
Essay

AI Patents and the Self-Assembling Machine

Dan L. Burk[†]

INTRODUCTION

Technological change is a constant in our society. Such change sometimes poses novel challenges for regulation,¹ but unfailingly elicits its dubious declarations regarding its revolutionary legal character. Artificial intelligence, or “AI,” seems to be the novel technology *du jour*, prompting a growing literature devoted to the problems such information systems pose for the existing legal order.² Much of this commentary surrounding AI systems partakes of what has been called the “magical worldview,”³ a somewhat breathless and overwrought perspective on technology that assumes AI systems are somehow transcendent or miraculously unprecedented in their qualities or applications.⁴ In the patent context, such hyperbole manifests itself in assertions that these technologies upend the patent system, defy long-established patent doctrines, or portend the end of innovation as we know it.⁵

[†] Chancellor’s Professor of Law, University of California, Irvine. My thanks to participants in the 38th Annual Congress of the International Association for the Advancement of Teaching and Research in Intellectual Property (ATRIP) for helpful discussion on a previous version of this paper. Copyright © 2021 by Dan Burk.

1. See Monroe Price, *The Newness of New Technology*, 22 CARDOZO L. REV. 1885, 1886–90 (2001).

2. See generally Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399 (2017) (reviewing the legal landscape related to AI systems).

3. See Madeline Clare Elish & danah boyd, *Situating Methods in the Magic of Big Data and AI*, 85 COMM. MONOGRAPHS 57, 62–63 (2017).

4. Cf. Lyria B. Moses, *Why Have a Theory of Law and Technological Change?*, 8 MINN. J.L. SCI. & TECH. 589, 596 (2007) (“Despite occasional statements that some new technology changes everything, legal problems stemming from technological change are relatively rare and quite specific.”).

5. See, e.g., Tabriz Y. Ebrahim, *Data-Centric Technologies: Patent and Copyright Doctrinal Disruptions*, 43 NOVA L. REV. 287 (2019); Ryan Abbott, *Everything Is Obvious*, 66 UCLA L. REV. 2 (2019); Shlomit Yanitsky Ravid & Xiaoqiong (Jackie) Liu, *When*

To be clear: I am fully convinced that AI technologies do pose profound and fundamental challenges for the law and for social policy, and I have discussed some of those challenges at length elsewhere.⁶ But I remain wholly unconvinced that these technologies pose such challenges either to intellectual property law in general, or to patent law in particular. The patent system is surprisingly adaptable, accommodating over its history a wide range of new technologies that display divergent characteristics.⁷ Accommodating the current acceleration of innovation in machine learning technologies seem to me far less of a challenge to conventional patent doctrines than was the advent of recombinant biotechnological inventions thirty-five years ago.⁸

Indeed, in many cases the solutions developed for the patenting of biotechnological inventions appear to provide ready answers to the concerns raised regarding patents and AI technologies. Consequently, I shall argue throughout this Essay that where solutions to AI patenting difficulties are not already apparent, doctrines accommodating patent law to biotechnology can either themselves be applied to AI patenting, or point the way to solutions for whatever small difficulties AI patenting now poses. Thus, far from challenging the existing order of patent law, the patent system is fully equipped to encompass AI innovation, with perhaps some minor doctrinal accommodations that are well within the policy lever discretion available to the courts and to the Patent Office.⁹

Consequently, in this Essay I offer a more sanguine and hopefully more sensible view of AI patents than has so far been advanced in the literature. I begin by examining and de-sensationalizing some of the supposedly revolutionary qualities of AI systems, particularly the trope surrounding AI revolution through “emergent” discovery. I then consider several of the patent doctrines that have been suggested as the sites of destabilization by AI technologies, showing that in fact patent law is entirely capable of accommodating AI innovation within

Artificial Intelligence Systems Produce Inventions: An Alternative Model for Patent Law at the 3A Era, 39 CARDOZO L. REV. 2217 (2018).

6. See Dan L. Burk, *Algorithmic Legal Metrics*, 96 NOTRE DAME L. REV. 1147 (2021).

7. See Dan L. Burk & Mark A. Lemley, *Policy Levers in Patent Law*, 79 VA. L. REV. 1575 (2003).

8. See, e.g., Dan L. Burk, *Biotechnology and Patent Law: Fitting Innovation to the Procrustean Bed*, 17 RUTGERS COMP. & TECH. L.J. 1 (1991) (discussing doctrinal challenges to patenting posed by early biotechnology).

9. See Burk & Lemley, *supra* note 7 (discussing the doctrinal policy levers that allow patents to accommodate new technologies).

those doctrines. I conclude with several observations related to the actual challenge that in fact faces AI patenting, which is the problem of causation within patent law. This is not a new challenge, and in particular is not an AI specific challenge, but is rather an existing doctrinal gap that has become apparent in recent patent cases, and which is highlighted by the advent of AI related innovation. AI innovation thus offers the opportunity to improve and refine patent doctrine.

I. THE IRRELEVANCE OF EMERGENT OUTPUTS

We should begin by recognizing that “artificial intelligence” is something of a misnomer. What is now being touted as “AI” is almost entirely, and perhaps altogether entirely, systems implementing machine learning routines.¹⁰ Such systems are not intelligent in any robust sense of the word; they lack any hint or expectation of encompassing “strong” AI with general cognitive abilities of the sort that humans (or even animals) routinely display.¹¹ There is at present no serious prospect of designing machines with such capabilities; as Fourcade and Healy have observed, computer science has given up on building machines that can think in favor of building machines that can learn.¹²

For that matter, even “learning” is something of a euphemism or metaphor in this context; the devices that are garnering attention for their astounding capabilities are primarily statistical optimization systems, capable of rapidly generating and testing statistical models against massive data sets, iteratively amending the fit of the model to the data until some specified parameter is met.¹³ But if we set aside the distracting anthropomorphisms, the technology is unquestionably impressive in the scope and breadth of its potential applications. It is highly adept at solving certain types of problems, particularly problems that lend themselves to predictive data modeling. Among its many applications are uses in generating innovative technical advances, such as designing new devices or developing new molecular structures.

