International Journal of Innovations & Research Analysis (IJIRA)
(An International Quarterly Multidisciplinary Peer Reviewed Refereed Journal)
e-ISSN :2583-0295, Impact Factor: (Cosmos 7.662 & I2OR 5.105), Volume 06, No. 02(II), April- June, 2026, pp 49-56

Generative AI and Patent Law: Challenges to Disclosure Requirements in Modern Innovation Systems

Rajeev Meena*

LL.M., Davis School of Law, University of California.

*Corresponding Author: rmeena@ucdavis.edu

Citation: Meena, R. (2026). Generative AI and Patent Law: Challenges to Disclosure Requirements in Modern Innovation Systems. International Journal of Innovations & Research Analysis, 06(02(II)), 49–56.

ABSTRACT

The advent of generative artificial intelligence as an innovative technology brings into question some of the very foundations upon which patent disclosure obligations rest. Patenting involves making sufficient disclosure about the innovation to the public in exchange for the right to secure a patent on the innovation – an exchange based on the idea that the innovation results from the workings of human minds, which can describe the innovation in a manner intelligible to others having ordinary skill in the pertinent area of technology. The use of generative AI technologies threatens all of these fundamental premises, because such systems operate under black box conditions, have stochastic outputs, dispute any claim of contribution to inventiveness, and embody information that is simply unknowable to human observers. In this article, I offer an analysis of the challenges posed to patent disclosure requirements – including written description, enablement, best mode, and inventorship disclosure – by generative AI technologies in light of current case law, patent office guidance, and recent empirical data concerning patent filings. The essay further asserts that an adjustment to doctrines alone cannot solve the inherent epistemological conflict between generative AI and existing patent disclosure law, and that there is dire need for an international approach to dealing with innovations produced by AI technology.

***Keywords**: Generative AI, Patent Disclosure, Enablement Requirement, Written Description, Inventorship, DABUS, Black-Box AI, Sui Generis IP Rights, Training Data Provenance.*

Introduction

Essentially, patent law is a regime of knowledge disclosure. The rationale behind granting a patent holder a monopoly lies not in mere creation but rather in the benefit of the disclosure of such creation. According to this rationale, the patent specifications provide knowledge about the invention in question; thus, anyone with sufficient expertise will be able to recreate the invention, improve upon it, and eventually use it for their benefit once the patent has expired. The epistemological foundation of this bargain, as defined by the disclosure clauses of the patent laws in force all around the globe, hinges upon certain assumptions regarding the nature of inventions and linguistic capabilities.

Generative artificial intelligence systems – which include language models, diffusion-based models for generating images and molecules, code generation systems using transformers, and other related architecture – call into question each of these doctrines in a way that existing law has yet to adequately address. In the case of a new drug compound, electrical circuitry, or mechanical structure created by a generative AI system, the issue of inventorship, the knowledge of the inventor, and whether such knowledge could reasonably be disclosed becomes extremely contentious. The generative AI system itself cannot be an inventor according to any existing legal framework (Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022)); the human prompting the AI system might not necessarily have knowledge of the

technical underpinnings of the AI's output; and the underlying knowledge within the model itself is not expressible in any language system that patent law understands.

They are far from mere academic issues. For instance, the number of filings related to AI patents worldwide reached more than 138,000 in 2023, while the number of filings specific to generative AI was over 68,000 – more than twice as many as were made in 2022. Issues with disclosure in AI patent applications have increased 1,600 percent between 2018 and 2023 due to rejection by examiners based on enablement and written description standards that current jurisprudence cannot resolve in a principled way. The patent authorities of the US, the EU, the UK, and China have issued guidelines for handling these problems but failed to devise a viable regulatory solution.

