AI-Generated Inventions: Implications for the Patent System

Gaétan de Rassenfosse
Adam Jaffe
Melissa Wasserman

May 2023

Innovation and Intellectual Property Policy
Working Paper series no. 22

Available at: https://ideas.repec.org/p/iip/wpaper/22.html

EPFL Working Paper Series
STiP lab

AI-GENERATED INVENTIONS: IMPLICATIONS FOR THE PATENT SYSTEM

GAÉTAN DE RASSENFOSSE[*]

ADAM B. JAFFE[†]

MELISSA WASSERMAN[‡]

This symposium Article discusses issues raised for patent processes and policy created by inventions generated by artificial intelligence ("AI"). The Article begins by examining the normative desirability of allowing patents on AI-generated inventions. While it is unclear whether patent protection is needed to incentivize the creation of AI-generated inventions, a stronger case can be made that AI-generated inventions should be patent eligible to encourage the commercialization and technology transfer of AI-generated inventions. Next, the Article examines how the emergence of AI inventions will alter patentability standards, and whether a differentiated patent system that treats AI-generated inventions differently from human-generated inventions is normatively desirable. This Article concludes by considering the

*. Associate Professor, College of Management of Technology, Ecole polytechnique fédérale de Lausanne, Switzerland.

†. Professor Emeritus of Economics, Brandeis University; Senior Research Associate, Motu Research, Wellington, New Zealand.

‡. Charles Tilford McCormick Professor of Law, Associate Dean for Research, University of Texas School of Law.

We thank participants at the UK IPO Research Seminar (online), the ISEG Research Seminar at Lisbon School of Economics and Management (Lisbon), the LUISS Law Research Seminar (Rome), the 6th WIPO Conversation on IP and Frontier Technologies: AI inventions (Geneva), the 16th Annual Conference of the EPIP Association (Cambridge), the Competition & IP in digital markets workshop (Brussels), the CIPCO-IGE discussion on AI and IP Law (Zurich), the IP Australia Knowledge creation to utilization Workshop (Canberra), the Oxviews Webinar (online), and the Digital Transformation in Business and Law, Southern California Law Review Symposium (Los Angeles) for useful comments.

larger implications of allowing patents on AI-generated inventions, including changes to the patent examination process, a possible increase in the concentration of patent ownership and patent thickets, and potentially unlimited inventions.

INTRODUCTION

AI-generated inventions—inventions autonomously created by AI software—are around the corner.[1] They have already surfaced in some applications, including genomic.[2] "Invention machines," as we will generically call them, will, in all likelihood, become more prevalent in the future with more and better data, methods, and computers. They will also fundamentally alter the innovation process, with inventions becoming cheaper and faster to produce—at least in some technological fields or for some types of inventions.

If the innovation process changes, so, perhaps, should the support schemes put in place to encourage it. Scholars have traditionally seen innovation activities as needing policy support with tools such as the patent system, grants, research and development ("R&D") tax subsidies, and prizes, among others.[3] It is not clear that the current policy toolbox is well adapted to this changing landscape.

One concrete question that has received a great deal of scholarly attention is whether AI-generated inventions can be protected by patents under existing intellectual property ("IP") laws.[4] The issue also received coverage

---

1. *See* Hiroaki Kitano, *Nobel Turing Challenge: Creating the Engine for Scientific Discovery*, 7 NATURE PARTNER JS.: SYS. BIOLOGY AND APPLICATIONS 1, 1–2 (2021).

2. *See* Ross. D. King, Kenneth E. Whelan, Ffion M. Jones, Philip G. K. Reiser, Christopher H. Bryant, Stephen H. Muggleton, Douglas B. Kell & Stephen G. Oliver, *Functional Genomic Hypothesis Generation and Experimentation by a Robot Scientist*, 427 NATURE 247, 247–51 (2004). Genomics is the study of genes, including interactions of those genes with each other and the environment. WILLIAM KLUG, MICHAEL R. CUMMINGS, CHARLOTTE A. SPENCER, MICHAEL A. PALLADINO & DARRELL KILLIAN, CONCEPTS OF GENETICS 46 (12th ed. 2019).

3. *See* Jakob Edler & Jan Fagerberg, *Innovation Policy: What, Why, and How*, 33 OXFORD REV. ECON. POL'Y 2, 2–6 (2017); Johan Schot & W. Edward Steinmueller, *Three Frames for Innovation Policy: R&D, Systems of Innovation and Transformative Change*, 47 RSCH. POL'Y 1554, 1554–55 (2018); Nicholas Bloom, John Van Reenen & Heidi Williams, *A Toolkit of Policies to Promote Innovation*, 33 J. ECON. PERSPS. 163, 163–65 (2019).

4. *E.g.*, RESEARCH HANDBOOK ON THE LAW OF ARTIFICIAL INTELLIGENCE at 411–537 (Woodrow Barfield & Ugo Pagallo eds., 2018) [hereinafter "RESEARCH HANDBOOK ON LAW OF AI"]; Ryan Abbott, THE REASONABLE ROBOT: ARTIFICIAL INTELLIGENCE AND THE LAW (2020); Marta Duque Lizarralde & Claudia Tapia, *Artificial Intelligence: IP Challenges and Proposed Way Forward*, 2022 PAT. LAW. 16, 16–21 (2022). *See, e.g.*, Dan L. Burk, *AI Patents and the Self-Assembling Machine*, 105 MINN. L. REV. HEADNOTES 301, 301–03 (2021); W. Michael Schuster, *Artificial Intelligence and Patent Ownership*, 75 WASH. & LEE. L. REV. 1945, 1946–52 (2018); Liza Vertinsky, *Thinking Machines and Patent Law*, *in* RESEARCH HANDBOOK ON LAW OF AI, *supra*, at 489; Shlomit Yanisky Ravid & Xiaoqiong (Jackie) Liu, *When Artificial Intelligence Systems Produce Inventions: An Alternative Model for Patent Law at the 3A Era*, 39 CARDOZO L. REV. 2215, 2217 (2018); Ryan Abbott, *I Think,*

from mainstream media when Professor Ryan Abbott's team from the University of Surrey filed patent applications, as part of the Artificial Inventor Project, designating an AI system as the inventor at several patent offices worldwide.[5] The applications were (so far) rejected by some patent offices (including in the United States, the European Patent Office, and the United Kingdom), but accepted by others (including in South Africa and Australia).[6] The issues posed in that case were whether an AI-generated invention can be patented and whether an AI system can be named as an inventor in a patent application. The patentability of AI-generated inventions is also high on the policy agenda, with the main patent offices actively discussing the issue.[7]

However, the question of the patentability of AI-generated inventions under current patent laws is too narrow a framing of the issue. The important question is whether and how the emergence of this new invention technology changes our judgment as to how the patent system can best operate to achieve its objectives. The fundamental aspects of patent laws have barely changed since the 1474 Venetian Patent Statute. Having resisted two industrial revolutions, it is not immediately apparent that the patent system must adapt to the digital revolution. However, whereas the previous industrial revolutions

---

Therefore I Invent: Creative Computers and the Future of Patent Law, 57 B.C. L. REV. 1079, 1079–83 (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); John Villasenor, *Reconceptualizing Conception: Making Room for Artificial Intelligence Inventions*, 39 SANTA CLARA HIGH TECH. L.J. 197, 199–203 (2022); KEMAL BENGI & CHRISTOPHER HEATH, *Patents and Artificial Intelligence Inventions*, *in* INTELLECTUAL PROPERTY LAW AND THE FOURTH INDUSTRIAL REVOLUTION 127, 127–30 (Christopher Heath, Anselm Kamperman Sanders & Anke Moerland eds., 2020).There is also a growing literature addressing whether AI generated work can be protected by copyright. *See, e.g.*, Daniel J. Gervais, *The Machine as Author*, 105 IOWA L. REV. 2053 (2020); Matthew Sag, *The New Legal Landscape for Text Mining and Machine Learning*, 66 J. COPYRIGHT SOC'Y 291, 291–92 (2019).

5. *See, e.g.*, AJ Willingham, *Artificial Intelligence Can't Technically Invent Things, Says Patent Office*, CNN (Apr. 30, 2020, 4:39 AM), https://edition.cnn.com/2020/04/30/us/artificial-intelligence-inventing-patent-office-trnd/index.html [https://perma.cc/625V-FUZK]; Leo Kelion, *AI System 'Should Be Recognized as Inventor'*, BBC (Aug. 1, 2019), https://www.bbc.com/news/technology-49191645 [https://perma.cc/ETP2-NXKN]; Angela Chen, *Can an AI be an Inventor? Not Yet.*, MASS. INST. TECH. TECH. REV. (Jan. 8, 2020), https://www.technologyreview.com/2020/01/08/102298/ai-inventor-patent-dabus-intellectual-property-uk-european-patent-office-law [https://perma.cc/7UKU-8DDE].