10. See Elish & boyd, *supra* note 3, at 63 (reviewing the trend toward machine-learning technologies).

11. *Id.* Admittedly, some legal commentators have chosen to focus on patents in the context of “strong” or “good old fashioned AI” (GOF AI). See, e.g., Clark Asay, *Artificial Stupidity*, 61 WM & MARY L. REV. 1187, 1198–1217 (2020).

12. Marion Fourcade & Kieran Healy, *Seeing Like a Market*, 15 SOCIO-ECON. REV. 1, 24 (2017).

13. Adrian MacKenzie, *The Production of Prediction: What Does Machine Learning Want?*, 18 EURO. J. CULT. STUD. 429, 435 (2015).

And therein lies the origin of the more sensational commentary regarding AI and patents: Much of the agitation regarding AI and patenting (and indeed, regarding AI and other areas of intellectual property) has arisen due to the so-called “emergent” properties of machine learning systems—that is, due to the unexpected and unprogrammed outputs generated by such systems.¹⁴ Machine learning systems may generate unpredictable outputs that were unforeseen and unforeseeable to their human developers. In many cases, the unpredictability of the outputs is taken as a sign that the system is in some way operating creatively or autonomously, because it generates results well beyond the capabilities of human cognition.¹⁵

As we shall see, the emergent nature of AI outputs is by no means an indicator of autonomous creation. But even assuming that AI systems in fact generate outputs that are truly emergent and therefore autonomous—a dubious assumption at best—this need not preclude them from coverage within the patent system, nor even pose any particularly new challenges to patent doctrine. The patent system has encountered and encompassed such technologies before. Such “emergent” outputs have long been the norm in the chemical and biological sciences, at least in the sense that “emergent” products with unforeseeable characteristics are generated by the technology once it has been set in motion. Organic synthesis, mutagenesis, cell transformation, cell fusions, and other complex reactions in biochemical systems frequently lead to unexpected but valuable products and by-products. Indeed, often the generation of just such unusual and unforeseen outputs is the goal of engaging such multifarious processes.

The Nobel Prize-winning technology for production of monoclonal antibodies offers a classic example.¹⁶ Monoclonal antibodies were produced by the development of artificially created cells called hybridomas, which are the fusion of two types of precursor cells: an antibody producing B-cell and a cancer cell. Cancer cells have the qualities of continuous growth and reproduction; this makes them fatal in the context of the body but allows them to grow continuously in an artificial laboratory setting where other, better regulated cells would reach senescence and die. Antibody producing cells each synthesize and secrete a single type of antibody, so that a hybridoma fusion of the two cell types continuously grows and produces one particular antibody in the laboratory. The combination of the qualities of the two

14. See Bruce Boyden, *Emergent Works*, 39 COLUM. J.L. & ARTS 377, 389 (2016).

15. See Calo, *supra* note 2, at 405.

16. See Georges Köhler & César Milstein, *Continuous Cultures of Fused Cells Secreting Antibody of Predefined Specificity*, 256 NATURE 495 (1975).

precursor cells creates a molecular factory that neither precursor could supply alone.

However, the exact type of antibody produced by the fusion of such cells was largely unforeseeable; some fusions work, others do not, and the process of cell fusion is designed to generate large numbers of hybridomas in the hope of finding one that secretes an antibody of interest. Among a sufficiently large number of fusions, an inventor can locate one or more with the desired qualities. Indeed, landmark patent cases involving hybridoma technologies often concerned the need to make hundreds of cell fusions before a hybridoma producing a desirable antibody product was obtained.¹⁷ The exact antibody product from any given fusion was unforeseeable, unpredictable, and one might say “emergent.” But those of skill in the art were able to identify antibodies with the proper characteristics once they were produced, and their emergentness was never considered an impediment to patenting.

II. ENCOMPASSING AI INNOVATION

The supposed impediments to patenting arising from the perceived autonomy and emergentness of AI outputs have encompassed several different aspects of patent doctrine. Inventorship is the perennial favorite topic of AI commentators, but non-obviousness, disclosure, and infringement have been considered as well. It is worth briefly reviewing and correcting the concerns arising in each of these areas, not only to rectify the misunderstandings that are rife in the current literature, but to identify certain common threads among the doctrinal concerns. Looking for such commonalities allows us to triangulate on the actual, legitimate patent issues posed by AI innovation.

A. INVENTORSHIP

The most appropriate doctrinal interrogation to begin our tour of AI patenting is undoubtedly the question of inventorship. A number of commentators have seized upon the emergent qualities of AI systems to question whether innovations generated using such systems can have a human inventor, and whether the AI itself should perhaps be considered the inventor.¹⁸ Indeed, the European Patent Office, the

17. See *In re Wands*, 858 F.2d 731 (Fed. Cir. 1988).

18. See Ebrahim, *supra* note 5, at 308; W. Michael Schuster, *Artificial Intelligence and Patent Ownership*, 75 WASH. & LEE L. REV. 1945 (2018); Liza Vertinsky, *Patents and Thinking Machines*, in RESEARCH HANDBOOK ON LAW AND ARTIFICIAL INTELLIGENCE 489 (Woodrow Barfield & Ugo Pagallo, eds., 2018); Yanisky Ravid & Liu, *supra* note 5; Ryan

United Kingdom Intellectual Property Office, and the United States Patent Office have all recently rejected test applications, bordering on publicity stunts, in which the claimed invention was credited to an AI system as inventor.¹⁹

This outcome is a fairly straightforward conclusion from patent doctrine, particularly in the United States. In many jurisdictions, the jurisprudence of inventorship is fairly sparse, presumably because it was never a question of much salience. Historically, most jurisdictions have granted patents on the basis of application filing date priority, without much inquiry into the process of invention.²⁰ But because the United States long granted patents on priority of invention, rather than application, American patent jurisprudence has a rich and robust fund of doctrine defining invention and inventorship. Even though U.S. patent law has recently moved to a modified first-to-file system, the concept of inventorship as developed over the previous 200 years remains deeply embedded in the current system, and offers fairly distinct (and resoundingly negative) answers on the question of AI inventorship.

