

# THE RELIABILITY RESPONSE TO PATENT LAW'S AI CHALLENGES

Arti K. Rai\*

*Pervasive AI use adds newfound importance to longstanding debates over patent timing and reliability. Patent claims on speculative ideas generated by AI, or even the infusion of speculative AI-generated ideas into the public domain, may defeat patent incentives for more careful research. Although challenges that AI use poses for patent validity requirements like human inventorship and nonobviousness have received more attention, reliability is equally important.*

*Indeed, as this Essay argues, the issues are linked. If requirements for inventorship and nonobviousness were adjusted to emphasize reliability, a human role could be preserved, and AI use would not necessarily threaten patents. Currently, as empirical evidence presented in this Essay shows, the fear of imperiling patents may be chilling normatively desirable transparency about such use.*

*The path forward requires embracing reliability throughout patent doctrine. In addition to changes to inventorship and nonobviousness doctrine, robust adoption of reliability requires fortification of the utility requirement for securing a patent and a parallel tightening of requirements for the types of information that can be used to thwart patent grants. Longer term, if cost barriers to innovation across fields fall dramatically, certain non-patent exclusivities may need to play the dominant incentive role. But for the time being AI can provide a powerful catalyst for bolstering a level of reliability the patent system should arguably have had all along.*

INTRODUCTION .....	2
I. SCIENTIFIC RELIABILITY AND THE PATENT SYSTEM .....	6
A. <i>The Normative Debate</i> .....	6
B. <i>Relevant Doctrine: Utility, Novelty, Inventorship, and Enablement</i> .....	8
C. <i>How AI Use Shines a Spotlight on Reliability</i> .....	13

---

\* Elvin R. Latty Distinguished Professor of Law, Duke Law School; Founding Faculty Director, Center for Innovation Policy; Faculty, Duke-Margolis Institute for Health Policy. Thanks to Ryan Abbott, Stuart Benjamin, Colleen Chien, Janet Freilich, Lidiya Mishchenko, Nicholson Price, Tejas Narechania, and participants in the University of California, Berkeley Law & Technology Colloquium for helpful comments. I am also grateful to Mike Hostetler of Wilson Sonsini for spurring this line of inquiry and to Duane Valz of Insitro, Daniel Reker of Duke University's Pratt School of Engineering and Eric Wagner of Duke University's Office for Translation and Commercialization for useful insights. Jill Baggett, Gina Campanelli, Brendan Genaw, Cliff Gilman, Ken Hong, and Annmarie Messing provided excellent research assistance. All errors are my own.

II. THE ILLUSTRATIVE CASE OF DDD.....	15
A. <i>Reliability and Patents in Traditional DDD</i> .....	16
B. <i>The Impact of AI Use</i> .....	18
C. <i>The Potential for AI Use to Improve Reliability</i> .....	21
D. <i>Failure to Disclose AI Use</i> .....	22
III. PATENT LAW AND POLICY IMPROVEMENTS.....	25
A. <i>Utility in Patents and Unpatented Prior Art</i> .....	25
B. <i>Protection Against Inventorship and Obviousness Challenges</i> .....	27
1. <i>Inventorship</i> .....	27
2. <i>Obviousness</i> .....	31
CONCLUSION .....	33
APPENDIX A.....	34

## INTRODUCTION

Artificial intelligence (AI)<sup>1</sup> is everywhere – applications span the spectrum from the public-facing large language models that dominate popular discourse to the specialized models of protein folding that won its developers a Nobel Prize in 2024.<sup>2</sup> In the legal realm, AI use affects virtually all doctrinal areas, perhaps none more so than intellectual property. Indeed, if AI use were to produce an across-the-board, dramatic reduction in the cost of innovation or creation relative to that of copying, pervasive use would undermine the foundational incentive rationale for intellectual property.

This Essay addresses the patent system. At least in the near term, AI use will not cut the ratio of innovation to copying cost in all contexts. Accordingly, legal incentives to innovate, including patent law, will still have a role to play.<sup>3</sup>

---

<sup>1</sup> In using the term AI, I focus on machine learning systems, developed using training data and learning algorithms, that output predictions, text, images, videos, designs or other technical or non-technical output. The complexity, and interpretability to humans, of the machine learning models developed from the training data can vary substantially. Systems that operate entirely through *ex ante* specification of rules by humans fall outside my focus. Additionally, I focus on the use of AI systems in producing innovative *output*, not on the innovation in AI systems themselves.

<sup>2</sup> See, e.g., *Nobel Prize in Chemistry 2024*, NATURE, (Oct. 9, 2024), <https://www.nature.com/collections/edjcfjihdi> (discussing award of 2024 Nobel Prize in Chemistry to Demis Hassabis and John Jumper for protein structure prediction). David Baker also shared the 2024 Chemistry prize for computational protein design.

<sup>3</sup> Additionally, at least in the near term, these legal incentives will target humans. This article does not address questions surrounding what it would mean for AI to be “incentivized” or even for AI to be considered a person, at least in the manner in which corporations are considered legal persons. See generally Elizabeth Pollman, *Reconceiving Corporate Personhood*, 2011 UTAH L. REV. 1629 (2011) (critiquing the legal and historical developments that led to corporations being considered persons within

Scholars have discussed the ways in which AI use could call into question compliance with doctrinal requirements in patent law, including human inventorship<sup>4</sup> and the requirement that the patented invention be “non-obvious” to the average scientist or technology in the area.<sup>5</sup> The U.S. Patent and Trademark Office (PTO) is actively confronting these and other specific doctrinal questions.<sup>6</sup> This Essay makes a broader argument, namely that AI use not only provides an opportunity to reassess longstanding policy debates that cut across doctrinal issues but also that the insights gleaned from such reassessment actually solve the patent puzzles AI raises.

At its core, patent policy must address the timing of when rights are allocated as well as how broad these rights should be. For decades, commentators have debated these fundamental questions. Some have argued that if patent rights are awarded too early or sweep too broadly, patent rewards will go to speculators who race to the patent office and stake broad claims rather than careful researchers who are more likely to be in the best position to bring the invention to market.<sup>7</sup> On this view, such misallocation will ultimately thwart innovation rather than promote it. A somewhat related commentary notes that even if inchoate ideas are not specifically covered by a patent right, public availability will represent patent-defeating prior invention (so-called “prior art”) that thwarts careful work by thwarting the ability to patent such work.<sup>8</sup>

Pitted against these concerns have been arguments that placing greater burdens on patent applicants would defeat an important

---

the meaning of the Constitution). For an illuminating discussion of challenges to traditional visions of personhood posed by advances in machine intelligence, see JAMES BOYLE, *THE LINE: AI AND THE FUTURE OF PERSONHOOD* (The MIT Press 2024).

<sup>4</sup> See, e.g., Ryan Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. REV. 1079 (2016); David L. Schwartz & Max Rogers, “Inventorless” Inventions? *The Constitutional Conundrum of AI-Produced Inventions*, 35 HARV. J.L. & TECH. 531 (2022).

<sup>5</sup> See, e.g., Ryan Abbott, *Everything is Obvious*, 66 UCLA L. REV. 2, 6–8 (2019) (discussing the impact the widespread use of inventive machines, such as computers equipped with artificial intelligence, will have on the current obviousness standard for patentability). In the context of patents that cover AI systems themselves (as contrasted with patents that cover outputs from AI use), scholars have also discussed questions of adequate patent disclosure. See, e.g., Mateo Aboy et al., *The Sufficiency of Disclosure of Medical Artificial Intelligence Patents*, 42 NATURE BIOTECHNOL. 839 (2024); W. Nicholson Price, *Big Data, Patents, and the Future of Medicine*, 37 CARDOZO L. REV. 1401 (2017). That set of disclosure issues is distinct from (though relevant to) disclosure issues arising from patents that cover outputs from AI use.

<sup>6</sup> See, e.g., Request for Comments Regarding the Impact of Proliferation of Artificial Intelligence on Prior Art, the Knowledge of a Person Having Ordinary Skill in the Art, and Determinations of Patentability Made in View of the Foregoing, 89 Fed. Reg. 34217 (Apr. 30, 2024); Inventorship Guidance for AI-Assisted Inventions, 89 Fed. Reg. 10043 (Feb. 13, 2024) [hereinafter PTO Inventorship Guidance]. While the analysis in this article focuses on patents and other innovation incentives, both the USPTO and the Copyright Office are also tracking AI use closely in the context of copyright law.

<sup>7</sup> See *infra* note 20 and accompanying text.

<sup>8</sup> See *infra* note 21 and accompanying text.

economic role for patents as early-stage innovation “prospects.”<sup>9</sup> On this view, patents are analogous to rights in tangible property: granting broad patents early in the R&D process provides incentives for efficient further innovation on the path to market, either by the patent owner itself or through licensing. And even for skeptics of prospect theory, the “startup/small firm” rationale for allowing early (though not necessarily broad) patenting – that for such firms inchoate patents can be important to attract the venture capital or large firm partnership necessary to conduct further experimentation – has empirical support.<sup>10</sup>

In recent years, U.S. patent law has evolved away from very broad rights that far exceed what the inventor has actually accomplished. However, it still allows early patenting of inchoate ideas. Perhaps most notably, although the so-called utility doctrine of patent law does invoke the concept of scientific reliability – even using a term, credibility, that should get to the heart of the issue<sup>11</sup> -- the doctrine generally fails to require that patent applicants demonstrate more than a highly cursory level of credibility for their claims about their invention.<sup>12</sup> Because this lax standard for credibility gives patent law’s use of the term a strained meaning, I will therefore use the term *reliability* to address the relevant normative goal.

This Essay argues that pervasive AI use makes the longstanding debate over reliability, both of patents and of prior art, an urgent issue. As a general purpose technology, AI use has the capacity to rapidly and cheaply generate inchoate ideas across all industries. Absent vigilance on the question, counterproductive racing to the patent office is likely to increase. Even if inchoate ideas are simply put into the public domain, the lax standards for what constitutes prior art may defeat the possibility of patent incentives for careful researchers. The net result will be diminished innovation.

On the positive side, proper use of AI systems could bolster the scientific reliability of patents. Indeed, because high-quality AI use will generally require not only human input but selectivity with respect to such input, attention to the quality of AI use could provide a defense against the challenges to human inventorship and non-obviousness that AI use itself raises. Although some tweaking of the inventorship and non-obviousness doctrines will be necessary, dramatic change of the sort that would require legislation should not be needed.

---

<sup>9</sup> Edmund Kitch first used the term “prospect” in this context. See Edmund Kitch, *The Nature and Function of the Patent System*, 20 J.L. & ECON. 265, 269-70 (1977). Scholars have developed extensions of, or variations on, aspects of prospect theory. See *infra* note 14. Although these scholars tend to draw a distinction between invention of the “prospect” and subsequent commercialization, these steps all exist on an innovation spectrum. Thus, the relevant inquiry is one of timing – that is, when on the spectrum the patent right should be allowed.

<sup>10</sup> See *infra* notes 15 - 16 and accompanying text.

<sup>11</sup> See *infra* note 32 and accompanying text.

<sup>12</sup> See *infra* notes 34 - 40 and accompanying text.

In order for patent adjudicators to determine whether the AI use in question is high-quality, however, they must (at a minimum) know about such use. Unfortunately, current doctrine provides few incentives for applicants seeking patents on AI-derived inventions to report such use. To the contrary, patent applicants may be concerned that disclosure could jeopardize inventorship and non-obviousness. To illustrate the problem empirically, I use as a case study the highly patent-sensitive context of drug discovery and development (DDD). As I demonstrate using an original dataset of drug patents likely derived through the use of AI, these patents say virtually nothing about how AI was used.