6. In Australia, the initial decision to accept the AI-inventor patent has been overturned by a five-judge panel. This decision can still be appealed to the highest court. *Commissioner of Patents v Thaler* [2022] FCAFC 62 (13 Apr. 2022) (Austl.), rev'd, *Thaler v. Commissioner of Patents* [2021] FCA 879 (30 July 2021) (Austl.) (holding inventor for a patent application must be a natural person).

7. *See, e.g.*, *Artificial Intelligence*, EUROPEAN PAT. OFF., https://www.epo.org/news-events/in-focus/ict/artificial-intelligence.html [https://perma.cc/BGR2-3KXC] (May 2, 2022); *Artificial Intelligence*, U.S. PAT. AND TRADEMARK OFF., https://www.uspto.gov/initiatives/artificial-intelligence [https://perma.cc/S36W-WQ37] (last visited Aug. 31, 2023); *Artificial Intelligence and Intellectual Property*, WORLD INTELL. PROP. ORG., https://www.wipo.int/about-ip/en/frontier_technologies/ai_and_ip.html [https://perma.cc/9LZR-AQSX] (last visited Aug. 31, 2023); *Artificial Intelligence and IP: Copyright and Patents*, U.K. INTELL. PROP. OFF. , https://www.gov.uk/government/consultations/artificial-intelligence-and-ip-copyright-and-patents [https://perma.cc/5K9N-KPVN] (June 28, 2022).

essentially concerned invention-driven changes in the organization of production, AI affects the invention process itself and, consequently, the incentives for innovation that are the focus of the patent system.

This Article takes a normative approach to how the patent system should handle AI-generated inventions. It also discusses implications for the patent system of invention machines. It draws on arguments from economic theory and evidence from empirical analyses of analogous situations. The focus is on *technical inventions* that would clearly and unambiguously meet the novelty, non-obviousness, and utility criteria if invented by a human. We are concerned with inventions that AI has fully and autonomously invented; we are not considering the use of AI as a mere tool in the invention process. However, we note that many of the points we raise apply to this broader issue as well. The fact that AI speeds up and lowers the cost of inventing does change the innovation incentives—and, perhaps, the way we should conceive the patent system.

This Article proceeds in four parts. Part I considers whether patent protection for AI-generated inventions is normatively desirable. Part II examines how invention machines could affect the patentability standards, especially the non-obviousness requirement. Part III argues against a differentiated patent system for AI-generated inventions vs. human-made inventions. Part IV discusses some systemic consequences of invention machines for patent systems and proposes potential solutions. The last Part offers concluding remarks.

I. SHOULD AI-GENERATED INVENTIONS BE PATENTABLE?

Artificial Intelligence is notoriously difficult to define but is commonly associated with the ability of a computer to learn. We utilize the term AI to refer to computer systems that can perform tasks that normally require human intelligence. AI is used in hundreds of ways all around us. Apple uses AI technology in its voice recognition software, Tesla in its self-driving technology, and Spotify and Amazon use AI to learn customer preferences. AI is used to identify the shape of proteins, which could lead to breakthroughs in drug discovery and development. AI chatbots like ChatGPT are poised to change the way students learn and study.[8]

AI, however, can also invent. Perhaps the most infamous AI-generated inventions include those associated with DABUS. DABUS is an AI system

8. ChatGPT and other natural language processing algorithms raise normative issues for copyright policy that are analogous to those considered here for patent policy. We do not consider AI-driven copyright policy issues herein because the incentive issues are different in the copyright and patent contexts.

developed by Stephen Thaler. According to Thaler, DABUS created inventions that Thaler did not conceive.[9] However, DABUS is far from the only AI system that has created inventions without human intervention, which rise to the level of inventor under current patent law.[10] Among other examples, AI-generated inventions currently include an AI-designed airplane cabin and an AI-designed race car chassis.[11]

In this Part, we address the fundamental economic question of whether society would be better off granting patent protection for AI-generated inventions instead of keeping them unprotected in the public domain. We do so by examining three canonical reasons for granting patent protection, the incentives to innovate, the incentive to commercialize inventions, and the ability of patents to encourage technology transfer. During our analysis, we assume that the invention machine autonomously creates patentable inventions at zero cost.

A. DO WE NEED PATENTS TO ENCOURAGE AI-GENERATED INVENTIONS?

The primary justification for the patent system is to provide incentives to innovate.[12] Patents enable inventors to recoup their research and development expenses by granting inventors the time-limited ability to exclude others from making, selling, or importing their inventions. By doing so, patents provide dynamic incentives for investments in new technologies.

Despite its primacy in theoretical discussions of the patent system, it is not immediately apparent that patents are needed to incentivize the act of inventing. Curiosity is a fundamental human trait, and exploration for its own sake is a widespread human activity. Inventions would undoubtedly occur in the absence of patents. It is possible that the incentive created by patents increases the rate of invention over its natural rate. This proposition is difficult to determine because we do not have good "natural experiments"

---

9. *See* Jared Council, *Can an AI System Be Given a Patent?*, WALL ST. J. (Oct. 11, 2019, 9:45 AM), https://www.wsj.com/articles/can-an-ai-system-be-given-a-patent-11570801500 [https://perma.cc/F3BX-2WKS] (stating with respect to two inventions that, according to a group associated with Thaler, he "didn't conceive of those two products and didn't direct the machine to invent them . . . .").

10. *See* Michael McLaughlin, *Computer-Generated Inventions*, 101 J. PAT. & TRADEMARK OFF. SOC'Y 224, 238–39 (2019). For other examples of AI-generated inventions, see Ben Hattenbach & Joshua Glucoft, *Patents in an Era of Infinite Monkeys and Artificial Intelligence*, 19 STAN. TECH. L. REV. 32, 32 (2015).

11. *See* McLaughlin, *supra* note 10, at 238–39.

12. *See* Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, *in* THE RATE AND DIRECTION OF INVENTIVE ACTIVITY: ECONOMIC AND SOCIAL FACTORS 609, 609 (Nat'l Bureau Econ. Rsch. ed., 1962).

comparing societies with and without patent systems.[13]

The issue of incentives to bring inventions to market plays out similarly for AI as for human-made inventions. With AI, the act of creating inventions moves away from a costly, time-consuming trial-and-error process towards an automated data-crunching task. This approach drastically reduces the cost and time of inventions, such that it costs nothing for the AI machine to produce an invention—bar the computing costs.[14] Producing inventions is cheap, and machines do not need to be incentivized.

However, producing the invention machines is presumably costly. Thus, the relevant question is whether these machines would be developed in a world where their output cannot be patented. In other words, would a patent on the invention machine itself provide enough of an incentive to create the machine, or would the machine's outputs also need to be patent eligible?[15] This question is difficult to answer, as the answer depends upon a number of factors, including the costs to produce an invention machine and the ability to monetize any invention the machine creates without patent protection. At the most, if innovators cannot secure the property of their AI inventions, there is limited financial incentive to produce invention machines in the first place. On the other hand, allowing every invention produced by an invention machine to be patentable seems like a windfall to the inventor. At some point, the reward will substantially outweigh the original incentive to innovate. As a result, it is unclear whether AI-generated inventions should be patentable based on the incentive to innovate alone.

B. DO WE NEED PATENTS TO ENCOURAGE THE COMMERCIALIZATION OF AI-GENERATED INVENTIONS?

Although it is uncertain whether we need patents on AI-generated inventions to maintain invention incentives, patents also play a critical role in invention commercialization. To be clear, we differentiate between "invention costs," which are assumed close to zero with the invention machine, and "commercialization costs," which are necessary to bring the invention to market—covering activities such as development, optimization of design,

13. *See* Eric Budish, Benjamin N. Roin & Heidi Williams, *Patents and Research Investments: Assessing the Empirical Evidence*, 106 AM. ECON. REV. 183, 183 (2016).

14. Whether there is some critical human input in the creation of inventions is an important consideration in the legal literature to establish that inventions are allowed patent protection. The distinction between AI-generated vs. AI-aided inventions (autonomy vs. automation) does not matter so much in the present discussion, where the cost and speed of creation carry more weight. If inventions are cheap and fast to come up with, one could argue that there is a priori no need to incentivize inventive activities.

15. *See* Deepak Somaya & Lav R. Varshney, *Embodiment, Anthropomorphism, and Intellectual Property Rights for AI Creations*, 2018 PROC. AAAI/ACM CONF. ON AI, ETHICS & SOC'Y 278, 278–283 (2018).

market research, scale-up of production, distribution, and the like.[16] In the particular but important case of pharmaceuticals and medical devices, human safety and efficacy testing also form part of commercialization costs.