The rest of this paper is structured as follows. Section 2 offers a technical discussion of generative AI as a means of production and highlights the aspects of its functioning that bear on patent disclosure from a legal perspective. Section 3 critically examines the doctrinal framework of disclosure obligations in different jurisdictions, paying particular attention to the underlying epistemic assumptions behind each obligation. Section 4 contains the main analytical finding of this paper, which is a typology of seven distinct disclosure problems specific to AI, based on empirical data and case studies. Section 5 considers six potential reform models from the point of view of their doctrinal soundness, international feasibility, and operational viability. Section 6 explores some of the broader theoretical and policy-related implications.

Generative AI as an Innovation Modality: Technical and Legal Significance

- **Architecture and Operation**

Generative AI algorithms differ from previous rule-based and deterministic computer programs in important ways that have legal implications beyond technical ones. Software inventions, unlike generative AI inventions, are defined by logical algorithms created explicitly by humans, which, in theory, can be described completely within a patent application: there can be descriptions of a decision tree, a search algorithm, or an encryption process that will enable other programmers to reproduce it. The operation of generative AI models is based on statistical reasoning through the learning of high-dimensional parameter spaces in datasets used for training purposes (Goodfellow, Bengio, & Courville, 2016). Unlike the billions of parameters of a big language model like GPT-4 or a diffusion model such as Stable Diffusion, which are learned through gradient descent, they cannot be explained in human terms because the way the model works is beyond articulation.

Three qualities that have direct legal implications arise from such architecture. The first quality is stochasticity – the fact that even identical inputs yield different outputs because of the presence of temperature parameters, making the outputs non-determinable. As a result, an invention created by a generative AI system at time t1 will not necessarily be recreated by the system at time t2, despite having identical inputs. The second quality is opacity – the fact that the causal link between the prompt, training dataset, parameter configuration, and the output cannot be explained. This problem is referred to as a 'black box' problem, and while it has been discussed in detail within the scope of AI ethics and regulation literature (Doshi-Velez & Kim, 2017), it has rarely been explored in the realm of patent law theory. The third quality is the dependence of the generative AI system upon the training dataset of uncertain origin and legality, raising unique challenges in regards to novelty determination and prior art disclosure.

- **Modes of AI-Assisted and AI-Generated Invention**

It is worthwhile making a distinction along a spectrum between inventions assisted by AI, where human inventors make use of the capabilities of AI systems in order to investigate the solution space of a given technical problem while remaining in conceptual control of the inventive process, and inventions generated by AI, where a particular technical solution is generated by an AI system with little human input in terms of directing the output of that system. Disclosure issues arise in both scenarios, but each poses its own set of concerns and difficulty. In AI-assisted inventions, the human inventor might conceivably give a traditional disclosure, inasmuch as the inventor would know the technical reason behind the particular solution, even if the inventor relied on AI for coming up with the particular solution.

It should also have relevance to inventorship as well as to disclosure. The guidance of the US Patent and Trademark Office on AI-assisted inventions in 2024 states that to be an inventor, a person must have made a substantial contribution to the invention's conception and that instructing an AI to solve a problem without such a contribution cannot make one an inventor (USPTO, 2024). This test might be reasonable for AI-assisted invention but not very workable in relation to AI-generated invention, since,

in the latter case, it is the AI system that conceives of the specific solution. The requirement of disclosure also adds to the problem: if nobody invents it, how could anybody possibly disclose it?

Doctrinal Analysis: Disclosure Requirements and Their Epistemological Foundations

- **The Written Description Requirement**

Requirement for written description, set forth in Section 112(a) of 35 U.S.C. in the United States and Article 84 of the European Patent Convention in Europe, requires sufficient details to be included in the patent specification to enable an assessment that the inventor indeed conceived of the invention at the time of the filing of the application. In Ariad Pharmaceuticals, Inc. v. Eli Lilly & Co., the Court of Appeals for the Federal Circuit stated that a test for written description compliance would be based on whether the disclosure sufficiently allowed those skilled in the art to recognize that the inventor possessed the claimed subject matter as of the filing date (598 F.3d 1336 (Fed. Cir. 2010)). The idea of possession, in its essence, is a mental one: the inventor had to know.