As a doctrinal matter, the path forward will require weaving reliability throughout patent law. This procedure should include fortifying the utility requirement and, in parallel fashion, fortifying requirements for the types of prior art that can thwart patent grants. Those who reported high-quality AI use would not only be well positioned to satisfy these fortified requirements, they could also be provided safe harbors against inventorship and non-obviousness challenges. Although this fortification would necessarily operate as an industry-sensitive standard rather than a bright-line rule, it could (as the Essay shows) be relatively administrable.

If and when we reach a stage at which reliable science does not require human input, but still needs substantial additional investment to reach the market, a major reason will likely be pre-market safety requirements enunciated in risk regulation schemes. In that case, commercialization incentives that directly track safety requirements may prove more useful than patents.

The reliability criterion highlighted by AI use also provides insight into claim scope, including the scope issue that recently caught the attention of the Supreme Court in the recent *Amgen v. Sanofi* case.<sup>13</sup> AI use shows the wisdom of rights that closely track reliability. Failure to emphasize reliability could allow speculators that raced to the patent office claiming all types of products or processes to preempt more careful researchers.

The Essay unfolds in three Parts. Part I reviews the longstanding patent policy debate over scientific reliability, discusses the limited emphasis patent doctrine places on reliability, and argues that pervasive AI use heightens the importance of reliability. Part II turns to empirical evidence I have gathered through a study of AI use in the highly patent-sensitive context of drug discovery and development (DDD). It discusses the DDD innovation process, AI use in the DDD process, and how AI use – and disclosure thereof – can be used to bolster reliability. Part III outlines patent law improvements that would promote reliability and simultaneously address concerns regarding human inventorship and non-obviousness. It also briefly addresses options for

---

<sup>13</sup> See, e.g., *Amgen v. Sanofi*, 598 U.S. 594, 614–16 (2023) (holding a patent to be invalid due to its failure to meet the so-called enablement standard for broad scope).

addressing incentives if and when the central barriers to market become regulatory regimes for addressing safety risks.

## I. SCIENTIFIC RELIABILITY AND THE PATENT SYSTEM

### A. *The Normative Debate*

The argument that patents should contain scientifically reliable information might seem unremarkable, even banal. For multiple reasons, however, the argument is not only contestable but has often been rejected. To begin, patents containing information that is on its face suspicious could be seen as harmless, as patents are a commercial instrument, and the most frequent consumers of patent information – investors and licensees – are sophisticated. On this view, suspicious patents will simply be ignored by such consumers.

Alternatively, to the extent that a patent is not on its face dubious, and even sophisticated parties may therefore find it challenging to evaluate, the reason is likely to involve timing. Specifically, the research in question may be promising but too early-stage to be reliable. In that case, a patent may be useful to promote subsequent investigation and development, either by the patent owner or through patent licensing.<sup>14</sup> Delaying patent availability until full reliability was established could hamper the requisite investment by those who feared the appropriation of such investment by competitors.

While supporters of development-based approaches to patenting have deployed largely theoretical arguments, other scholars have used empirical evidence to argue that early-stage patents on inchoate ideas are important for younger firms and startups that seek to attract financing.<sup>15</sup> Although these scholars have not specifically focused on

---

<sup>14</sup> A prominent proponent of broad patents on nascent invention is Edmund W. Kitch. See Kitch, *supra* note 9. Kitch has also argued (arguably in tension with concerns about underdevelopment) that without early patent grants we would have inefficient and duplicative “R&D races” to secure later-stage patents. See *id.* at 269–70. This aspect of Kitch’s theory has been squarely challenged by prominent real-world races to secure patents on early-stage information. See, e.g., Arti Rai, *Addressing the Patent Gold Rush: The Role of Deference to PTO Patent Denials*, 2 WASH. U. J.L. & POL’Y 199, 204, 212 (2000) (discussing claims to genetic information and to internet-based software). Indeed, some scholars have defended a modified version of prospect theory by noting that even if early rights allocation simply moves the timing of racing further up, it has the salutary effect of channeling inevitable racing in the direction of earlier patent grants and hence earlier patent expiration. John F. Duffy, *Rethinking the Prospect Theory of Patents*, 71 U. CHI. L. REV. 439, 443 (2004). However, early patent expiration could itself be a problem if the result was a failure to commercialize. Cf. Ted Sichelman, *Commercializing Patents*, 62 STAN. L. REV. 341, 363 (2010) (noting that many patents are never commercialized).

<sup>15</sup> See, e.g., Joan Farre-Mensa et al., *What is a Patent Worth: Evidence from the U.S. Patent “Lottery”*, 75 JOURNAL OF FINANCE 639, 642, 647–65 (2020) (concluding, based on their empirical study, that patents were a causal factor in greater startup access to venture

doctrinal issues associated with patent timing, an extension of the literature would counsel in favor of allowing early patenting. With such early patenting, cash-strapped young firms could attract the capital needed to develop their ideas further. The argument for early patenting is bolstered by suggestions in these studies that the first patent secured by the young firm is the most important for purposes of attracting capital.<sup>16</sup>

A contrasting view is offered by scholars who argue that the absence of a strong reliability requirement means “too many applications, too many patents . . . and increased assertion of patent rights.”<sup>17</sup> Rather than fostering development, speculative patents impose a tax on development by third parties.<sup>18</sup> Moreover, in keeping with transaction cost economics, the transaction costs of licensing speculative patents may be high.<sup>19</sup> Additionally, even speculative patents can represent prior invention that impedes subsequent patenting by those who do the work of developing inventions that are more fully vetted.<sup>20</sup> Absent the ability to secure patents, these entities may be reluctant to do such work.

Prior art can be created not only by prior patents but also by unpatented information discussed in publications, including patent publications. Various scholars have argued that in this context, the absence of a reliability requirement for prior art bars patents that should be granted, particularly in highly patent-sensitive areas like biopharmaceuticals and chemicals.<sup>21</sup>

---

capital). The “attracting capital” argument assumes of course that a relatively robust supply of young firms and startups are present, as this is important for a healthy innovation ecosystem.

<sup>16</sup> *Id.* at 658–64.

<sup>17</sup> Christopher A. Cotropia, *The Folly of Early Filing in Patent Law*, 61 U.C. L.J. 65, 71 (2009).

<sup>18</sup> *See id.* at 115–16 (describing the tax on innovation generally associated with patent trolls and subsequent licensing fees). Certain discussions of the economic costs created by poor patent “quality” illustrate similar themes. Many discussions of quality tend to focus, however, not on reforms to the early filing allowed by the current system but on the failure of granted patents to satisfy existing criteria of novelty and non-obviousness.

<sup>19</sup> Such transaction costs may be particularly high if multiple overlapping inputs must be licensed, with the potential for holdup by one of the rights owners. *See generally* Carl Shapiro, *Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting*, in 1 INNOVATION POLICY AND THE ECONOMY 119, 124–26 (2001) (explaining the holdup problem resulting from the need to obtain licenses). But they exist even in contexts with fewer patented inputs.

<sup>20</sup> *See, e.g.*, Mark A. Lemley, *Ready for Patenting*, 96 B.U. L. REV. 1171, 1195 (2016) (noting that the current system “encourag[es] ideas at the expense of those who take the time to develop and test their invention”).

<sup>21</sup> *See, e.g.*, Ben Roin, *Unpatentable Drugs and the Standards of Patentability*, 87 TEX. L. REV. 503, 569 (2009) (describing how drugs rendered unpatentable under the novelty and non-obviousness requirements of patent law are unattractive for commercialization and therefore unlikely to reach the public); Sean B. Seymore, *Rethinking Novelty in Patent Law*, 60 DUKE L.J. 919, 963–69, 976 (2011) (proposing a new framework for addressing the novelty requirement given complexities for patenting across unpredictable fields like biology and chemistry).

The normative literature thus makes clear that lax approaches to reliability may lead to errors of both overinclusion and underinclusion. Even so, in the absence of more definitive proof that reliability should be valued, patent doctrine has favored a lax approach. The only substantial doctrinal lever that has been used to bolster reliability, the enablement component of patent claim scope, does so only to some extent. The next section discusses current doctrinal levers that address reliability.

### *B. Patent Validity Doctrine*

In order to understand how current validity doctrine intersects with reliability, it is important first to understand some basic principles of how the patent process works. A patent application, which must be evaluated by the PTO for compliance with patent validity requirements before it can be granted, typically describes the prior art as well as the invention in a section known as the specification. The specification concludes with claims that define the “metes and bounds” of precisely what the applicant argues is its right. Even if a patent is granted by the PTO, its validity can be challenged post-grant, either in an Article III court or at an administrative tribunal within the PTO known as the Patent Trial and Appeals Board (PTAB).

In assessing compliance with the validity requirements of the patent statute, the PTO and courts determine whether a given patent claim, construed in light of the larger specification,<sup>22</sup> represents patent-eligible subject matter,<sup>23</sup> has utility,<sup>24</sup> and is novel<sup>25</sup> and nonobvious.<sup>26</sup> Additionally, under Section 112(a) of the patent statute, the specification must disclose information sufficient to “enable” the claim – that is, to show how to “make and use” the scope of the inventive territory claimed.<sup>27</sup> While enablement has historically been the most important component of adequate disclosure, the Court of Appeals for the Federal Circuit (which hears all appeals in patent cases) has, in recent decades, determined that Section 112(a) also contains a separate written description requirement. As construed by the Federal Circuit in cases like the court’s 2010 *en banc* decision in *Ariad v. Eli Lilly*,<sup>28</sup> written description tends to focus on disclosure of invention through description of structure, whether physical structure (in the case of most

---

<sup>22</sup> See *Phillips v. AWH Corp.*, 415 F.3d 1303, 1317 (Fed. Cir. 2005) (explaining how the PTO and courts assess the scope of claims).

<sup>23</sup> 35 U.S.C. § 101.

<sup>24</sup> *Id.* § 101.

<sup>25</sup> *Id.* § 102.

<sup>26</sup> *Id.* § 103.

<sup>27</sup> *Id.* § 112(a).

<sup>28</sup> *Ariad Pharm., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336 (Fed. Cir. 2010) (*en banc*). As a textual matter, this determination is based on a reference in 112(a) to a “written description” of the invention.

inventions) or sometimes algorithmic structure (in the case of software inventions).<sup>29</sup> A major goal is promoting notice.<sup>30</sup> Utility, novelty, non-obviousness, enablement, and written description are all evaluated from the perspective of the ordinary scientist or technologist working in the field – the so-called person having ordinary skill in the art (POSITA or PHOSITA). Proper inventorship is also a fundamental requirement of patent validity.<sup>31</sup>

Of these various requirements, the ones most closely linked to the current failure to value reliability are utility, novelty, nonobviousness, and inventorship. In contrast, courts have recently begun to construe the enablement requirement to place greater emphasis on reliability. But enablement cannot do the job alone. I address each requirement in turn.

Patent law's utility doctrine most directly regulates how much work an applicant has to do before filing a successful application. As noted in the Introduction, utility nominally requires "credibility."<sup>32</sup> Additionally, it requires demonstration of "specific and substantial" usefulness.<sup>33</sup> In principle, these terms could mean that the invention in question had been tested relatively rigorously or was the product of relatively rigorous experimental or analytical discovery processes. As commentators have discussed, however,<sup>34</sup> the PTO and the courts not only have interpreted utility as a low-threshold bar but also have said that the patent examiner bears the burden of showing that a given application lacks utility.<sup>35</sup> In patent parlance, credibility simply means

---

<sup>29</sup> See Kevin E. Collins, *An Initial Comment on Ariad: Written Description and the Baseline of Patent Protection for After-Arising Technology*, 2010 PATENTLY-O PATENT L.J. 60, 62–65 (discussing the structural emphasis in the court's decision and noting that, going forward, written description could play a role with respect to all claims similar to the patent statute's requirement for structure with respect to claims that use "means plus function" language under 112(f)).

<sup>30</sup> See Arti K. Rai, *Improving (Software) Patent Quality Through the Administrative Process*, 51 HOUS. L. REV. 503, 518–33 (2013) (discussing the role of written description as well as Section 112 requirements like definiteness in improving notice to competitors).