Recent history provides part of the answer to that question. Let us go back in time, to 1980, and call the invention machine a "public research organization" ("PRO"). The U.S. government used to retain title to inventions and license them only non-exclusively. As we now know, this situation led to many valuable inventions being left unused. According to a governmental report, at the time, "fewer than 5 percent of the 28,000 patents being held by federal agencies had been licensed," compared with 25–30 percent of the federal patents for which the government allowed companies to retain title to the invention.[17] Thus, many valuable inventions fell into oblivion.

The context changed with the Government Patent Policy Act of 1980, also known as the Bayh-Dole Act, which allowed PROs and universities to patent and exclusively license federally-funded inventions. Research on the effects of the Bayh-Dole Act shows that university patenting and licensing revenues increased after 1980, suggesting greater use of inventions.[18] Several countries in Europe adopted similar legislation, including Germany and Italy.[19]

This situation is known in economics as the free-good problem.[20] A free good has zero opportunity cost, and the textbook example is air, which everyone can freely consume. By its very nature, nobody can possibly sell a free good. The picture changes when one introduces scarcity. Consider Swissbreeze, a startup that sells "the best, most pristine and freshest Swiss canned air, gathered in the most beautiful and remote lake and mountain regions."[21] Swissbreeze's business model only works because not everyone has access to fresh air, let alone from the Swiss mountains. It is easy to imagine that

16. *Ted* Sichelman, *Commercializing Patents*, 62 STAN. L. REV. 341, 348–355 (2010).

17. U.S. GEN. ACCT. OFF., GAO/RCED-98-126, TECHNOLOGY TRANSFER: ADMINISTRATION OF THE BAYH–DOLE ACT BY RESEARCH UNIVERSITIES (1998).

18. *See* David C. Mowery, Richard R. Nelson, Bhaven N. Sampat & Arvids A. Ziedonis, *The Growth of Patenting and Licensing by U.S. Universities: An Assessment of the Effects of the Bayh–Dole Act of 1980*, 30 RSCH. POL'Y 99, 99 (2001); Jerry G. Thursby & Marie C. Thursby, *University Licensing and the Bayh–Dole Act*, 301 SCI. MAG. 1052, 1052 (2003); Scott Shane, *Encouraging University Entrepreneurship? The Effect of the Bayh-Dole Act on University Patenting in the United States*, 19 J. BUS. VENTURING 127, 127 (2004).

19. Dirk Czarnitzki, Wolfgang Glänzel & Katrin Hussinger, *Heterogeneity of Patenting Activity and its Implication for Scientific Research*, 38 RSCH. POL'Y 26, 28 (2009).

20 Wendy Gordan, *Fair Use as Market Failure: A Structural and Economic Analysis of the Betamax Case and its Predecessors*, 82 COLUM. L. REV. 1600, 1611 (1982).

21. Martha Cliff, Would You Pay £19 for a Bottle of Fresh Air? Swiss Company Sells Containers of Oxygen Collected in the Mountains to 'Clear Your Mind,' DAILYMAIL, https://www.dailymail.co.uk/femail/article-5294701/Would-pay-19-bottle-fresh-AIR.html [https://perma.cc/9Q59-8JRB] (Jan. 21, 2018, 11:54 AM).

wealthy consumers in Delhi, India, or Anyang, China—two of the world's most polluted cities—may want to pay a high price for a shot of fresh air. Fresh air in these cities is scarce, and breathing it has a high opportunity cost.

Only scarcity makes the business model of bottling and selling fresh air viable. By the same reasoning, only scarcity makes viable the business model of bringing an invention to the market. Put differently, the inability to secure exclusive rights to an invention limits firms' appetite for that invention. This fate was that of many PRO and university inventions before the Bayh-Dole Act. The need for investment to bring the product to market means that at least some level of scarcity (achieved with patent protection) is warranted.[22] Thus, patent protection may be necessary to ensure commercial opportunities for the output of invention machines and, consequently, for creating invention machines themselves.[23]

C. DO WE NEED PATENTS TO ENCOURAGE TECHNOLOGY TRANSFER?

The third rationale for granting patents is to enhance technology transfer. If an invention is not patented, inventors may keep the invention secret.[24]

---

22. *See* Benjamin N. Roin, *Unpatentable Drugs and the Standards of Patentability*, 87 TEX. L. REV. 503, 509–10 (2009). The present reasoning does not apply to inventions with zero commercialization costs, that is, inventions that can be directly implemented in products without further investment. In the absence of patent protection, firms in a competitive market would immediately adopt the invention, and consumers would absorb all the surplus. However, most inventions require some amount of investment to get them from concept to market. Yet, the corner case of inventions with zero commercialization cost is an interesting one because it suggests another argument in favor of patent protection: ensuring disclosure. There has been ongoing debate regarding the extent to which patents actually disclose helpful information. *See, e.g.*, Lisa Larrimore Ouellette, *Do Patents Disclose Useful Information?*, 25 HARV. J.L. & TECH. 545, 546–50 (2012) (summarizing the existing debate and arguing that benefits of disclosure are stronger than generally thought). We have assumed thus far that inventions are disclosed publicly. It is clearly the case for university and PRO inventions, but it will not necessarily be the case for AI-generated inventions. In the absence of protection, many ready-to-market inventions—but also inventions with non-zero commercialization costs—would be kept secret, severely limiting the diffusion of these inventions. *See, e.g.*, Daniel P. Gross, *The Consequences of Invention Secrecy: Evidence from the USPTO Patent Secrecy Program in World War II* at 2–3 (Harv. Bus. Sch., Working Paper, No. 19-090, 2019); Gaétan de Rassenfosse, Gabriele Pellegrino & Emilio Raiteri, *Do Patents Enable Disclosure? Evidence from the Invention Secrecy Act* (Mar. 26, 2020) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561896 [https://perma.cc/7YKN-MR7F]; Jeffrey L. Furman, Markus Nagler & Martin Watzinger, *Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program*, 13 AMERICAN ECON. J.: ECON. POL'Y 239, 241–242 (2021).

23. Such a situation will have admittedly a lower impact for "integrated" innovators, who are both invention creators and implementors. They may obtain high enough returns from commercializing their own AI-generated inventions.

By analogy with tangible goods, one might argue that patenting the machine and its output is inappropriate. One does not get a patent for a screw machine and additional protection for the screw it produces. This analogy is misleading as the economic appropriation of tangible goods is inherently different than that of intangible goods. The "public good" nature of knowledge calls for additional protection mechanisms.

24. One might object that secrecy creates scarcity, solving the free good problem. Indeed, nothing would prevent the owner of an invention machine from approaching would-be licensees or buyers to

Secrecy hampers transactions in markets for technology, as it hurts the search for a licensing partner. Secrecy reduces the search to a one-sided process, where only the owner has the ability to reach out to interested parties.[25] Furthermore, even if the owner of the invention identifies an interested party, contracting over the information is notoriously difficult. Once the owner discloses the information, the interested party may be able to take it without paying.

Patents help increase technology transfer in two ways. First, a patent helps enable a two-sided search process where licensees and licensors search for each other. Hegde and Luo provide evidence that the publication of U.S. patent applications 18 months after their filing date rather than at the time of the patent grant has sped up licensing transactions.[26] They attribute this effect to the patent system being a "credible, standardized, and centralized repository [that] mitigates information costs for buyers and sellers."[27] Second, patents may help solve the information disclosure paradox. Patent rights are legal title that protects buyers against the expropriation of the traded idea, including when searching for a licensing partner, which also facilitates technology transactions.[28]

Implicit in this argument is that a transfer must occur between invention producers and implementers. Such transfers are necessary in the case of PROs and universities, which produce non-market-ready inventions and do not commercialize products. However, owners of invention machines may very well implement the inventions themselves. In practice, many patented

---

transfer the secret inventions. However, secrecy is not always an adequate protection mechanism. *See* Edwin Mansfield, *Patents and Innovation: An Empirical Study*, 32 MGMT. SCI. 173, 176 (1986); Wesley M. Cohen, Richard R. Nelson & John P. Walsh, *Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)* 6 (Nat'l Bureau Econ. Rsch., Working Paper No. 7552, 2000). It offers no protection for inventions that can be easily reverse-engineered, with drugs being a notable example.

25. More generally, the option of keeping an invention secret is available by default for all inventions, patentable or not. Although secrecy is sometimes used in lieu of patent protection, we do not generally judge that the option of secrecy (or other possible appropriation methods) means that patents are not a valuable policy tool. We see no reason why AI inventions are different in this regard.

