INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS

Open Access, Refereed Journal Multi Disciplinary
Peer Reviewed



www.ijlra.com

ISSN: 2582-6433

DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

Copyright © International Journal for Legal Research & Analysis. All rights reserved.

ISSN: 2582-6433

ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

ISSN: 2582-6433

PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

ISSN: 2582-6433

PATENT DRAFTING IN THE AGE OF GENERATIVE AI: FOUNDATIONS, COMPARATIVE REQUIREMENTS, AND THE PROBLEM OF HALLUCINATION

AUTHORED BY - ISHAN ANAND & ARUNDHATI KUMARI

Abstract

Patent drafting is the stage at which an inventor's contribution is converted into a legal instrument and the quality of that instrument largely determines whether a patent can be granted, enforced and defended. This article examines the foundations of patent drafting: the historical development of the patent system, the comparative requirements for specifications and claims under the principal regimes (the United States, the European Patent Convention, the United Kingdom, China, India and the Patent Cooperation Treaty), and the anatomy of a well-constructed application. It then turns to the development that is now reshaping the discipline, namely the use of large language models to draft claims, specifications and prosecution responses. The article argues that generative AI leaves the substantive law of disclosure untouched but changes the economics and the failure modes of compliance. It develops a taxonomy of drafting hallucinations, maps each category onto existing doctrines (enablement, written description, sufficiency, added matter, definiteness, candour and inventorship), reviews the emerging technical research and regulatory guidance, and proposes a verification framework for AI-assisted drafting. A worked example in the annexes illustrates where unsupervised generation is most likely to fail. The article extends the author's earlier paper on the basics of patent drafting, published in this journal in May 2025.

Keywords: Patent drafting; generative artificial intelligence; hallucination; sufficiency of disclosure; inventorship; comparative patent law.

ISSN: 2582-6433

I. Introduction

A patent is a time-limited exclusive right granted by the State in exchange for the public disclosure of an invention.[1] The specification must teach the public how to make and use the invention, and the claims must delimit the territory from which the patentee may exclude others. Errors in that document are rarely curable after filing, because every major system forbids amendments that add matter not disclosed in the application as filed.[2] Drafting is therefore not clerical work that precedes the legal process; it is the point at which most of the legal outcome is fixed.

For most of the modern history of the patent system, drafting was a craft performed by human specialists: patent attorneys and agents who translated an inventor's disclosure into a document satisfying the formal and substantive requirements of one or more jurisdictions. Since 2022, large language models (LLMs) have entered that craft. Purpose-built platforms now offer to generate claim sets from invention disclosures, expand claims into full specifications, produce figure descriptions, check internal consistency and draft responses to examiners' objections. Their adoption has been rapid, although much of the available evidence on adoption comes from vendors.[3]

It is important to distinguish two roles that artificial intelligence plays in patent law. AI may be the *subject matter* of an invention, as when a machine-learning method is claimed; patent offices report that applications of this kind have grown dramatically.[4] Or AI may be a *tool* used to prepare an application for any kind of invention. This article is concerned principally with the second role, although the first is relevant where it changes what must be disclosed.

The article proceeds as follows. Part II sets out the historical and international foundations of the patent system. Part III compares the principal criteria for patentability and entitlement. Part IV analyses the anatomy of a specification and the comparative requirements that govern

[1] Agreement on Trade-Related Aspects of Intellectual Property Rights art. 33, Apr. 15, 1994, 1869 U.N.T.S. 299 [hereinafter TRIPS] (requiring a minimum term of twenty years from the filing date).
[2] *See* 35 U.S.C. § 132(a); Convention on the Grant of European Patents art. 123(2), Oct. 5, 1973, 1065 U.N.T.S. 199 [hereinafter EPC]; Patents Act 1977, c. 37, § 76 (UK); The Patents Act, 1970, § 59 (India).
[3] *See, e.g.*, Jeff O'Neill, *AI Patent Drafting Tools: Patent Bots vs. Other Solutions*, Patent Bots, https://www.patentbots.com/patent-drafting-tools-comparison (comparing five drafting tools); *Best AI Tools for Patent Drafting in 2026: 13 Platforms Ranked by Use Case*, Rework, https://resources.rework.com/tools/ai-tools/best-ai-tools-for-patent-drafting-2026 (reporting, on the basis of a Clarivate survey, a rise in AI adoption among IP professionals from 57% to 85% over two years). These are commercial or vendor-adjacent sources and the figures have not been independently verified.
[4] Bastian Best, *[Guest Post] Will the "Human in the Loop" Please Stand Up? Recap from Search and Examination Matters 2026*, IPKat (Feb. 11, 2026), https://ipkitten.blogspot.com/2026/02/guest-post-will-human-in-loop-please.html (reporting an EPO statement that AI-related applications have increased 45-fold since 2015).

ISSN: 2582-6433

it. Part V examines how generative AI has changed drafting practice and reviews the technical literature and the tools. Part VI develops a taxonomy of hallucination in patent drafting and maps it onto existing doctrine. Part VII reviews regulatory and judicial responses, including on inventorship. Part VIII proposes a verification framework, and Part IX concludes.

A. Research Questions, Hypothesis and Method:

A1. Research Questions:

1. What are the essential components of a patent application, and how do the principal jurisdictions differ in their requirements?
2. How has generative AI altered the practice of patent drafting, and what new risks, in particular hallucination, does it introduce?
3. What doctrinal and regulatory responses exist, and what verification practices follow from them?

A2. Hypothesis:

The working hypothesis is that the clarity, support and internal consistency of the language of an application remain determinative of the validity and enforceability of the resulting patent, and that generative AI increases rather than reduces the premium on rigorous human verification.

A3. Method Used:

The method is doctrinal and comparative. It draws on primary sources (statutes, treaties, judicial decisions and official guidance) and on secondary literature, including technical research on LLM-based drafting. No primary empirical data were collected. Developments are stated as at November 2026.

A4. Citation Followed:

Bluebook 21st edition citation guidance is followed throughout the paper.

B. Relationship to the Author's Earlier Paper

This article builds directly on, and substantially extends, the author's earlier paper, *Analysing the Basics of Patent Drafting*, published in this journal in May 2025.[5] The earlier paper introduced the origins of the patent system, the criteria for patentability and eligibility, the contents of a specification and the anatomy of a patent application. It was framed principally around the Indian Patents Act, 1970 and did not address the use of artificial intelligence in

[5] Ishan Anand, *Analysing the Basics of Patent Drafting*, 2(7) Int'l J. Legal Rsch. & Analysis 5 (May 2025) (ISSN 2582-6433) [hereinafter Anand, *Basics*].

ISSN: 2582-6433

drafting.

The present article retains the earlier paper's core subject, namely the foundations of patent drafting, and its worked example, which is reproduced in revised form as Annex A. It differs in four respects. First, it recasts the analysis comparatively, treating the Indian provisions as one regime among several rather than as the organising framework. Second, it restates the historical and statutory material in updated form and corrects points in the earlier treatment of subject-matter exclusions and specification requirements. Third, and most substantially, it adds an analysis of how generative AI has changed drafting practice, including a taxonomy of hallucination, a review of the technical research and of the regulatory guidance issued up to November 2026, and a proposed verification framework (Parts V to VIII). Fourth, it adds a new annex applying the taxonomy to the example application. Passages that restate the earlier paper's material have been rewritten and the author acknowledges the earlier paper as the foundation of Parts II to IV.

II. Historical and International Foundations

A. Origins

Commentators trace precursors of the patent to antiquity, but the institution is conventionally dated from fifteenth-century Venice.[6] The Venetian Patent Statute of 1474 granted inventors of new and ingenious devices an exclusive right for a fixed period, with the aim of attracting craftsmen and encouraging innovation. In England, the Statute of Monopolies of 1623 (21 Jac. 1, c. 3) curtailed the Crown's practice of granting monopolies while preserving a limited exception for the "true and first inventor" of a new manufacture for a term of fourteen years. The United States adopted its first patent statute in 1790, pursuant to the Intellectual Property Clause of its Constitution, also with a fourteen-year term.[7] Revolutionary France enacted its first patent law in 1791, grounding the right in a conception of the inventor's natural entitlement to the fruits of invention.[8]

These early statutes were transplanted, adapted and consolidated across the world through commerce, empire and later treaty. Many national systems began as local enactments modelled on the law of a trading or colonial partner. The Indian Act VI of 1856, for example, was modelled on the British Patent Law Amendment Act of 1852, and similar patterns can be traced

[6] G. Krishna Tulasi & B. Subba Rao, *A Detailed Study of Patent System for Protection of Inventions*, Nat'l Libr. Med., https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3038276 (last visited Apr. 28, 2024).
[7] Patent Act of 1790, ch. 7, 1 Stat. 109; U.S. Const. art. I, § 8, cl. 8.
[8] Loi du 7 janvier 1791 relative aux découvertes utiles et aux moyens d'en assurer la propriété à ceux qui seront reconnus en être les auteurs (Fr.).

