Inventorship and Disclosure in the Age of Algorithmic Discovery

Ashish Arora, Sharon Belenzon, and Larisa Cioaca

Forthcoming in *Bringing Medicines to Life: How Intellectual Property Enables Innovation in the Life Sciences* (eds. Jonathan M. Barnett and Bowman Heiden, Cambridge University Press, forthcoming 2026).

I. INTRODUCTION

Artificial intelligence (AI) is reshaping pharmaceutical discovery, raising new questions about intellectual property protection[1] and patentability.[2] Companies such as BenevolentAI, Atomwise, Insilico Medicine, and Exscientia have been using advanced machine learning (ML) methods to select promising drug candidates for further scrutiny. An inflection point arrived in 2020, when DeepMind's AlphaFold 2 demonstrated that predicting protein structure with high accuracy at an unprecedented scale was possible.[3] Investors have noticed: between 2019 and 2023, they channeled $10 billion into AI drug discovery and design startups.[4]

These modern AI systems, typically labeled as *deep learning*, differ from the machine learning (ML) tools that drug developers have relied on for years. Classical statistical and ML methods such as regression, random forests, and support vector machines rely on descriptors chosen by researchers, like molecular weight or hydrogen bond counts. Given identical inputs, these models produce identical outputs. In contrast, modern deep learning models for molecular data increasingly rely on graph neural networks, transformers, and hybrid architectures trained on extensive and diverse collections of labeled and unlabeled molecular sequences, graphs, and 3D structures.

1. World Intellectual Property Organization, *WIPO Conversation: Generative AI*, Accessed: 2025-05-06, 2024, https://www.wipo.int/edocs/pubdocs/en/wipo-pub-rn2024-2-en-wipo-conversation-generative-ai.pdf.

2. Gaétan De Rassenfosse, Adam B. Jaffe, and Melissa Wasserman, "AI-Generated Inventions: Implications for the Patent System," Available via HeinOnline, *Southern California Law Review* 96 (2022): 1453–1492; Shlomit Yanisky-Ravid and Xiaoqiong (Jackie) Liu, "When Artificial Intelligence Systems Produce Inventions: The 3A Era and an Alternative Model for Patent Law," Available at SSRN: https://ssrn.com/abstract=2931828 or http://dx.doi.org/10.2139/ssrn.2931828, *Cardozo Law Review* 39 (2018): 2215–2263.

3. John Jumper et al., "Highly Accurate Protein Structure Prediction with AlphaFold," *Nature* 596, no. 7873 (2021): 583–589, https://doi.org/10.1038/s41586-021-03819-2, https://doi.org/10.1038/s41586-021-03819-2.

4. CB Insights, *AI Drug Discovery & Design Startups Have Raised $10B Since 2019*, 2023, https://www.cbinsights.com/research/ai-drug-discovery-design-startups-funding-trends/.

Because the learning process is stochastic and path-dependent, two training runs starting from the same conditions can diverge after only a few passes through the data. Moreover, small changes in input or model configuration can guide the model into previously unexplored regions. Their outputs are often unpredictable *ex ante* and difficult to rationalize *ex post*, challenging traditional notions of interpretability.[5] Some have argued that these systems behave more like human inventors than rule-based robots.[6]

That resemblance collides with two requirements in current patent law: that an inventor be a natural person and the duty to teach how the invention works (i.e., *enablement*). The human inventorship issue was central in the DABUS litigation of 2020-2023. DABUS—the "Device for the Autonomous Bootstrapping of Unified Sentience"—was built by Stephen Thaler to generate inventions without human direction. Thaler insisted that DABUS, not he, was the true inventor and filed patent applications in multiple jurisdictions. Patent offices and courts rejected the filings.[7] Across leading patent regimes, an inventor must still be human.

Inventorship, however, is only part of the problem. Under 35 U.S.C. §112(a), a U.S. patent must "contain a written description of the invention" and must "enable any person skilled in the art to make and use" it without undue experimentation. U.S. courts have already invalidated method patents involving software or AI functionality whose disclosures offered little more than "run our algorithm." In Ibormeith IP, LLC v. Mercedes-Benz USA, LLC, the Court of Appeals for the Federal Circuit struck down a driver-fatigue detector that recited functional goals but disclosed no underlying code.[8] In EON Corp. IP Holdings LLC v. AT&T Mobility LLC, it voided network-management claims for the same reason.[9] The European Patent Office has followed suit, refusing applications such as T 0161/18 (2021) for an AI-based cardiac-output estimator because the disclosure did not enable a skilled person to reproduce the trained model.[10]

5. Zachary C. Lipton, "The Mythos of Model Interpretability: In Machine Learning, the Concept of Interpretability Is Both Important and Slippery," *Queue* 16, no. 3 (2018): 31–57.

6. Will Bedingfield, "The inventor behind a rush of AI copyright suits is trying to show his bot is sentient," Accessed: 2025-10-01, *WIRED*, August 2023, https://www.wired.com/story/the-inventor-behind-a-rush-of-ai-copyright-suits-is-trying-to-show-his-bot-is-sentient/.

7. Thaler v. Vidal, *43 F.4th 1207, 1210 (Fed. Cir. 2022)*, 2022, accessed May 23, 2025, https://www.wipo.int/wipolex/en/text/590863; Thaler v. Comptroller General of Patents, *UKSC/2021/0201 (UK Supreme Court 2023)*, UKSC 49, 2023, accessed May 23, 2025, https://www.supremecourt.uk/cases/uksc-2021-0201.html; European Patent Office Legal Board of Appeal, *J 8/20 and J 9/20 (DABUS)*, Decision of 21 December 2021, 2021, https://www.epo.org/boards-of-appeal/decisions/pdf/j200008eu1.pdf; Library of Congress, *Germany: Federal Court Rules AI Cannot be Recognized as an Inventor*, Published September 9, 2024, 2024, https://www.loc.gov/item/global-legal-monitor/2024-09-09/germany-federal-court-rules-ai-cannot-be-recognized-as-an-inventor/.

8. Ibormeith IP, LLC v. Mercedes-Benz USA, LLC, *732 F.3d 1376 (Fed. Cir. 2013)*, 2013, accessed May 23, 2025, https://www.courtlistener.com/opinion/1086341/ibormeith-ip-llc-v-mercedes-benz-usa/.

9. EON Corp., IP Holdings LLC v. AT&T Mobility LLC, *785 F.3d 616 (Fed. Cir. 2015)*, 2015, accessed May 23, 2025, https://law.justia.com/cases/federal/appellate-courts/cafc/14-1392/14-1392-2015-05-06.html.

10. Technical Board of Appeal European Patent Office, *T 0161/18 Decision of the Technical Board of Appeal – Insufficiency of disclosure – AI-based cardiac-output estimator*, Insufficient disclosure under Art. 83 EPC: training data for neural network not reproducible by skilled person, 2020, accessed June 21, 2025, https://www.epo.org/de/boards-

Using firm-level case studies, we show that companies are increasingly avoiding references to AI-generated content in patent applications.[11] Whenever possible, companies pursue separate method patents covering the use of AI in the discovery process, treating the model as a tool rather than a co-inventor. In addition, our case studies show a clear difference in disclosure strategies. Firms that rely on AI to identify promising solutions usually provide extensive experimental data to prove that the invention can be reproduced.

The remainder of this chapter proceeds as follows. Section II briefly surveys the literature on AI's impact on scientific discovery and invention. Section III unpacks the DABUS decisions from various patent offices. Section IV presents case studies of BenevolentAI, Atomwise, Insilico Medicine, and Exscientia. Section V discusses potential legal and policy reforms, while Section VI concludes.