<sup>31</sup> 35 U.S.C. §§ 116, 117, 118 (discussing various aspects of inventorship).

<sup>32</sup> See U.S. PATENT & TRADEMARK OFFICE, MANUAL OF PATENT EXAMINING PROCEDURE § 2107 (9<sup>th</sup> ed. Rev. 1, Jan. 2024) [hereinafter MPEP], available at <http://www.uspto.gov/web/offices/pac/mpep/index.html>.

<sup>33</sup> See *id.*; *In re Brana*, 51 F.3d 1560, 1568 (Fed. Cir. 1995) (holding that although usefulness for pharmaceuticals requires some research, clinical testing is not necessary to prove usefulness). The statutory foundation for utility is most directly found in 35 U.S.C. § 101, which requires that a patent cover a "useful" process or product. The PTO and the courts have also linked utility to Section 112(a) of the patent statute, which requires that the patent's disclosure "enable" the claimed invention – that is, that the disclosure show the person having ordinary skill in the art how "to make and use" the claimed invention. *Cf.* Jacob S. Skerkow, *Patent Law's Reproducibility Paradox*, 66 DUKE L.J. 845, 906 (2017) (arguing that a utility rejection should be limited to the situation where *no* species in a genus is useful whereas an enablement rejection could apply more broadly).

<sup>34</sup> See, e.g., Sean Seymore, *The Teaching Function of Patents*, 85 NOTRE DAME L. REV. 621, 639 (2010) (discussing the low-threshold utility bar and burden of proof allow).

<sup>35</sup> In fact the PTO's official manual of patent examination and procedure itself notes that rejections for lack of utility are "rare." MPEP § 2107.01.

that the invention not operate in a manner that is on its face fantastical or otherwise inconsistent with known scientific principles (e.g., a perpetual motion machine).<sup>36</sup> Although specific and substantial utility poses a slightly higher hurdle, the requirement is still quite lax: fulfilling the requirement simply requires that the invention have a use that is not a “throwaway” use (e.g. using a computer as a doorstop). Similarly, courts have held that patent applicants do not have to include any working examples of the claimed invention.<sup>37</sup>

On occasion, particularly in the context of the so-called “unpredictable” sciences of biology and chemistry, the PTO or the courts have required a somewhat higher level of rigor. In 1999, for example, the PTO issued guidance indicating that speculative patent applications on gene fragments that emerged from large-scale human genome sequencing technology would be rejected for lack of specific and substantial utility.<sup>38</sup> This guidance responded to concerns that gene fragment patents could create insuperable licensing cost hurdles for follow-on innovators.<sup>39</sup> In 2005, the Federal Circuit affirmed this guidance.<sup>40</sup> The gene fragment case study serves as an instructive precedent on which this Essay draws in Section III. This one-time guidance for gene fragments aside, however, the utility standard is low even in the unpredictable arts.<sup>41</sup>

Another key patent law doctrine that gives short shrift to reliability is novelty. Here the doctrine’s failure to value reliability works against the patent applicant. Specifically, the law surrounding what unpatented prior publications<sup>42</sup> constitute prior art that defeats novelty places even

---

<sup>36</sup> *Newman v. Quigg*, 877 F.2d 1575, 1582 (Fed. Cir. 1989) (holding a patent application claiming to cover an “Energy Generation System Having Higher Energy Output Than Input” to be unpatentable).

<sup>37</sup> Instead, patent applicants can (and often do) use so-called prophetic examples – examples of experiments that are merely hypothesized. See Janet Freilich & Lisa Larrimore Ouellette, *Science Fiction: Fictitious Experiments in Patents* 364 SCIENCE 1036 (2019) (discussing the challenges arising from the use of prophetic examples).

<sup>38</sup> For a detailed discussion of the history of this USPTO action, see Arti K. Rai, *Patent Validity Across the Executive Branch: Ex Ante Foundations for Policy Development*, 61 DUKE L.J. 1237, 1251–56 (2012).

<sup>39</sup> A prominent scholarly expression of concern was articulated in Michael Heller & Rebecca Eisenberg, *Can Patents Deter Innovation? The Anticommons in Biomedical Research*, 280 SCIENCE 698 (1998).

<sup>40</sup> See *In re Fisher*, 421 F.3d 1365 (Fed. Cir. 2005) (holding that the claims surrounding expressed sequence tags (ESTs) lacked a specific and substantial utility and failed to satisfy § 101).

<sup>41</sup> Part II discusses in detail case law and literature on experiments found in drug patents.

<sup>42</sup> In addition to ordinary publications, unpatented prior publications include disclosures in patent specifications that are not specifically claimed by the applicant as well as publications of patent applications that are abandoned after applicants fail to secure a grant. Lidiya Mishchenko, *Thank You for Not Publishing (Unexamined Patent Applications)*, 47 BYU L. REV. 1563, 1566–67 (2022). As Mishchenko notes, there is some evidence that examiners like to cite abandoned applications as prior art against subsequent applicants. *Id.* at 1586 (citing Christopher Cotropia & David Schwartz, *The*

less weight on technical trustworthiness than does the utility doctrine. While patents must show *some* use, however nominal, novelty-defeating prior art need not demonstrate *any* use.<sup>43</sup> Rather, it simply must be capable of being made – itself a fact that patent law presumes.<sup>44</sup> For example, in the chemical arts, a reference to a chemical X that discloses a structure for X as part of a laundry list of structures, can invalidate a subsequent patent application or patent on X. The reference can invalidate novelty even if the chemical was never made and has no known practical use.<sup>45</sup>

In contrast with novelty, the nonobviousness doctrine does allude to reliability, particularly in its axiom that an invention that is “obvious to try” does not become obvious unless the average scientific or technologist in the field would have some reasonable expectation of success. However, because success is defined in terms of whether the invention can be made and whether it has at least one utility, nonobviousness ultimately fails to stress reliability.

Patent law’s inventorship doctrine similarly skirts questions of scientific reliability. The disconnect between patent doctrine and reliability looms largest when it comes to case law that focuses on an inventor’s “conception.” Under the traditional formulation of conception, which is keyed to an idea in the mind of the inventor,<sup>46</sup> no actual creation of the invention is required.<sup>47</sup> The low reliability threshold persists even for versions of the inventorship doctrine prevalent in unpredictable technology such as biotechnology, which hold that conception is complete only at the point that it has been created

---

*Hidden Value of Abandoned Applications to the Patent System*, 61 B.C. L. REV. 2809, 2812 (2020)).

<sup>43</sup> *In re Hafner*, 410 F.2d 1403, 1426 (C.C.P.A. 1969).

<sup>44</sup> MPEP § 2122 (citing various Federal Circuit cases for the proposition that utility need not be disclosed in a prior art reference); MPEP § 2121 (stating that prior art references are presumed to show the researcher how to “make and use” the reference). I focus in this Essay on written prior art, as that is most relevant to the new issues raised by AI use.

<sup>45</sup> For a detailed discussion of the relevant case law, see Sean Seymore, *Rethinking Novelty in Patent Law*, 60 DUKE LAW JOURNAL 919, 946–53 (2011). For discussion of some normative consequences, see Benjamin N. Roin, *Unpatentable Drugs and the Standards of Patentability*, 87 TEX. L. REV. 503, 564–69 (2009) (arguing that promising compounds are routinely not pursued due to concerns about patentability and advocating a long FDA-administered exclusivity as an alternative). *Cf.* W. Nicholson Price, *The Cost of Novelty*, 102 COLUM. L. REV. 769, 783 (2020) (arguing that one cost of the current novelty standard is a bias against innovation that deepens knowledge of a particular area).

<sup>46</sup> *Townsend v. Smith*, 36 F.2d 292, 295 (C.C.P.A. 1929). For a recent invocation of this traditional formulation, see *Burroughs-Wellcome Co. v. Barr Lab’y, Inc.*, 40 F.3d 1223, 1227 (Fed. Cir. 1994).

<sup>47</sup> The traditional view of conception is consistent with the aspect of the so-called operability aspect of the utility, which simply requires that the patent not make claims affirmatively contrary to the laws of physics. *See, e.g.*, *In re Swartz*, 232 F.3d 862, 864 (Fed. Cir. 2000) (affirming the USPTO’s determination that a patent application purporting to cover cold fusion was not operable).

– so-called “reduction to practice.”<sup>48</sup> The fact that an invention can be created doesn’t mean that it’s likely to be scientifically or technologically promising.

One recent doctrinal trend that does have some potential for addressing reliability is the Federal Circuit’s often strict evaluation of enablement. Notably, in the 2023 case of *Amgen v. Sanofi*,<sup>49</sup> the Supreme Court unanimously agreed with the Federal Circuit’s approach. In so doing, it rejected prominent arguments, advanced by commentators who discussed the Federal Circuit case,<sup>50</sup> that this approach necessarily doomed all “genus” claims – that is, claims to a group of products or processes.<sup>51</sup>

The case involved functional claims to the genus of all antibodies that bind to one or more of 16 different sites on the PCSK9 protein. Through such binding, the antibodies may prevent PCSK9 from binding to LDL receptors. For individuals with high cholesterol, blocking PCSK9’s binding to LDL receptors is valuable, as the binding triggers a process that degrades such receptors and increases LDL cholesterol levels. Although Amgen had identified the amino acid sequence structure of 26 antibodies that performed the function of blocking binding, the scientific evidence indicated that millions more antibodies could fall within the genus. Moreover, although Amgen purported to have approaches for identifying these additional antibodies, the Court agreed with arguments that these approaches amounted to little more than “research

---

<sup>48</sup> See Dan L. Burk, *AI Patents and the Self-Assembling Machine*, 105 MINN. L. REV. HEADNOTES 301, 308 (2021) (making this point about biotechnology and comparing biotechnology systems to AI systems).

<sup>49</sup> *Amgen v. Sanofi*, 598 U.S. 594 (2023).

<sup>50</sup> See, e.g., Dmitry Karshtedt et al., *The Death of the Genus Claim*, 35 HARV. J.L. & TECH. 1, 13-71, 63 (2021).

<sup>51</sup> The Supreme Court’s approach – like that of the Federal Circuit before it – does not in fact doom genus claims. Instead, it allows initial narrow claims to the results of experimental forays that use trial-and-error. As technological improvement allows identification of more fundamental scientific principles that move beyond trial-and-error, patent applicants that file later in time will be able to claim more broadly. The more complex question of whether (and if so, when) patent owners should be able to take advantage of technology developed *after* the time of filing lies beyond the scope of this Essay. The issue has, however, generated extensive scholarly discussion over many decades. For a sampling of the commentary, see, e.g. Kevin Emerson Collins, *Enabling After-Arising Technology*, 34 J. CORP. L. 1083 (2009) (assessing the intersection of enablement with after-arising technology); Christopher A. Cotropia, “*After-Arising Technologies and Tailoring Patent Scope*,” 61 N.Y.U. ANN. SURV. AM. L. 151 (2005) (examining the doctrine of equivalents and its ability to give exclusivity over after-arising technology and after-arising equivalents); Timothy R. Holbrook, *Patent Disclosures and Time*, 69 VAND. L. REV. 1459 (2016) (analyzing the relationship between patent disclosure and technological improvement over time); Joshua D. Sarnoff, *Correcting Misunderstandings of Literal Infringement Scope Regarding After-Arising Technologies Protected by the Doctrine of Equivalents*, 53 AKRON L. REV. 767 (2019) (arguing that, in respect to literal infringement, claim scope does not expand over time to cover after-arising technologies).

assignments” that would require “painstaking experimentation” by the average antibody scientist.<sup>52</sup>

Cases like *Sanofi* foster reliability in the sense that they force inventors that purport to claim a broad genus functionally to do the scientific work necessary to identify which species actually achieve that function. In fact, even written description, which generally focuses on notice, can promote reliability by limiting functional claims to species the inventor has actually done sufficient work to identify. However, neither enablement nor written description necessarily addresses the reliability of relatively narrow claims already worded in structural terms. Assuming the structures can be made, and the patent applicant hypothesizes some use in the specification for one or more of the structures, the claim will not face validity problems based on inadequate disclosure.