26. Deepak Hegde & Hong Luo, *Patent Publication and the Market for Ideas*, 64 MGMT. SCI. 652, 652 (2017).

27. *Id.*

28. *See* Joshua S. Gans, David H. Hsu & Scott Stern, *The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence from Patent Grant Delays*, 54 MGMT. SCI. 982, 988 (2008); Gaétan de Rassenfosse, Alfons Palangkaraya & Elizabeth Webster, *Why Do Patents Facilitate Trade in Technology? Testing the Disclosure and Appropriation Effects*, 45 RSCH. POL'Y 1326, 1326 (2016). *But see* Michael J. Burstein, *Exchanging Information Without Intellectual Property*, 91 TEX. L. REV. 227, 235–46 (2012) (arguing that there is a range of ways in which to exchange information without patent protection). *See generally* Benjamin Mitra-Kahn, *Economic Reasons to Recognise AI Inventors*, *in* RESEARCH HANDBOOK ON INTELLECTUAL PROPERTY AND ARTIFICIAL INTELLIGENCE 376, 378 (Ryan Abbott ed., 2022) (arguing that recognizing AI inventors will facilitate technology transfer).

inventions are traded and licensed on markets for technology.[29] Using patent reassignment data, Serrano found that about 12–16 percent of U.S. patents are traded over their lifecycle,[30] while Ciaramella et al. found 12 percent of European patents in medical technologies are traded.[31] Furthermore, we may speculate that invention machines will exacerbate the division of innovative labor. Creating invention machines is costly, but producing inventions is cheap and fast, which may lead to more specialization (in other words, inventors versus implementers). In addition, the skills and capabilities required for creating invention machines differ drastically from those required to commercialize the inventions. If invention machines lead to a greater division of labor (where producers of inventions do not implement them), the issue of technology transfer will become particularly salient.

* *

In summary, under the traditional theory of incentives to innovate, it is uncertain whether AI-generated inventions should be patent eligible. AI makes inventing cheap, and AI machines do not need to be incentivized to invent. However, producing the AI invention machine is presumably costly. It is unclear whether these machines would be developed if their outputs cannot be patented. A stronger case for patenting AI-generated inventions is made under commercialization and technology transfer rationales for patents. Without protection, the output of the invention machine becomes more challenging to transfer and commercialize, which reduces the incentives to invent and develop such machines in the first place. Moreover, AI-generated inventions may result in the further stratification of labor markets, where producers of inventions do not commercialize them. This division of labor would make it more critical for AI-generated inventions to be patentable, as patents facilitate both technology transfer and commercialization. Of course, patents also impose costs on society, such as limiting competition and access to the invention. Thus, the benefits of allowing patents on AI-generated inventions should outweigh the costs. While we believe these arguments taken together make the uneasy case for allowing AI-generated inventions to be patented, we acknowledge that it is difficult to say so definitively. Before considering whether the patent system should treat AI-generated inventions differently, we discuss a potentially significant implication of AI systems for the patentability standards.

---

29. ASHISH ARORA, ANDREA FOSFURI & ALFONSO GAMBARDELLA, MARKETS FOR TECHNOLOGY: THE ECONOMICS OF INNOVATION AND CORPORATE STRATEGY 15–45 (2001).

30. Carlos J. Serrano, *The Dynamics of the Transfer and Renewal of Patents*, 41 RAND J. ECONOMICS 686, 693 (2010).

31. Laurie Ciaramella, Catalina Martínez & Yann Ménière, *Tracking Patent Transfers in Different European Countries: Methods and a First Application to Medical Technologies*, 112 SCIENTOMETRICS 817, 817–20 (2017).

II. THE EXISTENCE OF AI-GENERATED INVENTIONS AND IMPLICATIONS FOR THE NON-OBVIOUSNESS STANDARD

The emergence of inventions generated by AI systems also has implications for how we interpret patent validity. At any given time, there is an unknown but presumably large set of inventions that are makeable in the sense that humanity's underlying knowledge and technology base has advanced to the point where they are a feasible step beyond what has come before—an argument known as the "inevitability of inventions" at least since Ogburn and Thomas,[32] and Ihde.[33] Historically, the flow of patent applications from this unknown feasible pool has been determined by some combination of the contemporary socio-economic context, the breadth of human ingenuity, and the resources devoted to finding them. The addition of AI systems to the technology for fishing in this pool of potential inventions will likely significantly relax the latter two constraints. Human ingenuity will quite literally no longer be necessary, and the cost of exploration may be so dramatically reduced that resources available for inventing will be much less binding (perhaps almost irrelevant) as a constraint.

To begin, countries are not uniform in allowing a machine to be an inventor of a patent. Appeals courts in both the United States[34] and England[35] have held that machines cannot be inventors of patents. In contrast, Australia and South Africa allow machines to be inventors of patents.[36] Thus, we acknowledge that the patent acts of some countries, such as the United States, may need to be amended in order for machines to be inventors of patents. Assuming such reform efforts will occur, the rest of this Part examines how AI-generated inventions may affect the non-obviousness standard of patentability.

An invention is deemed obvious (and, therefore, not patentable) if the differences between what is claimed and what has been done before are such

---

32. William F. Ogburn & Dorothy Thomas, *Are Inventions Inevitable? A Note on Social Evolution*, 37 POL. SCI. Q. 83, 88 (1922).

33. Aaron J. Ihde, *The Inevitability of Scientific Discovery*, 67 SCI. MONTHLY 427, 427 (1948).

34. This conclusion seems to follow a straightforward interpretation of the Patent Act. The Patent Act defines an inventor as an "individual or, if a joint invention, the individuals collectively who invented or discovered the subject matter of the invention." 35 U.S.C. § 100(f). The Federal Circuit interpreted the term "individual" to be a natural person and that the term inventor, as used in patent statutes, does not include machines. Thaler v. Vidal, 43 F.4th 1207, 1211 (Fed. Cir. 2022). The Federal Circuit did in part by noting that the Patent Act refers to individual inventor in gendered pronouns as herself or himself, which would exclude a machine from comprising an individual. *Id.* at 1209.

35. Thaler v. Comptroller Gen. of Pats. Trade Marks and Designs, [2021] EWCA (Civ) 1374 (U.K.).

36. In Australia, the initial decision to accept the AI-inventor patent has been overturned by a five-judge panel. This decision can still be appealed to the highest court. *Commissioner of Patents v Thaler* [2022] FCAFC 62, rev'd, *Thaler v. Commissioner of Patents* [2021] FCA 879 (holding inventor for a patent application must be a natural person).

that it is obvious to a person having ordinary skill in the art ("PHOSITA") how to adapt existing technology to make the proposed invention.[37] The level of skill associated with the PHOSITA is critical in the non-obviousness inquiry. The PHOSITA is defined as an average person in a given field with "ordinary creativity, not an automaton,"[38] who has access to the same tools, skills, and knowledge base. The more skilled the PHOSITA, the more likely a new invention is obvious. Another key determinant of the obviousness inquiry is establishing what constitutes prior art, which references such as scientific articles may be used to determine whether an invention is obvious. The more prior art that can be considered, the more likely an invention is obvious. The emergence of AI systems for invention will likely have at least two ramifications for the obviousness inquiry.

First, we must confront the question of whether the PHOSITA includes AI systems. Said differently, if a proposed invention *could have been* adapted from existing technologies by a normally-skilled AI system, does that make the invention obvious and, hence, invalid? Currently, because most fields do not use AI, inventors do not have to disclose the use of AI to the Patent Office. Consider a scientist who decides to use neural networks to help come up with a new microchip design. The AI might help her calculate the ways that different materials can impact the microchip's operations. The new microchip may represent an improvement in the technology, but if an ordinary microchip inventor could have arrived at the same invention, then the new microchip would not qualify for a patent. However, suppose the AI assists in developing a novel microchip design that is beyond the skill of the ordinary microchip inventor to design. In that case, the invention may qualify for a patent. As more companies and inventors use AI to create new inventions, the legal standard will have to adapt. At some point, patent examiners will have to start assuming that a PHOSITA, which is a legal fiction that is presumed to know the relevant prior art, has access to AI, which will raise the bar for obviousness in the patent process.

Second, AI machines may alter the analogous art doctrine, which limits the prior art considered in an obvious inquiry to only prior art in the same field of the invention[39] or reasonably pertinent to the problem faced by the

37. Graham v. John Deere Co., 383 U.S. 1, 17 (1966).

38. KSR Int'l Co. v. Teleflex Inc., 550 U.S. 398, 421 (2007). The MPEP provides guidance on the level of ordinary skill in the art. *See* UNITED STATES PATENT & TRADEMARK OFFICE, MANUAL OF PATENT EXAMINING PROCEDURE § 2141.03 (9th ed. 2023); *see also* John F. Duffy & Robert P. Merges, *The Story of Graham v. John Deere Company: Patent Law's Evolving Standard of Creativity*, *in* INTELLECTUAL PROPERTY STORIES 109, 110 (Jane C. Ginsburg & Rochelle Cooper Dreyfuss eds., 2006) (noting that determining the appropriate level of ordinary skill for the nonobviousness standard "is one of the most important policy issues in all of patent law.").