II. RELATED LITERATURE

At the firm level, AI adoption correlates with stronger innovation performance. An empirical study found that while only ~5.8% of German firms had adopted AI by 2019, those firms accounted for 18.1% of all new-to-the-world product innovations (by sales) in 2018.[12] Theoretical models help explain why. Innovation can be modeled as a search over a combinatorial space where hypothesis quality matters.[13] While scientists traditionally rely on intuition and theory, AI can prioritize better hypotheses. Using AI for hypothesis generation can raise discovery rates, shorten search time, and raise the expected value of outcomes. However, AI's benefits depend on sufficient testing capacity. If experimentation and prototyping lag, even high-quality AI predictions cannot speed innovation. Some scholars caution against overestimating AI's standalone capabilities. Patrick Walters and Mark Murcko, for instance, find that generative AI in medicinal chemistry can produce compounds that closely resemble the training data.[14] This suggests that human creativity remains essential.

of-appeal/decisions/t180161du1.

11. The USPTO issued guidance in 2024 reminding patent lawyers to disclose any "material" part played by AI in creating an invention. (United States Patent and Trademark Office, "Guidance on Use of Artificial Intelligence-Based Tools in Practice Before the United States Patent and Trademark Office," Document No. 2024-07629, *Federal Register* 89, no. 71 [April 2024]: 25609–25614, https://www.federalregister.gov/documents/2024/04/11/2024-07629/guidance-on-use-of-artificial-intelligence-based-tools-in-practice-before-the-united-states-patent)

12. Christian Rammer, Gastón P. Fernández, and Dirk Czarnitzki, "Artificial intelligence and industrial innovation: Evidence from German firm-level data," *Research Policy* 51, no. 7 (2022): 104555, ISSN: 0048-7333, https://doi.org/10.1016/j.respol.2022.104555.

13. Ajay Agrawal, John McHale, and Alexander Oettl, "Artificial Intelligence and Scientific Discovery: A Model of Prioritized Search," *Research Policy* 53, no. 5 (2024): 104989, https://doi.org/10.1016/j.respol.2024.104989.

14. W. Patrick Walters and Mark A. Murcko, "Assessing the impact of generative AI on medicinal chemistry," *Nature Biotechnology* 38, no. 2 (2020): 143–145, https://doi.org/10.1038/s41587-020-0418-2.

A. AI in Scientific Discovery

AI is now part of everyday lab work and supports the entire discovery process, from creating hypotheses to planning experiments and examining data.[15] DeepMind's AlphaFold 2 solved the long-standing problem of predicting a protein's 3D structure from its amino-acid sequence.[16] This breakthrough now guides enzyme engineering and drug design. Generative models explore chemical space to suggest new compounds with therapeutic potential.[17] Recent methods go further, designing proteins that bind tightly to previously "undruggable" targets, such as disordered proteins without fixed structures.[18] In parallel, AI systems trained on large biological datasets are emerging as powerful research tools. For example, OpenCRISPR-1, the first fully AI-generated genome editor, matches the performance of natural Cas9 enzymes while reducing off-target effects and immune responses.[19] Arc Institute and Stanford University have reported the first viable genomes, for bacteriophages, created using generative AI.[20] Together, these advances show how AI is moving beyond prediction to generate functional biological components that can be directly used in research.

Beyond generating ideas, AI also designs better experiments.[21] AI-enhanced microscopes use real-time computer vision to decide which subcellular events to record, increasing useful image yield while reducing storage needs.[22]

When hypothesis generation and experimentation are linked by automated analysis, the result is closed-loop discovery. Early examples like Robot Scientists Adam and Eve generated hypotheses, ran experiments, analyzed results, and iterated with minimal human input.[23] Newer platforms like

15. Sebastian Musslick et al., "Automating the Practice of Science: Opportunities, Challenges, and Implications," *Proceedings of the National Academy of Sciences* 122, no. 5 (2025): e2401238121, https://doi.org/10.1073/pnas.2401238121.

16. Jumper et al., "Highly Accurate Protein Structure Prediction with AlphaFold"; Ewen Callaway, "'It will change everything': DeepMind's AI makes gigantic leap in solving protein structures," Published November 30, 2020, *Nature*, 2020, https://www.nature.com/articles/d41586-020-03348-4.

17. Kyle Swanson, Guang Liu, Daryl B. Catacutan, et al., "Generative AI for designing and validating easily synthesizable and structurally novel antibiotics," *Nature Machine Intelligence* 6 (2024): 338–353, https://doi.org/10.1038/s42256-024-00809-7.

18. James Urquhart, "AI method makes designer binders for 'undruggable' proteins," Accessed: 2025-08-04, *Chemistry World*, July 2025, https://www.chemistryworld.com/news/ai-method-makes-designer-binders-for-undruggable-proteins/4021911.article.

19. Jeffrey A. Ruffolo et al., "Design of highly functional genome editors by modelling CRISPR–Cas sequences," Published: 30 July 2025, *Nature* 625 (2025): XX–XX, https://doi.org/10.1038/s41586-025-09298-z, https://www.nature.com/articles/s41586-025-09298-z.

20. Niko McCarty, "AI Phages," Accessed: 2025-10-02, 2025, https://www.asimov.press/p/ai-phages.

21. Hanchen Wang et al., "Scientific discovery in the age of artificial intelligence," *Nature* 620, no. 7972 (2023): 47–60, https://doi.org/10.1038/s41586-023-06221-2.

22. For example, ZEISS has developed AI-powered microscopy systems that use real-time image analysis to guide which parts of a sample to image.

23. Ross D. King et al., "The Automation of Science," *Science* 324, no. 5923 (2009): 85–89, https://doi.org/10.1126/science.1165620; Kevin Williams et al., "Cheaper faster drug development validated by the repositioning of drugs against neglected tropical diseases," *Journal of the Royal Society Interface* 12, no. 104 (2015): 20141289, https:

AutoRXN in chemistry and A-Lab in photovoltaics go further, autonomously exploring reaction pathways or material compositions.[24]

B. AI in Invention

AI supports several stages of the innovation lifecycle: identifying market opportunities (via predictive analytics), generating design alternatives (through generative design), and optimizing production (via machine learning). In aerospace, generative AI can produce hundreds of aircraft part designs that meet performance criteria; engineers then select and refine the best.[25] This expands the search space and yields non-obvious designs, such as strong yet lightweight lattice structures for the aerospace industry.[26] AI aids the invention of medical devices (by optimizing configurations) and diagnostics (by combining signals in novel ways).[27] In energy and materials, AI helps invent solar cells and batteries by screening millions of material combinations.[28] In software, AI "invents" algorithms, such as using evolutionary methods to generate encryption schemes or neural architectures (as in AutoML).[29]

In drug development, generative models propose chemical structures, some of which have resulted in patented compounds entering clinical trials, as detailed in Section IV. Gaétan De Rassenfosse, Adam Jaffe, and Melissa Wasserma call such systems "invention machines" and predict they will become more common as algorithms, data, and computational resources improve.[30] Increasingly, firms explicitly describe AI contributions in method and process patent filings, where the invention centers on algorithmic steps, and in device or system patents, where AI is embedded

//doi.org/10.1098/rsif.2014.1289; Hector Zenil et al., "The future of fundamental science led by generative closed-loop artificial intelligence," *arXiv preprint arXiv:2307.07522*, 2023,

24. Jan P Unsleber et al., "High-throughput ab initio reaction mechanism exploration in the cloud with automated multi-reference validation," *The Journal of Chemical Physics* 158, no. 8 (2023); Nathan J. Szymanski et al., "An autonomous laboratory for the accelerated synthesis of novel materials," *Nature* 624 (2023): 86–91, https://doi.org/10.1038/s41586-023-06734-w.

25. Raymond Deplazes, *Autodesk and Airbus Demonstrate the Impact of Generative Design on Making and Building*, Autodesk News, November 2019, https://adsknews.autodesk.com/en-gb/news/autodesk-airbus-generative-design-aerospace-factory/.

26. Lawrence Livermore National Laboratory, *LLNL Researchers Unleash Machine Learning in Designing Advanced Lattice Structures*, https://www.llnl.gov/article/51656/llnl-researchers-unleash-machine-learning-designing-advanced-lattice-structures, Accessed June 21, 2025, 2024.

27. World Intellectual Property Organization, *The frontier of MedTech: AI-driven medical devices*, WIPO Global Health News, Accessed via WIPO website, July 2023, https://www.wipo.int/en/web/global-health/w/news/2023/news_0021.