ISSN: 2582-6433

in other Commonwealth jurisdictions.[9] What the early statutes shared was a quid pro quo: a temporary exclusivity in return for useful novelty and, increasingly, for disclosure.

B. The Modern International Framework

Four instruments shape contemporary drafting practice. The Paris Convention of 1883 established national treatment and the right of priority, under which an applicant who files in one member State has twelve months to file in others and claim the earlier date.[10] The Patent Cooperation Treaty (PCT) of 1970 allows a single international application to preserve the right to seek protection in over one hundred and fifty contracting States, deferring national-phase costs and decisions.[11] The European Patent Convention of 1973 created a centralised grant procedure that results in a bundle of national patents. And the TRIPS Agreement of 1994 set minimum standards: patents must be available in all fields of technology, subject to exclusions, and the applicant must disclose the invention in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art.[12]

TRIPS also gave developing countries transition periods, and its implementation drove significant legislative change in several of them. India's amendments of 1999, 2002 and 2005, which culminated in the reintroduction of product patents for pharmaceuticals, are one well-documented example among many.[13] The practical consequence for the drafter is that an application is now seldom written for one jurisdiction. A specification prepared for a PCT filing will be examined, sometimes years later, under the divergent standards of the United States, the European Patent Office (EPO), China and other offices, and must be drafted accordingly.

III. Patentability and Entitlement: A Comparative Sketch

A. Substantive Criteria

The core criteria are common in outline. The invention must be new (novelty), must not be obvious to a person skilled in the art (non-obviousness in the United States, inventive step elsewhere), and must be capable of industrial application or have utility.[14] The detail differs. The EPO assesses inventive step through the "problem-solution approach," which turns on the

[9] *History of the Indian Patent System*, Off. of the Controller Gen. of Pats., Designs & Trade Marks (India), https://ipindia.gov.in/history-of-indian-patent-system.htm (last visited May 24, 2024).
[10] Paris Convention for the Protection of Industrial Property arts. 2, 4, Mar. 20, 1883, 828 U.N.T.S. 305.
[11] Patent Cooperation Treaty arts. 3-5, June 19, 1970, 28 U.S.T. 7645, 1160 U.N.T.S. 231.
[12] TRIPS, *supra* note 1, arts. 27, 29.
[13] TRIPS, *supra* note 1, art. 65; The Patents (Amendment) Act, 2005, No. 15 of 2005, India Code (introducing product patent protection in all fields of technology from Jan. 1, 2005).
[14] *See* 35 U.S.C. §§ 101-103; EPC arts. 52(1), 54, 56, 57; Patents Act 1977 §§ 1(1), 2, 3, 4 (UK); The Patents Act, 1970, § 2(1)(j), (ja), (ac) (India); Patent Law of the People's Republic of China art. 22 (2020).

ISSN: 2582-6433

objective technical problem solved by the claimed features relative to the closest prior art, so that the technical problem and effect identified in the description have direct legal importance.[15] The United States applies a more open-textured obviousness analysis, in which the specification's statements about the prior art can be used against the applicant. The lesson for the drafter is that the description is not a neutral narrative but evidence in the examination of patentability.

B. Subject-Matter Exclusions

All major systems exclude some subject matter, but they do so on different grounds and with different vocabulary. In the United States, 35 U.S.C. § 101 is subject to judicial exceptions for laws of nature, natural phenomena and abstract ideas.[16] The EPC excludes discoveries, scientific theories, mathematical methods, aesthetic creations, schemes for mental acts, games and business methods, and programs for computers, in each case only "as such," and separately excludes inventions contrary to *ordre public* and methods of medical treatment.[17] The United Kingdom's list in section 1(2) of the Patents Act 1977 is similar to the EPC's.[18] China excludes scientific discoveries, rules and methods for mental activities, and methods of diagnosis or treatment, among others.[19] India's list in sections 3 and 4 of the Patents Act, 1970 is more extensive and includes the exclusion of "a computer programme per se," and of the mere discovery of a new form of a known substance unless it shows enhanced efficacy.[20] TRIPS permits exclusions for public order, medical methods, and certain plants and animals.[21]

These differences matter for AI-implemented inventions in particular. The Federal Circuit has held that claims to the application of generic machine-learning techniques to a new data environment, without a specific improvement to the technique itself, were ineligible under § 101.[22] The EPO, by contrast, applies the "any hardware" approach and asks whether the claimed subject matter makes a technical contribution, a framework now elaborated in its Guidelines for artificial intelligence and machine learning.[23] A drafter must therefore frame the same invention differently depending on the intended forum.

---

[15] Eur. Pat. Off., *Guidelines for Examination* pt. G, ch. VII, § 5 (Apr. 2026).
[16] *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208 (2014); *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66 (2012); *Ass'n for Molecular Pathology v. Myriad Genetics, Inc.*, 569 U.S. 576 (2013).
[17] EPC arts. 52(2)-(3), 53(a), 53(c).
[18] Patents Act 1977 § 1(2) (UK).
[19] Patent Law of the People's Republic of China art. 25 (2020).
[20] The Patents Act, 1970, §§ 3(d), 3(k), 4 (India).
[21] TRIPS, *supra* note 1, art. 27(2)-(3).
[22] *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025), *cert. denied* (U.S. Dec. 8, 2025).
[23] *Hitachi*, T 258/03, 2004 O.J. EPO 575; *Comvik*, T 641/00, 2003 O.J. EPO 352; Eur. Pat. Off., *Guidelines for Examination* pt. G, ch. II, § 3.3.1 (Apr. 2026).

ISSN: 2582-6433

C. Entitlement

The right to apply belongs in principle to the inventor or the inventor's successor in title. The details differ. Under U.S. law an assignee or obligated assignee may file as applicant, whereas the EPC allows any person to whom the right has passed to apply, and the United Kingdom and India follow broadly similar approaches.[24] What is common to all of them, and has become central since the development of AI systems, is the premise that the inventor is a natural person. That premise is examined in Part VII.

IV. The Anatomy of a Patent Specification

A. Comparative Requirements

The principal components of a specification (the description, the claims, the drawings and the abstract) are broadly common, but the standards governing them are not. Table 1 summarises the central requirements in the jurisdictions examined.

Table 1. Principal specification requirements in selected jurisdictions

Jurisdiction	Disclosure standard	Claim requirements	Best mode	Principal authorities
United States	Enablement of the full scope of the claims; written description showing possession	Particularly point out and distinctly claim; definiteness (reasonable certainty)	Required to be disclosed, but failure to do so is not a ground of invalidity	35 U.S.C. § 112; *Amgen Inc. v. Sanofi*, 598 U.S. 594 (2023); *Ariad Pharms., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336 (Fed. Cir. 2010) (en banc); *Nautilus, Inc. v. Biosig Instruments, Inc.*, 572 U.S. 898 (2014)

[24] *See* 35 U.S.C. § 118; EPC art. 58; Patents Act 1977 § 7 (UK); The Patents Act, 1970, § 6 (India).

ISSN: 2582-6433

Jurisdiction	Disclosure standard	Claim requirements	Best mode	Principal authorities
European Patent Convention	Sufficiency: invention disclosed clearly and completely enough to be carried out by a skilled person; added-matter bar	Clear, concise and supported by the description	Not required; at least one way of carrying out the invention must be described	EPC arts. 83, 84, 123(2); EPC R. 42(1)(e)
United Kingdom	Clear enough and complete enough to be performed by a skilled person	Clear, concise, supported by the description	Not required	Patents Act 1977 § 14(3), (5)
China	Clear and complete enough for a skilled person to carry out	Supported by the description; clear and concise	Not required	Patent Law of the P.R.C. art. 26
India	Full and particular description of the invention, its operation or use, and the method of performing it	Clear, succinct and fairly based on the matter disclosed; single inventive concept	Required: best method known to applicant	The Patents Act, 1970, § 10(4)-(5)
PCT	Disclosure sufficiently clear and complete to be carried out by a skilled person	Clear and concise; fully supported by the description	Not required (national phase governs)	PCT arts. 5, 6; PCT R. 5

Sources: statutory and treaty provisions cited in the final column. The table is a summary for orientation and omits many qualifications.