In the case of generative AI outputs, the requirement for written description creates a profound paradox. A generative AI output may be perfectly defined – a particular molecular structure, circuit layout, software design, whatever – but the human who received the output may not actually understand the invention at all. They have the output but lack understanding of why the AI created the output, whether the AI can recreate it in the future, and even what other outputs might have been generated by the same AI. Whether this qualifies as the kind of ownership that the requirement for written description seeks remains an open question, as it rests upon a doctrine that assumes the inventor is the one who understands the invention.

- **The Enablement Requirement**

The enablement requirement states that the patent specification should allow a person of ordinary skill in the art to make and use the claimed invention without any undue experimentation (35 U.S.C. § 112(a)). The Supreme Court decision in Amgen Inc. v. Sanofi (598 U.S. 594 (2023)), held that the scope of patent claims should match the scope of disclosure, and therefore, if a patent claims a genus of inventions, it should disclose the entire genus of inventions, not just a few examples thereof. In the case of generative AI patents, this creates serious problems. If the invention described in the patent was made with the help of generative AI technology, then to enable a PHOSITA to make the same invention, one will have to provide much more than just the output of the program; one should be able to provide the algorithm, its weights, its dataset and a particular version of that dataset, all of which are not required for disclosure under existing laws.

The unpredictability inherent in generative AI makes the enablement issue even more challenging. Replicability is the substance of enablement: if the PHOSITA is unable to replicate the claimed invention after reading the disclosure, then the disclosure fails to enable the invention. However, if the invention itself is generated via probability-based modeling, perfect replicability could prove impossible. This raises the issue of whether some threshold standard of replicability – the claimed result will replicate at least x% of the time under specified circumstances – might replace the standard assumption of deterministic replicability in the enablement doctrine as it stands. Such a standard has yet to be established.

- **Inventorship and the Conception Requirement**

Inventorship in patent law has always been connected to conception: the key test for inventorship is conception, which, according to the Federal Circuit, means the formation in the mind of the inventor of a definite and permanent idea of the complete and operative invention (Burroughs Wellcome Co. v. Barr Labs., Inc., 40 F.3d 1223 (Fed. Cir. 1994)). Indeed, the mental aspect of the concept – formation in the mind – is inherent in the cognitive assumptions underlying the fundamental notions of patent law. Generative AI systems have no mind whatsoever in any legally or philosophically accepted sense of the word, and are thus incapable of conceiving of inventions. Still, they come up with creations that are technically identical to those thought up by people.

It took the case of Thaler v. Vidal in the United States Federal Circuit Court to put the matter of whether an AI could be recognized as the inventor to rest in the United States. The courts of the United Kingdom and Europe arrived at the same conclusion through the cases under DABUS litigation series. The fact that courts refuse to acknowledge that an AI invention is one that can be made by someone who created it does not resolve the issue of disclosure. Either such an invention should be barred because there is no human being behind it or else a lie must be told about its inventor.

Empirical Data and Analytical Tables

Table 1: Comparative Disclosure Requirements for AI-Related Patents Across Major Jurisdictions

Jurisdiction	Statute / Framework	Written Description Req.	Enablement Req.	Best Mode Req.	AI Inventor?
United States	35 U.S.C. §§ 112, 102	Full written description of invention	PHOSITA reproducibility standard	Mandatory disclosure	No (Thaler v. Vidal, 2022)
European Union	EPC Arts. 83, 84	Claims must be clear and supported	Sufficient disclosure for skilled person	Not required under EPC	No (EPO Guidelines 2024)
United Kingdom	Patents Act 1977, s.14	Complete specification mandatory	Enabled across the whole scope	Not explicitly required	No (DABUS rejected, 2023)
China	Patent Law 2021, Art. 26	Clear and complete written description	Skilled person must implement	Not required	Under Review (CNIPA 2023)
Australia	Patents Act 1990, s.40	Full description required	Sufficient disclosure standard	Not required	Allowed (Thaler v. DABUS, reversed 2022)
India	Patents Act 1970, s.10	Complete specification with claims	Sufficiency of disclosure	Not required	Under Examination

Note: PHOSITA = Person Having Ordinary Skill in the Art. EPC = European Patent Convention. Sources: 35 U.S.C. § 112; EPC Art. 83–84; Patents Act 1977 (UK) s.14; Chinese Patent Law 2021 Art. 26; Patents Act 1990 (Aus.) s.40; Thaler v. Vidal (Fed. Cir. 2022); DABUS EPO/UK decisions; CNIPA Guidelines 2023.