28. Nathan Baker, *Unlocking a New Era for Scientific Discovery with AI: How Microsoft's AI Screened over 32 Million Candidates to Find a Better Battery*, Microsoft Azure Blog, January 2024, https://azure.microsoft.com/en-us/blog/quantum/2024/01/09/unlocking-a-new-era-for-scientific-discovery-with-ai-how-microsofts-ai-screened-over-32-million-candidates-to-find-a-better-battery/.

29. Matt Burgess, "How Google's AI taught itself to create its own encryption," Accessed June 21, 2025, *WIRED*, October 2016, https://www.wired.com/story/google-artificial-intelligence-encryption/.

30. De Rassenfosse, Jaffe, and Wasserman, "AI-Generated Inventions: Implications for the Patent System."

into diagnostic or therapeutic tools.[31] But the rise of AI-driven invention raises legal and policy questions. If AI contributes to an invention, how should inventorship credit be assigned?

III. LEGAL CHALLENGES: INVENTORSHIP AND BLACK-BOX INVENTIONS

AI-driven innovation raises two key legal issues: (1) whether an AI can qualify as an inventor and (2) how to meet disclosure requirements when inventions stem from "black-box" AI processes.

A. The DABUS Cases Across Jurisdictions

The DABUS case has become central to the global debate over AI and inventorship. DABUS, an AI system created by Stephen Thaler, was designed to autonomously generate inventions. In 2018-2019, Thaler filed patent applications in multiple countries for two inventions allegedly conceived by DABUS: a "neural flame" emergency beacon and a fractal-designed food container. Each application listed DABUS as the inventor, with Thaler as the applicant and assignee. His goal was to provoke a legal ruling on whether an AI could be recognized as an inventor. He argued that if a machine generated the invention, naming a human would misrepresent its origin and discourage investment in creative AI. Patent offices and courts reaffirmed that inventorship is *reserved for humans*, though their reasoning varied by jurisdiction. Table 1 summarizes the outcomes.

United States: The U.S. Patent and Trademark Office (USPTO) was among the first to rule on Thaler's DABUS applications. In April 2020, it rejected them because no human inventor was named, as required in the U.S. Patent Act (35 U.S.C. § 100 and § 115). The courts upheld the USPTO decision. In Thaler v. Vidal (Fed. Cir. 2022), the Federal Circuit held that under U.S. patent law, only natural persons can be inventors.[32] The Patent Act defines an inventor as an "individual" (35 U.S.C. §100(f)), and the court emphasized that Congress used personal pronouns like "himself or herself," never "itself." Policy arguments in favor of recognizing AI inventors could not override this clear statutory text. The court also noted that inventorship has previously been denied to corporations and governments, and placed AI in the same category.

European Patent Office (EPO): Although the EPC does not define "inventor," the EPO interprets it to mean a natural person with legal capacity. In January 2020, the EPO rejected Thaler's applications, citing two reasons: (i) an AI cannot satisfy Article 81 because inventors must be natural persons, and (ii) a machine lacks legal personality to transfer rights, creating an ownership

31. From 2013 to 2016, mentions of deep learning and neural networks in patent filings (not restricted to pharmaceuticals) grew at an average annual rate of 175% and 46%, respectively. World Intellectual Property Organization (WIPO), *WIPO Technology Trends 2019 – Artificial Intelligence* (Geneva, Switzerland: World Intellectual Property Organization, 2019)

32. Thaler v. Vidal, *43 F.4th 1207, 1210 (Fed. Cir. 2022)*.

Table 1: DABUS Patent Applications

Jurisdiction	Human Inventorship Required	Guidance/ Legislation Issued	DABUS Case Outcome
United States	Yes	Yes (2024 USPTO Guidance)	Rejected; USPTO and courts ruled only natural persons can be inventors
European Union	Yes	Yes (EPO & EU AI Act)	Rejected; EPO ruled only humans can be inventors
United Kingdom	Yes	No	Rejected; UKIPO and courts upheld human inventorship requirement
Germany	Yes	Yes (Court ruling)	Rejected; Federal Court ruled AI cannot be inventor
South Africa	No	No	Granted; First and only country to accept DABUS as inventor
Australia	Yes	No	Initially accepted, later reversed; Full Court of the Federal Court of Australia overturned earlier decision
Canada	Yes	No	Rejected; CIPO requires human inventors (informal guidance)
China	Yes	No	Rejected; CNIPA requires human inventors
Japan	Yes	No	Rejected; JPO requires human inventors
South Korea	Yes	No	Rejected; KIPO requires human inventors
India	Yes	No	Rejected; IPO requires human inventors
Brazil	Yes	No	Rejected; INPI requires human inventors
New Zealand	Yes	No	Rejected; inventors must be natural persons
Saudi Arabia	Yes	No	Filed; no public decision confirmed
Taiwan	Yes	No	Rejected; requirement for human inventors
Israel	Yes	No	Rejected
Singapore	Yes	No	Rejected

Notes: This table summarizes how different jurisdictions have treated inventions generated by DABUS, as of July 2025.

issue. Thaler appealed, but the Legal Board of Appeal upheld the decision in December 2021.[33] The EPO also rejected Thaler's request to leave the inventor field blank or list "none," insisting that a natural person must be named—even if that person was not the true conceiver. In effect, the EPO suggested that if an AI generated the invention, a human surrogate (e.g., the AI's owner or developer) must still be listed.

United Kingdom: The UK Intellectual Property Office and High Court rejected Thaler's applications, and in 2021, the Court of Appeal upheld those decisions. Thaler appealed to the UK Supreme Court, which unanimously ruled in December 2023 that, under the UK Patents Act 1977, inventors must be natural persons.[34] Writing for the Court, Lord Kitchin explained that inventorship is tied to someone who can hold and transfer rights, which only legal persons can do. The Court clarified that the case was not about whether AI-generated inventions are patentable, but about who can be named as an inventor.

Australia: Australia briefly broke ranks with the rest of the world. In July 2021, Justice Beach of the Federal Court issued a judgment accepting DABUS as an inventor (Thaler v. Commissioner of Patents, [2021] FCA 879). He reasoned that since "inventor" was undefined, it could mean "the entity that invents"—in this case, the AI. However, in April 2022, the Full Federal Court unanimously overturned the decision.[35] The appellate judges concluded that the Patents Act implied inventors must be natural persons, citing provisions about rights and assignments that only apply to humans.

Germany: In Germany, the DABUS applications resulted in a compromise. The German Patent and Trademark Office rejected DABUS as an inventor, and Thaler appealed to the Federal Patent Court (Bundespatentgericht). In late 2021, the court upheld the principle that inventors must be natural persons.[36] However, it allowed a workaround: while a human must be named as inventor, the applicant may note the AI's role in the application. In 2024, the Federal Court of Justice (Bundesgerichtshof) affirmed this solution[37].

South Africa and Other Jurisdictions: South Africa became the first country to grant a patent listing DABUS as the inventor in 2021. The South African Patent Office registers patents without substantive review. Since the DABUS application met formal requirements, it was granted by default. Beyond South Africa, patent offices in jurisdictions such as Canada, China, Japan, South Korea, India, and Brazil have maintained that inventors must be human. Some express this in formal guidelines; others treat it as the default interpretation of existing law. For example, Canada's IP

33. European Patent Office Legal Board of Appeal, *J 8/20 and J 9/20 (DABUS)*.
34. Thaler v. Comptroller General of Patents, *UKSC/2021/0201 (UK Supreme Court 2023)*.
35. Full Federal Court of Australia, *Commissioner of Patents v. Thaler*, [2022] FCAFC 62, decided 13 April 2022, 2022, https://www.wipo.int/wipolex/en/judgments/details/1529.
36. Federal Patent Court of Germany, *Thaler, In re DABUS*, Case 11 W (pat) 5/21, decided 21 December 2021, 2021.
37. Federal Court of Justice of Germany, *Thaler v. German Patent and Trade Mark Office*, Case X ZB 5/22, decided 11 June 2024, 2024, https://www.wipo.int/wipolex/en/text/592754.

Office rejected DABUS and later informally affirmed that only natural persons qualify as inventors.