Three observations follow. First, the disclosure standard is formulated differently but is directed to the same end: the skilled person must be able to perform the invention without

ISSN: 2582-6433

undue burden. Second, support for the claims in the description is a recurrent requirement, framed as written description in the United States, as clarity and support in the EPC, the United Kingdom, China and the PCT, and as "fair basis" in India. Third, the best-mode requirement is the principal outlier, being mandatory in the United States and India (though in the United States no longer enforceable as a ground of invalidity) and absent in Europe and China.[25] A specification written for global filing should satisfy the most demanding standard on each point.

B. The Components

Title and field. The title should identify the subject of the invention concisely. The statement of the technical field situates the invention for classification and search.

Background and statement of the problem. The background describes the state of the art and the problem the invention addresses. It must be drafted with care because characterisations of prior art may later be treated as admissions, and a problem statement that is too narrow may limit the way the claims are read. Where an application sets out the objects of the invention, they should be framed as the problems addressed rather than as results that every embodiment must achieve.

Summary. The summary states the invention in terms that correspond to the claims. Language such as "the present invention is" invites a court to treat a described feature as essential to every embodiment and to limit the claims accordingly.[26] Summaries are best drafted by reference to embodiments ("in one embodiment") and to the claims.

Drawings and their description. Drawings need not be provided for every invention but are almost always advisable for apparatus and process inventions. Reference numerals must correspond consistently between drawings, description and claims, an area in which automated consistency checking is now especially active (Part V).

Detailed description. The detailed description must describe at least one embodiment in sufficient detail to enable the skilled person to practise the invention, together with alternatives and variations that support the breadth of the claims. It is here that the sufficiency and support requirements bite. A claim to a broad genus supported by one narrow example runs the risk that the specification will be held not to enable the full scope claimed.[27] Where results are claimed, the conditions under which they are obtained must be disclosed. Consider an engine

[25] 35 U.S.C. §§ 112(a), 282(b)(3)(A).

[26] *Verizon Servs. Corp. v. Vonage Holdings Corp.*, 503 F.3d 1295 (Fed. Cir. 2007).

[27] *See Amgen*, 598 U.S. 594 (holding that enablement requires the specification to teach the full scope of the claimed invention).

ISSN: 2582-6433

claimed to deliver one hundred kilometres per litre of fuel. If some users obtain ninety kilometres per litre and others one hundred, the specification must state the operating conditions under which the claimed figure is achieved, or the claim will not be reproducibly performable. Where examples are predicted rather than actually performed, they must be presented as such; in United States practice, applicants are expected to distinguish working from prophetic examples.[28]

Claims. The claims define the scope of protection. They should contain independent claims of different scope and dependent claims that offer fall-back positions. Each term must have antecedent basis and be used consistently, and dependent claims must refer to the claim that introduces the elements on which they depend. Claims must relate to a single inventive concept or a group of inventions so linked as to form a single general inventive concept.[29] The claims are construed in light of the specification, which in the United States is the settled approach under *Phillips* and in Europe is embodied in Article 69 EPC and its Protocol.[30]

Abstract. The abstract serves as technical information only and does not define the scope of protection.[31]

C. Provisional and Priority Filings

Most systems allow an applicant to secure an early filing date with a preliminary document that does not need to contain claims. The United States provides for provisional applications, India for provisional specifications to be followed by a complete specification, and the Paris Convention priority system lets an applicant follow a first filing anywhere with later filings within twelve months.[32] The practical hazard is that a thin preliminary filing may fail to support the claims later presented, and priority is lost to the extent that the later claims are not supported by the earlier disclosure. Annex A presents an example application that illustrates the components discussed above, and Annex B applies review checkpoints to it.

V. Drafting in the Age of Generative AI

A. The AI-Assisted Drafting Workflow

Generative AI is now used at nearly every stage of the drafting and prosecution cycle. It is

[28] U.S. Pat. & Trademark Off., *Manual of Patent Examining Procedure* § 2164.02 (current ed.).
[29] *See* EPC art. 82; PCT R. 13; The Patents Act, 1970, § 10(5) (India).
[30] *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc); EPC art. 69 & Protocol on the Interpretation of Article 69 EPC.
[31] EPC art. 85; PCT art. 3(3); PCT R. 8.
[32] 35 U.S.C. § 111(b); The Patents Act, 1970, § 9 (India); Paris Convention, *supra* note 10, art. 4.

ISSN: 2582-6433

useful to separate the stages, because the risk profile differs across them.

At the *intake* stage, tools summarise inventor disclosures, technical papers and meeting notes into structured invention records. At the *search* stage, AI-assisted engines retrieve and cluster patent and non-patent literature, and several platforms generate preliminary novelty assessments. At the *claim drafting* stage, models propose independent and dependent claim sets from a disclosure, sometimes in multiple scopes. At the *specification* stage, models expand claims into a full description, generate background and summary text, and produce descriptions of figures. At the *quality-control* stage, tools check antecedent basis, reference numerals and term consistency. At the *prosecution* stage, they analyse office actions, map claim elements against cited references and draft responses. At the *foreign filing* stage, they translate and localise specifications.

The stages do not carry equal risk. Tasks that restructure or check text already written by a human, such as proofreading or consistency checking, produce errors that a reviewer can see by comparing the output with the input. Tasks that generate new technical content, such as expanding a claim into a description or proposing alternative embodiments, produce errors that cannot be detected by comparison because there is no source against which to compare them. The second category is where hallucination is most dangerous (Part VI).

B. The Tool Landscape

Practitioners can choose among three broad categories of tool. The first comprises general-purpose LLMs used through enterprise subscriptions or application programming interfaces. The second comprises purpose-built drafting and prosecution platforms, including Solve Intelligence, DeepIP, Patlytics and Patent Bots, which differ in whether they run in a browser or inside a word processor and in whether they extend from drafting into claim charting, invalidity analysis and examiner analytics.[33] The third comprises rules-based drafting aids, such as ClaimMaster, that check formal consistency without generating text, and prior art engines, such as IPRally and PatSnap's Eureka, that use machine learning to retrieve rather than to draft.[34]

[33] *See* O'Neill, *supra* note 3; *Best AI Patent Drafting Tools in 2026*, DeepIP (last updated Jan. 18, 2026), https://deepip.ai/blog/best-ai-patent-drafting-tools-in-2025; *8 Best AI Patent Drafting Tools in 2026 (Compared)*, Solve Intelligence (last updated Sept. 24, 2026), https://solveintelligence.com/blog/post/best-ai-patent-drafting-tools. Each of these is published by a vendor or a competitor and should be read as marketing material rather than as independent evaluation.

[34] *See* O'Neill, *supra* note 3 (distinguishing rules-based proofreading from generative drafting); *Top 7 Patent Prosecution Software Tools in 2026*, TechBullion, https://techbullion.com/top-7-patent-prosecution-software-tools-in-2026/ (describing PatSnap's Eureka as covering drafting, claim comparison and global search).

ISSN: 2582-6433

Patent offices have also become developers and users. The EPO has deployed a generative AI tool, the Legal Interactive Platform, trained on the EPC, its Guidelines and the case law of the Boards of Appeal and integrated into its online filing environment; it cites its sources, which users are expected to check.[35] The EPO's stated position is that AI assists and humans decide, a formulation that has an obvious corollary for applicants.

C. The Technical Research

The academic literature on automated patent generation predates the current generation of LLMs. An early line of work fine-tuned the GPT-2 language model on patent text to generate claims and other parts of patent documents.[36] More recent work has moved from demonstrating that models can imitate patent language to measuring whether the output is adequate.

