

Artificial Intelligence Inventions & Patent Disclosure

Tabrez Y. Ebrahim*

ABSTRACT

Artificial intelligence (“AI”) has attracted significant attention and has imposed challenges for society. Yet surprisingly, scholars have paid little attention to the impediments AI imposes on patent law’s disclosure

* Associate Professor of Law, California Western School of Law; Visiting Associate Professor, University of California, San Diego; Ostrom Visiting Scholar, Indiana University (Bloomington); Visiting Fellow, University of Nebraska (Lincoln); Nebraska Governance & Technology Center; Thomas Edison Innovation Fellow & Leonardo da Vinci Fellow, George Mason University Antonin Scalia Law School; Visiting Scholar, University of California, Los Angeles School of Law; Registered U.S. patent attorney; J.D., Northwestern University Pritzker School of Law; M.B.A., Northwestern University Kellogg School of Management; LL.M., University of Houston Law Center; Graduate Entrepreneurship Certificate, Stanford Graduate School of Business; M.S. Mechanical Engineering, Stanford University School of Engineering; B.S. Mechanical Engineering, University of Texas at Austin Cockrell School of Engineering.

I am grateful for helpful comments, feedback, and suggestions from Michael Risch, Ted Sichelman, Brenda Simon, Thomas D. Barton, Robert A. Bohrer, Shawn Miller, Lisa Ramsey, Anjanette Raymond, Daniel R. Cahoy, Sonia Katyal, Tejas Narechania, Jonathan Barnett, Eric Claeys, John Duffy, Sean O’Connor, Ashish Bharadwaj, Loletta Dardin, Charles Delmotte, H. Tomás Gómez-Arostegui, Taorui Guan, Devlin Hartline, Christa Laser, Daryl Lim, Kevin Madigan, Talha Syed, James Stern, Seth C. Oranburg, Agnieszka McPeak, Gregory Day, Nicole Iannarone, Emily Loza de Siles, Eric C. Chaffee, Robert F. Kravetz, Ashley London, Aman Gebru, Elizabeth I. Winston, A. Michael Froomkin, Mason Marks, Larry DiMatteo, Robert W. Emerson, Robert E. Thomas, Colleen M. Baker, Lawrence Trautman, George Cameron, David Orozco, Thomas Freeman, Christopher Guzelian, Daniel Herron, Michelle Romero, Tyler Smith, Brian Haney, Jihwang Yeo, Sikander Khan, Erica Pascal, Ryan Hsu, Kevin R. Tamm, and Daniel R. Peterson.

Thanks to the following forums for presenting this Article and their participants for insightful comments: 2020 Huber Hurst Research Seminar at University of Florida Warrington College of Business, The Junior #FutureLaw Workshop 4.0 at Duquesne University School of Law, Ostrom Workshop’s Colloquium Series at Indiana University (Bloomington), Junior Intellectual Property Scholars Association (JIPSA) at George Washington University School of Law, We Robot 2019 at University of Miami School of Law, PatCon 9 (The Annual Patent Conference) at University of Kansas School of Law, and the 9th Annual Patent Law Conference at University of San Diego School of Law. Thanks to the Academy of Legal Studies in Business (ALSB) Interdisciplinary Section for selecting this Article for its inaugural “Best Paper Award” at the 2020 ALSB Annual Conference and to ALSB members for thoughtful comments.

function from the lenses of theory and policy. Patents are conditioned on inventors describing their inventions, but the inner workings and the use of AI in the inventive process are not properly understood or are largely unknown. The lack of transparency of the parameters of the AI inventive process or the use of AI makes it difficult to enable a future use of AI to achieve the same end state. While patent law’s enablement doctrine focuses on the particular result of the invention process, in contrast, this Article suggests that AI presents a lack of transparency and difficulty in replication that profoundly and fundamentally challenge disclosure theory in patent law. A reasonable onlooker or a patent examiner may find it difficult to explain the inner workings of AI. But even more pressing is a non-detection problem—an overall lack of disclosure of unidentified AI inventions, or knowing whether the particular end state was produced by the use of AI.

The complexities of AI require enhancing the disclosure requirement since the peculiar characteristics of the end state cannot be described by the inventive process that produced it. This Article introduces a taxonomy of AI and argues that an enhanced AI patent disclosure requirement mitigates concerns surrounding the explainability of AI-based tools and the inherent inscrutability of AI-generated output. Such emphasis of patent disclosure for AI may steer some inventors toward trade secrecy and push others to seek patent protection against would-be patent infringers despite added *ex ante* costs and efforts. Utilitarian and Lockean theories suggest justifications for enhanced AI patent disclosure while recognizing some objections. Turning to the prescriptive, this Article proposes and assesses, as means for achieving enhanced disclosure, a variety of disclosure-specific incentives and data deposits for AI. It concludes by offering insights for innovation and for a future empirical study to verify its theoretical underpinnings.

Table of Contents

INTRODUCTION	149
I. FOUNDATIONS: HOW & WHY AI IS “ARTIFICIAL”	165
A. The “Magic” of AI	165
B. Transparency Viewed from a Taxonomy of AI	171
II. THEORIES & NORMATIVE ASSESSMENTS OF	
PATENT DISCLOSURE WITH AI.....	181
A. Theories of the Patent and Trade Secret	
Trade-off Decision with AI.....	182
B. Theories of Patent Disclosure	187
1. Theories by a Skeptic of Patent Disclosure.....	187
2. Theories by Proponents of Patent Disclosure.....	189
C. Normative Assessment with AI Disclosure	190

1. Disclosure for Patentability	191
2. Detection of Disclosure for Patent Examination	193
D. Theories for Justification of Patent Disclosure of AI-Generated Output	197
1. Utilitarian Theories for Adequate AI-Generated Output	199
2. Lockean Theories for Adequate AI-Generated Output	201
E. Toward an Enhanced Patent Disclosure & Responding to Objections	202
III. NORMATIVE IMPLICATIONS, PRESCRIPTIONS, & FUTURE DIRECTIONS	204
A. Normative Implications of Enhanced AI Disclosure	205
1. Greater <i>ex ante</i> Inventor Costs and Efforts for AI Inventors	205
2. Possible Move Toward Trade Secrecy for AI Inventions	207
B. Assessments & Prescriptions for Enhanced AI Disclosure	210
1. Establishing Disclosure Incentives for the Inventor	211
2. Establishing a Data Deposit as an Alternative Disclosure	214
C. Future Directions	218
CONCLUSION	220

INTRODUCTION

Artificial intelligence (“AI”) inventions are artificial.¹ There is much disagreement about what it means to be artificial. I define “artificial” broadly to mean that inventors use AI-based tools in the inventive process or AI-based tools invent autonomously to produce AI-generated output without human intervention. AI-based tools (run by

1. See Ryan Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. REV. 1079, 1083–85 (2019) (explaining that AI is characterized by autonomously-created inventions, can brainstorm inventions, and is self-assembling); Harry Surden, *Artificial Intelligence and Law: An Overview*, 35 GA. ST. U. L. REV. 1305, 1307 (2019) (describing AI as using technology to automate tasks that normally require human intelligence); Clark D. Asay, *Artificial Stupidity*, 61 WM. & MARY L. REV. 1187, 1190 (2020) (defining AI as computing systems that perform tasks that normally would require human intelligence); see also KAY FIRTH-BUTTERFIELD & YOON CHAE, *ARTIFICIAL INTELLIGENCE COLLIDES WITH PATENT LAW* 5 (Apr. 2018), <https://bit.ly/3eloC0N> (defining AI as “a computerized system exhibiting behavior commonly thought of as requiring intelligence” or “a system capable of rationally solving complex problems or taking appropriate action to achieve its goals in real-world circumstances”); Chris Smith et al., *The History of Artificial Intelligence*, U. WASH., Dec. 2006, at 4, <https://bit.ly/36xGwec> (defining AI as “a system which amplifie[d] people’s own knowledge and understanding”); Phillipe Aghion et al., *Artificial Intelligence and Economic Growth* 2 (Nat’l Bureau of Econ. Research, Working Paper No. 23928, 2017), <https://bit.ly/3goqSGp> (defining AI as “the capability of a machine to imitate intelligent human behavior [or] an agent’s ability to achieve goals in a wide range of environments”); Sean Semmler & Zeeve Rose, Comment, *Artificial Intelligence: Application Today and Implications Tomorrow*, 16 DUKE L. & TECH. REV. 85, 86 (2017) (defining AI as “the process of simulating human intelligence through machine processes”).

complex algorithms) and AI-generated output (the result of applying such complex algorithms), are rapidly proliferating in business² and becoming ubiquitous in society.³ The potential societal benefits and costs of AI-based tools and AI-generated output force re-examination of the basic tenets of patent law, since the exact parameters of AI-based tools and their use in the inventive process are not clear, and since it is difficult to reproduce the same end state in the future.⁴ Should patent grants to AI inventors be conditioned upon more transparent disclosure of AI? Does the black-box nature of AI jeopardize the normatively desirable transparency⁵ of patent disclosure? Are patents that cover

2. See James Bessen et. Al., *The Business of AI Startups* 2, 3, 26 (B. U. Sch. of L., L. & Econ. Research Paper, Working Paper No. 18-28, 2018), <https://bit.ly/36wmQXT> (suggesting that AI dramatically alters the economy and ways of conducting business where it is applied, enhances human capabilities in commercial applications, and changes competition and entry barriers for startups); Jason Furman & Robert Seamans, *AI and the Economy* at 1–2, 17 (Nat'l Bureau of Econ. Research, Working Paper No. 24689, 2018), <https://bit.ly/3epai7z> (explaining that AI is having a large effect on the economy, will increase productivity and economic growth, and could lower the cost of doing business); see also Roberto Moro Visconti, *The Valuation of Artificial Intelligence*, UNIVERSITÀ CATTOLICA DEL SACRO CUORE at 3–5 (2019), <https://bit.ly/2zvAsql> (noting that AI can be applied to a variety of business models, including business-to-business (B2B), business-to-consumer (B2C), back-office infrastructure, and business applications such as web search engines, e-mail Spam detectors, virtual personal assistants, or any business application where forecasting plays a fundamental role in decision-making or increases the value of data).