Global Status Quo: As of July 2025, no country has amended its patent laws to allow AI to be listed as an inventor. Most DABUS rulings reaffirm the traditional rule: inventors must be natural persons. Judges in the UK and Australia noted that AI's role in innovation is a policy matter for lawmakers. For now, the position is clear: AI can assist in the inventive process and generate patentable subject matter, but a human must be named as the inventor.

B. Disclosure and Enablement Challenges

1. Section 112: Disclosure Burden

Patent protection operates as a *quid pro quo*: in exchange for a limited-time monopoly, the applicant must *teach* how to practice the invention. The tradeoff is formalized in the U.S. Patent Act (35 U.S.C. § 112), which sets out three distinct but related disclosure requirements. First, the **written description** requirement ensures that the inventor was in possession of *every* claimed feature at the time of filing. Second, the **enablement** requirement demands that the patent disclose enough information for a person having ordinary skill in the art (PHOSITA) to make *and* use the invention without *undue experimentation*. Third, the **definiteness** requirement mandates that the claims clearly define the invention's boundaries, so others can understand what is covered and what is not.

All three requirements are related to the scope or breadth of claims. In drug discovery, inventors typically seek protection for a class of compounds because once a lead compound or antibody is shown to work, it is relatively straightforward to develop close substitutes.[38]

The amount of information needed to satisfy these requirements depends heavily on the predictability of the field. In more **predictable arts**—such as mechanical engineering or simple electronic circuitry—a single well-described embodiment may be sufficient. A skilled practitioner can often extrapolate from that example to make modest variations without difficulty. By contrast, in **unpredictable arts**—such as chemistry, biotechnology, or pharmaceuticals—small changes in molecular structure can lead to major differences in performance, efficacy, or safety. As a result, the USPTO cautions that "disclosure of a single species usually does not provide an adequate basis to support generic claims."[39]

In short, the less predictable the field, the heavier the disclosure burden. Inventors working

38. Once the underlying target and pathways to disease have been discovered, and an antibody capable of precisely binding to that target has been generated, it may be routine and not necessarily innovative, to manufacture similar antibodies that also precisely bind to that target and treat the same disease. Thus, a patent limited to a single antibody (i.e., an antibody defined by a deposit number, or by its specific protein or DNA sequence) may not prevent the commercialization of highly similar products. Jorge Goldstein, "Solutions to the Problem of Therapeutic Antibody Genus Claims," Summer Issue, *AIPLA Quarterly Journal* 52, no. 3 (2024): 513–

39. United States Patent and Trademark Office, *Manual of Patent Examining Procedure, Section 2164 - Enablement*, https://www.uspto.gov/web/offices/pac/mpep/s2164.html, 9th Edition, Revision 01.2024, 2024.

in complex or less predictable domains must offer more detailed support for their claims. This principle applies regardless of whether the invention was generated by a human or an AI system. However, patents are potentially more vulnerable if they incorporate "black-box" generative AI systems, where the *why* and *how* may be subordinated to *whether* the AI-generated invention works.

2. Opaque Algorithms, Opaque Disclosures

Note that there is no obligation to reveal whether inventors used trial and error, stumbled upon the invention by accident, or employed a sophisticated AI algorithm. For example, if an AI system designs a new drug molecule, the patent application doesn't need to say "we used a neural network to find this." It does need to describe the molecule (e.g., its structure or formula) and how to make or use it (such as synthesis steps, formulations, or dosages for treatment).[40] AI-generated outputs may lack the scientific reasoning or structure–activity rationale that typically supports patent claims. This opacity can create problems.

Concerns arise in three common scenarios. First, when the invention falls within an **unpredictable art**, where small changes can lead to large functional differences. Second, when claims are broad. For instance, when they take the form of a **functional genus**, such as "any molecule that inhibits Enzyme X," yet the patent specification fails to provide a unifying structural principle or a sufficiently representative range of examples. Third, when reproducing the invention (typically covering the process of discovery) would require **rerunning** a proprietary AI model or resorting to **trial-and-error screening**.

Two leading cases illustrate how courts apply § 112 in these contexts. In *Ariad v. Eli Lilly* (Fed. Cir. 2010), the court struck down broad claims covering methods for regulating a protein (NF-κB) because the inventors did not disclose the specific molecules or methods needed to carry them out.[41] Similarly, in *Amgen v. Sanofi* (U.S. 2023), the Supreme Court invalidated Amgen's claim to all antibodies that bind to a protein (PCSK9), despite the patent disclosing 26 examples. The Court emphasized that the patent failed to specify "some general quality [...] running through" the class (*of the claimed antibodies* that gives it "a peculiar fitness for the particular purpose" and reiterated that "the more one claims, the more one must enable."[42]

The USPTO echoed these concerns in its 2024 guidance on AI-assisted inventions. Here, it is important to distinguish between two types of patents: those claiming outputs of generative AI (such as novel molecules) and those claiming the AI tools themselves. In the former, applicants must still provide a specification that enables others to make and use the molecule; AI-assisted

40. United States Patent and Trademark Office, *Manual of Patent Examining Procedure, Section 2164 - Enablement*.

41. Inc. v. Eli Lilly & Co. Ariad Pharmaceuticals, *598 F.3d 1336 (Fed. Cir. 2010)*, 2010, accessed July 16, 2025, https://cafc.uscourts.gov/opinions-orders/08-1248.pdf.

42. Amgen Inc. v. Sanofi, *598 U.S. 594 (2023)*, 2023, accessed July 16, 2025, https://supreme.justia.com/cases/federal/us/598/21-757/.

discovery does not relax that burden. In the latter case, when the invention is an ML model or training method, the application must disclose the inputs, architecture, and training processes, not merely refer to a "neural network" module.[43] While both types raise enablement issues, our focus here is on the first: inventions generated *using* AI, rather than inventions *of* AI.

Consider an illustrative compound claim: "A compound of formula I, R^1–CO–NH–R^2, where R^1 is any C_1–C_{20} alkyl group and R^2 is any heteroaryl with one to three heteroatoms; the compound inhibits Kinase ABC with $IC_{50} \leq 10nM$." This claim covers millions of possible compounds, discloses only a few (which may not be fully representative of the class claimed), and provides no structure-activity relationship. An examiner would likely demand more data—diverse species, binding assays, synthetic routes—or force the applicant to narrow the claim to the disclosed examples.

These risks are not unique to AI, but they may be amplified in AI-generated claims. Traditional drug discovery often starts with a hypothesis that provides an intuitive explanation for why a compound might work. In contrast, generative AI can propose active molecules without offering any insight into why they bind or what features matter. If the AI functions as a black box, and the inventors cannot extract a meaningful structure-activity rationale, the resulting patent may fail both enablement and written description. Without interpretability or guiding principles, disclosing a few unexplained hits is unlikely to support a broad genus.

C. The Human Contribution Standard

The USPTO's 2024 guidance specifically asks whether each human inventor made a significant contribution to the conception of the invention.[44] The EU and UK authorities are implicitly asking a similar question. In response, companies have begun to document the inventive steps that involve humans when using AI.[45]

In practical terms, there are several stages in an AI-driven drug discovery pipeline where human creative input can be injected. These stages are now often highlighted in patent applications and internal records to establish inventorship.[46] Key points include:

43. See, for example, the decision in United States Patent and Trademark Office, *Ex parte Yassin Labyed and Andrzej Milkowski*, https://www.uspto.gov, Appeal No. 2021-002042, Application No. 15/716,444, Patent Trial and Appeal Board, United States Patent and Trademark Office, Alexandria, VA, 2022.

44. "A natural person who merely recognizes and appreciates the output of an AI system as an invention, particularly when the properties and utility of the output are apparent to those of ordinary skill, is not necessarily an inventor. However, a person who takes the output of an AI system and makes a significant contribution to the output to create an invention may be a proper inventor." U.S. Patent and Trademark Office, "Inventorship Guidance for AI-Assisted Inventions," *Federal Register* 89, no. 29 (2024): 10043–10049, https://www.federalregister.gov/documents/2024/02/13/2024-02623/inventorship-guidance-for-ai-assisted-inventions

45. The requirement for meaningful human input in AI-assisted inventions stems from the long-standing "significant contribution" test in Pannu v. Iolab Corp., 155 F.3d 1344 (Fed. Cir. 1998), which holds that inventors must contribute creatively to conception, not merely recognize or execute ideas.