Several findings are consistent across studies. First, general-purpose and domain-specific models differ in unexpected ways: a benchmarking study of six leading LLMs, whose authors describe it as the first to compare patent generation under a unified framework, reports earlier work finding that general-purpose models such as GPT-4 outperformed domain-specific patent models and that generating claims from detailed descriptions produced better results than generating them from abstracts.[37] Second, smaller open models can fail at the basic task. An evaluation of four open-source models (two from the LLaMA-2 family and two from the Mistral family) on biomedical claim generation found that they did not produce correctly styled claims.[38] Third, fine-tuning on patent data improves style while worsening faithfulness. The Pap2Pat benchmark, which pairs research papers with the patents later filed on them to test generation of full patent descriptions, found that fine-tuning produced more patent-like language but also more hallucination.[39] Fourth, existing tools remain short of professional quality. A 2025 study constructed a dataset of eighty human-authored and eighty AI-generated patents, the latter produced by two drafting tools, and applied a multi-dimensional LLM-based quality framework; it concluded that the tools had not achieved human-level quality, with significant gaps in figure-text alignment and in the alignment of claims and specification, particularly for inventions with many interacting components.[40] Fifth, multimodal approaches

[35] Best, *supra* note 4.
[36] Jieh-Sheng Lee & Jieh Hsiang, *Patent Claim Generation by Fine-Tuning OpenAI GPT-2*, 62 World Pat. Info. 101983 (2020).
[37] *PatentWriter: A Benchmarking Study for Patent Drafting with LLMs*, arXiv:2507.22387 (2025).
[38] *Evaluating Application of Large Language Models to Biomedical Patent Claim Generation*, World Pat. Info. (2025), https://www.sciencedirect.com/science/article/abs/pii/S0172219025000067.
[39] *Pap2Pat: Benchmarking Outline-Guided Long-Text Patent Generation with Patent-Paper Pairs*, arXiv:2410.07009 (2024).
[40] *Towards Automated Quality Assurance of Patent Specifications: A Multi-Dimensional LLM Framework*,

ISSN: 2582-6433

that condition generation on drawings improve specification drafting, whereas the evidence suggests that figures add little to claim generation, which depends primarily on precise modelling of inventive scope and legal structure.[41]

Two cautions apply to this literature. The first is that evaluation remains immature. The authors of PatentScore, a framework for evaluating generated claims, start from the premise that conventional text-generation metrics do not capture the legal and structural qualities that make a claim effective, and that evaluation frameworks specific to claims are needed.[42] The second is that most benchmarks use granted patents as the reference standard. Granted patents are a convenient proxy for good drafting, but they include documents that are narrow, over-broad or later invalidated, and a model that imitates the corpus will imitate its flaws. Most of the studies are also preprints that have not yet been subject to peer review.

Not all recent work is directed at measuring error after the fact; some is directed at designing it out. One promising example is a tool for pro se inventors presented at the 2026 International Conference on Artificial Intelligence and Law in Singapore by Ashwin Murthy and co-authors, who are Silicon Valley professionals. As described, the system begins with a deterministic, non-AI initial generation step. A multi-agent stage follows, in which several drafting agents, each given slightly different instructions, write the individual sections. A combining step then merges their drafts on the theory that a detail invented by only one agent is unlikely to survive the merger. Separate reviewer agents then check the merged draft, each seeing only part of the context, such as the original description of the invention. Where several reviewers independently fail to confirm a passage, the issue is surfaced to the user as a high-priority warning.

The novelty of the system lies in combining multi-agent system based drafting with deterministic search, with the deterministic methods run first to narrow the candidate set, which lowers cost and limits the scope for invented detail.

The design is significant for three further reasons. First, it applies the logic of the verification controls proposed in Part VIII, namely grounding in the disclosure and independent review, within the architecture of the tool rather than leaving them to the user. Deterministic stages have a particular attraction in this respect, because their outputs are fixed by their inputs and rules and they cannot invent content in the way a generative model can. Second, agreement across agents is a plausible filter for idiosyncratic inventions but offers no guarantee against

arXiv:2510.25402 (2025).

[41] *PatentVision: A Multimodal Method for Drafting Patent Applications*, arXiv:2510.09762 (2025); *Benchmarking Patent Drafting from Inventor-Style Disclosures*, arXiv:2608.21249 (2026).

[42] *PatentScore: Multi-Dimensional Evaluation of LLM-Generated Patent Claims*, arXiv:2505.19345 (2025).

ISSN: 2582-6433

errors that every agent shares, as may happen when the agents are instances of the same underlying model. Third, the system ends in a warning to a human rather than in an autonomous decision, which accords with the human-in-the-loop principle adopted by the offices (Part VII). Whether the architecture suppresses invented detail as effectively as its design suggests is an empirical question that remains to be answered.

D. What Changes and What Does Not

Generative AI changes the economics of drafting. A first draft that took days can be produced in hours, and the cost of producing additional embodiments, alternative claim sets and description of figures approaches zero. It also changes the drafter's role from author to editor and verifier, and the habits of attention that editing requires differ from those of authorship. The literature on human-automation interaction has long recognised that reviewers of fluent automated output tend to under-scrutinise it, and the risk is greater when, as in patent drafting, the output is long, repetitive and technically specialised.[43]

What does not change is the legal standard to which the finished document is held. An application drafted by a model is judged under the same sufficiency, support, definiteness and candour requirements as one drafted by hand. This asymmetry, in which production is cheap but liability is unchanged, is the central problem of AI-assisted drafting, and it is most acute where the model's output looks authoritative while being wrong.

VI. Hallucination in Patent Drafting

A. The Concept

In the technical literature, hallucination denotes generated content that is nonsensical or unfaithful to the source material.[44] Researchers distinguish *intrinsic* hallucination, which contradicts the source, from *extrinsic* hallucination, which adds content that cannot be verified from the source. The distinction carries over to patent drafting with specific legal meaning. An intrinsic hallucination is a statement in the specification or claims that contradicts or is unsupported by the inventor's disclosure. An extrinsic hallucination is an assertion about the world that has no basis in the disclosure, such as a non-existent prior art reference, an invented experimental result or a misstated rule of law.

[43] Raja Parasuraman & Dietrich H. Manzey, *Complacency and Bias in Human Use of Automation: An Attentional Integration*, 52 Hum. Factors 381 (2010).
[44] Ziwei Ji et al., *Survey of Hallucination in Natural Language Generation*, 55 ACM Computing Surveys, no. 12, art. 248 (2023).

ISSN: 2582-6433

B. Why Patents Are Fertile Ground

Several features of patent drafting make it unusually susceptible. The domain is highly technical and relatively narrow, so there is less training data than in general text. The documents are long, which strains a model's ability to maintain consistency over tens of thousands of words. The genre rewards expansive, formulaic language ("in some embodiments, the processor may be any suitable device"), which models produce readily whether or not it is true of the invention. And the instruction that a drafter gives a model, to broaden coverage and add alternatives, is an instruction to generate content beyond what the inventor said. Commentators have observed that techniques used to reduce hallucination in general-purpose models are less effective on complex technical prompts.[45]

The experience of adjacent fields confirms the concern. In the legal domain, a preregistered study of leading commercial legal research tools found that, although tools using retrieval-augmented generation hallucinated less than a general-purpose model, the products of two major vendors nevertheless produced hallucinated or incorrect answers in a substantial share of queries, between seventeen and thirty-three per cent, contrary to marketing claims that the problem had been eliminated.[46] An earlier study had profiled the frequency of legal hallucinations in general-purpose models.[47] Courts have sanctioned lawyers for filing briefs containing fabricated citations,[48] and the problem has reached intellectual property tribunals: a UK Intellectual Property Office decision in a trade mark matter has been cited as a reminder of the same danger.[49] Patent applications are not immune. A commentator has reported an Australian application that appears to have been produced by an AI system without evident human editing.[50]

C. A Taxonomy of Drafting Hallucinations

The following taxonomy is proposed as a working classification for practitioners and tribunals. It is organised by the nature of the error and the doctrine that the error engages. Table 2 summarises it.

---

[45] *Use of AI in the Patent Industry: The Spectre of Hallucination*, IPKat (Oct. 2025), https://ipkitten.blogspot.com/2025/10/use-of-ai-in-patent-industry-spectre-of.html.
[46] Varun Magesh et al., *Hallucination-Free? Assessing the Reliability of Leading AI Legal Research Tools*, 22 J. Empirical Legal Stud. 216 (2025).
[47] Matthew Dahl et al., *Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models*, 16 J. Legal Analysis 64 (2024).
[48] *Mata v. Avianca, Inc.*, 678 F. Supp. 3d 443 (S.D.N.Y. 2023).
[49] *See Use of AI in the Patent Industry*, *supra* note 45 (citing UKIPO decision O/0559/25 in a trade mark matter as a reminder of the dangers of AI-generated authorities).
[50] *Use of AI in the Patent Industry*, *supra* note 45 (discussing AU2023233168A1).