In sum, then, current patent doctrine tends to place limited emphasis on reliability. Of course, as discussed in Section I.A, the normative debate surrounding how to balance reliability with delayed patent availability remains unsettled. In the next Section, I address how the advent of AI use makes reliability a more important criterion.

### C. *How AI Use Shines a Spotlight on Reliability*

AI systems have the ability to draw upon disparate literature and generate large numbers of ideas virtually instantaneously. With respect to the patent system, commentators have begun to worry about an expansion of speculative ideas, whether in patent claims themselves or in unpatented disclosures.<sup>53</sup> Moreover, given the advent of AI-assisted patent drafting, speculative ideas may arise not only in the context of AI-assisted science but also in the patent drafting process.<sup>54</sup> Indeed, certain firms could develop business models in which they merged science and drafting to both file thousands of patent applications and put large numbers of ideas into the public domain, so as to preclude patenting by others.<sup>55</sup> Picking up on these concerns, a recent PTO Request for Comments asks whether AI-generated disclosures that have not have been reviewed by a human “should be afforded the same presumption that they are operative and enabled” as other prior art.<sup>56</sup>

---

<sup>52</sup> *Amgen*, 598 U.S. at 614.

<sup>53</sup> See John Villanesor, *Ten Thousand AI Systems Typing on Keyboards: Generative AI in Patent Applications and Preemptive Prior Art*, 26 VAND. J. ENVT. & TECH L. 375, 377-378 (2024) (discussing challenges created by AI-generated patents and by unpatented disclosure).

<sup>54</sup> Lisa L. Ouellette, Victoria Fang, and Nicholas T. Ouellette, *How Will AI Affect Patent Disclosure*, 26 NATURE BIOTECHNOLOGY 26 (2025) (focusing on AI-assisted patent drafting).

<sup>55</sup> Villanesor, *supra* note 53, at 391-396 (2024) (discussing viability of potential AI-based business models that either preempt patents or generate patents).

<sup>56</sup> Request for Comments Regarding the Impact of the Proliferation of Artificial Intelligence on Prior Art, the Knowledge of a Person Having Ordinary Skill in the Art,

The PTO has also released guidance informing those who practice before the agency that they must review all document filings that have been drafted with AI assistance for factual errors and misstatements of the law.<sup>57</sup>

Widespread use of poorly trained AI systems could also cause data flaws to spread more quickly than in prior eras. For example, if training data sets are based on flawed experimental results, and AI model developers do not address these flaws, output based on this training data will also be inaccurate. Training processes may also be problematic. AI researchers have been at the forefront of calling attention to such flaws, including failure to fully separate the data on which a model has been trained from the test data used for validation of the model.<sup>58</sup>

Concerns about flooding either patent offices or the public domain with unreliable ideas are certainly meritorious. Indeed, because the sheer quantity of ideas that AI use can generate has a quality all its own, AI creates an inflection point in the normative debate over reliability. The appropriate response, however, is not to create a special category of AI-generated prior art, or AI-generated patented invention, that is treated differently from human-developed technology. To the contrary, pervasive AI use should be treated as an opportunity to raise all prior art, and all patent applications, to a threshold level of reliability. Moreover, because (as discussed further in Part III), raising the floor of reliability will in many if not most cases require human intervention, an emphasis on reliability can help address the human inventorship and obviousness dilemmas raised by AI use.

Notably, AI systems that are properly deployed can themselves bolster reliability. AI systems may prove particularly fruitful in areas of science and technology, like drug discovery and development (DDD), where we have limited knowledge of fundamental principles. In these areas, AI techniques can be used to analyze large amounts of data on real world failures and success that is fed back into the R&D system to bolster subsequent earlier stage work.

Part II, to which the Essay now turns, uses as a case study the patent-sensitive area of drug discovery and development (DDD). It also discusses mechanisms by which AI use could promote virtuous feedback loops with respect to reliability. However, bolstering reliability through appropriate AI use requires at least some reporting

---

and Determinations of Patentability in View of the Foregoing, 89 Fed. Reg. 24217, 24219 (2024).

<sup>57</sup> Department of Commerce, USPTO, Guidance on Use of Artificial-Intelligence Tools to Practice Before the United States Patent and Trademark Office, 89 Fed. Reg. 25609, 25614-25615 (2024).

<sup>58</sup> Sayash Kapoor and Arvind Narayanan, *Leaking and the Reproducibility Crisis in Machine-Learning-Based Science*, 4 PATTERNS 100804 (2023) (discussing 17 fields, including many life science fields, with a focus on the challenge of appropriate data collection, sampling, and pre-processing).

of such use. Unfortunately, as Part II also discusses, DDD patents that likely involved AI use contain virtually no discussion of its use.<sup>59</sup>

## II. THE ILLUSTRATIVE CASE OF DDD

For several reasons, DDD presents a particularly fruitful case study of the intersection of AI use with patent reliability. First, although AI use may exacerbate certain problems in the DDD process, the process also contains ample room for improvement through AI use. Second, as the scholarly literature has shown,<sup>60</sup> DDD innovation is highly sensitive to patents. Not only does getting patent law incentives right matter a great deal in this context, but suggestions for improvement should translate into other areas of innovation where patent incentives matter.

Part II.A discusses reliability challenges in the traditional DDD process and how those challenges relate to patent doctrine governing drug molecule patents. Part II.B discusses the impact of AI use. Part II.C addresses how proper AI use could bolster reliability. Finally, Part II.D discusses the challenges to assessing reliability posed by the dearth of reporting of AI use in the patent record.

---

<sup>59</sup> Reporting is important even for output inventions that can be tested independent of discovery process. See Janet Freilich, *The Replicability Crisis in Patent Law*, 95 IND. L.J. 431, 448-451 (2020) (extensive discussion of scientific literature establishing tight link between methodological quality and reproducibility). Notably, a recent report from the National Academies on computational reproducibility states: “[R]esearchers should convey clear, specific, and complete information about any computational methods and data products that support their published results in order to enable other researchers to repeat the analysis . . . That information should include the data, study methods, and computational environment.” NATIONAL ACADEMIES OF SCIENCES, ENGINEERING, AND MEDICINE. REPRODUCIBILITY AND REPLICABILITY IN SCIENCE 7 (2019). NASEM adopts a definition of reproducibility that requires consistent results by independent investigators that use the same experimental input materials, methods, data, and code as the original investigators. It defines replicability separately, as “obtaining consistent results across studies aimed at answering the same scientific question, each of which has obtained its own data.” Because I do not engage formal epistemology, I use the terms synonymously. See, e.g., Leonard P. Freedman, Iain M. Cockburn, and Timothy B. Simcoe, *The Economics of Reproducibility in Preclinical Research*, 13(6) PLOS BIO e10020165, <https://doi.org/10.1371/journal.pbio.10020165> (using the terms reproducibility and replicability synonymously). See also Janet Freilich, *The Replicability Crisis in Patent Law*, 95 IND. L.J. 431, 448-451 (2020) (extensive discussion of literature establish tight link between methodological quality and reproducibility).

<sup>60</sup> See, e.g., Eric Budish, Benjamin N. Roin, and Heidi Williams, *Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials*, 105(7) AMERICAN ECONOMIC REVIEW 2044, 2044-2048 (2015) (demonstrating, using a theoretical model and original data, that because of the fixed 20-year patent term, firms conducting cancer trials focus on late-stage cancer treatments that reduce patent term less than cancer prevention or early-stage treatment); Ben Roin, *Unpatentable Drugs and the Standards of Patentability*, 87 TEXAS LAW REVIEW 503, 545-549 (2009) (collecting extensive qualitative evidence, including interview evidence, that firms screen promising molecules for patentability and drop molecules that cannot be protected by a strong patent).

### A. *Reliability and Patents in Traditional DDD*

In the biopharmaceutical arena, the expected strength and length of the core patent on the therapeutic molecule (the so-called “composition of matter” or product patent) forms the basis for go/no-go decisions on clinical development.<sup>61</sup> Product patents may be particularly critical for small molecule drug discovery,<sup>62</sup> where reverse engineering of the end product is straightforward<sup>63</sup> and incentives provided by the clinical data exclusivities conferred by the Food and Drug Administration (FDA) are limited to 5 years.<sup>64</sup> Accordingly, I focus on small molecules.<sup>65</sup>

The standard DDD process for small molecules comprises a series of stages that takes 10-15 years. Discovery often begins with laboratory identification of a biological target, typically a protein that appears to be implicated in a relevant biological pathway. Identification of this target is followed by high-throughput screening of the biological target against large libraries of molecules,<sup>66</sup> with the goal of identifying molecules that bind to the target (known as ligands or, more colloquially, “hits”).<sup>67</sup>

Screening for drug-target interactions can take place both physically and *in silico*, with the latter using computational models of both the target and the drug to predict binding. In either case, recent explosion in the size and chemical diversity of libraries has resulted in screening processes that often yield large numbers of hits.<sup>68</sup> From the hits, chemists must identify a small subset of promising candidates. These candidates can then be tested in cell-based *in vitro* assays as well as *in vivo* in animal models, with the overall goal of assessing efficacy as well as safety.<sup>69</sup>

After preclinical testing, the drug developer will typically file a product patent application covering the most promising candidates. Although this application is filed early in the innovation process, before any clinical testing, it is nonetheless considered the most economically important. Moreover, because the current system holds that any prior

---

<sup>61</sup> *See id.*

<sup>62</sup> Unlike large molecule biologics, which are generated in living systems, small molecules are generated through chemical synthesis.

<sup>63</sup> W. Nicholson Price and Arti K. Rai, *Manufacturing Barriers to Biologics Competition and Innovation*, 101 IOWA L. REV. 1023, 1033-1034 (2016).

<sup>64</sup> *Id.* at 1027.

<sup>65</sup> That said, the discussion below generally applies as well to therapeutic proteins.

<sup>66</sup> *See, e.g.,* Andreas Bender et al. *Which aspects of HTS are empirically correlated with downstream success?* 11 CURR. OPIN. DRUG DISCOV. DEVEL. 11, 32 (2008) (discussing small molecules). Though this Essay does not focus on biologics, biologics discovery also involves screening against molecule libraries. Marissa Mock et al., *AI Can Help to Speed Up Drug Discovery But Only If We Give It the Right Data*, 621 NATURE 467, 468 (2023) (depicting high-throughput screening of therapeutic proteins).

<sup>67</sup> *See* Bender et al. *supra* note 66, at 33.

<sup>68</sup> Anastasia Sadykov & Vsevolod Katritch, *Computing approaches streamlining drug discovery*, 616 NATURE 673, 674-76 (2023) (providing a comprehensive discussion of expansion in chemical space).

<sup>69</sup> For purposes of this high-level overview, I include within the category of safety and efficacy considerations surrounding appropriate pharmacokinetics and pharmacodynamics.

third-party competitor's disclosure of a similar molecular structure, no matter how uninformative, can be considered "prior art" that defeats patentability, firms may race to file.

After patent filing, the developer will apply to conduct clinical trials by filing an investigational new drug application (IND) at the FDA. Historically, at least for approved drugs with publicly available IND filing dates, this filing has generally occurred several years after patent filing.<sup>70</sup> One reason for delay may be that the developer, having raced to file the patent, has not yet done comprehensive *in vivo* testing. This hypothesis is supported by the fact that while FDA has historically required all IND applications to include animal pharmacology and toxicology studies,<sup>71</sup> drug patents don't always have evidence of such testing.<sup>72</sup>

On the other hand, the developer is not likely to want to wait too long after patent filing to file the IND. After patent filing, the 20-year exclusivity clock begins to tick, and patent term extensions that can be secured do not fully compensate for the delay occasioned by clinical trials. In fact, concerns about patent term reduction occasioned by long clinical trials are so acute that careful econometric research has shown that firms conducting research on cancer drugs appear to preferentially develop drugs for late-stage cancers (as contrasted with early stage cancers or prevention, where project duration would be longer).<sup>73</sup>

Because information on IND filings is secret, and even the date of IND filing is not revealed unless the drug is approved, systematic study of information contained in INDs, or even of the overall lag between patent filing date and IND filing date, is challenging. But regardless of what is contained in the IND, patent races may lead to the patent reward going to a firm that is not necessarily well placed to develop the drug, or to license the drug for development by other firms.