3. See Andrei Iancu, Remarks by Director Iancu at the Artificial Intelligence/Intellectual Property Considerations Event, USPTO (Jan. 31, 2019) [hereinafter Remarks by Director Iancu at the AI Event] (unpublished manuscript), <https://bit.ly/3gr6RyM>. The USPTO Director stated:

Today, AI is becoming ubiquitous in our society. For example, faster, more-powerful processors and chips now provide sufficient computing power to perform trillions of calculations per second. Very quickly, AI technologies are evolving from far-off dreams of science fiction to mainstream, everyday uses that take computers to new levels at awe-inspiring speeds.

Id.

4. See S. 2217, 115th Cong. (2017) (introducing Congress's first major steps toward comprehensive regulation of the AI tech sector as AI technologies become ubiquitous in society); H.R. 4625, 115th Cong. (2017) (same); see also Exec. Order No. 13859, 84 Fed. Reg. 3967 (Feb. 11, 2019) (introducing the United States' national strategy on artificial intelligence that aims to promote and protect AI technology and innovation for U.S. society); Roberto Verganti, Luca Vendraminelli & Marco Iansiti, *Innovation and Design in the Age of Artificial Intelligence*, 37 J. PROD. MGMT. 212, 215–19 (2020) (describing implications of design on innovation policies, which are based on the analysis of AI—the practices, principles, and theory of AI design—as an inherently decision-making technology that automates many tasks relating to learning and devising solutions in a number of business processes and virtually every industrial setting); Iain Cockburn, Rebecca Henderson & Scott Stern, *The Impact of Artificial Intelligence on Innovation* 7 (Nat'l Bureau of Econ. Research, Working Paper No. 24449, 2017), <https://bit.ly/3dTFImw>.

5. See Robert Brauneis & Ellen Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J.L. & TECH. 103, 107–08, 110, 132 (2018) (suggesting that a lack of transparency refers to “the use of algorithms that are highly dynamic or that use modeling

inventions developed by the use of AI-based tools being granted at an alarming enough frequency such that it causes serious problems to the patent system, and if so, how should patent law and the administration of the patent system evolve?

AI⁶ involvement in the invention process falls on a technological spectrum and is controversial for the patent system. At one end of the spectrum, an AI-based tool assists a human inventor without contributing to the conception of the invention; at the other end of the spectrum, an AI-based tool produces AI-generated output that would satisfy patentability if otherwise created by a human. However, controversy arises over whether patents' disclosure function is inhibited when an AI-based tool's operation is unexplainable⁷ (not discernible to the relevant audience) or the production of AI-generated output is inherently

which makes them difficult to interpret even when records are revealed" and implying that transparency refers to: disclosing predictive models or algorithms, opaque "black box" processes, or "black boxes"; additionally, suggesting that meaningful transparency is demonstrating the knowledge that is sufficient to approve or disapprove of the algorithm's performance); David Freeman Engstrom & Daniel E. Ho, *Artificially Intelligent Government: A Review and Agenda*, in ROLAND VOGL, *BIG DATA LAW* (forthcoming 2020) (stating that "transparency requires, at a minimum, a description of the decision's 'provenance,' including an accounting of its inputs and outputs and the main factors that drove it"); Alfred Früh, *Transparency in the Patent System: Artificial Intelligence and the Disclosure Requirement*, in RETHINKING PATENT LAW AS AN INCENTIVE TO INNOVATION (Žaneta Pacud & Rafał Sikorski eds., forthcoming 2019), <https://bit.ly/2LXDZjV>; Michèle Finck, *Automated Decision-Making and Administrative Law*, in MAX PLANCK INSTITUTE FOR INNOVATION AND COMPETITION RESEARCH PAPER SERIES 14 (Max Planck Inst. for Innovation & Competition Research, Research Paper No. 19-10, 2020) (stating that transparency can emerge where the learning dataset and source code are disclosed together within information regarding its design and parameters); Andrew Burt, *The AI Transparency Paradox*, HARV. BUS. REV. (Dec. 13, 2019), <https://bit.ly/2MRLEAB> (implying that transparency refers to a lack of opaque algorithms); Alex Engler, *The Case for AI Transparency Requirement*, BROOKINGS INST. (Jan. 22, 2020), <https://brook.gs/2zoSwCn> (suggesting that a lack of transparency refers to anonymity in the AI context).

6. I define artificial intelligence broadly, to include various umbrella terms such as "machine learning," "big-data analytics," "deep learning," "smart machines," "neural networks," "learning algorithms," and "reinforcement learning." For the purposes of this Article, I do not parse differences in the technological meanings of these terms or delve deeply into specific technological techniques. However, I briefly summarize the basic properties of AI in Part I of this Article. It is sufficient to note that AI is not a monolith.

7. See Arti K. Rai, *Machine Learning at the Patent Office: Lessons for Patents and Administrative Law*, 104 IOWA L. REV. 2617, 2625 (2019) (viewing "explainability on a spectrum with 'complete' explainability meaning that the algorithm's complete decisionmaking can be made fully understandable to the relevant human audience"); Finck, *supra* note 5, at 15 (defining "explainability" as the opening of the black box or an analysis of the computational logic that transforms an input into an output, as well as an examination of the system by comparing its inputs and outputs; also suggesting that explainability is narrower than transparency).

inscrutable⁸ (not understandable as to how, or from which method it was produced).⁹ Enablement, which provides the statutory form of disclosure, has presented problems in many technological domains, and AI presents a new challenge due to the opaqueness and its rapid use acceleration that necessitates an reevaluation of disclosure theory. While other AI-based patent law scholarship¹⁰ has assessed these issues as they pertain to

8. See Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* 1085, 1094 (2018) (defining “inscrutability” in this context as “a situation in which the rules that govern decision-making are so complex, numerous, and interdependent that they defy practical inspection and resist comprehension”); Katherine J. Strandburg, *Rulemaking and Inscrutable Automated Decision Tools*, 119 *COLUM. L.R.* 1851, 1851 (2019) (describing that inscrutability undermines the traditional function of explanation (in rulemaking) and suggesting that inscrutability refers to difficulty in assessing “whether decision criteria will generalize to unusual cases or new situations and heightens communication and coordination barriers between data scientists and new situations and heightens communication barriers between data scientists and subject matter experts”).

9. Consider as an example U.S. Patent No. 6,845,336 (Jan. 18, 2005), owned by Plaintiff Neochloris and asserted against Defendants Emerson Process Management LLP and CITGO Petroleum Corporation in *Neochloris, Inc. v. Emerson Process Mgmt, LLLP*, where the first patent claim states:

In a treatment facility adapted for use with water wherein pollutants are removed from the water and wherein waster exiting from the facility has various acceptable and predetermined effluent quality parameters; said facility including a plurality of operational sensors for operably determining process water quality conditions while the water is in the facility; the improvement comprising:

- a) a monitoring computer at a site remote from said facility and including software to receive data from said sensors, analyze the water quality conditions inputted by said sensors and predict effluent water quality and process upsets; said monitoring computer further including an artificial neural network module to determine solutions to actual and potential water quality and process upsets; and
- b) an internet interface operably connecting said computer to said sensors for transferring said process water quality conditions from said sensors to said monitoring computer and transferring said solutions from said monitoring computer to said facility.

Neochloris, Inc. v. Emerson Process Mgmt., LLLP, 140 F.Supp. 3d 763, 773 (N.D. Ill. 2015) (emphasis added).