ISSN: 2582-6433

Table 2. Taxonomy of hallucination in patent drafting and the doctrines engaged

Type	Illustration in a drafting context	Principal doctrinal exposure
1. Fabricated authority or prior art	A non-existent patent number or article is cited in the background or in an information disclosure statement; an office action response relies on a case that does not exist or says something it does not say	Duty of candour and good faith in the United States (37 C.F.R. §§ 1.56, 11.18); disciplinary and sanction exposure; risk of unenforceability if the misstatement is material and made with intent; burden on examiner of unverified IDS volume
2. Unsupported embodiments and invented data	A model supplies parameter ranges, materials or test results that the inventor never reported, or presents a predicted example as if it had been performed	Written description and enablement (35 U.S.C. § 112(a)); sufficiency (EPC art. 83); inconsistency with the working/prophetic distinction in MPEP § 2164.02; possible misrepresentation to the office
3. Technically inaccurate or inoperable description	A circuit, mechanism or algorithm is described in a way that cannot function, though the prose is fluent	Sufficiency and enablement; inoperability and utility challenges; loss of credibility of the whole specification
4. Terminological and structural drift	Different terms for the same element across the document; claim elements without antecedent basis; mismatched reference numerals; claims that no longer match the description after edits	Definiteness (35 U.S.C. § 112(b); *Nautilus*); clarity and support (EPC art. 84); disputes over claim construction (*Phillips*)
5. Scope inflation and boilerplate generalisation	Generic "any suitable" language and wider alternatives than the inventor possessed; statements that "the invention is" a feature that is merely one embodiment	Written description of the full genus (*Ariad*; *Amgen*); lack of support for broad claims; limiting effect of "present invention" statements (*Verizon*); added-matter exposure when later narrowing

ISSN: 2582-6433

Type	Illustration in a drafting context	Principal doctrinal exposure
6. Misstated law or procedure	A prosecution argument misstates the test for obviousness, the effect of a claim amendment or the content of a rule	Professional conduct rules; reasonable inquiry before signing; estoppel arising from arguments made to the office

Source: author's analysis. The categories are not mutually exclusive; a single fabricated example may engage types 2, 3 and 5 simultaneously.

Type 2 deserves emphasis because it is the least visible. Prophetic examples are not objectionable in themselves: the law does not require an applicant to have performed an experiment before filing, and a predicted embodiment grounded in sound science can be fully enabling. The difficulty is that a model has no way to distinguish a prediction grounded in science from a plausible fiction, and its output has the same cadence in either case. A reviewer who sees a table of measurements that "looks right" may not realise that no measurement was taken.

Type 4 has the strongest empirical support at present, since the quality-assurance study referred to in Part V found that tools struggled with figure-text and claim-specification alignment. The value of this finding for practitioners is that type 4 errors are amenable to mechanical checking, which can be automated in a way that types 2 and 3 cannot.

D. Why Detection Is Difficult

Three features make hallucinated content in a specification hard to catch. The first is fluency. A model's text is grammatical, uses the right register and is internally coherent. The second is the absence of an external check: a description of a novel invention has no ground truth other than the inventor's knowledge, which is the very thing the model lacks. The third is institutional. Patent offices examine against the application as filed and the cited art, and typically do not test whether a described embodiment would in fact work; examiners and courts will generally treat a specification as accurate unless a challenge is brought. Errors may therefore lie dormant until a patent is litigated years later, when an accused infringer will probe how the application was drafted and what the inventors actually conceived.[51]

A related concern is the effect of AI-generated disclosures on the prior art. Patent

[51] Kevin R. Davis, *Disclosure to Generative-AI Tools Can Create Patent Prosecution Risk*, Polsinelli (June 18, 2026), https://natlawreview.com/article/disclosure-generative-ai-tools-can-create-patent-prosecution-risk (observing that AI can turn a narrower technical contribution into broader and more polished language that makes the invention appear more complete than it was).

ISSN: 2582-6433

documents are favoured as prior art because they are drafted to meet enablement standards, and are therefore presumed to be enabling. Commentators have questioned whether that presumption should extend to the large volume of AI-generated disclosures, many of which may be speculative or hallucinated, and have proposed requirements such as proof of conception.[52] The question illustrates that hallucination is a problem not only for the drafter but for the system.

E. Confidentiality and Novelty

Hallucination is not the only risk of using generative tools. An invention disclosed to a publicly accessible AI system before filing may become available to the public, creating a risk to novelty under 35 U.S.C. § 102(a)(1) and Article 54(2) EPC, which turns on whether information has been made available to the public rather than on the technology used.[53] A confidential disclosure in a genuinely closed environment is materially different from entering the same information into a consumer product. The USPTO has also cautioned that vendors may retain or train on inputs and that using servers outside the country can raise national-security and foreign filing licence issues.[54]

VII. Regulatory and Judicial Responses

A. Professional Responsibility and Office Guidance

The regulatory response to AI-assisted drafting has so far been to apply existing duties rather than to create new ones. In the United States, the USPTO's 2024 guidance states that there is no general prohibition on using AI tools to prepare submissions and no general obligation to disclose their use, but it reminds practitioners that every paper must be signed by a natural person who thereby certifies that it has been reviewed and is accurate.[55] The guidance identifies the risk that AI tools will hallucinate, and states that verification extends to citations,

[52] Dennis Crouch, Patently-O (Apr. 2024), https://patentlyo.com/patent/2024/04/discerning-navigating-generated.html (discussing AI-generated prior art and the presumption of enablement); *The Patent Disclosure Trap: Will AI Undermine Biopharmaceutical Innovation?*, IPWatchdog (Oct. 20, 2026), https://ipwatchdog.com/2026/10/20/patent-disclosure-trap-will-ai-undermine-biopharmaceutical-innovation/ (raising concern that AI-generated hypothetical examples, some of which may be hallucinations, could be filed before the work needed to show conception).

[53] 35 U.S.C. § 102(a)(1); EPC art. 54(2); Davis, *supra* note 51.

[54] Guidance on Use of Artificial Intelligence-Based Tools in Practice Before the United States Patent and Trademark Office, 89 Fed. Reg. 25609 (Apr. 11, 2024) [hereinafter 2024 Practice Guidance]; *see* 37 C.F.R. § 11.106.

[55] 2024 Practice Guidance, *supra* note 54; 37 C.F.R. § 11.18; *see* Womble Bond Dickinson, *Patent Practitioners Take Note: (Even) if You Use AI, Read Before You File* (Feb. 12, 2024), https://www.womblebonddickinson.com/us/node/13034 (describing the Director's Feb. 6, 2024 memorandum to the Patent Trial and Appeal Board and Trademark Trial and Appeal Board).

ISSN: 2582-6433

technical information, claims of priority and factual assertions. It warns that AI-assisted search can enlarge information disclosure statements and that practitioners must not file one without verifying every reference. It also warns that the duty of disclosure in 37 C.F.R. § 1.56 applies to information material to patentability regardless of how it was found.[56] Where AI tools are used to draft a specification or drawings, the guidance states that practitioners must verify their technical accuracy.

The European Patent Office has taken a similar approach. Its Guidelines for Examination, in the edition that entered into force on 1 April 2026, introduced a dedicated section on the use of artificial intelligence, which provides that parties remain responsible for the content of their applications and submissions whether or not AI assisted in preparing them.[57] The Guidelines also elaborate the requirements for sufficiency of disclosure in the field of AI, consistent with Board of Appeal decisions that address sufficiency of disclosure for AI inventions.[58]

Offices are also regulating their own use. India's Patent Office issued guidelines on 7 August 2026 on the use of AI in examination, which permit cautious assistance for tasks such as prior art searching while reserving substantive judgment and accountability to human examiners and requiring awareness of hallucinations.[59] The common thread across offices is a "human in the loop" principle: AI may assist, but legal responsibility for what is filed or decided remains with a natural person.

B. Inventorship

The second area of response concerns who may be named as inventor. In the DABUS litigation, Dr Stephen Thaler filed applications naming an AI system as sole inventor in many jurisdictions and was unsuccessful in most. The Federal Circuit held that the U.S. Patent Act requires that an inventor be a natural person.[60] The UK Supreme Court reached the same

[56] 2024 Practice Guidance, *supra* note 54; 37 C.F.R. § 1.56(b).