If the FDA approves the IND, the candidate typically goes through three stages of clinical trials, the first for safety only, the second for an initial assessment of efficacy in a relatively small population, and the

---

<sup>70</sup> Interview with Professor Bhaven Sampat, Arizona State University, January 21, 2025. Professor Sampat is compiling a database of lags between patent filing and FDA filing for approved drugs

<sup>71</sup> <https://www.fda.gov/drugs/types-applications/investigational-new-drug-ind-application> (discussing requirements). In 2022, Congress passed legislation that allows the FDA to forego animal testing requirements. Meredith Waxman, FDA No Longer Needs to Require Animal Tests Before Human Drug Trials, *Science*, Jan. 10, 2023, <https://www.science.org/content/article/fda-no-longer-needs-require-animal-tests-human-drug-trials>. As discussed in Part II.C, the FDA is now moving towards allowing other types of data, including AI-powered computational models.

<sup>72</sup> See *infra* note 107 and accompanying text. In the life sciences context more generally, a study by Janet Freilich of patents and patent applications containing experimental results found that the majority did not disclose sample size or have any statistical analysis. Freilich, *supra* note 59, at 448-461 (2020) (finding, based on reading a random sample of 500 patents and patent applications, that preclinical experiments disclosed in patents lack key methodological indicia of reproducibility).

<sup>73</sup> Budish, Roin, Williams, *supra* note 60, at 2044-48.

third for a more substantial test of efficacy in a larger population.<sup>74</sup> Throughout these trials, the FDA continues to monitor for safety. After clinical trials have concluded, the FDA uses a risk-benefit analysis to determine whether the candidate should be approved.

According to some analysts,<sup>75</sup> the historical 90% failure rate for compounds that enter clinical trials can be attributed, at least in part, to insufficient testing of targets and drugs at the preclinical stage. Studies by prominent academic researchers,<sup>76</sup> and by industrial scientists<sup>77</sup> have put forward irreproducibility results ranging from 50 to 90%.<sup>78</sup> Using the conservative estimate of 50%, a group of economists suggested in 2015 that up to \$28 billion per year was being spent on irreproducible preclinical research.<sup>79</sup> Despite a decade of attention, the problem persists.<sup>80</sup>

The complexity of human biology, which in many cases cannot be parsed through isolated analysis of single targets,<sup>81</sup> clearly plays a role in the impasse. More generally, responsibility for adequate preclinical testing primarily rests with the FDA. But patent racing, combined with the need to file an IND soon after patent filing, cannot help.

## B. *The Impact of AI Use*

As the point about *in silico* screening in Part II.A suggests, computational approaches of various sorts have, for more than two decades, played a significant role in DDD. Prior uses of rules-based computation have been heralded as changing the DDD paradigm, with potentially useful results for reducing the overall cost of innovation, particularly in the preclinical context.<sup>82</sup> So the advent of data-based AI

---

<sup>74</sup> Congressional Research Service, Agata Dabrowska and Susan Thaul, How FDA Approves Drugs and Regulates Their Safety and Effectiveness, May 8, 2018, R41983.

<sup>75</sup> Sadyekov & Katritch, *supra* note 68, at 673 (arguing that failure rate is “largely attributable” to insufficient preclinical vetting).

<sup>76</sup> See C. Glenn Begley and John P.A. Ioannidis, *Reproducibility in Science: Improving the Standard for Basic and Preclinical Research*, 116 CIRCULATION RESEARCH 116, 118 (2015) (listing 16 studies that reviewed published preclinical research, with each of the studies showing high rates of irreproducibility in the specific area of preclinical research).

<sup>77</sup> C. Glenn Begley and L.M. Ellis, *Drug Development: Raise Standards for Preclinical Cancer Research*, 483 NATURE 531 (2012) (Amgen scientists); Florian Prinz et al., *Believe It or Not: How Much Can We Rely on Published Data on Potential Drug Targets*, 10 NATURE REVIEWS DRUG DISCOVERY 712 (2011) (Bayer scientists).

<sup>78</sup> Freedman et al., *supra* note 59, at 3.

<sup>79</sup> *Id.*

<sup>80</sup> See, e.g., Andreas Bender and Isidro Cortes-Ciriano, *Artificial Intelligence in Drug Discovery: What is Realistic, What Are Illusions?*, 26 DRUG DISCOVERY TODAY 511, 514 (2021).

<sup>81</sup> See *id.* at 515 (critiquing focus on “isolated mechanisms and targets”).

<sup>82</sup> Including by me. See Arti K. Rai, *The Information Revolution Reaches Pharmaceuticals: Balancing Innovation Incentives, Cost, and Access in the Post-Genomics Era*, 2001 U. ILL. L. REV. 173, 174-75 (2001) (suggesting that computational work in genomics could reduce drug innovation costs); Stephen M. Maurer et al., *Finding Cures for Tropical Diseases: Is Open Source an Answer?* 1(3) PLOS MED: e56 (2004) (proposing open source

systems, even AI systems based on sophisticated deep learning methods, arguably represents evolution rather than revolution.

That said, particularly for small molecules,<sup>83</sup> AI systems are playing a role at various preclinical stages that appears to move substantially beyond prior computer-aided drug discovery. In this section, I address the role these systems are playing at each stage. At the target identification stage, the boom in biomedical data sets and scientific literature profiling both healthy individuals and those with a particular disease provides opportunities for machine-learning-based prediction of new targets.<sup>84</sup> For example, use of machine learning on so-called knowledge graphs that associate data on expression of protein-encoding genes with disease data has resulted in the identification of novel genes associated with Alzheimer's disease.<sup>85</sup>

Once a target is identified, an important next step is identification of 3D structure. As the 2024 Nobel Prize in Chemistry recognized, protein structure prediction is a task at which AI systems have begun to excel. DeepMind's AlphaFold 3 (AF3), which debuted via a May 2024 *Nature* publication,<sup>86</sup> provides an example of a deep learning predictive model that may be particularly relevant to drug discovery. AF3 moves beyond protein structure prediction to prediction of protein interaction with other molecules, such as small molecule ligands, DNA, and RNA. The tool appears particularly good at predicting interactions between protein structure and small molecules.<sup>87</sup>

AI systems can also assist medicinal chemists in prioritizing hits that emerge from high-throughput screening (HTS). This important use was emphasized in a review article co-authored by 18 international experts, which noted that

[m]any parameters need to be considered in hit selection and subsequent optimization, including potency and selectivity at the desired pharmacological targets and potential off-target [effects] as well as the physicochemical characteristics that could be important in

---

computational approaches as a mechanism for reducing drug innovation costs in the area of neglected diseases).

<sup>83</sup> The use of AI in discovery of small molecules is further developed than for biologics, see Madura Jayatunga et al., *AI in Small Molecule Drug Discovery: A Coming Wave?*, 21 NATURE REVIEWS DRUG DISCOVERY 175, 175 (2022) (empirical study of AI use, noting that the authors "focused mainly on small molecule drug discovery, for which AI approaches are relatively more established.").

<sup>84</sup> Jessica Varnethevan et al., *Applications of Machine Learning in Drug Discovery and Development*, 18(6) DRUG DISCOVERY TODAY 463, 466-68 (2019).

<sup>85</sup> Catrin Hasselgren and Tudor J. Oprea, *Artificial Intelligence for Drug Discovery: Are We There Yet?*, 64 ANN. REV. PHARMACOL AND TOXICOL., 24-1, 24-5 (2024).

<sup>86</sup> The study is entitled Josh Abramson et al., *Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3*, 630 NATURE 493 (2024).

<sup>87</sup> See Derek Lowe, *AlphaFold 3 Debuts*, SCIENCE, May 10, 2024, <https://www.science.org/content/blog-post/alphafold-3-debuts> ("The paper presents comparisons between AF3 and tools such as Vina, Gold, and RoseTTAFold All-Atom, using several hundred protein-ligand structures that were reported after the training set cutoff dates, and it appears that AlphaFold 3 is far more accurate.")

drug pharmacokinetics and safety. Consequently, medicinal chemists typically face challenging multi-objective optimization (MOO) problems, with far more potential choices than are possible to explore systematically . . . Part of the appeal of applying AI in drug design thus lies in the potential to develop data-driven, implicit model-building processes to navigate vast datasets arising from HTS and to prioritize alternatives.<sup>88</sup>

Additionally, AI systems can assist in small molecule design. In contrast to traditional approaches to *de novo* design that rely on explicit rules regarding chemical transformation and assembly, generative AI systems are trained using existing successful designs.<sup>89</sup> Some industry analysts argue that while most large pharmaceutical firms work within a universe of approximately 10 million compounds, AI systems could be used to generate, and sift through, a number closer to the larger universe (estimated at between  $10^{33}$  and  $10^{60}$ ) of drug-like compounds.<sup>90</sup> While there is reason to be skeptical that size *per se* (as contrasted, for example, with molecular diversity) is a particularly valuable goal, it is also unlikely that 10 million compounds is the optimal number.

Machine learning (ML) models developed using data on quantitative chemical structure-activity relationships (QSAR) can be used to design small molecules with desired biological properties. Although the quality of these models differs depending on data availability with respect to the class of the biological target,<sup>91</sup> the ultimate aspiration is to abstract beyond particular targets and develop “broadly generalizable or even universal models”<sup>92</sup> of binding affinities and absorption, distribution, metabolism, and excretion-toxicity properties.

---

<sup>88</sup> Petra Schneider et al., *Rethinking Drug Design in the Artificial Intelligence Era*, 10 NATURE REVIEWS DRUG DISCOVERY, 353, 353 (2020).

<sup>89</sup> Varnethevan, *supra* note 84, at 468-70. More specifically, generative models are often estimated through an adversarial process in which two models are simultaneously trained. The first is a generative model G of the overall data distribution, and the second is a discriminative model D that estimates the probability that a sample came from the training data rather than G. The training goal for G is maximizing the probability of D making an error. See generally Ian Goodfellow et al., *Generative Adversarial Nets*, ADV. NEURAL INF. PROCESS SYSTEMS 2672 (2014). Or, as Goodfellow et al. colorfully put it, “The generative model can be thought of as analogous to a team of counterfeiters, trying to produce fake currency and use it without detection, while the discriminative model is analogous to the police, trying to detect the counterfeit currency. Competition in this game drives both teams to improve their methods until the counterfeits are indistinguishable from the genuine articles.” *Id.* at 2672. As with any other predictive AI, however, the discriminatory model is only as good as the data on which it is trained. As noted above, to develop good discriminatory models, failure data is essential.

<sup>90</sup> Will Douglas Heaven, *AI is Dreaming Up Drugs That No One Has Ever Seen*, MIT TECH. REV., February 15, 2023, <https://www.technologyreview.com/2023/02/15/1067904/ai-automation-drug-development/> (quoting the CEO of an AI-native firm that uses generative techniques).

<sup>91</sup> Sadybekov *supra* note 68, at 679.