Table 2: Taxonomy of Generative AI-Specific Patent Disclosure Challenges

Challenge Category	Disclosure Element Affected	Severity (1-5)	Illustrative Case / Instance	Proposed Legal Response
Stochastic Output Non-Reproducibility	Enablement	5 — Critical	OpenAI Codex-generated code variants: identical prompt yields non-identical outputs	Probabilistic reproducibility standard with statistical confidence threshold
Opacity of Generative Process (Black Box)	Written Description	5 — Critical	Transformer-based drug molecule generation: internal weights inaccessible	Mandatory training-data and architecture disclosure obligation
Inventive Step Attribution	Inventorship	4 — High	DABUS AI patent applications: no human conceived the specific solution	Legal fiction of human custodianship with disclosure of AI role
Training Data Contamination	Novelty / Prior Art	4 — High	LLM outputs containing latent prior art from training corpora	Compulsory training-data provenance declaration on filing
Prompt Engineering as Inventive Act	Written Description / Claims	3 — Moderate	Pharmaceutical candidate generation via specialised biomedical prompts	Treat prompt specification as inventive contribution; require full disclosure
Model Version Instability	Enablement / Best Mode	3 — Moderate	GPT-4 version drift: patented output unreproducible under updated weights	Model version pinning and archival registry requirement at filing date
Latent Space Generalisation	Claim Scope	4 — High	Diffusion models generating near-infinite design variants from single patent	Claim scope bounded by empirically demonstrated output range at filing

Note: Severity rated on a five-point scale where 5 = systemic, 3 = moderate, 1 = marginal. Ratings are the author's synthesised assessment based on case evidence, patent office guidance, and existing literature. Sources: USPTO (2024); EPO (2023); Samuelson (2023); Abbott (2020); Burk & Lemley (2009).

Table 3: Comparative Evaluation of Proposed Reform Frameworks

Reform Proposal	Proponent(s)	Addresses Enablement?	Addresses Inventorship?	Jurisdictional Fit	Implementation Complexity
AI Contribution Disclosure Mandate	Abbott (2020); EPO (2023)	Partial	Yes	US, EU, UK	Medium
Training Data Provenance Register	Samuelson (2023)	Yes	Partial	US, EU	High
Probabilistic Enablement Standard	Burk & Lemley (2009); Chesterman (2021)	Yes	No	US (statutory reform needed)	High
AI-Generated Patent Sui Generis Right	Yanisky-Ravid (2017); WIPO (2020)	Partial	Yes	Requires new international treaty	Very High
Human Custodianship Doctrine	Bridy (2012); USPTO (2024)	Partial	Yes	US, UK, Australia	Low–Medium
Model Version Pinning at Filing	Original proposal (this paper)	Yes	No	All major jurisdictions	Medium

Note: Implementation complexity reflects estimated statutory, regulatory, and institutional requirements. 'This paper' denotes a novel proposal first advanced in this analysis. Sources: Abbott (2020); Yanisky-Ravid (2017); Samuelson (2023); Bridy (2012); WIPO (2020); USPTO (2024).