39. *KSR Int'l Co.*, 550 U.S. at 417.

inventor.[40] Because an obviousness inquiry often involves combining multiple prior art references to render the invention unpatentable, the analogous art doctrine was adopted by courts to reflect the practical conditions facing an invention. An inventor likely would focus on this type of prior art when inventing. Adopting a "normally skilled" AI system as the PHOSITA could lead to a reconsideration of the analogous art doctrine. A normally skilled AI system may easily search the entire world of prior art (including patents and printed publications, but also technical blogs, standard documents, and other resources), and thus removing the analogous art limitation on the obviousness inquiry may reflect the practical realities of shifting the skilled artisan to a skilled AI system. Such removal would also result in raising the bar to patentability.

There are, however, some difficulties associated with trying to define a "normally skilled" AI system. Making the determination as to what represents an inventive enough leap for a person of ordinary skill is challenging enough; doing so for an AI machine may be even more challenging. To begin, it seems difficult to distinguish the AI system that did find the invention from the fictional one that could have. This problem does not arise with human inventors because we accept as a matter of course that each human is unique, and a given invention can come from one human's spark of genius without suggesting that any skilled human could have done it. Making this distinction for AI systems seems much harder.

A number of commentators argue that a PHOSITA AI system will place the bar for non-obviousness implausibly high, as a PHOSITA using AI can potentially create every invention—rendering "everything obvious."[41] However, as several commentators also note, this conception of AI currently is more science fiction than science,[42] in that AI only works within circumscribed attributes that humans input. Importantly, our piece is explicitly assuming AI-generated inventions. In such a scenario, it is important to keep in mind that AI systems likely would raise the non-obviousness bar, making patents harder to obtain in the future.

III. A DIFFERENTIATED PATENT SYSTEM?

The previous Part considers how AI inventions may affect the non-

40. *In re* Bigio, 381 F.3d 1320, 1325 (Fed. Cir. 2004).

41. Ryan Abbott, *Everything is Obvious*, 66 UCLA L. Rev. 2, 4–10 (2019). *See also* Tabrez Y. Ebrahim, *Data-Centric Technologies: Patent and Copyright Doctrinal Disruptions*, 43 NOVA L. REV. 287, 310 (2019). Notably, this would be true for an inventor who did not have access to AI. That is, once inventors in the field are assumed to have access to AI, this will raise the legal standard for nonobviousness across the board, including for those inventors in the field who do not have access to AI.

42. Burk, *supra* note 4, at 301.

obviousness standard. Assuming for the sake of the argument that AI-generated inventions are patentable, we turn now to considering whether we should treat AI-generated patents differently from patents on inventions generated by humans.

While the first Part of this Article makes the case for patent protection of AI-generated inventions, we have not yet addressed how strong such patent protection should be. At first, this problem seems a special case of sequential innovation with just one chain—that is, the invention machine and its inventions. Unfortunately, the vast literature on IP rights and sequential innovations is of little help. It usually assumes (1) that firms compete in the generation of follow-on inventions and (2) that follow-on inventions improve or complement, in some ways, the original invention.[43] In the present case, the same firm controls both the invention machine and the downstream inventions. Furthermore, downstream inventions are quite distinct from the invention machine itself.

It might be helpful to think of the invention machine and its offspring as one "mega invention." This mega invention is characterized by high fixed costs (the cost of producing the machine) and low marginal costs (the cost of producing one more invention using the machine). Taking such a perspective leads to an intuitive parallel with the existing literature on optimal patent strength. If we allow downstream inventions to be patented, the fractional nature of the mega invention implies that more valuable (or fruitful) mega inventions will receive stronger protection. Put differently, mega inventions associated with a larger offspring will receive a larger number of patents—and thus, broader patent protection. In that simple setup, the breadth of patent protection is proportional to the inventive potential of the mega invention. A priori, such a naturally differentiated breadth of protection may seem desirable.

However, simply allowing more patents to more fruitful mega inventions may not be the first best. This discussion naturally takes us back to the literature on optimal patent breadth.[44] From a theoretical perspective, optimal patent incentives will always depend on the incentive structure of the invention and investment processes, which clearly differ across technologies and markets. Thus, the first-best patent policy has to be a highly differentiated one, in which many aspects of the patent process and characteristics of

43. Suzanne Scotchmer, *Standing on the Shoulders of Giants: Cumulative Research and the Patent Law*, 5 J. ECON. PERSPS. 29, 29–30 (1991).

44. *See, e.g.*, Richard Gilbert & Carl Shapiro, *Optimal Patent Length and Breadth*, 21 RAND J. ECON. 106, 108–12 (1990) (providing conditions for optimal patent policy); Paul Klemperer, *How Broad Should the Scope of Patent Protection Be?*, 21 RAND J. ECON. 113, 120–24 (1990) (exploring the tradeoff between patent length and width).

patent protection differ for different kinds of inventions.[45] This route is sometimes encouraged in the policy literature, which argues in favor of "a more differentiated approach to patent protection that depends on specific characteristics of the inventions . . . ."[46]

In the present context, the first best might be a differentiated system for AI-generated and man-made inventions, reflecting the fact that the invention processes are intrinsically different. A differentiated system requires a sui generis IP right, as already pointed out by some scholars.[47] In practice, we do not and cannot implement first-best policies; political and institutional realities and myriad information and transaction frictions constrain actual policies.[48] At the most fundamental level, the theoretical argument for differentiated patent treatment assumes that it is costless to separate different types of inventions from each other. A patent policy that awards AI longer/shorter or stronger/weaker patents than other inventions would require an articulated set of criteria that determine whether an invention is "AI" or "not AI." If being "AI" resulted in less desirable treatment, we can be sure that applicants will figure out ways to characterize their inventions to meet the "not AI" criteria—and even more so if AI-generated inventions are deemed not patentable. We cannot know what fraction of truly-AI inventions would manage to escape the screen, but this positioning battle would inevitably waste resources and confuse the examination process. Recent history confirms this fear. In 2000, the United States Patent and Trademark Office ("USPTO") held that business method patent applications would be subject to a "second pair of eyes" review ("SPER"), unlike other patent applications.[49] Allison and Hunter show that the introduction of SPER led applicants of business method patent applications to write their applications so that they would not be subject to the extra review.[50]

A second problem with a differentiated patent system is that any differences in treatment would have to be introduced by statute, at least in the

45. *See* David Encaoua, Dominique Guellec & Catalina Martínez, *Patent Systems for Encouraging Innovation: Lessons from Economic Analysis*, 35 RSCH. POL'Y 1423, 1425 (2006); Angus C. Chu, *The Welfare Cost of One-Size-Fits-All Patent Protection*, 35 J. ECON. DYNAMICS & CONTROL 876, 877 (2011).

46. ORG. FOR ECON. COOP. AND DEV., PATENTS AND INNOVATION: TRENDS AND POLICY CHALLENGES 6 (2004), https://www.oecd.org/science/inno/24508541.pdf [https://perma.cc/6SNS-XUCU].

47. Deepak Somaya & Lav R. Varshney, *Ownership Dilemmas in an Age of Creative Machines*, 36 ISSUES SCI. & TECH. 79, 85 (2020); Alexandra George & Toby Walsh, *Can AI Invent?*, 4 NATURE MACH. INTEL. 1057, 1057–58 (2022).

48. *See generally* ADAM B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS (2004).

49 John R. Allison & Starling D. Hunter, *On the Feasibility of Improving Patent Quality One Technology at a Time: The Case of Business Methods*, 21 BERKELEY TECH. L.J. 729, 734 (2006).

50. *Id*.

United States. Patent policy already is a highly political process. If legislation treating AI inventions differently were to be passed, it does not require a high degree of cynicism to expect that the differentiation eventually ending up in the legislation might bear little relation to what was suggested by the first-best theoretical analysis of incentives.

Further, there is a danger that such discussion would open a bigger door: if AI patents are to be treated differently, other interests would be sure to jump in and argue that their patents should be treated differently. And in each case, the interests most affected by such differentiation would be those who expect to apply for the new special category. They have much more at stake in seeking favorable treatment than anyone has at stake in protecting the broader public interest. Opening the door to special treatment might well result in a series of differentiations in which particularly active groups get favorable treatment. Again, believing or hoping that theoretical results from welfare optimization would drive the differentiation seems naïve.