<sup>92</sup> *Id.*

Given this preclinical potential, what results have been achieved thus far? One 2024 review that utilized a broad list of AI-related keywords to retrieve data from the clinical trial database Pharmaprojects found a total of 165 use cases (164 investigational drugs and 1 approved drug) involving AI. The major categories of use were target identification (22%) and drug molecule discovery, whether through predictive use in platform screening or generative molecule design (76%).<sup>93</sup> The approved drug, stem cell therapy remestemcel-L, invoked a relative modest use of AI – Bayesian techniques to estimate the likelihood of achieving significant results for the primary endpoint at study completion. Another review from 2024 with a more specific focus on available information (including in Pharmaprojects) pertaining to “AI native” firms – that is firms that emphasize a central role for AI in their drug discovery techniques – found a total of 67 ongoing clinical trials in which AI was used at the preclinical stages.<sup>94</sup>

The goal, however, is not simply to generate candidate molecules but to generate candidates, and associated patents, that have indicators of reliability. The next section discusses the promise of AI use in achieving this goal.

### C. *The Potential for AI Use to Improve Reliability*

Ideally, AI systems would be used not simply to generate new candidate targets and molecules but also to improve the reliability of these molecules. For example, using predictive AI models to analyze clinical trial results, both failures and successes, could inform preclinical vetting with respect to safety and efficacy. In one exploratory study,<sup>95</sup> scientists from Abbvie used what they termed “highly curated” *in vitro* safety data on more than 1 million compounds to train a model to predict off-target effects that could compromise safety. They then used the trained model to predict off-target effects in 857 diverse small molecules with substantial clinical data -- 456 molecules that had failed clinical trials and 401 that had been approved by the FDA. The model correctly predicted over 70% of known off-target effects. More notably, over 50% of the model’s predictions for discontinued drugs had not been previously known. In another study, GSK scientists used machine

---

<sup>93</sup> Louise Druedahl et al., *Use of Artificial Intelligence in Drug Development*, 7(5) JAMA NETWORK OPEN e2414139, May 31, 2024, doi:10.1001/jamanetworkopen.2024.14139

<sup>94</sup> Madura Jayatunga et al., *How Successful Are AI-Discovered Drugs in Clinical Trials: A First Analysis and Emerging Lessons*, 29(6) DRUG DISCOVERY TODAY 1, 2 (2024) (discussing cases of AI-derived small molecules, AI-identified targets, and AI-derived biologics).

<sup>95</sup> Mohan Rao et al., *Novel Computational Approach to Predict Off-Target Interactions for Small Molecules*, 2 FRONTIERS IN BIG DATA 25, 25 (2019).

learning tools to correlate clinical data on drug target successes and failures with various different target features.<sup>96</sup>

The Abbvie and GSK studies illustrate the promise of using AI models, in conjunction with clinical results in humans, to create alternatives to traditional preclinical testing based on animal models, which not only raise ethical concerns but have well-known limitations in terms of translation to humans. Pursuant to legislation passed in 2022 that allows FDA to move away from reliance on animal data, the agency is now giving non-animal alternatives are well-deserved attention.<sup>97</sup>

The clinical trial data from which high-quality AI-based models would most usefully be built is, however, generally highly proprietary and accessible only to the large firms that typically take molecules through clinical trials. Indeed, even studies from large firms like Abbvie and GSK suffer from data limitations associated with the fact that the studies can rely only on that firm's own data on clinical successes and failures.

An option for improving reliability that would more accessible, including to small firms, is publicly funded creation of AI models for safety and efficacy that would be open for all to use. One promising example of such creation is FDA's safetAI initiative, which is developing public AI-based toxicity prediction models.<sup>98</sup> Going forward, even under-resourced small firms could use these models. Indeed, even if the FDA itself continued to require some amount of more expensive animal testing, firms could reference the models in their patent applications. Reference to such presumably trustworthy AI models would be an indicator of basic reliability.

Disclosure of AI use in patents could, however, run up against applicant concerns about compromising patentability requirements like human inventorship and non-obviousness. The next section discusses an original empirical inquiry into the disclosure issue.

#### D. *Failure to Disclose AI Use*

---

<sup>96</sup> Andrew Rouillard et al., *Systematic Interrogation of Diverse Omic Data Reveals Interpretable, Robust, and Generalizable Transcriptomic Features of Clinically Successful Therapeutic Targets*, 14(5) PLOS COMPUT. BIOL. e1006142, May 21, 2018, <https://doi.org/10.1371/journal.pcbi.1006142>

<sup>97</sup> Jeff Craven, FDA Seeks to Reduce Animal Testing Requirements for mAbs, Other Drugs, *Regulatory News*, April 11, 2025, <https://www.raps.org/news-and-articles/news-articles/2025/4/fda-seeks-to-reduce-animal-testing-requirements-fo>

<sup>98</sup> Interview with FDA's National Center for Toxicological Research, March 1, 2024. See also <https://www.fda.gov/about-fda/nctr-research-focus-areas/safetai-initiative> ("During the Investigational New Drug (IND) application submission process, FDA specifically reviews the safety of a submitted drug candidate before the sponsor can initiate any clinical trials. SafetAI is a collaborative initiative between CDER and NCTR to develop a suite of deep learning-based QSAR models with innovative approaches for various safety endpoints critical to regulatory science and the IND review. Currently, the initiative focuses on five key safety endpoints: hepatotoxicity, carcinogenicity, mutagenicity, nephrotoxicity, and cardiotoxicity.")

Identifying failures to disclose AI use of course requires some mechanism for identifying those firms that use AI in their DDD processes. Briefly, our team used the Pitchbook database, which tracks by industry sector – including in this case “drug discovery using artificial intelligence and machine learning” – firms that have received venture capital money. We combined Pitchbook results with publicly available promotional material and work done by prior researchers using the Pharmaprojects database.<sup>99</sup> Further details of how we identified “AI-native” DDD firms are contained in Appendix A. We then used USPTO databases to identify patents held by the firms.

Next, our team read the patents, excluding those that covered AI systems and other software, so as to isolate patents on AI-derived biological outputs such as molecules or methods of using molecules. Of the 135 output patents in the set as of December 2023, *only 4* mentioned AI use in any aspect of the disclosure. In cases involving molecules, the patent disclosures generally jumped immediately from the statement of the “problem” – a biological target in need of modulation – to the “solution” – a genus of molecules purportedly capable of modulating the target.<sup>100</sup>

Admittedly, one can connect the dots to some extent by looking at patents on AI systems held by some of the AI-native DDD firms. AI system patents, which contain some disclosure,<sup>101</sup> might be linked to outputs. In some cases, it might even be possible to examine peer-reviewed publications produced by the firms. That said, connecting the dots is a laborious and imperfect exercise.<sup>102</sup> For purposes of promoting reliability, a much better approach would involve affirmative disclosure in the output patent application itself.

One might argue that, as a doctrinal matter, when a patent claim covers a product, the patent statute’s language counsels against disclosure of the discovery process that led to the product. Responding

---

<sup>99</sup> Jayatunga, *supra* note 83.

<sup>100</sup> To be sure, disclosure of invention process also appears to be somewhat sparse even outside the AI context. My review of 50 drug molecule patents randomly sampled from a dataset developed and analyzed in Arti K. Rai et al., *Post-Grant Adjudication of Drug Patents: Agency and/or Court*, 37 BERKELEY TECH.L.J. 139 (2022), determined that the discussion of invention process in these patents was generally limited to mentions of prior art molecules that had been altered in various ways to make the claimed molecule. Even so, this is greater disclosure than was found in the AI-based patents. Moreover, in the case of more conventional processes, the general benefits and limitations of those processes are well known.

<sup>101</sup> Disclosure is far from comprehensive, however. See Mateo Aboy et al., *The Sufficiency of Disclosure of Medical Artificial Intelligence Patents*, 42 NATURE BIOTECHNOLOGY 839, 842-843 (2024) (finding that about two-thirds of a sample of 421 medical AI software patents disclose some AI architecture and training information, and approximately 30% disclose performance, but that very few disclose mathematical details or code listings).

<sup>102</sup> See Arti K. Rai et al., *Accountability, Secrecy, and Innovation Incentives in AI-Enabled Clinical Decision Software*, 7(1) JOURNAL OF LAW & BIOSCIENCES, <https://doi.org/10.1093/jlb/ljaa077> 1, 18-23 (2020) (engaging in this exercise in the context of patents on AI-enabled clinical decision software).

to a PTO Request for Comments on Inventorship and Artificial Intelligence that asked about disclosure of AI use in discovery, AI-native DDD firms pointed to the statutory provision stating that “patentability shall not be negated by the manner in which the invention was made.”<sup>103</sup> As these firms stressed, the 1952 Patent Act added this language in response to Supreme Court cases disparaging inventions that appeared to emerge from conventional discovery processes. Moreover, even despite this language, the danger of hindsight bias<sup>104</sup> could persist and requiring disclosure of discovery process would tempt patent examiners to unduly reject deserving patents.

Those skeptical of discovery process disclosure have also argued that disclosure requirements would be unduly onerous for both inventors and the patent office. In the specific context of AI use, they have noted that search prompts to AI models often have to be refined over time and requiring disclosure of every search step would be a waste of time for both applicants and patent examiners.<sup>105</sup> These skeptics have further suggested that disclosure might call into question human inventorship on any application that used AI.<sup>106</sup>

As discussed further in Part III, however, whether or not the PTO receives disclosures of AI use, any savvy defendant charged with infringement by a patent owner that discusses such use in its promotional materials is likely to raise issues of improper inventorship and obviousness. At least some patent applicants might prefer a route in which they could exchange *ex ante* disclosure in the patent application for protection against expensive contestation of inventorship and obviousness in a court setting. Additionally, as Part III discusses, the requisite level of disclosure need not be voluminous – the point is to promote basic scientific reliability, not to engage busy researchers, patent prosecutors, and examiners in tedious accounting exercises.

To be sure, disclosure of discovery process (AI-based or otherwise) isn't *necessary* to show reliability. For example, independent of discovery process, evidence that end product outputs have been tested relatively rigorously should suffice. In the DDD context, however, the evidence is not encouraging. In our sample, small molecule patents secured by AI-native firms have significantly less discussion of testing *in vivo* than patents secured by firms that employ conventional identification methods.<sup>107</sup>

---

<sup>103</sup> See, e.g., Comment from Atomwise Inc., Cellarity, Inc., Generate Biomedicines, Relay Therapeutics, Inc. and Tessera Therapeutics, Inc., <https://www.regulations.gov/comment/PTO-P-2022-0045-0046>. The main trade organizations that represent the biopharmaceutical industry – the Biotechnology Industry Organization and PhRMA – made similar points in their comments.

<sup>104</sup> Gregory N. Mandel, *Patently Non-Obvious: Empirical Demonstration that the Hindsight Bias Renders Patent Decisions Irrational*, 67 OHIO ST. L.J. 1391 (2006).

<sup>105</sup> Andrei Iancu and David Kappos, *New Patent Guidance on AI Could Quash Innovation*, WALL ST. J., July 11, 2024.

<sup>106</sup> *Id.*

<sup>107</sup> See Janet Freilich and Arti Rai, *What Patents on AI-Derived Drugs Reveal* (working draft, on file with author). The difference is significant at the 0.05 level.

In sum, the DDD case study illustrates that even in innovation contexts where the patent system plays a central role, it is not currently working to bolster the reliability of AI use. Instead, AI use may be exacerbating racing dynamics that misallocate rights. The next Part turns to proposals for law and policy improvements that draw upon the lessons of the case study to properly allocate, and safeguard, patent rights.

### III. PATENT LAW AND POLICY IMPROVEMENTS

Several law and policy improvements will be important for realizing the promise of AI use while simultaneously bolstering scientific reliability. These include: fortification of the utility standard for patentability; parallel fortification of the standard for what constitutes prior art; protection for those who disclose discovery process against at least two *ex post* challenges to patentability: first, a deficient showing of inventorship and, second, inappropriate use of this disclosure to make arguments regarding obviousness. Some adjustment to the current legal test for inventorship is also likely to be necessary.