The core of inventorship in American patent law has been conception of the invention. Conception in American patent law is classically defined as “the formation in the mind of the inventor, of a definite and permanent idea of the complete and operative invention, as it is thereafter to be applied in practice.”²¹ Indeed, conception of the invention is definitive of the act of invention. American patent law lauds and rewards the mental work of conception; once conception of the invention is completed, all that remains is to reduce the conception to practice. Reduction to practice is not the work of an inventor, it is rather the work of the “mere artisan” which can be done without inventive skill.²²

Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. REV. 1079 (2016); Liza Vertinsky & Todd M. Rice, *Thinking About Thinking Machines: Implications of Machine Inventors for Patent Law*, 8 B.U. J. SCI & TECH. L. 574, 581 (2002). See also Ben Hattenbach & Joshua Glucoft, *Patents in an Era of Infinite Monkeys and Artificial Intelligence*, 19 STAN. TECH. L. REV. 32 (2015) (considering inventorship of AI-generated claim texts).

19. James Nurton, *EPO and UKIPO Refuse AI-Invented Patent Applications*, IP WATCHDOG (Jan. 7, 2020), <https://www.ipwatchdog.com/2020/01/07/epo-ukipo-refuse-ai-invented-patent-applications/id=117648/> [<https://perma.cc/WFJ7-KC44>].

20. See Dan L. Burk, *From ‘First to Invent’ to ‘First to File’ – Changing Lanes in U.S. Patent Procedure?*, 42 IIC – INT’L. REV. INT. PROP. & COMPETITION L. 627 (2011).

21. *Townsend v. Smith*, 36 F.2d 292, 295 (C.C.P.A. 1929).

22. *Gunter v. Stream*, 573 F.2d 77 (C.C.P.A. 1978).

From a formalist perspective, the exclusion of AI systems, or any other automated system, from this definition of invention is very straightforward. Because the machine has no mind in which an idea of the complete and operative invention can be conceived, the machine by definition cannot be an inventor. Stated differently, the device that generates the innovative output has no mind in which conception can occur; if there is no conception, there can be no invention. One possible outcome could be that the artifact generated by automation has no inventor for patent purposes but is merely the product of physical circumstance. A novel molecule or device or process comes into existence as the result of AI machinations, but is never conceived, and so is ineligible for a patent.

This would not necessarily be a bad outcome from a policy perspective; patents are well understood to constitute not only a restraint on trade but a restraint on innovation. We tolerate patent constraints in the hope that imposing exclusivity costs in the short run will produce greater social welfare in the long run. As many others have observed, despite copious speculation around this hope there is to date little evidence that this trade-off in fact occurs, and increasingly strong evidence that it does not occur robustly or reliably across all technical sectors.²³ Where patents are not clearly providing a benefit, the justification for their costs is thin. In particular, there may be good reason to allow technical developments to pass immediately into the public domain if they are not developments that would have or could have been furthered by the promise of exclusive rights. Machines are not incented to innovation by pecuniary incentives, and arguably neither are human innovation initiators whose actions are sufficiently divorced from the outcomes of those mechanical actions.

However, existing patent law already has the formalist quandary well in hand. As an initial matter, patent doctrine has never hesitated to bestow inventorship, and subsequently patents, on serendipitous discoveries.²⁴ There is no labor or “sweat of the brow” rule to obtain a patent. Invention that occurs by accident or happenstance, and could not have been foreseen before actual instantiation of the invention, still merits the reward of a patent. Accidents and surprises receive the same reward as intention and deliberation. Inventors who expected a different outcome, or no outcome, are entitled to the same rights as inventors who meticulously planned and executed their research program. Inventors who generated a novel, useful, and nonobvious device

23. Mark A. Lemley, *Faith-Based Intellectual Property*, 62 UCLA L. REV. 1328 (2015).

24. Sean B. Seymore, *Serendipity*, 88 N.C. L. REV. 185, 190 (2009).

on their first try, with little or no work, receive the same consideration as inventors who succeeded only after laborious effort.

Serendipitous or unforeseen inventions qualify as inventions due to the doctrine of simultaneous conception and reduction to practice. In instances where the inventor does not or cannot imagine the form of the invention in advance of actually having reduced the invention to practice, patent law has held that the invention is conceived when it is recognized by the inventor.²⁵ This is a common occurrence in technologies such as in the case of “emergent” chemical or biotechnological systems mentioned above. New molecules are generated without explicit design or formulation by the inventor. At the moment it is perceived by a human, the invention is conceived in the mind of the inventor, although the molecule has already been previously reduced to practice, that is, physically instantiated.

There has of course never been any question of declaring a hybridoma or microorganism that produces a novel and non-obvious molecule an “inventor” or even a “co-inventor” of the substance it secretes, despite the fact that the molecule was unforeseen by the human who subsequently claims it. To the contrary, it is the human who recognizes the desirable qualities of the substances produced; at that moment the invention is conceived and qualifies for patenting. The same principle clearly applies to molecules or devices generated by machine learning systems; the output becomes an invention after its generation when perceived by the human operator or investigator, who comprehends its nature and use. Thus, there is no need to speculate or fantasize as to whether the AI is or should be considered an inventor.

B. OBVIOUSNESS

The patentability criterion of non-obviousness has also garnered considerable attention from AI commentators. Obviousness in American patent law is judged against the knowledge of a hypothetical person having ordinary skill in the art, or PHOSITA.²⁶ If the PHOSITA would have found the claimed invention obvious as of the date the patent application was filed, the invention is considered insufficiently innovative to constitute a patentable advance over the prior art, and no patent should issue.²⁷ Certain commentators suggest that this standard is thrown into doubt by the ability of AI systems to rapidly

25. *Amgen, Inc. v. Chugai Pharm. Co., Ltd.*, 927 F.2d 1200, 1206 (Fed. Cir. 1991).