Moreover, the creation of a sui generis right could distort the innovation ecosystem in unintended ways. Consider the case of a company that has a choice between allocating human pharmacologists and investing in an AI system to develop a new vaccine. It is not clear that we want to create a system whereby the firm decides to pursue one option over another depending on the type of right it will get at the end. The new vaccine should be produced in the most efficient manner, and IP rights should be neutral to this choice.

Finally, the creation of a differentiated patent system might run afoul of international treaty obligations under the Trade-Related Aspects of Intellectual Property Rights Agreement ("TRIPS"). TRIPS requires signatories to provide a minimum set of standards for all patents, such as the stipulation that the term of a patent must last at least twenty years from the filing date.[51] However, it may be possible to create new sui generis intellectual property rights for AI-generated inventions that do not violate TRIPS obligations if such rights are not conceived as patents.[52] A complete examination of this issue is beyond the scope of this Article.

Overall, although a differentiated system *might* be the first best solution, the *realpolitik* of the patent system suggests that developing a patent

---

51. Agreement on Trade-Related Aspects of Intellectual Property Rights, Apr. 15, 1994, 1869 U.N.T.S. 299, 33 I.L.M. 1197, sec. 5, art. 33, ¶ 1 [hereinafter "TRIPS Agreement"].

52. There is also an open question as to whether new intellectual property rights, such as database protection, violates TRIPs. The European Union created a new form of intellectual property rights with respect to database protection, which so far has survived TRIPs challenges. *See generally* Guido Westkamp, *TRIPS Principles, Reciprocity and the Creation of Sui-Generis-Type Intellectual Property Rights for New Forms of Technology*, 6 J. WORLD INTELL. PROP. 827 (2003).

policy specifically for AI inventions is not likely to improve public policy and may violate international obligations.

IV. TAKING INVENTION MACHINES SERIOUSLY

This Part examines the bigger-picture implications of allowing patents on AI-generated inventions. In particular, this Part argues that patents on AI-generated inventions may overwhelm the examination capacity of national patent offices, increase the concentration of patent ownership, increase patent thickets, and lead to unlimited inventions. This Part also begins to examine changes to patent practice that might be desirable in light of these potential implications.

A. THE EXAMINATION PROCESS

It is easy to see why invention machines pose significant challenges to the functioning of the patent system. The first challenge is a potential backlog at patent offices that would come with a patent application explosion. Examining patent applications is (currently) a labor-intensive, time-consuming task. If inventing becomes cheap and fast, patent offices may not keep up with the increasing demand for examination.[53] The "global patent warming" of the mid-2000s,[54] which put the U.S. and European patent systems under strain, might look pale in comparison. Pendency could reach excessively long delays, which is detrimental to welfare.[55]

The obvious policy response is that patent offices must also use AI to speed up the examination process. Currently, a third-party contractor with the U.S. Patent & Trademark Office ("PTO" or "Agency") utilizes AI to classify new patent applications so that they route to patent examiners with the right technological expertise.[56] The U.S. PTO has also considered

---

53. *Cf.* George & Walsh, *supra* note 47, at 1059–60 (making a similar point).

54. *See generally* Bronwyn H. Hall & Rosemarie Ham Ziedonis, *The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979–1995*, 32 RAND J. ECON. 101 (2001) (documenting the rise of patenting in the semiconductor industry); Joseph Straus, *Is There a Global Warming of Patents?*, 11 J. WORLD INTELL. PROP. 58 (2008) (examining the reasons behind the surge in patent application filings); Jérôme Danguy, Gaétan de Rassenfosse & Bruno van Pottelsberghe de la Potterie, *On the Origins of the Worldwide Surge in Patenting: An Industry Perspective on the R&D-Patent Relationship*, 23 INDUS. & CORP. CHANGE 535 (2014) (same).

55. *See* Alfons Palangkaraya, Paul H. Jensen & Elizabeth Webster, *Applicant Behaviour in Patent Examination Request Lags*, 101 ECON. LETTERS 243, 243 (2008); Warren K. Mabey, Jr., *Deconstructing the Patent Application Backlog: . . .A Story of Prolonged Pendency, PCT Pandemonium & Patent Pending Pirates*, 92 J. PAT. & TRADEMARK OFF. SOC'Y 208, 237–46 (2010); Lily J. Ackerman, *Prioritization: Addressing the Patent Application Backlog at the United States Patent and Trademark Office*, 26 BERK. TECH. L.J. 67, 67–68 (2011); Stuart J. H. Graham & Galen Hancock, *The USPTO Economics Research Agenda*, 39 J. TECH. TRANSFER 335, 341 (2014).

56. U.S. DEPT. COM., U.S. PATENT & TRADEMARK OFFICE, PTOC-016-00: PRIVACY IMPACT ASSESSMENT FOR THE SERCO PATENT PROCESSING SYSTEM (PPS) 1 (2018); Serco Processes 4 Millionth

incorporating AI to improve prior art searching of patent examiners.[57] AI holds great potential for improving the search process associated with patent examination as well as locating relevant passages in the prior art, mapping them to elements of the current application's claims, and hence suggesting potential rejections. Admittedly, AI may not be as helpful in reviewing patent applications on newer subject matters where inventors are just developing new patentable technologies.

Moreover, it seems unlikely that legislators will authorize a fully autonomous examination, that is, the automatic granting of traditional patent rights without a human in the loop. Some human intervention in the patent examination process may be necessary to satisfy a patent applicant's due process rights or administrative law's reason-giving requirements under current law.[58] Moreover, effectively keeping up with the increase in patent numbers requires patent offices to adopt AI tools as sophisticated as those of the most advanced applicants, which does not seem likely.[59] Because the need for human intervention puts a hard constraint on examination time, it is safe to assume that, on balance, pendency most likely will increase.[60]

The U.S. PTO has some experience with an increased onslaught in patent applications in the past. In the 1990s, the Agency experienced a torrential rise in the number of patent applications filed on express sequence tags ("EST") or small fragments of DNA.[61] The U.S. PTO estimated that it would take a single examiner over 90 years and cost the Agency upwards of 20 million dollars to review the EST patent applications in its queue. As a result, then U.S. PTO Commissioner Bruce Lehman considered several possible

---

Patent Application for U.S. Patent and Trademark Office, PR NEWSWIRE (Nov. 15, 2018), https://www.prnewswire.com/news-releases/serco-processes-4-millionth-patent-application-for-us-patent-and-trademark-office-300751330.html [https://perma.cc/GM86-EWPT] ("Since 2006, Serco has performed classification and other analysis services through awarded contracts including Pre-Grant Publication (PGPubs) Classification Services, Initial Classification and Reclassification (ICR) Services, and Full Classification Services (FCS) contracts.").

57. U.S. Patent & Trademark office, Patent-End-To-End Search Artificial Intelligence Capability Request for Information & Notice of Vendor Engagement at 3, Aug. 25, 2023, available at file:///C:/Users/feeneymr/Downloads/PTAG+RFI+AI+Search-Final.pdf.

58. Although the case law is far from settled on this matter. *See, e.g.*, Arti K. Rai, *Machine Learning at the Patent Office: Lessons for Patents and Administrative Law*, 104 IOWA L. REV. 2617, 2625–29 (2019); Aziz Z. Huq, *A Right to a Human Decision*, 106 VA. L. REV. 611, 661–71 (2020).

59. *Cf.* Rai, *supra* note 58, at 2638 ("To the extent that the Al-assisted search used by the Patent Office does not account for potentially rapid change in the average skill of practitioners itself spurred by AI, it will fall short.").

60. Interestingly, one might say that invention machines will reduce the demand for scientists and engineers. The pool of redundant inventors could then be hired by patent offices to examine the inventions of the very machines that took their job. For a modern example of machine slavery, see MODERN TIMES (United Artists 1936).

61. This rise in patent applications was due to changes in technology that made the sequencing of DNA easier. *See* Eliot Marshall, *Patent Office Faces 90-Year Backlog*, 272 SCIENCE 643, 643 (1996).

changes to combat the growing backlog of DNA patent applications, including requiring patent applicants to do more work themselves or contract out the research for searching the prior art.[62] Patent offices can consider these same approaches with respect to AI-generated patent applications.