Here, the lack of explainability refers to the “artificial neural network module” (an AI-based tool), which utilizes an AI-method (an artificial neural network), and which the district court found akin to a black box without any limitations on the system. As their reasoning suggests, the interaction between the artificial neural network’s operation would not be discernible to the relevant audience (even after reading the patent’s detailed description). See *id.* Here, inscrutability refers to one’s inability to understand how or by what method the water-quality information (AI-generated output, or data produced using an AI-based tool) was produced. See *id.*

10. See Daniel Gervais, *Exploring the Interfaces Between Big Data and Intellectual Property Law*, 10 *J. INTELL. PROP., INFO. TECH. & ELECTRONIC COM. L.* 22, 23 (2019) (reviewing the application and reach of IP rights to AI and big data, including ways to adapt IP rights, analysis of infringement of IP rights, the issue of right in AI and big data itself, and incentives for additional rewards through an assessment of patent, copyright, *sui generis* database rights, data exclusivity, and trade secrets); Brian S. Haney, *AI*

inventorship¹¹ and non-obviousness,¹² this Article addresses the prevailing gap and understudied phenomenon of AI disclosure. It claims that AI fundamentally challenges disclosure in patent law (“patent disclosure”), which has not kept up with rapid advancements in AI technology, and seeks to invigorate the goals that patent law’s disclosure function is thought to serve for society.

To understand the theories, normative assessments, and prescriptions of this Article, some preliminary examples are helpful. Businesses are increasingly identifying and utilizing some form of AI in developing or delivering goods and services. Even those whose business is driven primarily by customer experience rather than technology are leveraging AI. AI is infiltrating virtually all business sectors and

Patents: A Data Driven Approach, 19 CHI.-KENT J. INTELL. PROP. (forthcoming 2020); Brian S. Haney, *Patents for NLP Software: An Empirical Review*, 22 N.C.J.L. & TECH. (forthcoming 2020); Ana Ramalho, *Patentability of AI-Generated Inventions: Is a Reform of the Patent System Needed?*, INST. INTELL. PROP., FOUND. FOR INTELL. PROP. JAPAN, 2018, at 1, 2, <https://bit.ly/3eq7O8K> (assessing international harmonization of the inventive step in response to AI-generated inventions); Hyunjong Ryan Jin, *Think Big! The Need for Patent Rights in the Era of Big Data and Machine Learning*, 7 N.Y.U. J. INTELL. PROP. & ENT. L. 78, 81 (2018) (assessing whether patents are needed in an era of AI and big data, while determining how to apply patent protection in AI); Charlotte A. Tschider, Presentation at Intellectual Property Scholars Conference at DePaul University College of Law: Patenting Artificial Intelligence (Aug. 8, 2019), <https://bit.ly/38HXAYT>.

11. See W. Michael Schuster, *Artificial Intelligence and Patent Ownership*, 75 WASH. & LEE L. REV. 1945 (2018) (utilizing a law and economics approach by employing the Coase Theorem to analyze how the United States Patent & Trademark Office should grant AI patents); Shlomit Yanisky Ravid & Xiaoqiong Liu, *When Artificial Intelligence Systems Produce Inventions: The 3A Era and an Alternative Model for Patent Law*, 39 CARDOZO L. REV. 2215, 2218, 2223, 2232–39 (2018) (responding to challenges arising from patent law’s inventor-identification requirement by using a “Multiplayer Model” to show that AI presents a realm with multiple stakeholders with varying interests); Robin Feldman & Nick Thieme, *Competition at the Dawn of Artificial Intelligence*, at 3 (U.C. HASTINGS SCH. L., LEGAL STUD., Research Paper No. 298, 2018) (proposing, in response to AI, that “inventorship could be assigned to the humans, to both humans and the computers, exclusively to the computers, or to no one at all”); Abbott, *supra* note 1 (suggesting that creative computers that utilize AI should be considered inventors based on the Patent and Copyright Clause of the United States).

12. See Ryan Abbott, *Everything Is Obvious*, 66 UCLA L. REV. 2, 8 (2018) (suggesting that as AI technology improves, the bar for non-obviousness will be raised, and when taken to the extreme, the non-obviousness doctrine will lose importance since inventions themselves will become obvious to AI-based machines); Patric M. Reinbold, *Taking Artificial Intelligence Beyond the Turing Test*, WIS. L. REV., Apr. 15, 2020, at 1 (proposing a standard for AI obviousness inquiries that “accounts for the inventor’s objectives, access to big and deep data, and knowledge of the existing datasets to control the form and operation of the machine learning resulting in AI-assisted inventions”). See generally Liza Vertinsky, *Thinking Machines and Patent Law*, in 18 RESEARCH HANDBOOK ON THE LAW OF ARTIFICIAL INTELLIGENCE (Woodrow Barfield & Ugo Pagallo eds., 2018) (labeling AI as a “thinking machine” to suggest that patent law faces a conceptual disconnect challenge, doctrinal challenges, and practical issues that necessitate patentability policy responses, including whether the person-having-ordinary-skill-in-the-art standard should be modified to include such “thinking machines”).

performing numerous functions, and has been popularized by its use in Amazon Go's cashier-less grocery store,¹³ Google Deep Mind's AlphaGo system's win over a world-champion "Go" player,¹⁴ IBM's Watson's win in *Jeopardy!*,¹⁵ and numerous virtual assistants, including Apple's Siri and Amazon's Alexa.¹⁶ Examples of the pervasive use of AI are well-known by the media and consumers, pose new types of challenges in commerce, and raise new scholarly debate on how AI will impact the law.¹⁷ The proliferation of AI in the invention process and at the United States Patent & Trademark Office (USPTO) may be less known but presents similar challenges for patent law and for society.¹⁸

Consider two case studies as patent-law application examples that demonstrate the urgency for a fresh look at disclosure theory in patent law and necessitate a response to the problems posed by the lack of

13. See Vishrut Shivkumar & Rishab Mehta, *Amazon Go: The Future of Retail*, 3 INT'L. J. ACAD. RES. & DEV. 646, 646 (2018) (describing the use of AI, specifically deep learning computer vision combined with sensor fusion for Amazon's cashless, automated supermarket as "a type of store which has fast checkouts and few employees [with] quick service and advanced consumer technology [where] anyone can lift a product off [the store's] shelves, walk over to the checkout counter which is automated, scan their phones at the scanner and an online transaction takes place").

14. See David Silver et al., *Mastering the Game of Go with Deep Neural Networks and Tree Search*, 529 NATURE 484, 484–89 (2016) (stating, "[t]he game of Go has long been viewed as the most challenging of classic games for artificial intelligence due to its enormous search space and the difficulty of evaluating board positions and moves," and further explaining how "[t]his is the first time that a computer program has defeated a human professional player in the full-sized game of Go, a feat previously thought to be at least a decade away"); *AlphaGo*, DEEPMIND, <https://bit.ly/3gq26Wx> (last visited Jan. 7, 2020) ("AlphaGo is the first computer program to defeat a professional human Go player, the first to defeat a Go world champion, and is arguably the strongest Go player in history.").

15. See Louise Beltzung, *Watson Jeopardy! A Thinking Machine*, VIENNA U. TECH., Dec. 2013, at 1, 2, 4–7 (describing IBM's Watson win in *Jeopardy!*, where Watson's use of AI resulted in its victory over two human *Jeopardy!* champions over a span of three days, as an example of the AI domain); David Ferrucci et al., *Building Watson: An Overview of DeepQA Project*, AI MAGAZINE, Fall 2010, at 59, 77 (explaining that IBM Research built an AI system capable of "performing at human expert levels in terms of precision, confidence, and speed" in real time on *Jeopardy!*).

16. See Daniel Castro & Joshua New, *The Promise of Artificial Intelligence*, CTR. FOR NEW DATA INNOVATION, Oct. 2016, at 2, 3.

17. See Giuseppe Contissa et al., *Towards Consumer-Empowering Artificial Intelligence*, INT'L JOINT CONF. ON ARTIFICIAL INTELLIGENCE ORG., 2018, at 5150, 5150–51 (suggesting that the current and pervasive use of AI challenges consumer law since "consumer data [is] continuously collected by on-line and off-line consumer behavior tracking . . . [to] elicit further information about consumers through profiling [leading to] . . . attempts to influence their behaviour").

18. See Feroz Ali, *Digitalised Invention, Decentralised Patent System: The Impact of Blockchain and Artificial Intelligence on the Patent Prosecution*, ARTIFICIAL INTELLIGENCE & INTELL. PROP., June 1, 2020, at 1–2, 16–18, 22 (asserting that artificial intelligence has transformed the representation of the invention before the patent office and has influenced the standards of patent prosecution by introducing a digital representation of the invention that changes the foundation of the entire patent system).

transparency of AI. These examples demonstrate how AI is rapidly destabilizing patent law's disclosure theory—first, an AI-based tool for computational experimentation¹⁹ to generate AI-generated output of nanomaterial compounds and, second, an AI-based tool for predicting AI-generated output of optimized cooking recipes for personalized nutrition. The first case study stems from the turn of the twentieth century when chemical companies handled research and development for new compounds in essentially the same manner as they did at the turn of the nineteenth century; scientists and technicians would spend countless years experimenting in the laboratory, collecting data from the experiments, and making changes to reactions in hopes of making a particular compound. This process was labor-intensive and costly, and, while it entailed some scientific intuition, it involved much trial and error. Computational chemistry companies, such as Schrödinger,²⁰ and other innovative startups identified and solved this problem by computationally designing chemicals with AI-based tools to produce AI-generated output that could be claimed in the form of optimal performance, properties, and structures of compounds without running physical experiments.²¹ A patent examiner assessing patent claims based on the AI-generated output would be unable to determine that an inventor utilized an AI-based tool without adequate disclosure but could grant the inventor an early patent grant.²²

The second case study concerns AI-based tools that have moved beyond optimization to iteratively improve parameters to enable a personalized solution, which is more difficult and involves many more variables. Digital health companies, such as Yummly,²³ and innovative startups are developing AI-based tools that can be patented and can generate AI-generated output of personalized recipes tailored to a particular individual's diet, allergens, and biological characteristics. However, the rise in available computing power and sources of data, the scale and plethora of new use cases, and the ability to use developments

19. See Tabrez Y. Ebrahim, *Computational Experimentation*, 21 VAND. J. ENT. & TECH. L. 591, 601–03 (2019) (describing techniques such as (1) molecular dynamics simulations and machine learning, which allow for the determination of reactivity and kinetics of chemical compounds; (2) high-precision, pre-laboratory-experiment simulations and predictions of chemical properties and reactions; and (3) construction of structure-activity relationships).