26. 35 U.S.C. § 103 (2012).

27. *Graham v. John Deere Co.*, 383 U.S. 1 (1966).

generate new constructs in the chemical, mechanical, and electrical arts.²⁸ The concern regarding AI is that a PHOSITA using machine learning systems can rapidly and easily create multitudes of innovations, meaning that the PHOSITA can in a sense foresee, or at least easily generate all potential advances, rendering “everything obvious.”²⁹

We should note at the outset that if AI in fact assured such an outcome, this would be cause for celebration. The obviousness doctrine in patent law is well understood to constitute a policy lever that aligns risk with reward.³⁰ The more risky or uncertain a particular line of research, the less likelihood there is of success in that endeavor, the greater reward is necessary to induce an innovator to pursue that line of research.³¹ The non-obviousness doctrine ensures that technical achievement in the face of uncertainty, where there is no assurance of success, is more likely to merit the exclusive rights of a patent.³² If AI systems indeed made “everything obvious,” assuring technical success without risk, we could do away with the cumbersome and costly mechanism of patent rewards, and enjoy the utopia of research certainty that AI ushered in.

Of course, AI does no such thing; imagining that it does so is a symptom of the “magical worldview.”³³ Predictive analytical systems are not Borge’s fictional Library of Babel, encompassing the totality of knowledge.³⁴ To the contrary, machine learning systems find only what humans design them to find, within statistical parameters that humans must specify.³⁵ Indeed, AI outputs are so copious and non-discriminating that humans must specify which algorithmic outcomes are sufficiently “interesting” to merit inclusion in the pool of viable

28. See Brenda M. Simon, *The Implications of Technological Advancement for Obviousness*, 19 MICH. TELECOMM. & TECH. L. REV. 331 (2013).

29. See Abbott, *supra* note 5; see also Ebrahim, *supra* note 5, at 309 (speculating on the impact of AI technologies on non-obviousness).

30. DAN L. BURK & MARK A. LEMLEY, *THE PATENT CRISIS AND HOW THE COURTS CAN SOLVE IT* 131–32 (2009).

31. See Robert P. Merges, *Uncertainty and the Standard of Patentability*, 7 HIGH TECH. L.J. 1 (1992).

32. *Id.*

33. See *supra* note 3 and sources cited therein.

34. See Jorge Luis Borges, *The Library of Babel* in *LABYRINTHS: SELECTED STORIES AND OTHER WRITINGS* 51 (James E. Irby & Donald A. Yates eds., tr. James E. Irby, 1962). Borges’ fictional library contained every variation on a text, meaning that it not only contained all possible coherent texts, but a nearly infinite number of gibberish texts, and so required (much like predictive analytics) sentient human curation to locate anything meaningful.

35. TREVOR HASTIE, ROBERT TIBSHIRANI & JEROME FRIEDMAN, *THE ELEMENTS OF STATISTICAL LEARNING: DATA MINING, INFERENCE, AND PREDICTION* 38 (2d ed. 2009).

results.³⁶ Moreover, innovation is not simply invention, but the full set of investments to develop a new product. Generating new inventions in many fields, such as pharmaceuticals, is already a trivial exercise; deciding which to pursue for development remains the primary source of innovative uncertainty.

As with any human technology, the availability of new tools—whether a screwdriver, an ultracentrifuge, or a predictive algorithm—will make some inventive tasks easier than they were before the particular tool was available. Inventions that would have been impossible become possible with new foundational implements. This is quintessentially a sign of the “progress of the useful arts” that patents are intended to encourage.³⁷ But the ease or difficulty of producing a technological advance has never been the question for determining patentable obviousness—recall that serendipitous, accidental, even effortless inventions may be patented.³⁸ The question for obviousness is rather whether the particular outcome could have been foreseen, predicted, or routinely executed by those of ordinary skill in that technology.

The confusion on this point is in part semantic. Obviousness has been defined by the courts as whether the PHOSITA would have a “reasonable expectation of success” in obtaining the claimed invention as of the critical date.³⁹ As a reciprocal proposition, courts have often observed that patent obviousness was not synonymous with being “obvious to try” a particular inventive combination of elements.⁴⁰ Originating in the chemical arts, the phrase “obvious to try” indicated that unexpected results could come out of ostensibly obvious combinations in chemistry and the other “unpredictable arts” and so merit a patent.⁴¹ But over time the United States Court of Appeals for the Federal Circuit turned this logic on its head, turning the phrase “obvious to try” into a mantra for permissive standards of non-obviousness.⁴² Ultimately, the Supreme Court in *KSR Int’l Co. v. Teleflex Inc.* intervened to point out that some inventive steps that are “obvious to try” do indeed result in obvious outcomes.⁴³ Specifically, the Court

36. Luciana Parisi, *Critical Computation: Digital Automata and General Artificial Thinking*, 36 THEORY, CULTURE & SOC’Y 89, 109–10 (2019); MacKenzie, *supra* note 13, at 438.

37. U.S. CONST. art. I, § 8, cl. 8.

38. See Seymore, *supra* note 24, at 190.

39. *In re O’Farrell*, 853 F.2d 894, 903 (Fed. Cir. 1988).

40. Merges, *supra* note 31, at 42.

41. *In re Tomlinson*, 363 F.2d 928, 931 (C.C.P.A. 1966).

42. *In re Kubin* 561 F.3d 1351, 1358–59 (Fed. Cir. 2009).

43. 550 U.S. 398, 421 (2007).

observed that if trying a finite number of solutions to a commonly encountered problem results in anticipated success, the outcome is not innovative, but likely the product of “ordinary skill and common sense.”⁴⁴

Given the vast exploratory power of AI systems, it is tempting to conclude that the employment of such systems in research inevitably yields a “reasonable expectation of success.” The AI system can parse vast datasets that are incomprehensible to the unaided human mind, can detect patterns in the data that are invisible to unaided human analysis, and can predict previously unimagined outcomes that fit to the patterns that it finds. Use of such systems will become routine, and common sensical, and in hindsight any innovation generated by an AI tool may seem to have been obvious in the sense that it was “obvious to try finding a solution via an AI system.”