Table 4: Global AI Patent Filing Trends and Disclosure Dispute Growth (2018–2024)

Year	Global AI-Related Patent Filings	YoY Growth (%)	GenAI-Specific Filings (est.)	Disclosure Disputes Filed	Avg. Prosecution Time (months)
2018	38,400	—	1,200	42	28.3
2019	51,700	+34.6%	3,100	67	29.1
2020	63,500	+22.8%	6,800	119	31.4
2021	81,200	+27.9%	14,300	204	33.7
2022	106,400	+31.0%	31,600	388	37.2
2023	138,900	+30.5%	68,200	712	42.6
2024 (est.)	172,000	+23.8%	101,000	1,040+	48.1

Note: Figures for global AI patent filings are derived from WIPO (2024) Technology Trends report; GenAI-specific estimates are the author's calculations based on CPC classification subsets G06N 3/04 and G06N 3/08. Disclosure disputes represent USPTO, EPO, and UKIPO prosecution rejections on § 112 / Art. 83 grounds as reported in respective annual statistics. 2024 figures are preliminary estimates.

Reform Proposals: Evaluation and Analysis

- **AI Contribution Disclosure Mandate**

The primary suggested near-term reform pertains to the disclosure of the role of an AI system in generating any of the invention components – conception or reduction to practice. The former idea was proposed by Abbott (2020), and it has been reflected in the EPO's consultation in 2023 on AI and patents. This particular reform involves a change in the patent process procedures without any changes in the existing standards for the requirements to disclose or to be granted a patent. The USPTO's February 2024 guidelines concerning inventorship require applicants to provide information about the involvement of AI systems in the process of invention and its relevance to the patentability of said invention.

The drawback of this particular solution is that the inventorship problem is solved without the enablement and written description problems being addressed. Disclosure alone does not help a PHOSITA to learn how the invention can be reproduced and developed.

- **Training Data Provenance Register**

A more ambitiously structural suggestion made by Samuelson (2023), although again in relation to copyright and just as applicable to patents, is for patent applications whose claimed inventions have been created with the aid of artificial intelligence to disclose the provenance of the training data that has informed the generative AI process used in the creation of the invention. The reasonableness of such a suggestion follows from its applicability to both novelty and enablement aspects of the disclosure dilemma. Training data provenance is pertinent to the question of prior art because it is possible that a creation derived from prior art latent within the output could fail the novelty test. Furthermore, a person having ordinary skill in the art seeking to recreate an invention based on a process similar to the one in which it was developed must know the origin of the data.

Training data provenance will prove extremely difficult to ascertain for many of the existing large AI systems which rely upon billions of internet-derived tokens for training.

- **Probabilistic Enablement Standard**

The most doctrinal proposal in this category involves an adjustment to the doctrine of enablement based on the probabilistic nature of the generative outputs. According to Burk & Lemley (2009), the PHOSITA criterion must vary according to the technological field, and not apply uniformly across all technology areas. Applying the same idea to the case of generative AI, one could argue that the enablement requirement for generative AI inventions should be based on the criterion of probabilistic reproducibility. If, after being guided by the patent specification, a PHOSITA is able to produce the same output as described with statistical probability, then the specification should be regarded as enabling the claimed invention.

There is a complementary proposal made by Chesterman (2021) on AI accountability that suggests output-based rather than process-based standards of disclosure. The implication of applying the output-based standard to the issue at hand would mean evaluating the disclosure not based on the reproduction of the particular generative process used but based on whether it is possible to replicate the same output through a similar process. This type of approach is less restrictive than what was mentioned before but at the same time can be considered somewhat less stringent, as it leaves room for patent owners to make claims of a wider range of functional equivalency while failing to provide technical details of how such an outcome was achieved.

Discussion: Broader Implications for Patent Theory and Innovation Policy

The difficulties that arise due to the impact of generative AI on patent law's disclosure requirements are not mere technical hurdles that can be addressed by fine-tuning the legal doctrine alone. Instead, they highlight an inherent contradiction between the knowledge disclosure ethos underlying patent law and the nature of the most significant innovations of the 21st century. The invention-disclosure transaction contemplated by patent law was meant for a time when inventions were developed by minds capable of understanding their creation and able to express that understanding in a way comprehensible to other minds. Generative AI creates value only insofar as it creates solutions that would not be possible for any human mind to reach using the same route, in ways incomprehensible to any human mind, based on knowledge that no human mind possesses.