#### A. Utility in Patents and Unpatented Prior Art

Working with the PTO (which could issue initial guidance, as it did in the case of gene fragments),<sup>108</sup> the Federal Circuit should consider tweaking the utility standard. Patent applicants should have to make the case, whether through analytical reasoning, rigorous testing of the end product, disclosure of sound discovery methods, or a combination of the above that the ordinary scientist or technologist in the field would find their invention reliable.

To be sure, this proposed change is far from bulletproof. As a multi-factor standard that would have some teeth, it is more time consuming to administer than the current, largely non-existent, test. Additionally, patent examiners are highly constrained in the time they can devote to reviewing patent applications.<sup>109</sup> Even so, utility fortification along the lines of what happened with gene fragments is a plausible requirement. Moreover, a consistent threat of judicial implementation by courts should spur applicants to offer more credible evidence of utility *ex ante*.

In many if not most cases, fortification of the utility standard would not be commercial-level usefulness. Particularly for innovation that is risky from the standpoint of human or environmental health and safety, risk regulators, not the PTO, are institutionally best suited to make determinations regarding commercial marketing. Additionally, in cases of substantial risk to safety, the expense associated with proving commercial-level usefulness will likely be comparably substantial. To

---

<sup>108</sup> See *supra* note 38 and accompanying text.

<sup>109</sup> See Michael Frakes & Melissa F. Wasserman, *Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents*, 99(3) REVIEW OF ECONOMICS AND STATISTICS 550, 550 (2017).

promote investment in proving safety, patents will need to be granted *before* the invention begins to receive scrutiny by the risk regulator. However, as I discuss further below, the more that the patent applicant undertakes to show that it has appropriately tested its AI systems, and/or appropriately tested output from the systems, the more it should be protected from challenges regarding inventorship or obviousness, whether at the PTO or in the courts. A heightened utility standard would also mitigate challenges associated with patent racing leading to applicants having fewer years left on their patent post-issuance. Although such challenges pose a particular problem in the biopharmaceutical industry,<sup>110</sup> they are also a more general issue.

Some concrete examples that focus on specific cases of AI use in DDD can help clarify the contours of a fortified utility standard. As discussed in Part II, determining these use cases requires going outside the patent record. However, AI-native firms that are eager to license not simply to license specific molecules but also their broader AI-based discovery platforms do sometimes publish in the peer-reviewed literature. For these firms, one can identify more precisely how AI is being used.

Consider, for example, the case of the AI-native firm In Silico Medicine, which advanced into Phase II clinical trials a molecule targeting the TNIK protein kinase, with the goal of treating idiopathic pulmonary fibrosis. Although the information is not disclosed in its molecule patent, which was first filed in 2019,<sup>111</sup> In Silico Medicine published online in 2024 a peer-reviewed Nature Biotechnology article detailing its research methodology.<sup>112</sup> The paper contains a two-page, relatively high-level, disclosure of how the In Silico PandaOmics AI platform was used to identify the target and how its generative AI Chemistry 42 platform was used to identify molecules to modulate the target.<sup>113</sup> It also describes in summary terms a training data set and model validation details. This type of disclosure could, together with the *in vitro* testing actually contained in the patent disclosure, suffice to meet a heightened utility requirement.

Notably, even under a fortified utility standard, a patent applicant that disclosed more about discovery process along the lines suggested above might not need to engage in expensive *in vivo* testing (as In Silico Medicine apparently did not before filing its patent). In this way, patents on AI systems vetted by humans, and disclosure of such vetting, might be accessible even to under-resourced small firms. And to the extent that any of the testing could rely on well-vetted, public AI models

---

<sup>110</sup> See Budish et al., *supra* note 60 and accompanying text.

<sup>111</sup> The kinase in question is the NCK-interacting protein kinase TNIK, and is claimed in U.S. Patent 11795160 (“Kinase Inhibitors”).

<sup>112</sup> The print publication came out in 2025. Feng Ren et al., *A Small Molecule TNK1 Inhibitor Targets Fibrosis in Preclinical and Clinical Models*, 43 NATURE BIOTECHNOLOGY 63 (2025).

<sup>113</sup> *Id.* at 64-65.

of toxicity, of the sort that the FDA is attempting to develop,<sup>114</sup> the applicant's path would be easier still.

Just as the scientific standards for patents should be raised, so should the standards for what constitutes prior art. While the standards need not be precisely identical – a slight preference for the public domain over patents may weigh against precise parallelism – the advent of AI substantially reinforces the need to rethink the prior art inquiry. One straightforward change would require that a disclosure should not count as prior art unless the disclosure identifies a specific, scientifically plausible *use* for the scientific or technical item disclosed. Additionally,

## B. *Protection Against Inventorship and Obviousness Challenges*

More rigorous discovery and testing processes – and disclosure thereof – should be useful not only for compliance with a fortified utility standard but could also provide patentees *ex ante* protection against subsequent validity challenges relating to inventorship and obviousness.

### 1. *Inventorship*

Of the various topics at the intersection of AI and patents, none has received more attention than whether AI use could imperil the patent validity requirement of human inventorship. Indeed, some policy entrepreneurs and scholars have argued that conventions about human inventorship should be dispensed with, and AI systems should simply be recognized as inventors. The most notable example of this argument comes from the 2022 case of *Thaler v. Vidal*.<sup>115</sup> In that case, Stephen Thaler, a developer of AI systems and his attorney, law professor Ryan Abbott, argued that one of Thaler's systems, DABUS (Device for the Autonomous Bootstrapping of Unified Science) should be deemed the inventor of the “neural flame” and the “fractal container” that were claimed in the patent applications that Thaler filed.

On this view, recognition of AI system inventorship – with subsequent assignment of the invention to a human or corporate owner – squarely addresses the obstacle to patentability that will arise (and in Thaler's view has already arisen) when no human can claim to have contributed to the conception of an invention. Ryan Abbott has also noted the collateral benefit that allowing AI inventorship would encourage disclosure of AI use.<sup>116</sup>

Giving AI systems inventor status would, however, not only require Congressional action but might also test the constitutional boundaries of congressional power. On the statutory point, the Federal Circuit's *Thaler v. Vidal* decision correctly notes that the current patent statute

---

<sup>114</sup> See TAN *supra* note 98.

<sup>115</sup> 43 F.4th 1207 (Fed. Cir. 2022).

<sup>116</sup> Ryan Abbott, *Everything is Obvious*, 66 UCLA L.REV. 1, 6 (2019).

repeatedly refers to inventors and co-inventors as “individuals”<sup>117</sup> and that the Supreme Court has held that individuals are presumptively human beings. On the constitutional point, the Court has in past cases determined that Congress has broad power under the Constitution’s Intellectual Property Clause.<sup>118</sup> But even under an expansive view of congressional power, legislation stating that non-humans could be inventors would raise constitutional concerns. Particularly under the text, history, and tradition test to which the current Supreme Court hews, it is not clear that constitutional use of the term “inventor” encompasses non-humans.

Both normatively and constitutionally, a more compelling approach would leverage this Essay’s reliability framework. Just as those concerned about AI trustworthiness from the standpoint of improving safety, mitigating bias, and promoting other important values have emphasized the importance of humans in the loop, here the human in the loop would refine and test scientific reliability. In the context of DDD, for example, those who created and tested bespoke AI systems that generated credible targets could properly be deemed inventors. So could those who took the results from the system and refined them further. In contrast, the creators of a general purpose system, even a high-powered generative AI model with hundreds of billions of parameters, would not be inventors of products that emerged through more specialized training of those so-called foundation models. Nor would those who did nothing more than put generic queries into either a bespoke or general system.

Fostering a human role in promoting reliability is superior not only to deeming AI systems inventors but also to a proposal that Congress deem humans to be inventors of AI-generated invention by analogy to the patents of importation that existed in England at the time of the Constitutional Convention.<sup>119</sup> As the authors of this proposal themselves conclude, the historical record at the time of the Framing and the early Republic does not fully support congressional power to create patents of importation.<sup>120</sup>

Even the modest change proposed here takes the definition of inventorship sufficiently far from the courts’ historical common law gloss that the change will have to come from either Congress or the courts. Not only does the PTO lack rulemaking authority over patent validity questions but after the Supreme Court’s 2024 overruling of the *Chevron* doctrine, legal interpretations by agencies no longer receive

---

<sup>117</sup> *Thaler*, *supra* note 115, at 1211-1213(citing 35 U.S.C. § 100(f) and 100(g) of the patent statute).

<sup>118</sup> *See, e.g.*, *Eldred v. Ashcroft*, 537 U.S. 186, 199-208 (2003).

<sup>119</sup> *Schwartz & Rogers*, *supra* note 4, at 556-560.

<sup>120</sup> *Id.* at 560 (concluding, after a detailed historical discussion, that “[a]t worst, the question of whether the Constitution permits patents of importation and, by extension, patents on AI-produced inventions, is historically ambiguous.”)

deference even when promulgated via rulemaking.<sup>121</sup> Even so, given the delays that are inevitable in either Congressional or court actions, some action by the first mover in patent law – the PTO – is necessary.<sup>122</sup>

PTO guidance on inventorship in the context of AI use, issued in February 2024 pursuant to an Executive Order that directed the agency to publish guidance on the question,<sup>123</sup> (“Inventorship Guidance”) starts us down the path outlined above. The Executive Order and the guidance mirror earlier executive branch initiatives that took the lead on important patent policy issues, with courts subsequently concurring with the executive branch.<sup>124</sup>

According to the Inventorship Guidance, inventorship requires that a natural person(s) make a “significant” contribution to invention.<sup>125</sup> In the agency’s view, simply presenting a problem to an AI system may not be sufficient to show significant contribution to the AI system’s output.<sup>126</sup> That said, “a significant contribution could be shown by the way the person constructs the prompt in view of a specific problem to elicit a particular solution from the AI system.”<sup>127</sup> This framing appropriately recognizes the case-by-case nature of whether “prompt engineering” can itself constitute a type of inventiveness.

Additionally, according to the agency, a person who takes an AI system output and conducts further experimentation may make a significant contribution.<sup>128</sup> As one of the illustrative examples used by the agency recognizes (“Developing a Therapeutic Compound for Treating Cancer”)<sup>129</sup>, this particular principle may be quite relevant to drug discovery. Specifically, the agency’s example includes a variation (“Scenario 1”) in which the agency concludes that two scientists, who

---

<sup>121</sup> The 2024 Supreme Court decision in question was *Loper Bright v. Raimondo*, 603 U.S. 369 (2024). In addition to *Loper Bright*, the Court decided several other cases during the 2023-2024 term that substantially curtail administrative power. For a short discussion of what the Court’s administrative law cases from that term might mean for patent law, see Arti K. Rai, *Patent Puzzles After the Supreme Court’s 2024 Administrative Law Cases*, <https://patentlyo.com/patent/2024/07/administrative-rulemaking-discretion.html>

<sup>122</sup> See generally John Golden, *Working Without Chevron: The PTO as Prime Mover*, 65 DUKE L.J. 1657-1699 (2016).

<sup>123</sup> Inventorship Guidance, *supra* note 6, at 10044 (citing October 30, 2023 “Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence”).

<sup>124</sup> See, e.g., Arti Rai, *Patent Validity Across the Executive Branch: Ex Ante Foundations for Policy Development*, 61 DUKE L.J. 1237-1281 (2012); Tejas Narechania, *Patent Conflicts*, 103 GEO. L.J. 1483 (2014-2015). As discussed in these articles, the USPTO and other parts of the executive branch sometimes work in unison and sometimes do not.

<sup>125</sup> Inventorship Guidance, *supra* note 123, at 10046,

<sup>126</sup> *Id.* at 10046 (“A natural person who only presents a problem to an AI system may not be a proper inventor or joint invention of an invention identified from the output of the AI system.”)

<sup>127</sup> *Id.*

<sup>128</sup> *Id.* (“[A] person who takes the output of an AI system and makes a significant contribution to the output may be a proper inventor.”)