But this semantically loose application of the language of obviousness misapprehends the nature of the obviousness standard. “Obvious to try using an AI system” is not necessarily synonymous with statutory patent obviousness. Specifically, the Federal Circuit has opined that, under *KSR*, an innovative outcome could be “obvious to try” but not obvious under the patent statute in situations where the discovery was arrived at by varying all parameters or trying each of numerous possible choices without guidance or direction in the prior art as to which variations would be successful.⁴⁵ Similarly, the court opined that a non-obvious result might emerge from exploring a new technology or field where the prior art gives only general guidance as to or how to obtain the claimed result.⁴⁶ Although the court did not have AI tools in mind when advancing those definitions of non-obvious but obvious to try research, they fit astonishingly well to the uses of AI in innovative development—indeed, the former situation envisioned by the Federal Circuit literally describes the functioning of machine learning AIs.⁴⁷

The standard for obviousness is not now, nor has ever been gauged by the expectation that one of ordinary skill could expect to find *some* solution to the problem faced by the inventor. The standard is rather whether one of ordinary skill would as of the date of filing the application have found *the* solution the inventor found—the particularly claimed invention. As in the hybridoma example described

44. *Id.*

45. *In re O’Farrell*, 853 F.2d 894, 903 (Fed. Cir. 1988).

46. *Id.* See also *Kubin* 561 F.3d at 1359 (endorsing the “obvious to try” distinctions articulated in *O’Farrell*).

47. Mackenzie, *supra* note 13, at 435; Elish & boyd, *supra* note 3, at 16.

above, the question is not whether those of skill in the art could obtain some monoclonal antibody, but whether the prior art would have rendered obvious the one actually obtained. The question in AI assisted innovation will not be whether the person of ordinary skill would have reasonably anticipated some type of success; the question will be whether the person of ordinary skill would have anticipated the type of success ultimately claimed in the patent. Within the universe of solutions available at the time the patent application was filed, some solutions might be predictable and so legally obvious, but others may be unexpected or unpredictable, and potentially non-obvious. The question is always which category the claimed invention would have fallen into, not whether the method employed to obtain it was widely or routinely available.⁴⁸

C. DISCLOSURE AND CLAIMING

The patent statute requires disclosure of the invention by means of a written description that is sufficiently detailed to “enable” one of ordinary skill to make and use the claimed invention.⁴⁹ This requirement has been suggested to pose an impediment either for patents claiming AI systems or for patents claiming inventions that engage AI systems as a means of production.⁵⁰ Machine learning systems are frequently characterized by “black box” inscrutability of their constitutive processes. Because such systems iteratively re-program themselves, altering their own code and characteristics toward a particular result, their operations and inner workings may sometimes be opaque or incomprehensible to experts in the field. The ability of the system to essentially modify and re-program itself without requiring explicit coded instructions from a human programmer is considered to be an important advantage of machine learning.⁵¹ But this also is said to create a problem for patent disclosure; if the internal characteristics of the AI “black box” cannot be described, it may not be amenable to the enablement and written description disclosure requirements.

However, assuming that such inscrutability is indeed a fair characterization of AI systems, it need not preclude patenting. This is a

48. See *In re Durden*, 763 F.2d 1406, 1410–11 (Fed. Cir. 1985) (holding that the obviousness of a process and the obviousness of the product of that process are separate inquiries that must be decided on a case by case basis); see also Dan L. Burk, *The Problem of Process in Biotechnology*, 43 HOUSTON L. REV. 561, 572–75 (2006) (discussing the “Durden problem” in biotechnology).

49. 35 U.S.C. § 112(a) (2012).

50. Ebrahim, *supra* note 5, at 310.

51. Parisi, *supra* note 36, at 92 (noting that machine learning is in effect the inverse of computer programming).

problem that the patent system has confronted before, again in the biotechnology context, when accommodating patent disclosure of innovative genetically modified organisms or altered cell lines. Like the current crop of AI systems, biological organisms constitute inventions whose internal processes are often poorly characterized or are not fully understood. They are essentially “black boxes,” producing valuable biologics such as antibodies or hormonal proteins via complex metabolic processes that are largely unknown and opaque. Hybridoma cell fusions, for example, continuously produce monoclonal antibodies that are directed to a particular molecular target or “epitope” — the exact parameters of the targeted epitope, and the affinity or binding qualities of the molecule are the result of fortuitous molecular combinations that might be difficult to reproduce. Similarly, many biological processes required starting materials that were uniquely synthesized or difficult to obtain, making it difficult to enable their future use from a patent description.

The advent of such biological innovation during the 1980s presented the patent system with the quandary of how to satisfy enablement for biomaterials when the inventor could not effectively access essential components of the invention; it seemed impossible to describe in a textual exposition how to make and use such unique or incomprehensible cells, microorganisms, or starting materials. This quandary was solved by permitting public deposit of such materials. Patent applicants can and do put unique, rare, or difficult to describe biomaterials on deposit in an international repository, where they are held for the use of those who might practice the teachings of the patent disclosure.⁵² The patent document directs users to the repository in lieu of a textual explanation of the materials. The patent purposes of enablement, facilitating those of skill to make and use the claimed invention, is effectuated by public access to the necessary materials.

There seems to be in principle no reason that a similar approach could not be taken with “black box” algorithms that are difficult to describe or to replicate. If the AI system cannot be readily described in the text of the patent document, a working system might be placed on deposit where it would be accessible to the public; if this satisfies enablement for biotechnology, it would surely do so for machine learning. As in the case of biological inventions, some type of repository would have to be established with fees for curation and maintenance of the deposits. In many cases, it is likely that the training data used to prime the AI system should be put on deposit as well, as this would

52. *Enzo Biochem, Inc. v. Gen-Probe Inc.*, 323 F.3d 956, 965 (Fed. Cir. 2002); *In re Argoudelis*, 434 F.2d 1390, 1392–93 (C.C.P.A. 1970).

constitute a “starting material” for producing the resulting device. Indeed, there may be cases where the difficulty in replicating the claimed AI system is a difficulty in replicating the training data set, in which case perhaps only the training data need be put on public deposit.