[57] Eur. Pat. Off., *Guidelines for Examination* pt. General, § 5 (Apr. 2026); *see Five Key Changes to the EPO Guidelines*, Wilson Gunn (May 14, 2026), https://www.wilsongunn.com/resource/news/five-key-changes-to-the-epo-guidelines; *2026 EPO Guidelines: How the EPO Is Approaching AI in Practice*, Marks & Clerk (Mar. 19, 2026), https://www.mondaq.com/uk/patent/1761606/2026-epo-guidelines-how-the-epo-is-approaching-ai-in-practice.

[58] Eur. Pat. Off., *Guidelines for Examination* pt. F, ch. III & pt. G, ch. II, § 3.3.1 (Apr. 2026); *see, e.g.*, T 1669/21, T 161/18, T 1191/19 (Eur. Pat. Off. Bd. App.) (identified in the EPO programme for *Search and Examination Matters 2026* as addressing sufficiency of disclosure in AI inventions); *see also* Joel Todd & Felix Hall, *Updated EPO Guidelines on AI*, EIP (May 20, 2024), https://www.eip.com/uk/case-reports/why-us-companies-should-encourage-close-collaboration-between-their-us-and-european-attorneys-on-ai-patent-applications (describing the March 2024 update on disclosure of training data).

[59] *India's Patent Office Sets Guardrails for AI-Assisted Examination*, Mondaq (2026), https://www.mondaq.com/india/patent/1849792/indias-patent-office-sets-guardrails-for-ai-assisted-examination.

[60] *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022), *cert. denied*, 143 S. Ct. 1783 (2023).

ISSN: 2582-6433

conclusion under the Patents Act 1977 and rejected the argument that ownership of the machine conferred a right to the invention, while noting that it was not deciding whether inventions made by a person with the assistance of AI were patentable.[61] The EPO's Legal Board of Appeal likewise held that a designated inventor must be a person with legal capacity, while observing that patentable inventions are not limited to human-made inventions.[62] The Full Federal Court of Australia reversed a first-instance decision that had accepted an AI inventor.[63] Germany's Federal Court of Justice held in 2024 that only natural persons can be inventors but accepted an amended filing in which the human applicant was named as inventor and DABUS was identified as the means by which the invention was arrived at.[64] South Africa, which does not substantively examine patents, granted a patent naming DABUS.[65]

The more consequential question is how much human contribution is needed when AI assists. In February 2024 the USPTO issued guidance that applied the *Pannu* factors for joint inventorship to AI-assisted inventions. On 28 November 2025, it rescinded that guidance in its entirety and replaced it with a statement that the same conception-based standard applies to all inventions, with no separate or modified standard for AI-assisted inventions.[66] Commentators noted that the revision gives little practical guidance for the harder scenarios it leaves open.[67]

For drafting, the inventorship rules produce a practical consequence that is easy to overlook. If a model proposes claim language that goes beyond what the named inventors conceived, the application may recite subject matter that no named inventor conceived, raising questions of correct inventorship. The drafting file should therefore be capable of showing what the human inventors contributed. A drafting practice in which a model generates the inventive content and a human merely approves it invites the objection that the named inventors did not conceive what is claimed.

---

[61] *Thaler v. Comptroller-General of Patents, Designs & Trade Marks* [2023] UKSC 49 (appeal taken from Eng.).
[62] J 8/20 and J 9/20 (Eur. Pat. Off. Legal Bd. App. Dec. 21, 2021); *see* Mayer Brown, *Can AI Be an Inventor? The US, UK, EPO and German Approach* (Jan. 2024), https://www.mayerbrown.com/en/perspectives-events/publications/2024/01/can-ai-be-an-inventor-the-us-uk-epo-and-german-approach.
[63] *Commissioner of Patents v Thaler* [2022] FCAFC 62 (Austl.).
[64] Bundesgerichtshof [BGH] [Federal Court of Justice] June 11, 2024, X ZB 5/22 (Ger.); *see Germany Federal Court Holds AI Cannot Qualify as Inventor*, Stones Law, https://www.stoneslaw.net/germany-federal-court-holds-ai-cannot-qualify-as/.
[65] *AIPPI 2024: Global Debate over AI-Generated Patents*, Asia IP Law (Oct. 20, 2024), https://asiaiplaw.com/sector/patents/aippi-2024-global-debate-over-ai-generated-patents.
[66] Revised Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54636 (Nov. 28, 2025) (rescinding Inventorship Guidance for AI-Assisted Inventions, 89 Fed. Reg. 10043 (Feb. 13, 2024)); *see* Mateo Aboy & Kathleen Liddell, *Revised USPTO Guidance on Inventorship for AI-Assisted Inventions: A Pro-Innovation Pivot Away from Pannu Factors*, 21 J. Intell. Prop. L. & Prac. 272 (2026).
[67] *See* Osha Bergman Watanabe & Burton LLP, *USPTO Issues New Guidance on AI-Assisted Inventorship, Rescinding Its 2024 Guidance* (Dec. 11, 2025), https://www.mondaq.com/unitedstates/patent/1717494/uspto-issues-new-guidance-on-ai-assisted-inventorship-rescinding-its-2024-guidance.

ISSN: 2582-6433

C. Drafting About AI

Finally, the rise of AI as subject matter has itself changed drafting practice. Applications for machine-learning inventions must describe the model architecture, training data and the technical effect in enough detail to satisfy Article 83 EPC and § 112(a), and must frame the claims so as to survive eligibility scrutiny under § 101 and the technical-character test at the EPO. The "black box" description, in which a claim recites that a neural network determines a result without saying how, is increasingly vulnerable.[68] The irony is that the drafter who relies on a model to write such a description may reproduce, with fluent language, exactly the generality that the examiner is trained to reject.

VIII. A Verification Framework for AI-Assisted Drafting

The preceding analysis suggests that responsible use of generative AI in drafting does not require abstention but does require a disciplined process. The following controls are proposed. They are based on the doctrinal analysis above and on the guidance of the offices, and they have not been tested empirically.

First, preserve a record of human conception. The file should show what the inventors disclosed before AI was used and what they approved afterwards. This protects the correctness of inventorship and enables the drafter to answer later questions about who conceived what.

Second, work in a confidential environment. Unpublished inventions should be entered only into systems that are contractually closed, that do not train on customer data and whose data location is acceptable under foreign filing and export rules. Consumer tools should not be used for unpublished inventions.

Third, ground generation in the disclosure. The model should be instructed to use only the information in the inventor's disclosure, to flag gaps rather than fill them, and to mark any content it has added. A traceability matrix that links each limitation of each claim to the passage of the specification and the disclosure that supports it makes unsupported content visible. Multi-agent architectures of the kind discussed in Part V.C can operationalise this control, although they do not replace human review.

Fourth, run mechanical consistency checks. Antecedent basis, term consistency and reference-numeral correspondence can and should be checked by rules-based software and not left to a generative model. These checks address the type 4 errors identified above.

Fifth, verify every authority. Every cited patent, publication, case and rule must be checked

[68] *See* Best, *supra* note 4 (reporting EPO speakers' emphasis on sufficiency of disclosure for AI inventions).

ISSN: 2582-6433

against a primary source before filing. This is a requirement of the signing certification in United States practice and should be treated as a universal minimum.

Sixth, control data and examples. Every example, measurement and parameter range must be confirmed by an inventor, and each example should be identified as actual or prophetic. No number should appear in a specification unless a human can say where it came from.

Seventh, require technical review of each embodiment. A domain expert, preferably an inventor, should read every described embodiment for operability. This addresses the type 3 errors, which no software check will catch.

Eighth, audit scope. The breadth of each independent claim should be compared with what the inventors possessed and what the description supports across the full scope claimed, with particular attention to boilerplate alternatives that the inventors did not contemplate.

Ninth, document the process. A log of tools used, prompts given, outputs accepted and human edits made supports a defensible position if the drafting process is later questioned in litigation or before a disciplinary body.

Tenth, evaluate tools rather than trust them. A firm should test a tool on a gold set of matters with known answers, measure its error rate across the types in Table 2, and treat vendor claims of reliability sceptically, as the legal research study cited above shows that claims of "hallucination-free" operation can be overstated.

IX. Conclusion

Patent drafting has always required the translation of an inventor's contribution into a document that can bear legal weight. The comparative analysis in Parts III and IV shows that the principal systems converge on a common requirement, namely that the specification be clear, complete and sufficient to support the claims, even though they differ in detail. Generative AI has not altered that requirement. What it has altered is the cost of producing text and the character of the errors that text contains. The characteristic failure of an AI-assisted specification is not an obvious mistake but a plausible, fluent and unsupported statement that no one noticed.