20. See Schrödinger Announces Closing of Initial Public Offering, BUSINESS WIRE (Feb. 10, 2020, 4:05 PM), <https://bwnews.pr/3cbU4gg>.

21. See Ebrahim, *supra* note 19, at 595–96.

22. See Michael Abramowicz, *The Danger of Underdeveloped Patent Prospects*, 92 CORNELL L. REV. 1065, 1071 (2007) (suggesting that embryonic inventions that result in early-granted patents present great concerns for underdevelopment).

23. See Ingrid Lunden, *Whirlpool Acquires Yummly, The Recipe Search Engine Last Valued at \$100M*, TECHCRUNCH (May 4, 2017, 9:54 AM), <https://tcrn.ch/3c4M2Gc>.

in algorithms to perform complex computations and provide customization neither enable a person having ordinary skill in the art (“PHOSITA”) to practice the invention²⁴ nor prove that the inventor was in possession of the invention at the time of the patent application.²⁵ In both case studies, patents benefit the innovative AI companies and their inventors but also present new challenges for the theoretical justifications and economic theories of the patent system. Specifically, the use of AI in the process of inventing increases the scale of and lack of transparency in inventing, possibly resulting in premature patent grants or ones where the inventor arguably was not in possession of the invention.

The traditional justifications²⁶ for patents are based on *ex ante* labor theory and reward theory.²⁷ The utilitarian economic theory for patents is based on incentives, such that innovation would be at suboptimal levels

24. See 35 U.S.C. § 112(a) (2018); *Janssen Pharmaceutica N.V. v. Teva Pharms. USA, Inc.*, 583 F.3d 1317, 1320–23 (Fed. Cir. 2009); Sean B. Seymore, *Foresight Bias in Patent Law*, 90 NOTRE DAME L. REV. 1105, 1115 (2015); Sean B. Seymore, *Heightened Enablement in the Unpredictable Arts*, 56 UCLA. L. REV. 127, 137 (2008); Sean B. Seymore, *Making Patents Useful*, 98 MINN. L. REV. 1046, 1066–67 (2014); Sean B. Seymore, *The Enablement Pendulum Swings Back*, 6 NW. J. TECH. & INTELL. PROP. 278, 279 (2008).

25. See *Amgen Inc. v. Sanofi*, 872 F.3d 1367, 1373 (Fed. Cir. 2010); *Ariad Pharm., Inc. v. Eli Lilly & Co.*, 598 F.3d 1336, 1351 (Fed. Cir. 2010); MPEP (9th ed. Rev. 8, Jan. 2018) § 2163 (explaining that a patent applicant must show possession of the claimed invention, which can be done with such “descriptive means as words, structures, figures, diagrams, and formulas that fully set forth the claimed invention” or by a “description of an actual reduction to practice” or by demonstrating that the invention was “ready for patenting”); Alan L. Durham, *Patent Scope and Enablement in Rapidly Development Arts*, 94 N.C. L. REV. 1099, 1105 (2016); Timothy R. Holbrook, *Patent Anticipation and Obviousness as Possession*, 65 EMORY L.J. 987, 990–91 (2016); Timothy R. Holbrook, *Possession in Patent Law*, SMU L. REV., Winter 2006, at 123, 158.

26. See M. Du Bois, *Justificatory Theories for Intellectual Property Viewed Through the Constitutional Prism*, PER/PELJ, Mar. 16 2018, at 1, 2, 21 (suggesting that “where the dissemination of information is promoted by the grant of intellectual property rights, the reward theory seems appropriate” but recognizing that some newer theories may be justified by unconventional forms of intangible property interests); Adam Karbowski & Jacek Prokop, *Controversy Over the Economic Justifications for Patent Protection*, PROCEEDIA ECON. & FIN. 5, 2013, at 393–94 (pointing out, while reviving the debate on the economic justifications for patent protection, that the main justifications of an inventor’s right to patent protection include “natural law, reward in the form of a monopoly, incentives created by the monopoly profits, and compensation for revealing the secret”); Wendy Lim, *Towards Developing a Natural Law Jurisprudence in the U.S. Patent System*, SANTA CLARA COMPUTER & HIGH TECH. L.J., May 2003, at 559, 561; A. Samuel Oddi, *Un-Unified Economic Theories of Patents—The Not-Quite Holy Grail*, 71 NOTRE DAME L. REV. 267, 269 (1996).

27. See Reto M. Hilty et al., *Intellectual Property Justification for Artificial Intelligence 4* (MAX PLANCK INST. FOR INNOVATION & COMPETITION, Research Paper No. 20-02, 2020) (suggesting that the labor theory allows people to own property rights based on the fruits of their labor and that the reward theory refers to providing enough reward to inventors to recover their innovation costs).

absent the incentives.²⁸ Scholars have introduced other justifications and economic reasoning for the patent system,²⁹ and each has shed a new perspective on the consequences of the scope, duration, or enforcement of patents, as well as commercial influence. The traditional and more modern justifications and theories for the patent system³⁰ are not mutually exclusive and, surprisingly, scholars have paid little attention to the role of disclosure in the patent system's treatment of AI inventions. As technological advances in information and data have made it easier for inventors to mask and obfuscate their innovations,³¹ and as patent law has doctrinally moved away from its historical origins in the physical world,³² disclosure theory has been brought to the forefront.³³ As the USPTO grants more patents covering AI-based tools³⁴ as a policy goal,³⁵

28. See Jay P. Kesan, *Economic Rationale for the Patent System in the Current Context*, 22 GEO. MASON L. REV. 897, 898–99 (2015).

29. See J. Janewa Osei-Tutu, *Humanizing Intellectual Property: Moving Beyond the Natural Rights Property Focus*, 20 VAND. J. ENT. & TECH. L. 207, 210 (2017) (introducing a human rights framework for intellectual property law); Peter Lee, *Towards a Distributive Agenda for U.S. Patent Law*, 55 HOUS. L. REV. 321, 323, 325 (2017) (suggesting that distributive justifications, such as providing access to inventions and serving marginalized communities, are consonant with and advance the normative foundations of the U.S. patent system); Shlomit Yanisky-Ravid, *The Hidden Thought Flourishing Justification of Intellectual Property Laws: Distributive Justice, National Versus International Approaches*, 21 LEWIS & CLARK L. REV. 1, 2, 17–18 (2017) (suggesting that intellectual property law embodies distributive justice principles and justifications by proposing that intellectual property and distributive justice are “born together” and can “get along” in harmony).

30. See generally ROBERT P. MERGES, *JUSTIFYING INTELLECTUAL PROPERTY* (2011) (providing foundation for justifying intellectual property by describing theories by Locke and Kant, considering distributive justice and redistribution, and tying together disparities and bridging different foundational viewpoints).

31. See Bronwyn H. Hall & Meagan MacGarvie, *The Private Value of Software Patents*, 39 RES. POL'Y 994, 1003–05 (2010) (summarizing critique of information and software patenting, including being of low quality, lacking adequate prior art, not including source code implementation, and being vague and broadly worded).

32. See Christopher A. Cotropia, *Physicalism and Patent Theory*, 69 VAND. L. REV. 1543, 1545 (2016).

33. See Ali, *supra* note 18, at 29–30 (stating that the use of AI and machine learning to produced “digitalized invention[s]” will allow for disclosure to “not [be] used strategically and done by the user”).

34. See DEAN ALDERUCCI ET AL., CARNEGIE MELLON UNIV., *MAPPING THE MOVEMENT OF AI INTO THE MARKETPLACE WITH PATENT DATA 3* (2019), <https://bit.ly/2ARSnru>; Michael S. Borella et al., *Artificial Intelligence-based Patents: Perspectives for Practitioners and Patent Owners*, MCDONNELL BOEHNEN HULBERT & BERGHOF LLP: SNIPPETS, 2019, at 8, 8, <https://bit.ly/2ZtxqW> (demonstrating with Figure 1 “AI/ML-based Patent Publications by Year (USPTO)” rapidly increasing patent activity in AI technologies).