46. Ropes & Gray LLP, *Patentability Risks Posed by AI in Drug Discovery*, https://www.ropesgray.com/en/insights/

Designing and Training the AI System: Humans design how the AI will operate. This includes choosing or developing the model architecture, curating the training data, setting training parameters, and defining the criteria the AI will optimize. For example, scientists might decide to train a generative model on a specialized dataset of compounds or they might craft a custom scoring function that guides the AI toward molecules with the desired drug-like properties. If those choices are non-obvious and crucial to the outcome, they contribute to the conception of the invention. Indeed, companies now frequently emphasize these contributions in patent filings; for example, an application might state that "the inventors designed a novel neural network architecture and trained it on 5,000 kinase inhibitors to generate candidate molecules with property X." In contrast, if a team simply uses an off-the-shelf AI tool with minimal tweaking, it becomes harder to claim any inventive contribution in how the tool was used.

Selecting and Filtering AI Outputs: AI systems often generate many potential solutions. The act of choosing which outputs to pursue is another juncture for human ingenuity. Researchers may filter AI outputs by eliminating molecules with undesirable features (such as reactive groups, likely toxicity, and synthesizability) or prioritize ones aligned with known medicinal chemistry heuristics or disease biology insights. For instance, if an AI lists 10,000 candidate compounds for a given target, the human team might recognize that one of the lower-ranked compounds is pharmaceutically more promising than the top scorers and choose to develop that one. That decision relies on human experience and intuition.

Experimental Validation and Refinement: Once an AI suggests a molecule, scientists synthesize the compound or a set of top candidates, test them in laboratory assays, and then refine or optimize the structure based on the results. This process can involve numerous rounds of chemical modifications to improve potency, reduce toxicity, enhance solubility, and so on. For example, the AI might suggest a molecule that binds to a target protein but has poor "drug-like" qualities. The human chemists might then modify the molecule that retains the AI-suggested core functionality but is more soluble or stable, thus making it a practical drug candidate. That modification is a creative leap by the humans, inspired by but not provided by the AI.

Final Decision-Making ("Go/No-Go"): Even the ultimate decision to file a patent application can be considered as part of the inventive process in an AI context. Simply deciding "let's patent this" is not an inventive act, but the deliberation behind it can involve creative evaluation. If a team evaluates numerous AI-generated ideas and then consciously selects one to patent because they recognize a unique advantage or a novel mechanism of action in it, that recognition is an exercise of human intellect. Such contributions, when combined with more concrete contributions, can help demonstrate that humans were in possession of the invention and contributed to its conception.

alerts/2024/10/patentability-risks-posed-by-ai-in-drug-discovery, Accessed: 2025-07-20, October 2024.

IV. INDUSTRY RESPONSES: ADAPTING R&D AND IP STRATEGIES TO AI INVENTORSHIP RULES

In this section, we examine how several leading AI-driven drug discovery companies have adapted their approaches to align with the "human inventorship" and disclosure requirements.

A. BenevolentAI

BenevolentAI is a UK-based biotechnology company that merges advanced AI with biomedical science to accelerate drug discovery. Founded in 2013 under the name Stratified Medical, the company rebranded as BenevolentAI in 2016, went public on the Euronext Amsterdam in 2022, and went private through a merger with Osaka Holdings in early 2025. BenevolentAI built an AI-powered platform that leverages vast biomedical data to generate new therapeutic hypotheses. Central to this platform is a comprehensive biomedical knowledge graph: a curated network of scientific literature, molecular data, and clinical information. The AI system mines this knowledge graph to identify non-obvious connections, linking diseases with potential drug targets or existing drugs with new disease indications. BenevolentAI's "Retrieve & Explain (R2E)" system provides researchers with not just predictions but also the underlying rationale. The explanation (such as the key literature or data points linking a drug to a disease) can be used to substantiate inventiveness and non-obviousness. If an AI tool suggests a result, the human operators are expected to verify its significance and design follow-up experiments or analyses to bolster the finding.

BenevolentAI includes both domain scientists (chemists, biologists, physicians) and computational scientists (AI/ML specialists) as co-inventors on its patents for drug discoveries. For example, a recent BenevolentAI patent application for novel compounds (WO-2024062250-A1) lists AI expert Mike Rawling as co-inventor. For each project, BenevolentAI scientists maintain detailed research logs (electronic lab notebooks and project reports) recording how an AI suggestion was generated, which human experts reviewed and interpreted the suggestion, and what rational judgment was applied. This human "reasoning layer" (e.g., why a particular molecule was selected from an AI-ranked list, or how a disease linkage proposed by AI was biologically contextualized) is captured in writing as part of the inventive step. Such documentation could serve as evidence in any inventorship dispute that a named human inventor indeed contributed intellectually to the invention. BenevolentAI's leadership has echoed the importance of this approach. Gareth Jones, Vice President of IP at BenevolentAI, noted that while AI might generate ideas, a human inventor is ultimately required to "exert a decisive influence" on an invention's creation. The company's inventorship committees (which review each patent filing) now include AI specialists to explain how they used the AI tools and to confirm their inventive contributions.

BenevolentAI's patent drafting strategy also carefully balances disclosure with protection of

its AI methods. In BenevolentAI's patent applications, the text focuses on the drug or biological insight itself, without explicitly mentioning that AI or machine learning played a role. For example, patent WO-2024062250-A1 merely states that "an imidazole-NH-pyrazine motif has been identified with good Chk1 activity, ... results in increased utility in the treatment of cancers, particularly CNS cancers," without reference to the AI-driven workflow behind the finding. However, a companion paper, published after the patent application was filed, describes the way in which BenevolentAI combined scientific knowledge and machine learning to create, synthesize, test, and select the patented compounds.[47]

A search on Google Patents reveals that BenevolentAI has filed 108 patent applications with the USPTO, of which 45 had been granted as of September 15, 2025. Six of the granted patents cover AI-based discovery methods, while the remainder pertain to therapeutic compounds, none of which mention the use of AI.

B. Atomwise

Atomwise Inc. is a biotechnology company that combines AI with human medicinal chemistry expertise for structure-based drug discovery. Founded in 2012, Atomwise developed AtomNet, the first deep learning convolutional neural network for structure-based small-molecule design. Atomwise's AtomNet model uses protein structural data to evaluate vast numbers of chemical compounds *in silico*, predicting which molecules are likely to bind a given target with high affinity. By 2015, the AtomNet model was demonstrated to identify potential drug leads rapidly, famously discovering two promising drug candidates for the Ebola virus in under a day, a task that would traditionally take years. In a partnership with Hansoh Pharma, Atomwise used AtomNet to virtually screen over 12 billion compounds in just two weeks, yielding ~200 candidates for a challenging cancer target.

After an initial virtual screen yields promising compounds, medicinal chemists at Atomwise and its partners scrutinize them for synthetic feasibility, and other features, notably ADMET (absorption, distribution, metabolism, excretion, and toxicity) properties, as well as novelty. The selected compounds are synthesized and experimentally tested by chemists. If a compound shows activity but needs improvement in potency, selectivity, or pharmacokinetics, the medicinal chemists suggest structural modifications. AI can be invoked again to predict the impact of those changes or to search for alternative scaffolds. For a given lead compound, the system's predictions (for example, highlighting why one analog binds better than another) help the chemist propose tweaks to improve activity. In short, AtomNet provides direction and insight, but human experts drive the decision-making and inventive leaps.

Through its Artificial Intelligence Molecular Screen (AIMS) program, Atomwise offers free

47. Michael J Rawling et al., "Discovery of a Potent, Selective, and Brain-Penetrant Checkpoint Kinase 1 Inhibitor, BEN-28010, for the Treatment of Glioblastoma," *Journal of Medicinal Chemistry* 68, no. 9 (2025): 9101–9125.

virtual screening to academic investigators. Rather than simply providing a list of in silico hits, Atomwise actually ships physical samples of the top 72 predicted molecules to the researchers for testing. This results in co-invention. Of the 10 patent applications filed by Atomwise with the USPTO, four cover aspects of Atomwise's proprietary AI platform, while the remaining six are for new compounds co-invented with university researchers. Only two of the patents on new compounds mention AI.