In addition to an enabling description, the patent statute also requires that a patent applicant provide claims that “distinctly point out and particularly claim” the relevant features of the invention.⁵³ Textually claiming novel and nonobvious AI systems might again seem to present a challenge if the characteristics of the system constitute an opaque or impenetrable “black box.” It might seem that if the characteristics of the AI system cannot be so described, claiming the system may be impossible. But patent law has, yet again, long encompassed claims to inventions that elude capture in a detailed textual claim. This problem has arisen before, in areas such as biotechnology or materials science, but patent practice has evolved to accommodate the inventor of such materials. In such cases, the claimant may employ “product by process” claiming, which specifies a particular product in terms of the process that produces it, rather than in terms of the product itself. AI systems that have opaque characteristics might be claimed in this fashion, relying on the process of designing and training the AI as the characteristic by which it is best described.

D. INFRINGEMENT

Commentators examining patent issues related to AI systems have also tended to fret about problems of infringement. The concern again stems from the supposedly emergent and autonomous nature of AI innovation; perhaps an AI acting autonomously might infringe existing patents, by generating or implementing infringing technology, without the knowledge or even the awareness of humans who deployed the AI. Direct patent infringement is generally acknowledged to be a strict liability offense,⁵⁴ requiring no mental state or scienter for a violation, merely the act of making, using, selling, offering for sale, or importing the claimed invention.⁵⁵ If the AI is acting truly autonomously, this seems to make the AI a patent infringer. If the AI’s activities are attributable to a human actor, then it seems problematic or unfair that the human actor could become an unwitting infringer based on the outputs of an unfathomable “black box.”

53. 35 U.S.C. § 112 (b) (2012).

54. See Robert P. Merges, *A Few Kind Words for Absolute Infringement Liability in Patent Law*, 31 BERKELEY TECH. L.J. 1, 3 (2016).

55. 35 U.S.C. § 271(a) (2012).

Once again, this is not a new or unusual problem for the patent system, which has previously encountered the problem of inventions that appear to autonomously infringe other patents—specifically, in the case of living organisms that grow, reproduce, and mutate into forms that might unexpectedly infringe existing claims. The secreted products or offspring of biological organisms might well read on the claims of existing enforceable patents, but this does not make the generative *organisms* infringers. Under the proper circumstances, their human owners or users might be infringers. Indeed, the question was directly confronted by the Supreme Court in its opinion dealing with infringement of genetically modified plants, *Bowman v. Monsanto*.⁵⁶

The genetically modified plants in *Bowman* were covered by patents owned by Monsanto. Bowman, the alleged infringer, was growing the patented plants from seed on his farm. Bowman had purchased the seeds he planted from a local grain elevator, meaning that he was not in privity of contract with Monsanto, having not purchased anything from them.⁵⁷ Additionally, Monsanto's exclusive rights covering the use of Bowman's seeds were exhausted by the prior authorized sale of the seeds to the grain elevator. Although Bowman had obtained seeds free of either contractual or patent encumbrance, Monsanto sued Bowman for having *made* new seeds by the process of cultivation—Monsanto argued that even if its exclusive rights in the seeds that were planted was exhausted, its exclusive patent right to *make* the subsequent generations of the claimed invention arising out of those seeds was not exhausted by the sale of the progenitor seeds.⁵⁸

Bowman argued that he could not be an infringer, having not “made” the newly grown GMO seeds—rather, he argued that the seeds constituted a self-replicating invention, which had essentially “made” themselves.⁵⁹ This argument was summarily brushed aside by Justice Kagan, writing for a unanimous Court, on grounds of causality.⁶⁰ Although the Court did not couch its refutation in those terms, it observed that the plants were by no means spontaneously or autonomously grown. Bowman was responsible for the gestation and maturation of the patented seeds by virtue of having planted, watered, and cultivated the plants that produced them; consequently, Bowman

56. 569 U.S. 278 (2013).

57. *See id.* at 282.

58. *Id.* at 284–85.

59. *Id.* at 288.

60. *Id.* at 288–89. Justice Kagan referred to Bowman's unsuccessful attempt to avoid responsibility as “the blame-the-bean defense.” *Id.* at 288.

was responsible for having made and used the claimed invention under the statute, in violation of the patent holder's exclusive rights.⁶¹

Clearly the same will be true for AI systems—they do not appear spontaneously; they do not train themselves; they do not supply their own electrical power, turn themselves on, or generate any kind of output without human initiation. The fact that the output may be unexpected or unpredictable does not remove human causality from their operation. Humans may be unwitting infringers, whether the infringement occurs directly by their own hand or indirectly through planting a seed or activating a device. As the *Bowman* case demonstrates, a human who engages a device that infringes an existing patent need not intend or even know that a patent will be infringed—direct patent infringement is a strict liability offense.

To the extent that the imposition of inadvertent liability seems problematic, the likelihood of accidental infringement liability is limited. As a practical matter, the patent holder probably will not notice or care about inadvertent infringement unless the infringement becomes commercial or widespread. Additionally, a number of commentators have suggested that a scienter requirement either exists inherently in the patent infringement statute, or could easily be grafted on to avoid injustice in appropriate situations.⁶² And, of course, although it is quite possible that existing emergent technologies, in the form of living organisms, might infringe without human intervention—by mutating, or producing outputs that fall within existing patent claims—there are essentially no practical consequences. If in fact infringement occurs with no human involvement, in the wild for example, then there is simply no infringer. Such spontaneous infringement may well occur with some frequency, but no one notices or cares.