Contracting out the research, however, would have the same problems as addressed above. That is, any contractor likely would need access to AI tools as sophisticated as those of the most advanced applicants. An alternative may be to require patent applicants on AI-generated applications to conduct their own patentability search and identify the most relevant prior art when they submit their applications to patent offices. Shifting the prior art search on the applicant would ease the burden on the patent offices as well as harness the most up-to-date AI search tools.[63] Moreover, the common refrain against requiring more search efforts of patent applicants—that such efforts would increase the cost of patenting and hence reduce patenting efforts for cost-conscious applicants—has less force for AI-generated inventions.[64] Given that invention machines presumably have processed and screened the prior art for coming up with the invention, it would be reasonably straightforward to identify the closest prior art. Nonetheless, shifting the patentability search to the applicant has its own set of drawbacks. Applicants, whose incentives may arguably cut against doing an exhaustive search, may find ways to game the search process.[65]

Other work-sharing options may also ease the administrative burden associated with a rapid influx of AI-generated patent applications. The U.S. PTO has patent work-sharing arrangements with foreign intellectual property offices to improve patent examination efficiency. Patent work-sharing permits patent offices to collaborate in the examination of commonly filed patent applications, reducing inefficiencies that patent offices experience when doing largely duplicative research into questions relating to patentability.[66] The most famous of these work sharing efforts occurs through the

---

62. In the EST context, the Agency successfully lobbied for an elevated utility standard with respect to EST—which required the patent applicant to describe the function and utility of the gene that the EST comprised. *In re* Fisher, 421 F.3d 1365, 1370–71 (Fed. Cir. 2005).

63. The current duty of candor whose breach can lead to a charge of inequitable conduct attempts to harness applicants' knowledge. 37 C.F.R. § 1.56(a) ("Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this Section.").

64. John M. Golden, *Proliferating Patents and Patent Law's "Cost Disease*," 51 HOUS. L. REV. 455, 494 (2013).

65. *Cf.* Jeffrey M. Kuhn, *Information Overload at the U.S. Patent and Trademark Office: Reframing the Duty of Disclosure in Patent Law as a Search and Filter Problem*, 13 YALE J.L. & TECH. 90, 112–119 (2010) (documenting that examiners receive too much information on prior art disclosure from patent applicants that examiners cannot process the information and often ignore it).

66. Mabey, *supra* note 55, at 231.

Patent Prosecution Highway ("PPH") programs, in which the partial examination of an application in one office can result in the expedited review of that application in another office.[67] Various reports suggest that PPH results in faster and cheaper reviews of patent applications.[68] Nevertheless, in the fiscal year of 2021, the 6,000 patent applications filed under PPH are minuscule in comparison to the 650,000 patent applications filed at the U.S. PTO.[69] As a result, work-sharing efforts seem unlikely to do much to combat the increase in filings associated with AI-generated patent applications.

A more radical approach might be to, in effect, aggregate examinations of patents produced by the same AI invention machine. Applicants could apply to have a specific AI algorithm certified as reliably generating novel and non-obvious inventions. Subsequent applications that could be shown to be the output of certified machines would be presumed valid and granted patent protection.[70] This two-track system does not necessarily imply a differentiated patent system since the nature of the patent right granted is the same across both tracks. Further, it does not seem to introduce the problem of people gaming the system to qualify for or avoid special treatment. An applicant could submit an "invention machine" for approval. Examiners would not need to determine whether the submitted "machine" meets some definition of AI; they would need only to determine whether or not it reliably produces inventions that meet the standards for patentability.

B. MARKET IMPACTS

The second challenge of cheap and fast inventions is the potential effects on the markets for innovation. This Section identifies two potential market impacts of allowing patents on AI-generated inventions. First, AI-generated inventions could result in an increase in the concentration of patent ownership. Owners of invention machines would have the opportunity to amass vast patent portfolios, possibly conferring on them strategic advantages over their rivals.[71] Along this line, Professors Choi and Gerlach

67. Toshinao Yamazaki, *Patent Prosecution Highways (PPHs): Their First Five Years and Recent Developments Seen from Japan*, 34 WORLD PAT. INFO. 279, 279 (2012) (providing an overview of PPH programs); U.S. PAT. & TRADEMARK OFF., PERFORMANCE AND ACCOUNTABILITY REPORT 105 (2021), https://www.uspto.gov/sites/default/files/documents/USPTOFY21PAR.pdf [https://perma.cc/HB76-XGY8].

68. *See* Yamazaki, *supra* note 67, at 280–82 (claiming PPH benefits in terms of speed of "patent acquisition," increased allowance rates, and reduced costs).

69. U.S. PAT. & TRADEMARK OFF., *supra* note 67 (in the fiscal year 2021, 5,821 patent applications were filed under PPH while over 650,000 patent applications were filed in total at the U.S. PTO).

70. To guarantee the quality of the certification, machines could be checked regularly and major changes to the algorithms would trigger re-examination. Examiners could also randomly select some AI-generated inventions at regular intervals and examine them.

71. Hall & Ziedonis, *supra* note 54, at 108–10; Gideon Parchomovsky & R. Polk Wagner, *Patent Portfolios*, 154 U. PA. L. REV. 1, 72–74 (2005).

have shown that an increase in one firm's patent portfolio unambiguously reduces the rival firm's incentives to develop a new product. One could also think of more severe chilling and blocking effects.[72]

Second, a market-related issue of a burst of inventions is an exacerbation of the problem of patent thickets, namely overlapping and fragmented patent rights.[73] Intertwined patent rights increase litigation risks for innovators, and the transaction costs associated with clearing these rights may become prohibitively expensive. This is especially true in industries in which many patent-protected technologies are necessary to manufacture a single product, such as a smartphone.

Relatedly, increased market concentration of patenting and patent thickets could also lead to the emergence of a new genre of patent assertion entities ("PAEs"), taking hold-ups and nuisance settlements to new heights. The leading critique of PAEs is that they assert weak or invalid patents against product manufacturers to extract nuisance settlements, which in turn stunt innovation.[74] While there is no reason to think that AI-generated inventions are inherently of lower quality than human-generated inventions, the rise of patenting fueled by AI-generated inventions could lead to more overlapping patent rights and could decrease the costs of amassing vast patent portfolios. Product manufacturers who face patent thickets often settle through cross-licensing agreements. This process is not possible for PAEs as they do not produce any products or services that could potentially infringe anyone else's patents. Thus, an increase in patent thickets and a decrease in barriers to amassing vast patent portfolios may create tantalizing opportunities for PAEs.

The adverse welfare effects of vast patent portfolios and patent thickets suggest that rewarding machine-made inventions with as many patents as

---

72. Jay Pil Choi & Heiko Gerlach, *A Theory of Patent Portfolios*, 9 AM. ECON. J.: MICROECONOMICS 315, 315–16 (2017).

73. Carl Shapiro, *Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting*, *in* 1 INNOVATION POLICY AND THE ECONOMY 119, 119–22 (Adam B. Jaffe, Josh Lerner & Scott Stern eds., 2000); Rosemarie Ham Ziedonis, *Don't Fence Me In: Fragmented Markets for Technology and the Patent Acquisition Strategies of Firms*, 50 MGMT. SCI. 804, 804–06 (2004).

74. Ashley Chuang, Note, *Fixing the Failures of Software Patent Protection: Deterring Patent Trolling by Applying Industry-Specific Patentability Standards*, 16 S. CAL. INTERDISC. L.J. 215, 232 (2006) ("Because of a patent troll's approach to generating revenue, a troll's charges of infringement and litigation can often be baseless and thus clog the legal system."); Spencer Hosie, *Patent Trolls and the New Tort Reform: A Practitioner's Perspective*, 4 I/S: J.L. & POL'Y FOR INFO. SOC'Y 75, 78 (2008) ("Perhaps the most common refrain in the patent debate is that plaintiffs will bring frivolous cases to extort unjustified settlements."); Sannu K. Shrestha, *Trolls or Market-Makers? An Empirical Analysis of Non-Practicing Entities*, 110 COLUM. L. REV. 114, 119 (2010) ("One of the most prominent criticisms against NPEs is that they acquire weak and obscure patents and use them to pursue 'baseless' litigation."); Robert P. Merges, *The Trouble with Trolls: Innovation, Rent-Seeking, and Patent Law Reform*, 24 BERKELEY TECH. L.J. 1583, 1603–04 (2009) (discussing allegations that NPEs file suits on weaker patents).

inventions produced may offer too large a reward. Considering that invention machines have high fixed costs and low marginal costs, there must be a point at which the machines are generating large numbers of very low value inventions. Past this point, additional patents have value to their owners only through the market power generated by a larger portfolio.[75] This optimal threshold is private information and varies across invention machines.

One could imagine several mechanisms to limit patent portfolios' strength. The suggestion above of creating the applicant option to have an invention machine certified as producing patentable inventions likely would exacerbate the portfolio market power and patent thickets problem, but it also offers potentially incentive-compatible ways to limit such market power. Patents granted through this route could bear limitations such as a shorter validity period or forced availability under Fair, Reasonable, and Non-Discriminatory ("FRAND") clauses—although FRAND clauses come with their own set of challenges.[76] However, as noted in Section 3 these limitations would need to be carefully crafted so as not to violate international treaty obligations under the TRIPS Agreement. Other options exist, such as increasing application fees with the size of the assignee's patent portfolio or for each new invention produced by the same machine.