<sup>129</sup> <https://www.uspto.gov/sites/default/files/documents/ai-inventorship-guidance-chemical.pdf> (“Cancer Compound Example”).

use laboratory chemistry to improve an AI-generated structure for a compound are properly named inventors.<sup>130</sup>

Another guiding principle put forward by the agency states that the designer of an AI system that solves a specific problem may be an inventor of the output generated by the system.<sup>131</sup> Scenario 2 in the cancer compound example, which would confer inventorship of an AI output molecule to the designers of a neural network “Molecule Optimizer” that specifically generates molecules with favorable pharmacological properties, illustrates this principle.<sup>132</sup>

The agency’s vision on what should constitute human inventorship in the context of AI use is generally sound. The guidance should provide applicants confidence that disclosure of AI use, including disclosure sufficient to contribute to utility, will not imperil inventorship. In fact, after many years of scholarly advocacy for the goal of greater patent reliability, it may even provide some concrete incentives – i.e., protection against subsequent inventorship challenges – towards this goal.

That said, the guidance fails to take sufficient affirmative steps towards incentivizing disclosure of AI use, as would be necessary to make such disclosure one mechanism for fortifying reliability. With respect to an affirmative duty of disclosure, the guidance states that this duty arises in the first instance only if the applicant has information “that raises a prima facie case of unpatentability due to improper inventorship.”<sup>133</sup> But very few applicants would presumably bother with an application in which they laid out the prima facie case for an inventorship challenge to their own application.<sup>134</sup> As currently drafted, the guidance represents a missed opportunity on questions of incentivizing reliability through protection against inventorship challenges.<sup>135</sup>

---

<sup>130</sup> *Id.* Scenario 1 posits that the contributions made by Marisa and Naz rise to the level of inventorship.

<sup>131</sup> Inventorship Guidance, *supra* note 123, at 10049.

<sup>132</sup> Cancer Compound Example, *supra* note 129. Notably, in a nod towards the importance of analyzing actual clinical trial results to train upstream AI models, Scenario 2 posits that the designers of the hypothetical molecular optimizer have trained the model on compounds that succeeded in clinical trials.

<sup>133</sup> Inventorship Guidance, *supra* note 123, at 10049.

<sup>134</sup> Perhaps not surprisingly, then, the agency believes that “[a]t this time, to meet their duty of disclosure, applicants rarely need to submit information regarding inventorship” and that the guidance should not “have a major impact on applicants’ disclosure requirements.” *Id.*

<sup>135</sup> The guidance does instruct examiners to inquire into inventorship in cases where they “have a reasonable basis to conclude that one or more named inventors may not have contributed significantly to the claimed subject matter.” *Id.* at 10050. Whether examiners would have the time or inclination to make such inquiries, even in contexts where a cursory internet search would reveal the AI focus of the applicant firm, is unclear. Moreover, even if the examiner made an inquiry, a response about the role that AI played would appear not in the patent itself but in the less public-facing prosecution history.

A preferable approach would incentivize disclosure of AI use by stating that, for those patent applicants that demonstrated a sufficient quantum of human involvement, examiners would put into the examination record a statement of appropriate inventorship. Courts could then treat that statement as a safe harbor against subsequent challenges to inventorship.<sup>136</sup>

Such a safe harbor could prove quite valuable. Even in ordinary challenges to inventorship, where the patentee may be directed to correct inventorship by adding an inventor, the prospect of having to share inventorship with additional parties is one that patentees do not<sup>137</sup> welcome.<sup>138</sup> In the case of AI use that could be used to invalidate a patent entirely, a safe harbor should be even more attractive. Moreover, patents that benefited from the safe harbor would have greater reliability, with the consequence that a safe harbor option would benefit not only the patentee but society as a whole.

Further, the work required to invoke the safe harbor would not need to be especially burdensome. For example, in the *In Silico Medicine* example above, the two page summary of methodology contained in the *Nature Biotechnology* article, the human contribution to each of the many AI models used sequentially, and to the sequencing, is clearly explained. Such a description would both suffice for showing inventorship and would also bolster confidence in the reliability of the output.

## 2. *Obviousness*

Another prominent concern about pervasive use of AI systems is that their use will render all inventions obvious. After all, a presumption of even publicly available AI assistance from foundation models raises the skill of the hypothetical “person having ordinary skill in the art” (PHOSITA) who serves as the benchmark for obviousness determinations.<sup>139</sup> Additionally, to the extent that AI-driven search engines make more prior art readily accessible and therefore potentially

---

<sup>136</sup> To be sure, the patent statute states that granted patents are “presumed valid.” 35 U.S.C. 282. Moreover, in the 2011 case *Microsoft v. i4i*, 564 U.S. 91 (2011), the Supreme Court interpreted the statutory language to mean that those who challenge patents in district court must prove invalidity by clear and convincing evidence. Even so, an examiner’s statement of proper inventorship should provide additional assurance for patentees. Just as the *i4i* Court held that a challenger’s evidence of prior art that the USPTO did not review should help the challenger discharge its clear and convincing evidence burden, so too should courts be more likely to find improper inventorship when the issue of inventorship has not been reviewed at the USPTO.

<sup>137</sup> See TAN *supra* note 106.

<sup>138</sup> In part this is because absent assignment of all inventors’ rights to a single owner (as happens when the inventors are all employees of the same firm), additional inventors are by default co-owners, with independent rights to license.

<sup>139</sup> Gaetan de Rassenfosse et al., *AI-Generated Inventions: Implications for the Patent System*, 96 S. CAL. L. REV. 101, 112 (2023). See also Haixin Lin et al., *AI Can Almost Generate Evidence of Patent Obviousness*, Law 360, March 13, 2023.

“analogous” within the meaning of patent case law on obviousness, the PHOSITA may be presumed to have access to substantially more prior art.<sup>140</sup>

Until general purpose AI models are able to generate scientifically reliable responses to general queries, however, obviousness should not pose a critical problem. As the PTO’s Inventorship Guidance indicates, in most if not all cases of AI use in sophisticated research, at least some of the relevant systems will have been trained and refined based on at least some skill sets and data proprietary to the applicant. Under black letter patent law, so long as a patent applicant is filed before those data and skill sets are used to market an output invention, proprietary skill sets and data cannot be used to deny the applicant a patent.<sup>141</sup> Additionally, as discussed earlier, while concerns about hindsight bias<sup>142</sup> that might result from clear explanations of methodology are certainly plausible, such disclosure could also show the ways in which humans and machines worked together in non-obvious ways to make the idea reliable.

At some future time, general purpose AI systems may produce scientifically reliable responses to general queries across all fields. In that case both human inventorship and non-obviousness could be called into question. But in that case patents will presumably be necessary only if the invention’s safety must be carefully evaluated from a risk-benefit perspective. These are generally spheres where risk regulation agencies require further pre-market testing before an invention can be deployed.

In those circumstances, it will be important to explicitly embrace incentives on the basis of business risk and cost. A large number of commentators have provided creative accounts of how *ex ante* innovation rights, or *ex post* rewards, that worked outside the current patent system might be invoked or developed.<sup>143</sup> Moreover, many risk regulation schemes already provide data exclusivity and trade secrecy protections for those firms that do the work of generating the safety data necessary to take the product to market. Commentators have discussed those schemes – which are available not only for all types of biopharmaceutical data but also for other types of safety data – and mechanisms by which they could be improved.<sup>144</sup>

---

<sup>140</sup> De Rassenfosse et al., *supra* note 139, at 113.

<sup>141</sup> This longstanding precedent was unanimously reaffirmed in 2019 by the Supreme Court in *Helsinn v. Teva*, 139 S. Ct. 628 (2019).

<sup>142</sup> See *supra* TAN 104.

<sup>143</sup> See generally Sichelman, *Commercializing Patents*, *supra* note 14 (discussing “commercialization” patents that could be given to those who committed to bringing a public domain invention to market); Daniel Hemel & Lisa Larriomore Ouellette, *Innovation Policy Pluralism*, 128 YALE L.J. 544 (2019) (discussing non-IP innovation incentives such as tax preferences, grants, other government funding, and prizes as well as creative combinations of IP with such non-IP incentives).

<sup>144</sup> For a survey of different federal regulatory schemes that provide protection from competition based on safety data submission, see Yaniv Heled, *Regulatory Competitive Shelters*, 76 OHIO ST.L.J. 1, 27-54 (2015). For an early, extensive discussion focused on

## CONCLUSION

This Essay has shown that pervasive AI use makes the longstanding patent law debate over reliability, both of patents and of prior art, an urgent issue. Patent claims on inchoate ideas, or a profusion of such ideas in the public domain, may defeat patent incentives for responsible research. The challenge is particularly pressing in highly-patent sensitive contexts like the drug discovery and DEVELOPMENT context. That said, proper use of AI systems, properly disclosed, could bolster the scientific reliability of patents. Moreover, the human input necessary for high-quality AI use could provide a powerful defense against the challenges to human inventorship and non-obviousness that AI use raises.

The path forward will require fortifying the utility requirement and, in parallel fashion, fortifying requirements for the types of prior art that can thwart patent grants. Those who report high-quality AI use will not only be well positioned to satisfy these fortified requirements, they will also be protected against inventorship and non-obviousness challenges. The reliability criterion highlighted by AI use also shows the wisdom of court decisions limiting the scope of patent claims.

If we reach a stage at which relatively reliable science is obvious, but still requires substantial investment to reach market, a major reason will be pre-market safety requirements enunciated in risk regulation schemes. In that case, commercialization incentives more specifically tailored to risk regulation schemes will be more useful than patents.

---

FDA risk regulation as a form of innovation incentive, *see* Rebecca Eisenberg, *The Role of the FDA in Innovation Policy*, 13 MICH. TELECOMM. TECH. L. REV. 345 (2007),

## APPENDIX A

*Search Methodology for “AI-Native” Firms’ Patents*

We began by using the Pitchbook market research database.<sup>145</sup> As of August 2023, Pitchbook listed 190 companies in the category of “AI-powered drug discovery,” a subcategory of Healthcare. It also listed 408 companies in the industry of “drug discovery” with a vertical of “Artificial Intelligence & Machine Learning.” There was a significant, but not perfect, overlap between the two lists. The union of these two search lists yielded 508 firms.

We first used Pitchbook to narrow the list of 508 firms to firms that had 3 or more patents. The vast majority of these firms had portfolios that clearly covered different modes of discovery. By relying on public-facing marketing materials, and the extent to which the marketing materials suggested exclusive reliance on AI for therapeutic outputs, we further narrowed the list down to 28 “AI-native” drug discovery firms. We added to these 28 firms an additional 12 firms not among the 28 that a 2022 review of Pharmaprojects data<sup>146</sup> identified as an “AI -native” firm that had assets in clinical trials (for a total of 40). Although this set of firms is likely underinclusive (and may be overinclusive, particularly because of “AI-washing” on the part of some firms), it should not have systematic biases.

We used PTO databases to identify all U.S. patents held by these firms. We then read the titles and abstract to screen for patents that covered potential “outputs” of AI use. The largest set of outputs was patents on small molecule or biologic compounds (claimed either directly through compound claims or indirectly through method claims to a genus of molecules). We also included claims to new methods of use and formulations – AI use may occur in these output contexts as well.

---

<sup>145</sup> Pitchbook uses web crawlers to scan the internet and capture relevant financial information from news articles, regulatory filings, websites, press releases and more. While its database provides reliable information on many individual companies and can be useful for analyzing trends in an industry, its database is not comprehensive and is likely under inclusive. One inherent blind spot is very new companies in stealth mode. Another stems from less comprehensive reporting requirements for private companies. We supplemented our Pitchbook analysis with analysis of compound patents held by firms that the peer-reviewed literature reported to be “AI-native” firms.

<sup>146</sup> Madura Jayatunga et al., *AI in Small Molecule Drug Discovery: A Coming Wave?*, 21 NATURE REV. DRUG DISCOVERY 175 (2022), Supplementary Materials.