35. See Tom Simonite, *Despite Pledging Openness, Companies Rush to Patent AI Tech*, WIRED (July 31, 2018, 7:00 AM), <https://bit.ly/2ZxAJ6Q>; ALDERUCCI ET AL., *supra* note 34, at 4–5 (pointing out policy considerations based on the impact of AI on labor demand and wages, future wages, and employment); Remarks by Director Iancu at the AI Event, *supra* note 3 (“AI has significant implications for intellectual property law,

and since AI is fundamentally transforming patentability³⁶ with complex and nonintuitive algorithms³⁷ to produce AI-generated output, society needs inventors to provide greater disclosure of AI in patents. More precisely, nonintuitive data relationships between AI-based tools and AI-generated output should compel AI inventors to strive for greater patent disclosure.³⁸

Disclosure refers to patents being awarded as the quid pro quo for disclosing the invention through enablement and written description.³⁹ Patent disclosure is a theory that underlies the statutory requirements of demonstrating enablement and written description, and an invention must satisfy both the enablement and written description requirements to meet patentability.⁴⁰ A patent examiner may reject patent claims that, due to a

the economy, and America's position as the global innovation leader. . . . Not surprisingly, AI is changing the landscape of intellectual property policy.”).

36. See generally Ben Hattenbach & Joshua Glucoft, *Patents in an Era of Infinite Monkeys and Artificial Intelligence*, 19 STAN. TECH. L. REV. 32 (2015) (suggesting that AI-generated inventions raise important patent law questions on patent protection without human intervention, while assessing whether technologies invented by AI techniques are worthy of patent protection); Charles Duan, *Patenting Artificial Intelligence Inventions*, R STREET (Nov. 13, 2019), <https://bit.ly/30CUJp0> (providing normative recommendations for Congressional policy changes for AI patenting, in response to the USPTO's request for comments on AI inventions); Telephone Interview with Sikander Khan, Attorney, Haynes Beffel & Wolfeld (Oct. 3, 2019) (addressing patentability concerns in patent law practice raised by a number of technology clients for whom AI is a critical and core part of business); FIRTH-BUTTERFIELD & CHAE, *supra* note 1 (suggesting patentability implications and changes are paramount in the face of AI's rapid technological changes in order for the U.S. patent system to achieve its objectives); Ramalho, *supra* note 10 (suggesting that the substitution of human ingenuity in the invention process or generating inventions without much human input requires recalibrating patentability).

37. See W. Nicholson Price II & Arti K. Rai, *Clearing Opacity Through Machine Learning*, 106 IOWA L. REV. (forthcoming 2020).

38. See DISCOVER #OXFORDAI: THE UNIVERSITY OF OXFORD GUIDE TO ARTIFICIAL INTELLIGENCE 5–6 (U. of Oxford, 2d ed. 2019); see also DAVID FREEMAN ENGSTROM ET AL., GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES 75 (2020); Roman V. Yampolskiy, Unexplainability and Incomprehensibility of Artificial Intelligence (June 20, 2019) (unpublished manuscript), <https://bit.ly/3dQSB0G>.

39. See Jeanne C. Fromer, *Patent Disclosure*, 94 IOWA L. REV. 539, 546 (2009) (explaining that the technical layer in the patent specification describes the invention, that the written description “ensures that the inventor is in possession of the claimed invention,” and that enablement requires the patent applicant to “demonstrate in the specification to ‘any person skilled in the [relevant] art [how] . . . to make and use the [invention] without undue experimentation”); Lisa Larrimore Ouellette, *Do Patents Disclose Useful Information*, 25 HARV. J.L. & TECH. 531, 537 (2012) (referencing the Federal Circuit, which has reaffirmed that the written description requirement is separate from enablement, and that “the test for sufficiency [of written description] is whether the disclosure of the application relied upon reasonably conveys to those skilled in the art that the inventor had possession of the claimed subject matter as of the filing date”).

40. See Joseph Scott Miller, *Enhancing Patent Disclosure for Faithful Claim Construction*, 9 LEWIS & CLARK L. REV. 178, 179 (2005).

lack of enablement or written description for an invention, are not reproducible by a person having ordinary skill in the art. While scholars have recognized that patent disclosure has an interrelated nature of teaching others and limiting patent claim scope,⁴¹ another skeptic has argued that patent disclosure is an ancillary feature that conflicts with incentives to invent and commercialize.⁴² Previous scholarship debated the justifications of disclosure in patent law, but both viewpoints agree that disclosure could be improved in the patent system and would increase social welfare.⁴³ Information and software-technology inventors need to disclose more accurate information about the underlying invention⁴⁴ and have been criticized for their use of vague “patent-speak” jargon.⁴⁵ Rather than focus on the debate surrounding the justification for disclosure in the patent system and lament the inherent opacity⁴⁶ problems in disclosing information and software technologies,⁴⁷ this Article asks: given existing U.S. interests in patenting AI in general⁴⁸—

41. See Jason Rantanen, *Patent Law’s Disclosure Requirement*, 45 LOY. U. CHI. L.J. 369, 370–71 (2013).

42. See Allan Devlin, *The Misunderstood Function of Disclosure in Patent Law*, 23 HARV. J.L. & TECH. 401, 404 (2010).

43. See Colleen V. Chien, *Contextualizing Patent Disclosure*, 69 VAND. L. REV. 1849, 1860–61 (2016) (quoting surveys that demonstrate that patent disclosure disseminates technical information in society); Devlin, *supra* note 42, at 411–12; Fromer, *supra* note 39, at 597 (explaining that “disclosure should increase social welfare because inventors provide the information they already possess and outsiders need not try to learn about the invention on their own”); Miller, *supra* note 40, at 196.

44. See Hall & MacGarvie, *supra* note 31, at 9–11 (describing the critique of software patentability, explaining that the perception of software patents is that the disclosure function is badly served since they rarely include the source code implementation and are often vaguely and broadly worded).

45. See Robin Feldman, *Plain Language Patents*, 17 TEX. INTELL. PROP. L.J. 289, 292–93 (2009) (characterizing patent law’s word interpretation as an “arcane code,” “downright incomprehensible under common sense notions of language,” and “code-like communication”); Devlin, *supra* note 42, at 403 (characterizing patents in the information technology industry as being notorious for their vague language).

46. See Simon Chesterman, *Through a Glass, Darkly: Artificial Intelligence and The Problem of Opacity* 1–26 (Nat’l Univ. of Sing., Working Paper No. 2020/011, 2020) <https://bit.ly/30BQavc> (suggesting that as computer programs become more complex, the ability to understand a given output diminishes).

47. See generally Sylvia Lu, *Algorithmic Opacity, Private Accountability, and Corporate Social Disclosure in the Age of Artificial Intelligence*, 23 VAND. J. ENT. & TECH. L. 1 (forthcoming 2021) (summarizing the conflict between algorithmic opacity and democratic transparency in incentives for information disclosures of AI algorithms).

48. See Andrei Iancu, AI Policy Update, USPTO (Feb. 6, 2020) [hereinafter Iancu AI Policy Update] (“One of the [USPTO’s] top priorities is to ensure that the United States maintains its leadership in . . . AI. To that end, the USPTO has been actively . . . promot[ing] the predictability and reliability of IP rights relating to AI technology”); Exec. Order No. 13859, 84 Fed. Reg. 3967 (Feb. 11, 2019); NATIONAL SCIENCE & TECHNOLOGY COUNCIL, SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE, THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN: 2019 UPDATE, at iii (June 2019) (specifying that the U.S. has made “concerted efforts to promote and

or more specifically with AI-based tools and of AI-generated output—should society want to enforce more robust AI disclosure requirements, and if so, how? More specifically, as AI inventions magnify disclosure problems of information and software technologies with insufficient transparency, are our current disclosure requirements sufficient for patentability of AI-based tools, and are changes in patent examination necessary in response to calls for stronger disclosure of AI-generated output?

The USPTO has recognized the importance of patent disclosure as AI patent filings have rapidly increased and comprise a significant portion of *identified* inventions, as reflected in the USPTO Director’s comments and disclosure-related question: “We know that . . . AI . . . is included in more than one-third of all *identified* inventions. . . . What level of detail is necessary in a patent disclosure as to the structure and function [of the underlying inner workings]?”⁴⁹ Moreover, Brenda Simon and Ted Sichelman’s 2019 article, *Data-Generating Patents*,⁵⁰ introduced the term “data-generating patents” to refer to patents on technologies that generate valuable data by their operation or use while exploring the use of patents and trade secrets as complements. The authors recognized that patent disclosure requirements are not always rigorous, but their analysis did not explore the relation between data generation, such as by AI-based tools, and adequate disclosure of their

protect AI technology and innovation,” with eight strategic priorities that include a desire to: (1) “[m]ake long-term investments in AI research”; (2) “[d]evelop effective methods for human-AI collaboration”; (3) “[u]nderstand and address the ethical, legal, and societal implications for AI”; (4) “[e]nsure the safety and security of AI systems”; (5) “[d]evelop shared public databases and environments for AI training and testing”; (6) “[m]easure and evaluate AI technologies through standards and benchmarks”; (7) “[b]etter understand the national AI R&D workforce needs”; and (8) “[e]xpand public-private partnerships to accelerate advances in AI”); ALDERUCCI ET AL., *supra* note 34, at 3–4 (alluding to the national strategic interest of AI by providing evidence and statistics that “the U.S. is the overwhelmingly dominant source of AI invention” and noting that “U.S.-based inventors account for an overwhelming majority of granted AI patents”); Simonite, *supra* note 35.