C. Insilico Medicine

Founded in 2014, Insilico Medicine has built an AI platform for drug design that includes tools for target identification, molecule generation, and clinical trial prediction. A core component is Chemistry42, launched in 2020, which connects state-of-the-art generative models (including generative adversarial networks, variational autoencoders, flow-based models, evolutionary algorithms, and transformer-based language models) with medicinal chemistry expertise.[48] It uses 42 pre-trained generative algorithms to propose novel molecular structures, then evaluates and optimizes them via customizable reward functions and physics-based simulations. Molecules are dynamically scored for drug-like properties, including potency, metabolic stability, synthetic accessibility, ADME profiles, and selectivity.

Chemistry42 integrates human feedback in the AI generation cycle. For example, the platform's reward function can be adjusted based on human-provided criteria, and the generative models will prioritize molecules that experts deem promising. A flagship achievement of Insilico's AI platform is the discovery and development of INS018_055 (also code-named ISM001-055), a novel therapeutic molecule for idiopathic pulmonary fibrosis (IPF). In 2019, Insilico's biology AI proposed new fibrosis targets. Insilico's team identified TRAF2- and NCK-interacting kinase (TNIK) as a promising fibrosis target—one implicated in aging and multiple fibrosis pathways but not previously pursued for IPF.[49] After selecting TNIK, Insilico deployed Chemistry42 to generate molecules capable of inhibiting this kinase. The AI system produced many virtual compounds and evaluated them against desired parameters. After multiple AI-human iterative cycles, 80 candidate molecules were synthesized and tested. One showed potency, and after a round of medicinal chemistry optimization to improve its solubility and safety profile, Insilico nominated INS018_055 as its development candidate for IPF. By early 2022, a full Phase I trial was underway.

48. Insilico was among the first to apply generative adversarial networks (GANs) to drug discovery; the company published a pioneering GAN-based molecule generation method in 2017.

49. It also filed for multiple patents covering the biomarkers for aging, namely US-10665326-B2, "Deep proteome markers of human biological aging and methods of determining a biological aging clock"; US-11260078-B2 "Method of treating senescence with multi-stage longevity therapeutics"; US-202000286625-A1 "Biological data signatures of aging and methods of determining a biological aging clock", US-20190034581-A1 "Deep transcriptomic markers of human biological aging and methods of determining a biological aging clock"; US-20220005552-A1 "Methylation data signatures of aging and methods of determining a methylation aging clock".

The entire cycle from project initiation to start of Phase I took only 30 months and around $2.6 million—vastly less time and cost than traditional estimates (3–6 years and hundreds of millions of dollars). In 2023, it advanced to Phase II. The U.S. FDA granted Orphan Drug Designation to INS018_055 for IPF, reflecting its novelty and potential. Insilico's case demonstrates that AI-generated hypotheses (a novel target, a novel molecule) can succeed in practice, though only with extensive human-led testing at each stage to confirm the AI's predictions.

Insilico Medicine's intellectual property strategy has evolved to protect both its AI technology and the novel drugs coming out of it. First, it secures patents on its AI platform and algorithms (machine learning models and training methods) such as US-11403521-B2 (granted in August 2022), which covers a "mutual information adversarial autoencoder" used in generative molecule design. Second, it secures composition-of-matter patents on the AI-designed drug candidates themselves. An example is US-11530197-B2 (granted in December 2022), titled "Analogs for the treatment of disease," which covers a series of small-molecule inhibitors for fibrotic diseases, including the lead compound corresponding to INS018_055. The patent contains extensive experimental data to support these inventions. For instance, the patented TNIK inhibitors are characterized by biochemical assays (showing TNIK enzyme inhibition), cell-based fibrosis models (demonstrating anti-fibrotic activity), metabolic stability tests, and even efficacy in animal models of lung and skin fibrosis.

Of the 71 patents applied for at the USPTO, 28 cover aspects of in-silico drug discovery and design, and biomarkers used in the process. The remainder relate to molecules for treating diseases. Notably, only 12 of the 43 molecule patents explicitly reference AI or in-silico discovery, even though many inventors named on the methods patents also appear on the remaining compound patents.[50]

D. Exscientia

Exscientia is a UK-based "pharmatech" company founded in 2012 by Prof. Andrew Hopkins as a spin-out from the University of Dundee with the vision of revolutionizing small-molecule drug discovery through artificial intelligence. The firm has grown rapidly: it was listed on NASDAQ in October 2021 at a valuation of about $2.9 billion, and by 2022, it had 450 staff across multiple sites (Oxford, Vienna, Dundee, Boston, Osaka). Exscientia has built robotic labs to fully integrate automation into chemistry and biology workflows.

50. The invention summary in US-20180125865-A1 notes that "*A recently-developed approach to large-scale transcriptomic data analysis, called in silico Pathway Activation Network Decomposition Analysis (iPANDA) has been applied to identify pathway signatures of senescent cells in various tissues and pathway signatures of known senolytic drugs. ... this information was utilized for obtaining a list of prospective protein targets. ... several in silico approaches were applied including:1) drug-target interactions analysis ... 2) structural similarity to known senolytics ... 3) transcriptomic response profile similarity to known senolytics, 4) transcriptomic response profile scoring of various aging tissue datasets and metformin perturbations, 5) deep neural network (DNN)-based model for data integration, and 6) structural analogs search for top candidates.*"

Similar to other companies in this space, Exscientia's platform leverages multiple AI techniques. It uses deep learning models trained on large chemical and biological datasets to predict target activity, ADME/toxicology properties, and other pharmacological profiles of virtual molecules. The system also incorporates an evolutionary or generative algorithm to propose new compounds by applying learned transformations to starting molecules. Exscientia has emphasized using a database of medicinal chemistry transformations, so that the AI's suggestions are biased towards compounds that are structurally reasonable.

Another facet is target identification and validation. Exscientia's Precision Target AI module can analyze large corpora of scientific publications and genomic data to identify novel drug targets or target-disease links. Once a target is selected, the Precision Design module takes over to generate candidate molecules, and Precision Experiment uses robotics to test them. Exscientia has built capabilities for high-throughput automated synthesis and screening, aiming for a fully automated pipeline from idea to experimental result. Exscientia claims drug design cycle times (from project start to candidate nomination) of under 12 months, significantly shorter than the multi-year industry norm.

Exscientia was the first company to announce a small-molecule drug invented using AI to enter clinical trials, DSP-1181. That milestone came in early 2020. Since then, Exscientia's pipeline of AI-generated drug candidates has expanded steadily.

- DSP-1181, a long-acting serotonin 5-HT_1A receptor agonist for obsessive-compulsive disorder (OCD), was developed jointly with Sumitomo Dainippon, and entered a Phase I study in Japan in January 2020. The exploratory design phase was completed in less than 12 months, compared to $\tilde{4}.5$ years on average using conventional methods. The exploratory chemistry phase required synthesizing only $\tilde{3}50$ molecules to arrive at DSP-1181. Only a handful of close analogs were ultimately claimed in the patent (US-10800755-B2, assigned exclusively to Sumitomo Dainippon), suggesting the program homed in on an optimal scaffold efficiently.

- EXS-21546 (A2A receptor antagonist for cancer), Exscientia's first internal drug candidate, targeting the adenosine A2A receptor in immuno-oncology. This molecule, co-invented with Evotec, entered a Phase 1 clinical trial in the UK in December 2020. Its discovery was remarkably efficient: about 163 compounds were made and tested, of which 46 representative examples (28% of total) were disclosed in the patent application as evidence of the structure-activity landscape (US-11786528-B2).[51]

- DSP-0038, the second collaboration molecule with Sumitomo, is a dual-target compound

51. It is also a first-in-class agent, suggesting that AI can find innovative chemical matter. However, this molecule failed in clinical trials. https://www.echemi.com/cms/1896509.html

(a 5-HT_1A agonist and 5-HT_2A antagonist) designed to treat psychosis in Alzheimer's disease. Exscientia's AI had to achieve potent activity on two CNS targets while avoiding off-target receptors like dopamine D2. DSP-0038 entered Phase 1 trials in the U.S. in 2021. According to Exscientia, it took less than 12 months of design to produce this candidate as well.

