jiplp/jpac013
Chesterman S. Good models borrow, great models steal: intellectual property rights and generative AI. Policy and Society. 2025;44(1):23-37. doi: 10.1093/polsoc/puae006
Lucchi N. ChatGPT: A Case Study on Copyright Challenges for Generative Artificial Intelligence Systems. Eur j risk regul. 2024;15(3):602-624. doi: 10.1017/err.2023.59
Pasetti M, Santos JW, Corrêa NK, de Oliveira N, Barbosa CP. Technical, legal, and ethical challenges of generative artificial intelligence: an analysis of the governance of training data and copyrights. Discov Artif Intell. 2025;5(1):193. doi: 10.1007/s44163-025-00379-6
Kim K. Korean Copyright Issues in Text Data Mining for Generative AI. Journal of AI Law and Regulation. 2024;1(1):64-76. doi: 10.21552/aire/2024/1/8
Lentner GM. Generative AI and copyright in the EU and the USA: a view from international investment law. Journal of Intellectual Property Law & Practice. 2026;21(2):121-128. doi: 10.1093/jiplp/jpaf086
Chopra P, Sony RAL, Chopra S. Generative AI, Copyright and Personality Rights: A Comparative Legal Perspective. LIDA. 2025;6(3):23-51. doi: 10.17323/2713-2749.2025.3.23.51
Stockton-Brown M. Inking Cultures: Authorship, AI-Generated Art and Copyright Law in Tattooing. Int J Semiot Law. 2023;36(5):2037-2065. doi: 10.1007/s11196-023-10024-z
Gray JE, Suzor NP. Playing with machines: Using machine learning to understand automated copyright enforcement at scale. Big Data & Society. 2020;7(1):2053951720919963. doi: 10.1177/2053951720919963
Xiang A. Fairness & Privacy in an Age of Generative AI. stlr. 2024;25(2). doi: 10.52214/stlr.v25i2.12765