49. Andrei Iancu, Remarks by Director Iancu at the International Conference on AI-Emerging Technologies and IP, USPTO (July 16, 2019). The USPTO Director further stated:

AI published applications grew by 400% in the past decade. At the USPTO, AI technologies are part of about 26% of annual patent filings, which is a 34% increase in the share of AI patent filings since 2005. And we have doubled the number of examiners at the USPTO reviewing AI applications.

Id.

50. See Brenda M. Simon & Ted Sichelman, *Data-Generating Patents*, 111 NW. L. REV. 377, 377–78 (2019) (introducing the term “data-generating patents” to refer to “patents over inventions that generate unique data from users” and that with such “inventions involving technologies that by design generate valuable data through their operation or use,” the result being that the patentee also effectively “enjoys market power over the data generated by the invention”).

application to produce AI-generated output. Additionally, the USPTO and the U.S. Supreme Court⁵¹ have recognized the importance of patent disclosure, but these views only address the importance of greater and clearer disclosure of *identified* inventions. Their views do not address the challenge of AI-based tools and AI-generated output that could provide *unidentified* inventions⁵² (inventions based on AI-generated output but that appear as if they were invented by humans without use of AI-based tools), which, as this Article subsequently explains in more depth, stems from the transparency challenge of AI.

This Article argues that the lack of transparency in AI-based tools that produce AI-generated output results in *unidentified* AI inventions. In particular, it asserts that explainability and inscrutability have normative ramifications for the adequacy of disclosure of the inner workings of AI-based tools and the inventive method of AI-generated output. There are underappreciated benefits of patent disclosure of AI inventions, but they require evaluation of the marginal costs relative to the marginal benefits. In exploring patent disclosure of AI inventions, this Article also analyzes the related issues of whether current patent examination at the USPTO should evolve from the standpoint of AI capabilities⁵³ and detection of adequate disclosure.⁵⁴ The issue of the USPTO's patent examination of AI inventions is largely derivative of the normative ramifications of AI disclosure required for inventors and stems from the possibility of patent applicants flooding the USPTO with *unidentified* inventions.⁵⁵ The normative implications of patent disclosure of AI inventions are a response to the proliferation of AI inventions and the question of how society should develop a better understanding of interaction of patent law with AI. This Article analyzes whether and how a lack of AI transparency undermines the interrelated teaching and scope-limiting purposes of patent disclosure⁵⁶ and will stretch patent examination to its breaking point both conceptually and practically. On a prescriptive level,

51. See *Bonito Boats, Inc. v. Thunder Craft Boats, Inc.*, 489 U.S. 141, 151 (1989); *Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 481 (1974).

52. See *infra* Section I.B (defining “unidentified inventions” as inventions based on AI-generated output but that appear as if they were invented by humans without use of AI-based tools).

53. See Tabrez Y. Ebrahim, *Automation & Predictive Analytics in Patent Prosecution: USPTO Implications & Policy*, 35 GA. ST. U. L. REV. 1185, 1185, 1189 (2019) (prescribing that “the USPTO should develop a counteracting artificial intelligence unit in response to artificial-intelligence proliferation[,]” in part, “based on economic efficiency views [that] would impact fairness, time, and transparency policy considerations”).

54. See Ebrahim, *supra* note 19, at 650.

55. See Ebrahim, *supra* note 53, at 1236, 1239, 1242.

56. See Rantanen, *supra* note 41, 370–71, 373–74 (explaining that a reason for patent law's disclosure requirement is to “limit the maximum scope of patent claims [by] establish[ing] the outer boundaries of what the applicant might claim”).

it argues that disclosure of the inventive method, such as the use of AI tools for production of AI-generated output, is wholly appropriate and advances the patent system's objectives. In so arguing, it sketches the contours of an enhanced disclosure for AI inventions.

Turning wholly to the prescriptive, this Article proposes recommendations for improving the patent system's promotion of disclosure of AI inventions in multiple contexts, including disclosure incentives and a deposit requirement. It asserts that enhanced AI disclosure could be achieved through patent policy calibrations, starting with voluntary disclosure and continuing with increased disclosure incentives. Such incentives range from: (1) prioritized examination to (2) reduced maintenance fees to (3) greater patent terms to enable longer patent protection to (4) a working model requirement as a prerequisite for a complete patent application. In important ways, the adequacy of disclosure of AI inventions depends on one's normative vision of sufficient quid pro quo. Policymakers may regard incentivizing greater disclosure of an AI method of invention as posing greater costs for patent examination and for the inventor, yet it may prevent a socially harmful tragedy of the anticommons⁵⁷ with unclear competing rights to upstream training data⁵⁸ and algorithms that could prevent downstream innovation.⁵⁹

57. See Michale A. Heller & Rebecca S. Eisenberg, *Can Patents Deter Innovation? The Anticommons in Biomedical Research*, 280 *SCIENCE*, May 1, 1998, at 698, 698 (defining a “tragedy of the anticommons” as when “a resource is prone to underuse when multiple owners can each have a right to exclude others from a scarce resource and no one has an effective privilege of use”); Michael Heller, *The Tragedy of the Anticommons, A Concise Introduction*, 76 *MODERN L. REV.*, Jan. 2, 2013, at 1, 6 (specifying that the tragedy of the anticommons thesis is “when too many people own pieces of one thing, nobody can use it . . . [and the result is] wasteful underuse. . . . [C]ooperation breaks down, wealth disappears, and everybody loses”); Michael A. Heller, *The Tragedy of Anticommons: Property in the Transition from Marx to Markets*, 111 *HARV. L. REV.* 621, 622 (1998) (defining a “tragedy of the anticommons” as “[w]hen too many owners hold such privileges of use[] [such that] the resource is prone to overuse” and when “multiple owners are each endowed with the right to exclude others from a scarce resource, and no one has an effective privilege of use[,]” and further suggesting that “difficulties of overcoming a tragedy of the anticommons suggest that policymakers should pay more attention to the content of the property bundles”); see also note 59 and accompanying text.

58. See Phillip Hacker, *A Legal Framework for AI Training Data*, 13 *L., INNOVATION & TECH.* (forthcoming 2021) (stating that training data is of fundamental importance for the development of AI applications). For a more detailed understanding of the term “training data,” see ROBERT MUNRO & QAZALEH MIRSHARIF, *THE ESSENTIAL GUIDE TO TRAINING DATA 10 (Figure Eight)* (ebook) (explaining training data for AI as, “[AI] machines learn examples (a.k.a. training data); training data is an important—if not more important—than the algorithm itself[,]” and that “high quality and high quantities of training data are the surest way to improve models; training data needs labeling to be truly useful; and training data labeled by humans is the most accurate way to do this”);

This Article advances legal scholarship on AI and makes a number of novel contributions. First, it unravels the black box of AI by providing a detailed description and taxonomy of AI in patenting and includes graphical representations that identify a typology with resulting policy arguments. Second, it introduces the notion of unidentified inventions in the form of AI-generated output and associated normative inquiries concerning patent or trade secret protection for innovators. Third, it applies theories, justifications, and normative assessments of disclosure to AI to explain the benefits of enhanced AI patent disclosure, while responding to objections. Fourth, it prescribes incentivizing enhanced AI patent disclosure and distinguishes procedural reform at the USPTO and legislative reform. Fifth, it concludes with a description of a future proposed empirical study to verify the theoretical, normative, and prescriptive underpinnings of this Article. Sixth, and as applied to broader AI legal scholarship, the underlying problematic social effects of AI identified in this Article extend beyond patent disclosure and into similar forces that scholars have investigated with copyright law⁶⁰ and trademark law.⁶¹

This Article unfolds in three parts. Part I briefly introduces AI technology through a taxonomy comprising of AI-based tools and AI-generated output. Furthermore, Part I introduces AI's transparency

Training and Test Sets: Splitting Data, GOOGLE MACHINE LEARNING CRASH COURSE, <https://bit.ly/3gz6x1a> (last visited Apr. 28, 2020).

59. The tragedy of the anticommons refers to a type of coordination breakdown in which there is single resource with numerous rightsholders that prevent others from using that single resource, such that it would frustrate a socially desirable outcome. In other words, there are competing rights, such as with patents, so that people underuse scarce resources, since there are too many owners that block each other. As it relates to AI, the single resource is training data; if too many people owned pieces of the training data, then nobody would use algorithms used to apply the training data. The result would be wasteful underuse since the ownership rights and regulatory controls would be too fragmented. The analogy of the tragedy of the anticommons is particularly pronounced with AI as explained in this Article, since the application of AI-based tools is “hidden” and would result in AI-generated output that would be hard to spot.