- EXS-4318 Exscientia's collaboration with Bristol Myers Squibb (BMS, originally Celgene), a selective PKC-θ kinase inhibitor for immunology/inflammation disorders, entered Phase 1 trials in 2023 under BMS's sponsorship. The lead compound was co-invented with Evotec, and the patent was co-assigned to Exscientia and BMS (US-20240262821-A1). This program again underscored design speed: EXS-4318 was identified within 11 months of project start and was the 150th compound synthesized in that series.

Exscientia has pursued an active patenting strategy to protect both its AI technology and the novel compounds arising from that technology. Through 2024, Exscientia (and its IP holding subsidiaries, such as Exscientia AI Ltd.) had 138 patent filings worldwide, with 20 patents granted across jurisdictions. These belonged to 30 distinct patent families.

Exscientia's patent filings also cover two broad categories: (1) novel compounds and (2) AI-driven discovery methods and tools. The OCD drug DSP-1181 is protected by a family of patents, including US-10800755-B2, which claims a small set of specific compounds (three key examples) and their use in treatment. The patent specifications are notably data-rich: the DSP-0038 patent (US-10745401-B2) contains nearly 200 exemplified structures with assay results, but only three specific molecules were claimed in the granted U.S. patent. These extensive datasets demonstrate that the inventors possessed and tested a significant number of compounds, satisfying the written description requirement by showing representative breadth of the claimed chemical space. They also enable others to practice the invention by providing synthetic routes and biological results for many analogs.

Importantly, these patents make no mention of AI tools used in the discovery, though the teams of inventors contain inventors named on AI method patents by Exscientia. By contrast, the AI itself (e.g., the Centaur Chemist software) is not listed.

Of the 19 patent applications filed with the USPTO, 10 claim molecules—only one of which explicitly mentions AI (US-20230086120-A1, "Monolayer of PBMCs or bone-marrow cells and uses thereof")—while the remaining applications cover AI methods. One patent application (US-2024029834-A1) titled "Drug optimisation by active learning" describes an AI-driven iterative method for selecting compounds to synthesize in order to refine a predictive model. Another patent (US-20200013486-A1) covers techniques for generative molecular design using evolutionary algorithms and a database of chemical transformations. Exscientia's filings often include flowcharts

and formulae explaining the model's workings; for instance, the active learning application (US-20230335228-A1) details how to calculate a "subset score" based on feature frequencies and outlines the criteria for selecting compounds to synthesize. This level of detail is crucial for meeting disclosure requirements for algorithmic inventions.

E. Statistical Analysis of Office Actions

Figure 1 tracks 208 patent applications by BenevolentAI, Atomwise, Insilico Medicine, and Exscientia between 2001 and 2025[52]. Overall, 85 applications (41%) mentioned use of AI, compared to 123 (59%) that did not. As previously noted, the firms sought protection for both *discoveries* (molecules, targets, compounds; 134 applications, 64%) and *methods of invention* (algorithms, models, computational approaches; 74 applications, 36%). While discovery applications rarely mentioned use of AI (15 of 134, 11%), method applications almost always did (70 of 74, 95%), indicating that AI was central to the inventive process. There was substantial alignment between applicants and examiners in referencing AI. In only one case did the application omit mention of AI while the examiner referenced it in office action documents. Conversely, there were three cases where the application mentioned AI, but the examiner did not.

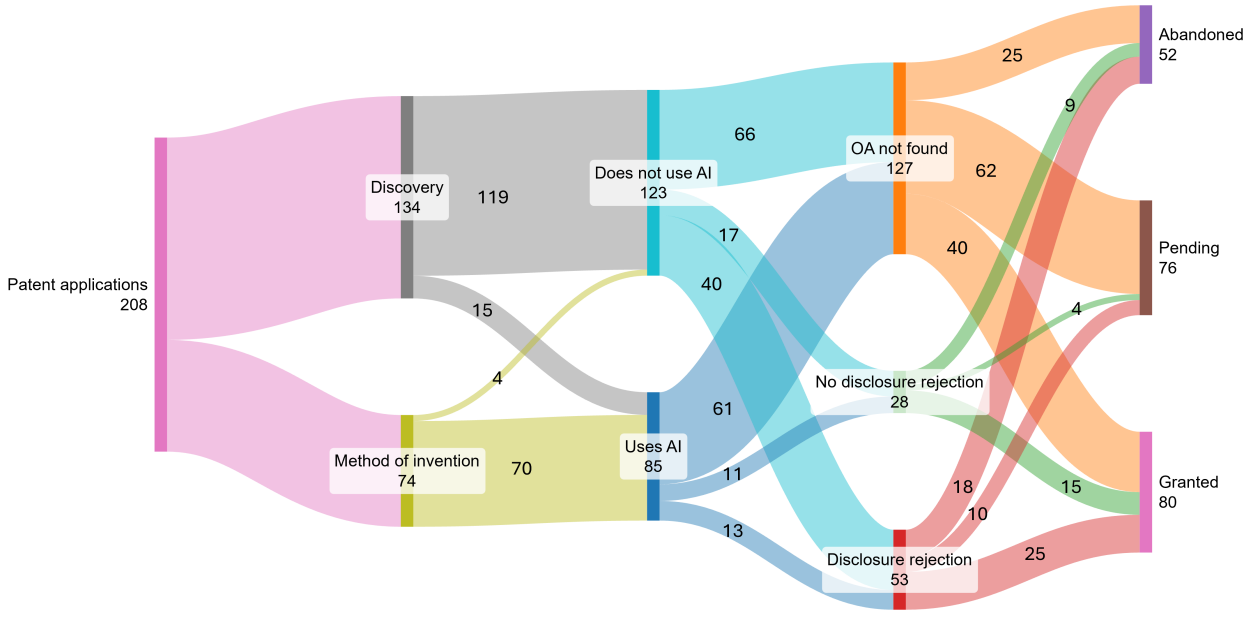

Figure 1: Patent Examination Process and Outcomes

Patent examination process and outcomes for 208 patent applications assigend to BenevolentAI, Atomwise, Insilico Medicine, and Exscientia (2001-2025). Flow widths are proportional to the number of applications at each stage, illustrating the examination actions and outcomes across different types of patent applications.

Of the 208 applications, 127 (61%) had no office action (OA) documents available for download

52. Patent data are current as of September 15, 2025.

from the USPTO.[53] Among the 81 applications (39%) with available OAs, *none* faced inventorship rejections under 35 U.S.C. § 115. This indicates that the four AI-based drug discovery firms systematically ensured a human inventor was listed, effectively avoiding the legal challenge that derailed the DABUS applications.

Among applications with available OAs, 53 (65%) received disclosure and enablement rejections under 35 U.S.C. § 112(a) or 112(b). While examiner feedback on clarity or sufficiency is common for first-filed applications, a two-thirds rejection rate indicates substantive issues with the breadth of claims made. Perhaps counterintuitively, applications referencing AI faced disclosure rejections at a much lower rate (13 of 85, 15%) than non-AI applications (40 of 123, 33%). This could reflect that AI applications involve more novel technical domains where current scientific understanding is limited, leading applicants to make narrower, more cautious claims that are less likely to exceed disclosure requirements.

Disclosure rejections are typically not fatal. Among the 53 applications receiving such rejections, 25 (47%) were eventually granted, 18 (34%) were abandoned, and 10 (19%) remain pending. Applications without disclosure rejections fared slightly better: 54% were granted, 32% abandoned, and 14% pending. Across all 208 applications, 52 (25%) were abandoned, 76 (37%) are still pending, and 80 (38%) were ultimately granted.