Putting conditions on patents from invention machines that potentially reduce the value of the patents would, again, introduce greater differentiation into the system. But this could perhaps be incentive-compatible rather than wasteful. It will be up to the applicants to decide whether to seek approval of an invention machine, and if they have an approved machine, whether to submit each new invention as a product of the machine or as a standard application. The machine route will yield faster but less valuable patents, while the standard route will yield slower but more valuable patents. In principle, these tradeoffs could be calibrated to limit the market power of vast portfolios while still affording appropriate incentives to patent the best inventions. Nevertheless, a differentiated system would still suffer from the political economy concerns set forth in Part III.

While it seems *a priori* desirable to limit the strength of AI-generated patent portfolios, the best mechanism to achieve this aim is unclear and deserves a careful theoretical investigation.

---

75. Alfonso Gambardella, Dietmar Harhoff & Bart Verspagen, *The Economic Value of Patent Portfolios*, 26 J. ECON. & MGMT. STRATEGY 735, 735–36 (2017).

76. Michael A. Carrier, *Why Is FRAND Hard?*, 2023 UTAH L. REV. 931, 932–53 (2023) (describing eight reasons why FRAND licensing is challenging).

C. UNLIMITED INVENTIONS?

Finally, even if the flood of inventions from AI is not all patented, the democratization of invention machines could still have systemic consequences for the patent system. Owners of such machines might not patent their inventions but generate a vast amount of prior art. This prior art would naturally form part of the literature used to assess the non-obviousness of inventions, implicitly raising the bar to obtain patents in these areas—perhaps to a point where it would be extremely challenging to obtain patents in a given area.[77] Firms may want to flood a technological area with prior art to ensure freedom to operate.[78] This practice could essentially impose patent-free technological zones with unknown consequences on product development and commercialization. Such situation would have similar consequences to allowing an AI-augmented PHOSITA. The issue would not be that the AI-augmented PHOSITA could have produced the invention, but an acknowledgement of the fact that a large pool of prior art exists that renders the invention obvious.

Taking this argument a step further (and maybe too far), suppose AI got so skilled at invention that invention itself became essentially irrelevant. Imagine a world where in some sense every invention that could possibly be made at a point in time was known to everyone, or knowable to anyone who cared at very low cost. At this point, there would be no need to provide any incentive for people to invent; indeed it would become somewhat unclear what it even meant to invent something. But there may still be a social need to provide incentives for people to invest in commercializing inventions, as argued above.[79]

To make this consideration concrete, consider the (admittedly artificial) hypothetical case in which every chemical compound that might have therapeutic benefits to humans was known or knowable, so no one could meaningfully "invent" a new drug. But it still costs millions to test the drug in humans. We would want companies to pay to run those tests, but they would not do so if anyone could then sell the drug because it was proved safe and effective. In that world, we might want to give companies some kind of exclusive right to test and then market new drugs. But we couldn't use first to

---

77. One such initiative is already under way. *See* ALL PRIOR ART, http://allpriorart.com [https://perma.cc/4RFE-8SQL] (last visited Sept. 6, 2023). A key question is whether the disclosures by the AI would be enabling.

78. Firms did something similar with DNA gene fragments before the law required that for a DNA gene fragment to be patentable, the utility of the underlying gene must be identified.

79. Unless we had AI that, without cost, could tell us exactly how to adapt, manufacture, scale-up, and market a new product. We have trouble imagining how this would work, but it would be silly to rule it out *ex ante*.

file as the criterion to determine who got that right. One could imagine a different kind of examination system, where companies made proposals for developing products out of the pool of available inventions, and were somehow evaluated on how much they proposed to invest and/or how good their development plan was. But that sounds hard. To economists, an obvious solution would be to auction the rights. The development of a particular invention out of a publicly-known pool is somewhat like a slice of electromagnetic spectrum in a given geographic area. We want someone to use it, but we don't want more than one entity to use it, so we auction it off.

We raise these possibilities neither to say that we know that AI will get that good, nor to suggest that we have done any careful analysis of the merits of public auctions for invention development rights. Rather, we only want to suggest that if AI becomes extremely successful at invention, we will need to think about potentially radical changes to innovation policy.

CONCLUSION

Patent law has traditionally adapted slowly to the changing environment. In 2004, the U.S. National Research Council issued a report entitled "A Patent System for the 21st Century."[80] The report addressed issues that had plagued the U.S. patent system for decades or more, including questionable patent quality, impediments to disseminating information through patents, and international inconsistencies.[81] Some inconsistencies, such as the United States' first-to-invent principle compared to the rest of the world's first-to-file principle, existed since the Patent Act of 1790. Many of the issues discussed in the report have not yet come to the fore. While they could materialize sooner than expected, the legislator is unlikely to act faster than expected. We hope that the patent system will be ready for the 22nd century by discussing these issues now.

In our view, some form of IP protection for AI-generated inventions is likely desirable. However, the nature of the IP regime is unclear and deserves in-depth theoretical and empirical examination. Regardless of whether AI-generated inventions are patentable, if AI radically reduces the cost and increases the production rate for inventions, it will have implications for the patentability standards that will have to be addressed. In addition, AI-generated inventions will have significant implications for the patent ecosystem more generally. A large increase in the rate of generation of patentable ideas will potentially overwhelm the examination process (if AI-generated

80. NAT'L RSCH. COUNCIL OF THE NAT'L ACADS., A PATENT SYSTEM FOR THE 21ST CENTURY (Stephen A. Merrill, Richard C. Levin & Mark B. Myers eds., 2004).

81. *Id.*

inventions are patentable), make patents unavailable in wide swaths of technology (if AI-generated inventions are not patentable but saturate the prior art), and increase the concentration of patent ownership and the likelihood of patent thickets.

We have proposed a series of potential solutions to these problems. We do not claim that any of our proposed solutions are the best. We note also that AI-generated inventions have the potential to exacerbate the problem of increasing market power from highly concentrated patent portfolios, and that certifying invention machines might make this problem worse. Our hope is that this Article illustrates a need to seriously consider the protection of AI-generated inventions and that creative solutions do exist, but those solutions may have complex ramifications that should be thought through. In addition, these solutions also require global cooperation to harmonize legislations. Meanwhile, some concrete steps may already be implemented, such as a change in disclosure requirements. By forcing patent applicants to disclose the extent of the involvement of AI in the invention process, it becomes possible to track AI-generated inventions. This step is necessary to quantify the phenomenon and empirically study its effects.

The pressure for changes in the system that AI-generated inventions may create is also an opportunity. The structure of our current system is essentially the result of historical accident. As noted, it is difficult to measure the consequences of the system, or of specific aspects of the system, because we do not have natural experiments that allow us to test one practice against another. If changes are to be made in response to these new pressures, they should be structured initially to provide explicitly for quantified evaluation of the effects of new policies and procedures, potentially including structures such as randomized control trials that isolate the causal effect of specific changes.[82]

There is little doubt that confronting the implications of AI playing a role in the invention process is now on the agenda, and is likely to become more and more important. This paper's focus on one set of issues should not be taken to mean that these issues are the main challenges facing tomorrow's patent system. Nor does it mean that there are no other ways of modernizing the patent system.[83] But AI is a rapidly evolving set of technologies, and the

82. For an example of the use of an RCT to measure the effect of a change in patent examination procedure, see Nicholas A. Pairolero, Andrew Toole, Peter-Anthony Pappas, Charles DeGrazia & Mike Teodorescu, *Closing the Gender Gap in Patenting: Evidence from a Randomized Control Trial at the USPTO* 2–5 (U.S. Pat. and Trademark Off., Econ. Working Paper No. 2022-1, 2022).

83. For example, proposals to "decentralize" the patent system using distributed ledger (a.k.a. blockchain) technologies may very well be an important component of a 22nd-century patent system. Lital Helman, *Decentralized Patent System*, 20 NEV. L.J. 67, 68–71 (2019); Gaétan de Rassenfosse & Kyle Higham, *Decentralising the Patent System*, 38 GOV'T INFO. Q. 1, 1 (2021). In the context of a burst of

longer we delay determining how the innovation system should respond, the more likely we are to see socially undesirable consequences.

---

inventions, a "block-chained" patent system can mitigate the transaction costs associated with intertwined patent rights. A license to an antecedent patent, essential to the use of a new invention, could be executed automatically by means of a smart contract under set conditions, should the owner of antecedent patent allow it.