60. See Dan L. Burk, *Algorithmic Fair Use*, 86 U. CHI. L. REV. 283, 283, 288 (2019) (examining algorithmic mediation of copyright exceptions and exploring the implications of incorporating legal standards into algorithms, and arguing that social action of algorithmic systems may become the legal standards that they seek to implement); Mark A. Lemley & Bryan Casey, *Fair Learning* 1–64 (Mar. 23, 2020) (unpublished manuscript), <https://bit.ly/2ZLVqff> (arguing that AI should be able to use training databases regardless of whether the contents of those databases are copyrighted); Gervais, *supra* note 10, at 22, 25–27, 38 (suggesting that copyright's traditional role is challenged by AI, given the unstructured nature of non-relational databases).

61. See Sonia Katyal & Aniket Kesari, *Trademark Search, Artificial Intelligence, and the Role of the Private Sector*, 35 BERKELEY TECH. L.J. (forthcoming 2020) (discussing the impact of AI on private trademark search engines and their economic and legal implications and demonstrating that AI fundamentally transforms the trademark ecosystem through an empirical investigation of various search engines).

problem by describing how patent law's disclosure requirement is challenged by the concepts of explainability of an AI-based tool's operation and inscrutability of AI-generated output.

Part II introduces disclosure's function in the patent system. It then describes the normative debate between disclosure proponents and a skeptic and over disclosure's effects on social goals. The normative balance will be challenging if, as some believe, AI technology advances such that AI displaces humans from the inventive process with an accelerating pace of innovation toward artificial general intelligence.⁶² Part II argues that AI, unlike other forms of software, lacks reproducibility due to its dynamic black-box⁶³ nature, thus creating a disclosure deficit. Furthermore, Part II provides a normative assessment of greater AI patent disclosure on patentability and for patent examination. Part II ends with theoretical justifications of AI patent disclosure and introduces normative reasoning for enhanced AI disclosure.

Part III explores mechanisms for enhancing disclosure of AI-based tools and AI-generated output. It proposes prescriptions to fix the disclosure function of patents—including disclosure incentives and a data deposit requirement to provide patent protections. In so doing, it

62. See FOTIOS FITSILIS, IMPOSING REGULATION ON ADVANCED ALGORITHMS, at v (2019) (defining “artificial general intelligence” as “the broader goal of reaching human-like intelligence, rather than a single technology, or bunch of technologies . . . through the inception of intelligent algorithms”); John Linarelli, *Advanced Artificial Intelligence and Contract*, UNIFORM L. REV. (SPECIAL ISSUES ON TRANSNAT'L COM. L. & TECH./DIGITAL ECON.) (forthcoming 2019), <https://bit.ly/3iOr6YA> (introducing artificial general intelligence, or AGI, as the displacement of the human in the innovation process by explaining that “AI cognitive architecture (code or otherwise) [that] will become something roughly analogous to ‘gene’ and AI evolution will occur independently of the initial human intervention of AI creation,” but further noting that “AGI does not yet exist [but] AGI will come to exist in the future”); Simon Deakin & Christopher Markou, *Is Law Computable? From Rule of Law to Legal Singularity* 4 (Univ. of Cambridge Faculty of L. Research Paper 2020), <https://bit.ly/3gOjUKh> (suggesting that “the AGI (artificial general intelligence) hypothesis is that a machine can be designed to perform any ‘general intelligent action’ that a human is capable of”).

63. See W. Nicholson Price II, *Artificial Intelligence in Health Care: Applications and Legal Issues*, 14 SCI. TECH. LAW, Fall 2017, at 10, 10 (describing AI as relying on “algorithms [that] may be best described as *black-box*”) (emphasis added); see also JACK BELZER ET AL., *Models of Learning Systems*, ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY 40 (Allen Kent et al. eds., 1978); James J. Anton et al., *Policy Implications of Weak Patent Rights*, in 6 INNOVATION POLICY AND THE ECONOMY 1, 6 (2006) (defining a “black-box invention” as one “for which the added performance is obvious when the product or service is observed, but the means—the magic ingredients—by which the performance is achieved cannot be readily discerned or reverse engineered”); *Black-Box Medicine: Legal and Ethical Issues: A Health Policy and Bioethics Consortium*, PETRIE-FLOM CTR. FOR HEALTH L. POL'Y, BIOTECH., & BIOETHICS AT HARV. L. SCH. (Feb. 8, 2019, 12:00 PM), <https://bit.ly/2XfwCuo> [hereinafter *Black-Box Medicine*] (describing the “black box” of AI algorithms as opaque computational models that make decisions).

assesses benefits and costs for the inventor, patent examination, and society. Part III also builds a foundation for a more holistic AI innovation law and policy framework by showing how the prescriptions would impact the decision between patent protection and trade secrecy. Furthermore, Part III argues that greater disclosure of AI-based tools and AI-generated output is essential for growth of AI innovation and explains how maximizing AI disclosure would promote aggregate welfare.

I. FOUNDATIONS: HOW & WHY AI IS “ARTIFICIAL”

To understand the legal frameworks and normative implications of inventions covering AI-based tools or application of them, it is vital to understand the technological forces that have led to the contemporary use of AI in patent law. This Part begins by describing how AI is “artificial” by unraveling its “black box,”⁶⁴ which serves as the technological force for examining why AI produces “artificial” inventions. By providing a taxonomy of AI, this Part demonstrates how AI-based tools are used in the inventive process to produce AI-generated output without human intervention and why such “artificial” inventing raises new considerations for patent disclosure. A descriptive account of the typologies of AI illuminates an understanding of its legal and normative significance, as is seen in Parts II and III, respectively.

A. *The “Magic” of AI*

AI is not “magic,” but much has been written about AI as a sort of magic and with futurist tones. The speculative discussion of AI is fraught with unsupported assumptions and thus requires definition and demystification. AI is difficult to define⁶⁵ but generally refers to the use of computing systems for automating tasks that would normally require human intelligence.⁶⁶ The varying definitions of AI diffused the concept into meaningless buzz,⁶⁷ and the definition also changed with time due to rapid technological developments.⁶⁸

64. See Nicholson Price, *supra* note 63, at 10; *Black-Box Medicine*, *supra* note 63.

65. See Mark A. Lemley & Bryan Case, *You Might Be a Robot*, 105 CORNELL L. REV. 1 (forthcoming 2020); Jonas Schuett, A Legal Definition of AI, at 1 (Sept. 4, 2019) (unpublished manuscript), <https://bit.ly/3eARORC> (stating that the term “artificial intelligence” is highly ambiguous and has a vast spectrum of definitions to argue that “there is no definition of AI which meets the requirements for legal definitions”).

66. See Asay, *supra* note 1, at 1187; Surden, *supra* note 1, at 1307 (describing AI as using technology to automate tasks that normally require human intelligence).

67. See Foster Provost & Tom Fawcett, *Data Science and Its Relationship to Big Data and Data-Driven Decision Making*, 1 BIG DATA at 51, 52 (2013).

68. See H.R. 4829, 115th Cong. (2018). Congress defined “artificial intelligence” as anything that can:

(A) think like humans (including cognitive architectures and neural networks);

AI is implemented by software technology,⁶⁹ which raises the question as to what distinguishes AI from run-of-the-mill software. AI performs different functions from software since AI can imitate intelligent behavior with computer programs.⁷⁰ AI has been referred to as being a “black box,”⁷¹ a “thinking machine,”⁷² and “a learning system.”⁷³

-
- (B) act like humans (such as passing the Turing test using natural language processing, knowledge representation, automated reasoning, and learning);
 - (C) think rationally (such as logic solvers, inference, and optimization);
 - (D) act rationally (such as intelligent software agents and embodied robots that achieve goals via perception, planning, reasoning, learning, communicating, decision-making, and acting); or
 - (E) automate or replicate intelligent behavior.

Id.

See also Phillippe Aghion et al., *Artificial Intelligence and Economic Growth* (Nat’l Bureau of Econ. Research, Working Paper No. 23928, 2017) (defining AI as “the capability of a machine to imitate intelligent human behavior [or] an agent’s ability to achieve goals in a wide range of environments”); Roger Parloff, *Why Deep Learning is Suddenly Changing Your Life*, FORTUNE (Sept. 28, 2016; 5:00 PM), <https://bit.ly/3hoiffg> (defining “modern artificial intelligence” as “a vast range of technologies—like traditional logic and rules-based systems—that enable computers and robots to solve problems in ways that at least superficially resemble thinking”); FIRTH-BUTTERFIELD & CHAE, *supra* note 1, at 5 (defining AI as “a computerized system exhibiting behavior commonly thought of as requiring intelligence” or “a system capable of rationally solving complex problems or taking appropriate action to achieve its goals in real-world circumstances”); Semmler & Rose, *supra* note 1, at 86 (defining AI as “the process of simulating human intelligence through machine processes”); Smith et al., *supra* note 1, at 4 (defining AI as “a system which amplifies people’s own knowledge and understanding”).