The doctrinal tools to address this problem already exist: enablement, written description and sufficiency, added matter, definiteness, candour and inventorship. What is needed is a drafting practice that applies them deliberately, using AI for tasks where its output can be checked against a source, and applying human judgment where it cannot. Offices on both sides of the Atlantic, and increasingly elsewhere, have adopted the principle that the person who files the document is answerable for it. The hypothesis advanced at the outset is therefore supported:

ISSN: 2582-6433

the quality of a patent's language remains determinative of its fate, and the arrival of generative AI increases the premium on rigorous human verification.

Three avenues for further research are suggested. First, empirical work is needed to measure hallucination rates by category in commercial drafting tools, using legally meaningful rather than purely textual metrics. Second, courts and offices will need to consider whether and how AI-generated disclosures should qualify as prior art. Third, the effect of AI assistance on the correctness of inventorship in granted patents deserves systematic study as the first cohort of AI-assisted patents reaches litigation.

Annex A: Illustrative Patent Application

Note. This annex reproduces, in revised form, the example application that appeared as Annexure A to the author's earlier paper (Anand, Basics, supra note 5). Typographical corruption has been corrected, grammar and punctuation standardised, and the dependency of claim 3 corrected (see Annex B). Technical content is unchanged.

Title

VEHICULAR ACCIDENT AVOIDING MECHANISM FOR FOGGY WEATHER

Field of the Invention

The present invention relates to automotive lighting and to external directional and warning light arrangements. In particular, it relates to a system for avoiding multi-vehicle accidents, especially in foggy weather, by automatically switching OFF the parking lights when the turn indicator lights are switched ON, without driver intervention, and, once the turn is complete, switching the turn indicator lights OFF and the parking lights ON.

Description of the Related Art

Driving a four-wheeled vehicle is not new, but driving in extreme fog presents particular difficulties. To avoid accidents in dense fog, drivers commonly use the headlamps, fog lamps and horn, reduce speed and, most importantly, use the parking lights. Because a parking light blinks, it is more likely to be noticed than a constant light source such as a headlamp. Parking lights also serve this purpose well during heavy rainfall.

US Publication No. 20140029277 describes a method of controlling lighting functions in a vehicle. The vehicle includes an operating device for switching between lighting functions, in which activation of one lighting function depends on the actual speed of the vehicle when the driver manually switches from a currently active lighting function to another. Publication No. WO2005042303 describes a module for generating high-intensity daytime running lights on a vehicle using existing warning and turn signal lamps, while retaining the vehicle's internal turn

ISSN: 2582-6433

signal bulb-out detection system. The same vehicular collision-avoidance system may also serve as a fog and inclement-weather lighting system that uses the vehicle's built-in factory directional lamps which, once switched ON, operate continuously.

In the existing arrangements, the parking lamps remain ON at the location of a turn, and there is no way for other road users to distinguish between a vehicle that is using its parking lamps for better visibility in fog and a vehicle that is about to turn. When a driver switches ON the left or right turn indicator while the parking lights are already ON, no new information is conveyed to the vehicle following. As the first vehicle turns with both its parking lights and turn indicator ON, the second vehicle, which does not know the significance of the lights, is likely to strike the first. This is a common cause of accidents in which a rear vehicle hits a front vehicle from behind while the latter is turning on a foggy or rainy day.

To avoid these problems, and to make driving in foggy or rainy weather safer as far as accidents due to parking lights are concerned, the present invention provides a system that automatically switches OFF the parking lights when the turn indicator lights are switched ON, without driver intervention, and that, when the turn is complete, switches the turn indicator lights OFF. Because the parking lights were never switched OFF manually, they return to their ON state.

Objects of the Invention

The principal object of the invention is to provide a system for avoiding multi-vehicle accidents in foggy or rainy weather by automatically switching OFF the parking lights when the turn indicator lights are switched ON, without driver intervention, and switching the parking lights ON again once the turn is complete.

Another object is to provide a system for regulating the parking lights and turn indicator lights that avoids multi-vehicle accidents and makes driving in foggy or rainy weather safer.

A further object is to provide a system for regulating the parking lights and turn indicator lights that can be used with a mechanically controlled vehicle with no electronic control, a digitally controlled vehicle, and a mixed-control vehicle having both mechanical and electronic control.

Summary of the Invention

The invention provides a system for avoiding multi-vehicle accidents in foggy or rainy weather by automatically switching OFF the parking lights when the turn indicator lights are switched ON, without driver intervention, and, once the turn is complete, switching the turn indicator lights OFF. As the parking lights were never switched OFF manually, they return to their ON state.

In one embodiment, the system is implemented in a mechanically controlled vehicle with no electronic control, a digitally controlled vehicle, or a mixed-control vehicle having both

ISSN: 2582-6433

mechanical and electronic control.

In another embodiment, where the vehicle has only an electric switch and no electronic control to regulate the supply of electricity, as in a mechanically controlled vehicle, a mechanical or electro-mechanical switch is placed at the turn switch. When the turn indicator is moved from its normal OFF position, the power supply to the parking lamp switch is cut. When the turn is complete, the turn indicator switch returns to its normal OFF position (an existing feature), which completes the parking light circuit and restores power to the parking lights.

In a further embodiment, in a digitally controlled vehicle in which most functions are controlled by an Electronic Control Unit (ECU), a modification of the existing programming of the ECU automatically switches OFF the parking lights when the turn indicator is not at its normal OFF position.

In yet another embodiment, for a vehicle with mixed control, it is first determined which type of control is used for the parking lights. The ECU can then be reprogrammed, or a mechanical switch can be placed in series with the wiring of the parking lights, as described for the two preceding embodiments.

Brief Description of the Drawings

The appended drawings illustrate only typical embodiments of the invention and are therefore not to be considered limiting of its scope, for the invention may admit of other equally effective embodiments.

Fig. 1 illustrates a switch in series with the parking light circuit, controlled by the turn indicator lever, in the case of mechanically controlled vehicles.

Fig. 2 illustrates a block diagram of the programming logic for the ECU.

Fig. 3 illustrates a modified standard turn lever with a limit switch regulating power to the parking lamp relay.

Description of the Preferred Embodiments

The present invention provides a system for avoiding multi-vehicle accidents and for making driving in foggy or rainy weather safer by automatically switching OFF the parking lights when the turn indicator lights are switched ON, without driver intervention. When the turn is complete, the turn indicator lights are switched OFF and, as the parking lights were never switched OFF manually, they return to their ON state. The invention can be implemented in a mechanically controlled vehicle with no electronic control, a digitally controlled vehicle and a mixed-control vehicle.

Mechanically controlled vehicle with no electronic control. If the vehicle has only an electric switch (Item 3, Fig. 1) to regulate the supply of electricity, a new mechanical or electro-

ISSN: 2582-6433

mechanical switch (Item 2, Fig. 1) is placed at the turn switch. When the turn indicator (Item 4, Fig. 1) is moved from its normal OFF position, the power supply to the parking lamp switch is cut. As a result, the parking lamps go OFF the moment the turn indicator is switched ON. When the turn is complete, the turn indicator switch returns to its normal OFF position (an existing feature), which completes the parking light circuit and restores the power supply to the parking lights.

Digitally controlled vehicle. In a modern vehicle, most functions are controlled through an ECU. All switches and engine functions are controlled through the ECU, which contains a computer to regulate the vehicle's functions. A modification of the existing ECU programming is required so that the parking lights are switched OFF automatically when the turn indicator is not at its normal OFF position. The workflow is shown in Fig. 2. The ECU continuously reads all the switches. When the driver has switched ON both the parking lights and the turn indicator, the ECU disconnects the power supply to the parking lights, while the turn indicator functions as usual. When the turn indicator goes OFF, the ECU restores power to the parking light circuit.

Mixed-control vehicle. In a vehicle with both mechanical and electronic control, it must be determined separately which type of control is used for the parking lights and which for the indicator. Depending on the type of switching used for the parking lights, either the ECU can be reprogrammed or a mechanical switch can be placed in series with the wiring of the parking lights, as described for the two preceding cases.