The involvement of AI systems as generative tools changes nothing—humans who use screwdrivers, soldering irons, centrifuges, or HPLC columns may also unwittingly infringe existing patents. This likely goes unremarked and unnoticed unless the activity progresses to a commercial stage that draws the attention of the relevant patent owner. AI system users will need to be cognizant of their freedom to operate, and of the patents their activity may potentially infringe just as any other innovators must be.

61. *Id.* at 289.

62. See Stephen Munzer, *Plants, Torts, and Intellectual Property*, in *PROPERTIES OF LAW: ESSAYS IN HONOR OF JIM HARRIS* (Timothy Endicott, Joshua Getzler & Edwin Peel, eds., 2006); Paul J. Heald & James C. Smith, *The Problem of Social Cost in a Genetically Modified Age*, 58 *HASTINGS L.J.* 87 (2006).

III. COMPREHENDING CAUSATION

The question of inadvertent patent infringement, and the holding of causal responsibility in the *Bowman* case, brings us effectively full circle in our consideration of AI patenting, back to the questions confronted regarding AI inventorship. We began by considering assignment of responsibility for innovation generated by means of AI, and in discussing of the *Bowman* opinion we end by considering responsibility for infringement committed by means of an AI. In each case the supposed autonomy of the device in operation seems to pose a conundrum for the patent system, but in each case closer consideration reveals human operators to whom the law applies, negating the question of whether the law can or should apply to a device deployed by humans.

Excluding the human causal origins of AI activity is tempting because machine learning systems are inventions that can be used to generate other inventions—possibly novel inventions that might be deserving of a patent, or non-novel inventions that infringe on existing patents. But this quality is in fact not unusual. As the *Bowman* case illustrates, the technology of self-assembling and self-replicating machines is one of the oldest human technologies, forming the foundation for all human civilizations—by this I refer to human cultivation and breeding of other living organisms as crops and livestock. The fact that the crops in *Bowman* made new copies of a patented invention did not absolve their human overseer from infringement responsibility. Neither is the act of novel assembly, or even the occurrence of design that is implemented in assembly, evidence of invention. Viruses routinely self-assemble themselves and reconfigure their own assembly instructions. But viruses are not inventors.

Asserting that AI tools are either inventors or infringers is equally absurd and can only be based on ignoring the human hand at work behind the AI. Such algorithms, far from operating consciously or even autonomously, constitute tools that are deployed by researchers in order to design and develop new inventions. Humans routinely use tools to construct other tools; manipulation of a simple screwdriver, a hammer, or a drill makes possible the construction of follow-on devices that could not have been built or devised without such aids. AI systems are the latest development on this trajectory, in the form of machines that build other machines; possibly replicating themselves or possibly designing and building other machines.

The degree of automation incorporated into such machines by their designers does not change this fundamental relationship. Humans are in essence cyborgs, adopting and wielding technical

prostheses to extend the capabilities of their hands, senses, and minds.⁶³ A hammer or drill operates in tandem with the human hand and eye to construct other artifacts, many of which are themselves also tools. So, too, the statistical optimization systems we have dubbed “machine learning” or “AI” systems leverage rapid digital processing, cheap data storage, and massive data sets to extend human cognitive capabilities—searching for correlative patterns that humans unaided might not find at all, or likely would not find within a human lifespan. There is nothing new in the use of tools to design, develop, or construct other human artifacts, and no one has ever endorsed the bizarre notion that a screwdriver, a power drill, or an arc welder is either singly or in combination with a human the “inventor” of an artifact that it was used to construct.

One might object that AI systems differ substantially from screwdrivers or other simple tools in their autonomy; a screwdriver clearly needs a human hand to wield it while an AI device can operate on its own, independently. But one sees such autonomous operation, or apparent autonomous operation, in any number of modern industrial settings. For example, automobile construction and other factory operations are routinely conducted by robotic systems that assemble, solder, weld, drill, and fasten the components of industrial products.⁶⁴ The functions of such robots appear to be executed autonomously. No human is obviously controlling them, the way a human hand must obviously grip and control a standard screwdriver. But no one would seriously argue that the robotic mechanisms of a factory floor are “inventors” of the automobiles or other products they fabricate, nor would they be “infringers” of patents whose claims read on the assembly line products, despite the apparent autonomy of the assembly robots.

This is because the apparent autonomy of the assembly line is largely a fallacy of framing. Factory robots appear autonomous in isolation, just as a screwdriver or drill would seem autonomous if one (rather absurdly) were to ignore the human hand holding it. Viewing

63. See Clive Lawson, *Technology and the Extension of Human Capabilities*, 40 J. THEORY SOC. BEHAVIOR 207 (2010) (reviewing arguments regarding technology as amplifying or supplementing bodily functions); see also Donna Haraway, *The Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century* in SIMIANS, CYBORGS, AND WOMEN: THE REINVENTION OF NATURE 149, 150 (1991) (“The cyborg is a condensed image of both imagination and material reality, the two joined centers structuring any possibility of historical transformation.”).

64. See MIKE WILSON, IMPLEMENTATION OF ROBOT SYSTEMS: AN INTRODUCTION TO ROBOTICS, AUTOMATION, AND SUCCESSFUL SYSTEMS INTEGRATION IN MANUFACTURING 75–100 (2015) (describing typical operations of factory robots).

the factory robots in a narrow frame gives the illusion that they are operating autonomously, but only because consideration of the machines is so myopic as to overlook the extensive human organization behind the machines' operation. The human control, input, and supervision of factory robotics is not apparent on the factory floor itself but becomes clear when the frame of reference is expanded to include the array of activities occurring at a spatial distance, off the manufacturing floor, or at a temporal distance, before or after the occurrence of the robotic action.

The manifest human guidance behind a factory assembly robot leaves no reason to believe that a machine that assembles other machines is an inventor, or an infringer, but anything other than a tool of the inventing or infringing humans who employed it. Yet when it comes to machine learning systems, the objection that might be raised against this logic is that AIs have "emergent" qualities, that their output is not pre-determined or specified in advance in the way that the output of an automobile factory robot would be. The activity of the factory floor robot is programmed and known (if it is not, then presumably something has gone terribly wrong). AI systems, on the other hand, are portrayed as developing unexpected and undetermined outputs, at least within the parameters specified in their design. For example, pharmaceutical design AIs may generate molecular structures and species that are unforeseen by the humans who set them in motion.