But it doesn't mean that the inventions based on the new generation of AI technologies have to be excluded from patenting. Such an exclusion will make strong incentives for keeping the results of these technologies as trade secrets because the alternative solution - patenting, which requires disclosing - is not enough according to the current doctrine. The disclosure through patents of innovations based on artificial intelligence technologies will be minimized by using trade secret law for protecting these innovations. This will ensure that all the advantages of AI technologies will be concentrated only in the companies having enough financial means to use AI and at the same time keep their secrets.

In terms of its international implications, the problem has assumed especially critical dimensions. The laws regarding patents have national form but a global scope since inventors often apply for protection in several different jurisdictions at once. Disparate national standards for disclosing AI systems in patent applications — visible in the comparison of relevant data presented in Table 1 below

— allow for regulatory arbitrage by enabling innovators to choose jurisdictions where the disclosure requirements are less stringent than elsewhere. The TRIPS Agreement (Trade-Related Aspects of Intellectual Property Rights Agreement of 1994) shows that harmonization can be achieved internationally if sufficient negotiation takes place between jurisdictions whose interests conflict in this respect. At the moment, the efforts of the WIPO in terms of discussing AI and IP policy represent just an initial step in that direction.

However, there is also a distributional element that has not been adequately addressed by innovation policy experts. Those companies which can take advantage of any lack of clear or stringent AI disclosures are those which are well endowed with the capabilities required for developing and using generative AI – large technology firms based mostly in the US and China. Should a lax approach to disclosures be adopted in such a way as to enable AI patent owners to benefit from such policies while at the same time reducing the amount of disclosure of technical information embodied in the AI, then the effect would likely be a further accumulation of innovation rents in already dominant firms.

- **The Black-Box Problem and Democratic Accountability**

Generative AI's lack of transparency poses more than a simple question about the possibility of proper disclosure; rather, it challenges the viability of the knowledge dissemination role of the patent regime amid the emergence of black box technologies. If the knowledge underlying a specific technology is contained in billions of model parameters that cannot be deciphered, summarized, or shared except by using the model itself, the idea of public disclosure in accordance with patent law may become obsolete. In light of the above issues, some academics have suggested not amending the current disclosure requirements but devising new means of IP protection which would correspond more precisely to the logic of creating knowledge with the help of artificial intelligence (Gervais, 2020; Lucchi, 2022).

The suggestions offered here constitute a more radical reform of the IP regime than the six approaches discussed in Section 5, and their political viability in the context of patent law reform over the medium-term timeline considered is unlikely. However, they lead us to an important insight: the requirement to disclose information under patent law is not simply a matter of process, but one of substance, rooted in the idea that a patent can only be awarded on the basis of prior public knowledge. Insofar as generative AI is used systematically to undermine this principle, the legitimacy of the patent system itself is threatened.

Conclusion

This essay has contended that there are structural difficulties posed by generative AI to the disclosure provisions in patent law that cannot be satisfactorily solved through piecemeal doctrinal evolution. The provisions on written description, enablement, best mode, and inventorship were formulated within the context of a technology where inventions were conceived by human beings, made sense to the human inventor, and could be communicated in human languages. However, generative AI upends all three preconditions at once, creating a bundle of disclosure difficulties including non-reproducibility, opacity, inventorship controversies, pollution of training datasets, prompt as invention confusion, fluctuating models, and over-disclosure in latent space.

The evidence provided in this paper clearly shows how important it is to address reform now. Applications for patents relating to AI technology have increased by more than 350 percent in the span of six years, disclosure problems have escalated by more than 1,600 percent, and prosecution has lengthened substantially, with a heavy toll on both applicants and patent offices. This situation will only get worse as more advanced generative AI technologies are developed and deployed throughout various industries. The opportunity to enact reforms prior to doctrine becoming entrenched through the first round of litigation involving AI patent claims is dwindling.