Limit-switch embodiment (Fig. 3). The invention comprises an extra limit switch having a rating equal to that provided for the parking lamps (typically 5 A, 12 V). In a standard turn lever, a mechanism powers the lamps of one side only; pushing the lever left or right completes the electrical contact of the respective indicator bulbs. In the modified turn indicator mechanism, a limit switch with a compression spring keeps the electrical contact switch in an outward position. The power to the relay of the parking lamps is re-routed through this limit switch. The power fed to the parking lamp relay is connected to the Common (Com) pin of the limit switch, and the power received from the battery or the ECU is fed to the Normally Connected (NC) pin. Power is therefore ordinarily available to the parking lamp relay. When the driver operates the turn indicator, the compression spring is pressed and a ball-ended member moves into the body to press the lever of the limit switch. When the lever is pressed, the connection between Com and NC is opened and Com is connected to the Normally Open (NO) pin. No power from the battery is then fed to the relay that powers the parking lamps. When the turn is complete and the lever returns to its mean position, Com is again connected to NC. The parking lamp circuit is thus broken during the turn.

ISSN: 2582-6433

ECU feedback embodiment (Fig. 2). In digitally controlled vehicles, the invention automatically switches OFF the parking lights when the turn indicator is not at its normal OFF position. As shown in the flow diagram of Fig. 2, if all functions are working, a voltage signal from the standard turn-indicator housing is collected from the output pin of the housing. When the driver triggers the turn indicator, power goes to the respective left or right direction indicator lights, and this power is fed back to the ECU as a feedback signal. Because the power available at the turn indicator is too large for the ECU, it is reduced using a standard electrical resistance to an analogue voltage range of 2 to 3 volts. When the turn indicator is at its home position, no feedback signal is received; when it is triggered, a voltage signal is fed to the ECU as feedback. Irrespective of the direction of operation (left or right), when the ECU receives a voltage pulse on the feedback pin, it immediately stops power to the parking lamp relay, so that the parking lamps are switched OFF. When the turn is complete and the lever returns to its normal position, the feedback voltage vanishes. A zero-volt feedback signal causes the ECU to restore power to the parking lamp relay. At any given time, therefore, only one of the turn indicator and the parking lamps is operative. For a vehicle with mixed control, the type of control used for the parking lamps and for the indicator must be checked separately, and either the ECU can be reprogrammed or a mechanical switch can be placed in series with the wiring of the parking lights, as already described. If the ECU has no port available to receive feedback as an input, the system designed for mechanically controlled vehicles can be used.

Speed indication using parking lamps. The system also provides a speed-limit indication using the parking lamps. An experienced driver will often switch ON the parking lamps when in an extreme hurry (for instance, when using a personal vehicle as an ambulance in an emergency). Other road users need to be able to assess the speed of such a vehicle, since many accidents arise from misjudgement of vehicle speed by others on the road (for example, pedestrians crossing, vehicles at turns and overtaking vehicles). In an ECU-controlled vehicle, the present speed of the vehicle is known, and, on the basis of legal speed limits, a flag is raised by the engine ECU or the speedometer. When the flag goes high because the speed limit has been exceeded, the parking lights are switched ON by the ECU without driver intervention. This attracts the driver's attention, so that a driver who was speeding unnecessarily will try to limit the vehicle's speed. It also signals to others on the road that the vehicle is in some trouble or hurry, so that they will try to keep a safe distance from it. The chances of an accident are therefore reduced.

If any newer type of control for vehicle lighting systems comes into existence, a similar approach may be adopted, so that clarity between the status of the parking lights and the turn

ISSN: 2582-6433

indicator lights can be assured and accidents avoided. Numerous modifications and adaptations of the system will be apparent to those skilled in the art, and it is intended by the appended claims to cover all such modifications and adaptations that fall within the true spirit and scope of the invention.

Claims

1. A system for controlling parking lights and turn indicator lights of a vehicle, particularly in foggy and rainy conditions, the system comprising a control unit having a feedback circuit and relays, and providing speed-limit indication using the parking lamps, the system being configured to switch OFF the parking lights automatically, without driver intervention, when the turn indicator lights are switched ON, and to switch the parking lights ON again when the turn indicator lights are switched OFF.
2. The system of claim 1, wherein the control unit receives the present speed of the vehicle and, based on a legal speed limit, a flag is raised by an engine ECU or a speedometer.
3. The system of claim 2, wherein, when the flag goes high because the speed limit is exceeded, the parking lights are switched ON by the control unit.
4. The system of claim 1, implemented in a mechanically controlled vehicle with no electronic control, a digitally controlled vehicle, or a mixed-control vehicle having both mechanical and electronic control.
5. The system of claim 1, wherein, in a mechanically controlled vehicle with no electronic control, a mechanical or electro-mechanical switch is placed at the turn switch to control the switching of the parking lights and the turn indicator lights.
6. The system of claim 1, wherein, in a digitally controlled vehicle, a modification of the existing programming of the ECU causes the ECU to switch OFF the parking lights automatically when the turn indicator is not at its normal OFF position.
7. The system of claim 1, wherein, for a vehicle having mixed control, the ECU is reprogrammable, or a mechanical switch is placed in series with the wiring of the parking lights.
8. The system of claim 1, wherein, when the driver actuates the turn indicator, power is supplied to the respective left or right direction indicator lights and is fed as a feedback signal to the control unit, the feedback signal being stepped down using a standard electrical resistance.
9. The system of claim 1, wherein, when the turn indicator is at its home position, no feedback signal is received, and, when the turn indicator is actuated, a voltage signal is fed back to the control unit, which, on receiving the voltage pulse at a feedback pin, immediately

ISSN: 2582-6433

stops power to the parking lamp relay, whereby the parking lamps are switched OFF.

10. The system of claim 9, wherein, when the turn is complete and the lever returns to its normal position, the feedback voltage vanishes and a zero-volt feedback signal causes the control unit to restore power to the parking lamp relay.

Abstract

A system automatically switches OFF the parking lights of a vehicle when the turn indicator lights are switched ON, without driver intervention, and, when the turn is complete, switches the turn indicator lights OFF while the parking lights, which were never switched OFF manually, return to their ON state. The system avoids multi-vehicle accidents and makes driving safer in foggy or rainy weather. It can be implemented in a mechanically controlled vehicle with no electronic control, a digitally controlled vehicle and a mixed-control vehicle.

Annex B: Review Checkpoints Applied to Annex A

The following table applies the taxonomy in Part VI to the application in Annex A. The first three rows identify issues that a careful human reviewer should catch in the existing text. The last three illustrate, hypothetically, what an unsupervised AI expansion of the same application might add. The hypothetical additions are not part of the application.

Table 3. Review checkpoints for the example application

Type	Observation	Recommended action
4. Terminological and structural drift (existing)	Claim 3 originally depended from claim 1 but referred to "the flag," which is introduced only in claim 2. Without a corrected dependency, "the flag" lacks antecedent basis.	Make claim 3 depend from claim 2. Run an automated antecedent-basis check on every claim set.
4. Terminological drift (existing)	The description uses "parking lights," "parking lamps," and "parking lamp relay" interchangeably, and refers to "turn indicator," "turning lights" and "direction indicators." Claim 1 originally ended with the turn indicator lights regaining their ON status, which contradicts the description (the parking lights are restored).	Adopt one term per element and define it. Correct the closing limitation of claim 1 to refer to the parking lights.

ISSN: 2582-6433

Type	Observation	Recommended action
Claim structure (existing)	Claim 1 combines two distinct features (parking-light suppression during turns and a speed-limit indication using parking lamps), which narrows it. The speed feature appears again in claims 2 and 3.	Consider restating claim 1 to the first feature alone and moving the speed feature to dependent or separate independent claims, subject to the unity requirement.
2. Unsupported embodiments (hypothetical AI addition)	A model asked to "expand the ECU embodiment" adds that the ECU communicates over a CAN bus with a 10 ms debounce interval and a configurable threshold.	Reject unless the inventor confirms each detail. Unsupported detail cannot later be relied on without added-matter objection, and may mislead.
5. Scope inflation (hypothetical AI addition)	A model adds that the system "may be applied to any vehicle, including aircraft and watercraft, using any suitable light source and sensor."	Delete or have the inventor confirm. A generic extension unsupported by an embodiment invites a written description or support objection and may not be enabled.
1. Fabricated authority (hypothetical AI addition)	A model adds to the Related Art a further prior art citation with a plausible but non-existent publication number.	Verify every reference against a patent database before filing; remove unverifiable references; consider whether any real, material art has been omitted.

Source: author's analysis of Annex A. Rows labelled "hypothetical" illustrate failure modes and do not describe actual tool output.