But, as discussed before with reference to biotechnology, patent law is already well familiar with biological machines that engage "emergent" processes and that assemble other devices—perhaps not fully autonomously, but by means of "black box" processes that humans do not fully understand. A machine that generates unforeseen outputs by design is by no means disengaged from human innovation; to the contrary, it is entirely the tool of human innovation. The specifics of an AI's output, or that of a hybridoma, may be unforeseen, but the general expectation of some output, and some output within a particular field, was surely intended in the machine's deployment by humans.

Thus, the apparent creativity of complex automated systems is an illusion of their opacity, and their apparent autonomy is an illusion of causative distance. As Carys Craig and Ian Kerr remind us in the copyright context, no AI system in fact functions autonomously.⁶⁵ There

65. Ian Kerr & Carys Craig, *The Death of the AI Author*, 52 OTTAWA L. REV. 31 (2021).

is always a vast array of human designers, programmers, trainers, and users behind any AI. Someone maintains the system, updates its software, makes certain it is plugged in to the electrical grid. Conceptually erasing all of those human personnel from the innovative process mistakenly valorizes and romanticizes the machine.⁶⁶

In the case of the screwdriver or the drill, human direction and control are typically manifest and proximate. The human hand may be less apparent in the case of machines that are used to tool or assemble other machines. Over the past several decades, human intervention has become causally attenuated, so that the human prime mover may appear to have disappeared from the system. As mechanisms become more complex, human causation becomes more distant, but not necessarily less proximate—there is no obvious human agency directing the robots that assemble automobiles on a factory floor, yet we have no doubt that a human agency arranged, orchestrated, and initiated their activity. Human involvement in the chemical sciences is in some sense even more attenuated; human manipulation is not present or required in each reactive step. Rather, humans create the entropic conditions for reactions, and having initiated them, allow them to proceed. Biotechnological inventions encompass vast arrays of such reactions.

Patent inventorship and patent infringement may therefore be regarded as questions of causality—how far up the chain of causality must one look to find the human initiator? It may be that in some cases the human is sufficiently far removed that there is no inventor or infringer; if Bowman receives as a gift patented seeds that are caught by the wind and grow as sports in his field, perhaps he is no longer the proximate cause of making the claimed invention even if he is a cause in fact. Such matters of causality are old and familiar (if perhaps intractable) problems across the domains of law, spanning criminal culpability, tort liability, and a host of other areas.⁶⁷ Infringement causality is clearly undertheorized in patent law, as some commentators such as Amy Landers and Dmitry Karshtedt have begun to recognize.⁶⁸ Those cracks in patent causality have begun to show in recent

66. *Id.*

67. *Cf.* Dan L. Burk, 36 *Views of Copyright Authorship*, by Jackson Pollock, 58 HOUSTON L. REV. 263 (2020) (reviewing common causality questions across legal specialties).

68. *See* Amy Landers, *Proximate Cause and Patent Law*, 25 B.U. J. SCI & TECH L. 330 (2019); Dmitry Karshtedt, *Causal Responsibility and Patent Infringement*, 70 VAND. L. REV. 565 (2017).

decisions involving multiple actors and divided infringement, and will perhaps be salient in some future scenarios involving AIs.

But the law of patent *invention* has never had a causality requirement, at least not a physical causality requirement, because physical reduction to practice has never been a requirement for patentable invention. Quite the contrary, patent law recognizes “constructive reduction to practice” in the filing of an enabling disclosure in a patent application.⁶⁹ Such “paper patents” disclose and claim inventions that have never been physically built or deployed; they exist only as the textual descriptions supplied by the inventor. The supposition behind this practice is that a detailed and enabling description is as good as a physical embodiment; once the mental work of invention has been done, those of ordinary skill can do the mundane work of implementation.⁷⁰

The result is that doctrines such as simultaneous conception and reduction to practice provide immediate answers to concerns regarding AI generated inventions, although such answers may not be fully satisfying. If self-assembling machines become sufficiently attenuated from human prime movers, we may no longer deem the human to be a proximate cause of the invention’s reduction to practice, but recognition of the invention in the mind of a human will always yield an “inventor.” Where infringement is concerned, attenuation between machine and human may reach a point that we are reluctant to assign liability to the human mover. But assessing when and where that point occurs requires a theory of infringement causation that patent law currently lacks. The commentary on AI innovation thus reveals a gap in patent doctrine, but not the gap that AI commentators thought they had identified.

CONCLUSION

AI systems promise to enable a host of technical advances, enormously enhancing human innovative capabilities. Patent law is fully equipped to encompass such innovative activity, as it has done in the past. The qualities of AI innovation that have been said to challenge the tenets of patenting are based on a category mistake, confusing the tool with the user. No one would seriously assert that the drill or screwdriver—much less the hydrocarbon cracking refinery or chromatographic column that is used to construct an invention—is an

69. *Burroughs Wellcome Co. v. Barr Labs, Inc.*, 40 F.3d 1220, 1228 (Fed. Cir. 1994).

70. *Townsend v. Smith*, 36 F.2d 292, 295 (C.C.P.A. 1929); *Feldman v. Aunstrup*, 517 F.2d 1351, 1355 (C.C.P.A. 1975).

inventor, or even a co-inventor with the human deploying the equipment. Neither would anyone seriously assert that these devices are patent infringers if their outputs happen to coincide with the claims of an existing patent. The fact that recent commentary assigns such status to our latest and most complex tools—AI systems—indicates not some transcendent property of AI technology, but rather that the patent system has a latent deficit in doctrines of causality. Rather than continued hysteria over the emergent properties of AI, focusing on patent law's unfinished business with causality offers the more fruitful path forward.