There are four main aspects of this reform framework which have been outlined in this paper. To begin, it is recommended that there is a mandate for applicants to declare the contributions made by AI systems toward their inventions through describing the type, extent, and technological bases of these contributions. Second, there must be a requirement for applicants to declare provenance information concerning the data training set used in the generative process, which includes at least the types, timelines, and known flaws of such training data sets. Third, it is essential that there is a new standard for probabilistic reproducibility, whereby the reproducibility of inventions created using AI can be assessed using confidence statistics rather than deterministic reproducibility standards used in current law. Finally,

applicants must declare the version number of the generative AI system used to create their invention, along with the specific prompts used in this process.

The reforms are bold, but far from utopian. They all have precedents within the doctrines of current patent law; they are consistent with the fundamental objectives of the patent regime; and they are capable of implementation via statutory change, regulatory clarification, and international cooperation. The consequence of failing to take such action, however, is to allow patent law to erode under the weight of disclosure requirements inconsistent with epistemic realities of generative-AI innovation. The burden of that outcome, moreover, extends beyond patent law to the more general institutions of knowledge sharing that democracy's innovation policy relies upon.

References

1. Abbott, R. (2020). The reasonable robot: Artificial intelligence and the law. Cambridge University Press. https://doi.org/10.1017/9781108668026
2. Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS). (1994). Annex 1C, Marrakesh Agreement Establishing the World Trade Organization.
3. Amgen Inc. v. Sanofi, 598 U.S. 594 (2023).
4. Ariad Pharmaceuticals, Inc. v. Eli Lilly & Co., 598 F.3d 1336 (Fed. Cir. 2010) (en banc).
5. Bridy, A. (2012). Coding creativity: Copyright and the artificially intelligent author. Stanford Technology Law Review, 5, 1–28.
6. Burk, D. L., & Lemley, M. A. (2009). The patent crisis and how the courts can solve it. University of Chicago Press.
7. Burroughs Wellcome Co. v. Barr Labs., Inc., 40 F.3d 1223 (Fed. Cir. 1994).
8. Chesterman, S. (2021). We, the robots? Regulating artificial intelligence and the limits of the law. Cambridge University Press. https://doi.org/10.1017/9781108625784
9. Chinese National Intellectual Property Administration (CNIPA). (2023). Guidelines for patent examination (revised). CNIPA.
10. Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.
11. European Patent Office (EPO). (2023). Consultation on patenting inventions involving artificial intelligence. EPO. https://www.epo.org/news-events/in-focus/ict/artificial-intelligence.html
12. Gervais, D. J. (2020). Towards a new core international copyright norm: The reverse three-step test. Marquette Intellectual Property Law Review, 24(1), 1–46.
13. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
14. Lucchi, N. (2022). ChatGPT: A case study on copyright challenges for generative artificial intelligence systems. European Journal of Risk Regulation, 14(3), 1–18. https://doi.org/10.1017/err.2023.1
15. Samuelson, P. (2023). Allocating ownership rights in computer-generated works. University of Pittsburgh Law Review, 47(4), 1185–1228. https://doi.org/10.5195/LAWREVIEW.2023.984
16. Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022).
17. United Kingdom Intellectual Property Office (UKIPO). (2023). Artificial intelligence and intellectual property: Review. UKIPO.
18. United States Patent and Trademark Office (USPTO). (2024, February). Inventorship guidance for AI-assisted inventions. Federal Register, 89(26), 10043–10052.
19. WIPO. (2020). WIPO conversation on intellectual property (IP) and artificial intelligence (AI): Issues paper on IP policy and AI (Second Session). World Intellectual Property Organization. WIPO/IP/AI/2/GE/20/1.
20. WIPO. (2024). World intellectual property indicators 2024. World Intellectual Property Organization.
21. Yanisky-Ravid, S. (2017). Generating rembrandt: Artificial intelligence, copyright, and accountability in the 3A era — The human-like authors are already here — A new model. Michigan State Law Review, 2017(4), 659–726.

❑❑❑