Table 2 presents nonparametric analyses of disclosure and enablement rejection rates for examined patent applications by AI usage, comparing 5-year periods before and after the USPTO's DABUS rejection decision on April 22, 2020.[54] It highlights that after DABUS, patents using AI were more likely to raise disclosure and enablement objections. As shown in Table 3, they were also less likely to be granted.

Table 2: Disclosure and Enablement Rejection Rates for Examined Patent Applications

Group	Filed before DABUS (Apr. 22, 2015-Apr.21, 2020)	Filed after DABUS (Apr. 22, 2020-Apr. 21, 2025)	Difference
Does not use AI	0.600 (15/25)	0.333 (7/21)	−0.267
Uses AI	0.267 (8/30)	0.294 (5/17)	+0.027
Difference			+0.294

Notes: This table reports disclosure rejection rates for examined patent applications by AI usage, comparing periods before and after the USPTO's DABUS rejection decision on April 22, 2020.

Table 4 presents difference-in-differences estimates of the effect of the DABUS decision on patent outcomes. All examined applications are included, not just those filed in the 5-years before/after DABUS. Columns 1 and 2 focus on disclosure and enablement rejections. Before DABUS,

53. It is unclear whether those cases had no office actions issued during prosecution or whether the documents were simply unavailable for public download. Notably, 62 of those cases (49%) remain pending.

54. *Examined* patent applications are those that have have received at least one recorded Office Action.

Table 3: Grant Rates for Examined Patent Applications

Group	Filed before DABUS (Apr. 22, 2015-Apr.21, 2020)	Filed after DABUS (Apr. 22, 2020-Apr. 21, 2025)	Difference
Does not use AI	0.560 (14/25)	0.667 (14/21)	+0.107
Uses AI	0.267 (8/30)	0.294 (5/17)	+0.027
Difference			−0.080

Notes: This table reports grant rates for examined patent applications by AI usage, comparing periods before and after the USPTO's DABUS rejection decision on April 22, 2020.

AI-related applications were approximately 30 percentage points less likely to receive disclosure rejections compared to non-AI applications, though this effect is not statistically significant at conventional levels (Column 2). The positive and significant interaction term "Uses AI × After DABUS" (0.329, $p<0.05$) indicates a substantial reversal after the DABUS decision. Specifically, AI-related applications experienced a relative increase of approximately 33 percentage points in disclosure rejection probability post-DABUS. The net effect suggests that AI applications' initial advantage was eliminated and potentially reversed following DABUS.

Columns 3 and 4 examine granted patents. The pre-DABUS effect of AI use is positive but statistically insignificant, indicating no clear advantage or disadvantage. However, the interaction term "Uses AI × After DABUS" is negative and substantial, though not statistically significant at conventional levels. This suggests that AI-related applications became approximately 32 percentage points less likely to be granted after DABUS (Column 4). Additionally, the "After DABUS" main effect is negative and significant (-0.384, $p<0.05$), indicating an overall decline in grant rates of approximately 38 percentage points in the post-DABUS period.

Overall, these patterns highlight two key points: first, the firms successfully navigated the human-inventor requirement that challenged DABUS; second, disclosure and enablement rejections increased post-DABUS for patents on AI (research and discovery tool) and patents using AI for discovery.

V. FUTURE DIRECTIONS: LEGAL AND POLICY REFORMS

Should patent law be revised to address inventions with minimal human input? And if so, how? Several ideas are being discussed in legal and policy circles:

Maintaining Status Quo (Human-Inventor Requirement) with Greater Clarity: The core policy question is unsettled. Do patents exist *primarily* to incentivize inventors, implying that only humans need patents, or to facilitate the disclosure and commercialization of new technologies?[55]

55. The U.S. Supreme Court has long emphasized patents both as a "reward" for inventors (see Graham v. John Deere Co., 383 U.S. 1, 1966) and as a means to promote "progress of science and useful arts" (U.S. Constitution, Art.

Table 4: Difference-in-Differences Estimates of the Effect of DABUS on Patent Disclosure Rejection and Grant Rates

Dependent variable:	(1) Disclosure rejection	(2) Disclosure rejection	(3) Granted	(4) Granted
Uses AI	-0.386*	-0.295	-0.026	0.177
	(0.137)	(0.160)	(0.165)	(0.227)
Uses AI x After DABUS	0.330*	0.329*	-0.315	-0.318
	(0.140)	(0.139)	(0.173)	(0.180)
After DABUS	0.110	0.109	-0.380*	-0.384*
	(0.111)	(0.113)	(0.136)	(0.141)
Discovery		0.103		0.232
		(0.197)		(0.169)
Filing Year FE	Yes	Yes	Yes	Yes
Mean DV	0.36	0.36	0.55	0.55
Observations	144	144	144	144
Adjusted R-squared	0.06	0.05	0.10	0.10

Notes: This table reports difference-in-differences (DiD) regressions estimating the effect of the DABUS decision on patent disclosure rejection and grant rates. "Uses AI" indicates patents exposed to the DABUS rejection decision, "After DABUS" is an indicator for periods after April 22, 2020, and "Uses AI × After DABUS" is the interaction capturing the DiD effect. All models include filing year fixed effects. Standard errors (in parentheses) are clustered by filing year. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

If the latter, then excluding AI-generated inventions may discourage innovation.

Even if the human inventor requirement remains, clarifications may be needed. For example, legislatures or patent offices could codify the "significant human contribution" test so that it is uniformly applied. They could require that, if AI was used, the application must state the nature of human contribution. The USPTO's 2024 guidance is a step in this direction. Under this approach, AI is still treated as a tool, and any evolution in law would focus on helping applicants and examiners deal with AI-assisted inventions.

Allowing Non-Human Inventorship in Exceptional Cases: Another view, led by academics like Ryan Abbott (DABUS project), is that if an AI truly devises an invention by itself, the law should not deny patents on a technicality.[56] The argument here is pragmatic: denying such patents may discourage AI use in discovery. A pharmaceutical company might not invest in an AI that finds drugs if they could not be patented due to the lack of a human inventor. Abbott and others have suggested that patent law could be amended to say, "if an invention is autonomously generated

I, §8).

56. Ryan Abbott, "I Think, Therefore I Invent: Creative Computers and the Future of Patent Law," *Boston College Law Review* 57, no. 4 (2016): 1079–1126, https://bclawreview.bc.edu/articles/566.

by AI, it may still be patentable; the person who owns/operates the AI is deemed to be the owner of the patent." There is precedent in patent law (for example, if an inventor dies, their estate can still file a patent, etc.). No country has adopted this yet, though discussions continue.

New IP Rights for AI-Generated Innovations: Some propose that inventions without a human inventor could get a special kind of IP right with shorter terms of other limits. However, this idea complicates the IP landscape. It shows up in policy debates, especially in Europe, where *sui generis* database rights and other tailored IP exist.

Global Harmonization Efforts: International bodies like the World Intellectual Property Organization (WIPO) are working to harmonize approaches to AI and IP. A patchwork of AI inventorship rules is undesirable. The ongoing WIPO Conversation on AI and IP may eventually lead to recommendations, such as a uniform definition of inventor in the age of AI. If and when a clearly valuable invention (e.g., a cure for a disease) is unpatentable solely because an AI conceived it, pressure might mount to change the law.

VI. CONCLUSION

The rise of AI in drug discovery is challenging traditional patent rules around inventorship and disclosure. Courts around the world have ruled that only humans can be inventors. This means companies using AI must show that people, not machines, are making key inventive contributions. AI-based biotech firms like BenevolentAI, Atomwise, Insilico Medicine, and Exscientia have adjusted their IP strategies to emphasize human oversight and creativity. They treat AI as a powerful tool, not a replacement for inventors.

By presenting AI as a "co-pilot" or "assistant," companies are aligning with legal requirements while continuing to integrate AI into discovery pipelines. Best practices have emerged: document the human-AI interaction, assemble multidisciplinary inventor teams, and draft patents to foreground human contribution. These are practical steps to ensure patentability under current law. Still, patent law is being tested. If AI can invent without human help, should we rethink who gets credit? Some suggest naming the AI's owner; others propose bigger changes. As AI advances, the rules may need to evolve.