69. See Andres Guadamuz Gonzalez, *The Software Patent Debate*, 1 J. INTELL. PROP. L. 1, 2 (2006) (explaining that software is not only source code, but to be able to operate in a computer, software has to be translated into object code by a process of compilation).

70. See Joost N. Kok et al., *Artificial Intelligence: Definitions, Trends, Techniques, and Cases*, in ARTIFICIAL INTELLIGENCE 1, 1–2 (2009), <https://bit.ly/3chsnD3>. The following definition of AI is based on The New International Webster’s Comprehensive Dictionary of the English Language, Encyclopedic Edition:

An area of study in the field of computer science. Artificial intelligence is concerned with the development of computers able to engage in human-like thought processes such as learning, reasoning, and self-correction.

The concept that machines can be improved to assume some capabilities normally thought to be like human intelligence such as learning, adapting, self-correction, etc.

The extension of human intelligence through the use of computers, as in times past physical power was extended through the use of mechanical tools.

In a restricted sense, the study of techniques to use computers more effectively by improved programming techniques.

Artificial intelligence, NEW INT’L WEBSTER’S COMPREHENSIVE DICTIONARY OF THE ENG. LANGUAGE (Encyc. ed. 1998).

71. See Nicholson Price, *supra* note 63, at 10; *Black-Box Medicine*, *supra* note 63.

72. See Liza Vertinsky & Todd M. Rice, *Thinking About Thinking Machines: Implications of Machine Inventors for Patent Law*, 8 B.U. J. SCI. & TECH. L. 574, 576–77 (2002) (discussing the growing use of computers to augment human capabilities and

A subset of AI is machine learning,⁷⁴ which is comprised of algorithms that provide new insights without being programmed to do so.⁷⁵

Scholarship and media anthropomorphize AI as an autonomous actor⁷⁶ with human cognitive functions,⁷⁷ but AI is not divorced from human control, and its intelligence is based on exceeding human capacity.⁷⁸ In fact, AI operates differently from the ways that humans sense, learn, reason, and act.⁷⁹ While some scholars have debated whether AI's unique features merit inventor status,⁸⁰ the European Patent Office⁸¹ and the USPTO have determined that AI cannot be an inventor.⁸²

replace human operators, as well as their effects on the invention process that cannot be easily accommodated within the current patent system).

73. MARCEL DEKKER, *ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY* 11, at 24 (1978) (defining a “learning system” as “any system which uses information obtained during one interaction with its environment to improve performance during future interactions”).

74. Juan C. Mateos-Garcia, *The Complex Economics of Artificial Intelligence* 6 (Dec. 19, 2018) (unpublished manuscript), <https://bit.ly/2ZTRDge> (stating that the artificial intelligence system combines machine learning algorithms).

75. See Carlton E. Sapp, *Preparing and Architecting for Machine Learning*, GARTNER RESEARCH (Jan. 17, 2017), <https://gtnr.it/2MOoIYs> (further defining “machine learning” as extracting knowledge or patterns from a series of observations; describing that data is fed into a machine learning system, which uses the data to fit to algorithms to solve a problem or derive an insight).

76. See Alicia Solow-Niederman, *Administering Artificial Intelligence*, S. CAL. L. REV. (forthcoming 2020).

77. See Catalina Goanta et al., *Back to the Future: Waves of Legal Scholarship on Artificial Intelligence*, in *TIME, LAW, AND CHANGE*, Oct. 23, 2019, at 1, 16.

78. See Finck, *supra* note 5, at 4 (noting that the capacity of AI has replaced human decision-making in some circumstances).

79. See Mirjana Stankovic et al., *Exploring Legal, Ethical and Policy Implications of Artificial Intelligence* 5 (L., Justice, and Dev., Working Paper 2017).

80. See Erica Fraser, *Computers as Inventors—Legal and Policy Implications of Artificial Intelligence on Patent Law*, 13 *SCRIPTED* at 306, 324 (2016); Ryan Abbott, *The Artificial Inventor Project*, *WIPO MAGAZINE* (Dec. 2019), <https://bit.ly/30CIROI> (suggesting that “no natural person, as traditionally defined, qualifies as an inventor”); Michael McLaughlin, *Computer-Generated Inventions*, 101 *J. PAT. & TRADEMARK OFF. SOC'Y* 224, 239–40 (2019) (proposing the establishment of categories of inventorship based on representing various degrees of human intervention that may take place throughout the inventive process); Abbott, *supra* note 1, at 1081 (arguing that AI has created patentable inventions and that, but for the judicial characterization of invention as a mental act, computers independently meet the requirements for inventorship); Yanisky Ravid & Liu, *supra* note 11, at 2221–22 (summarizing various proposals made by scholars and introducing an alternative approach for addressing inventions made by artificial intelligence).

81. See *EPO Publishes Grounds for Its Decision to Refuse Two Patent Applications Naming a Machine as Inventor*, EUROPEAN PAT. OFF. (Jan. 28, 2020), <https://bit.ly/2XnxZ9u> (“[T]he EPO considered that the interpretation of the legal framework of the European patent system leads to the conclusion that the inventor designated in a European patent must be a natural person[.]” and “[m]oreover, the designation of an inventor is mandatory as it bears a series of legal consequences, notably

AI is a form of statistical inference⁸³ that identifies correlations within datasets⁸⁴ to *imitate* human cognition and decision making.⁸⁵ More specifically, AI is different from traditional statistical analysis techniques where humans set up the analysis, specify which input variables in the data set to consider, and decide how those variables put together would yield an estimate of the outcome variable. By contrast, in AI, humans do not specify how input variables are put together, and instead, an algorithm tries many possible combinations of variables to optimize a function, and in so doing, “learns” how to make a more accurate prediction.⁸⁶ Deep learning,⁸⁷ which is the state-of-the-art of AI and machine learning, emulates the biological brain with the use of neural networks to process large amounts of data.⁸⁸

to ensure that the designated inventor is the legitimate one and that he or she can benefit from rights linked to this status”).

82. See U.S. Patent Application No. 16/524,350, Decision on Petition, 6 (filing date July 29, 2019) (unpublished) (reasoning that the denied patent application named a machine, “(Invention generated by artificial intelligence),” as the inventor, and . . . current statutes, case law, and USPTO regulations and rules limit inventorship to natural persons,” and “U.S. patent law does not permit a machine to be named as the inventor in a patent application”); Dennis Crouch, *USPTO Rejects AI-Invention for Lack of a Human Inventor*, PATENTLY-O (Apr. 27, 2020), <https://bit.ly/2TyWJKx>.

83. See generally Christopher Eldred, *AI and Domain Knowledge: Implications of the Limits of Statistical Inference*, BERKELEY ROUNDTABLE ON THE INT’L ECON., Oct. 2019 (describing AI’s fundamental nature as a form of statistical inference that serves as an effective use for many kinds of problem domains).

84. See Angela Daly et al., *Artificial Intelligence Governance and Ethics: Global Perspectives 5* (June 28, 2019) (unpublished manuscript) (describing AI as detecting patterns in data and making predictions on the basis of such datasets, which requires identification of correlations within the datasets).

85. See Osonde A. Osoba & Paul K. Davis, *An Artificial Intelligence/Machine Learning Perspective on Social Simulation 2* (RAND Corp., Working Paper 2018).

86. See Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1, 14–15 (2019).

87. See Maryam M. Najafabadi et al., *Deep Learning Applications and Challenges in Big Data Analytics*, 2 J. BIG DATA 1, 7 (2015) (asserting that deep learning is more powerful than statistics, since it can conduct difficult optimization tasks, extract non-local and global relationships and patterns in data, and can apply relational and semantic knowledge); Saptarshi Sengupta et al., *A Review of Deep Learning with Special Emphasis on Architectures, Applications, and Recent Trends*, 10 IEEE TRANSACTIONS at 2, 5 (Mar. 2019) (explaining that artificial neural networks, which are utilized to perform deep learning, refers to networks of interconnected hidden layers that are capable of approximating a function that represents patterns in the data from multiple passes of data through a learning procedure). See generally MICHAEL NIELSEN, *NEURAL NETWORKS AND DEEP LEARNING* (2019) (explaining that unlike algorithms in other software technologies, for which problems are broken into many small and precise tasks that a computer can easily perform, deep learning learns from observation of data to provide a solution to the problem at hand); TERRENCE J. SEJNOWSKI, *THE DEEP LEARNING REVOLUTION* (2018) (describing the motivation for the field of deep learning, neural network architectures, and deep learning’s impact now and in the years to come).

88. See Brian S. Haney, *Deep Reinforcement Learning Patents: An Empirical Survey 1*, 4 (2020) (unpublished manuscript), <https://bit.ly/2zpwMQs> (noting that deep

The benefits of AI over human capabilities stem from complexity, efficiency, scale, and time. AI can detect patterns from a massive amount of data and make predictions⁸⁹ that may be undetectable to the human brain. Additionally, AI is more sophisticated than other forms of software because of its ability to improve models to yield better predictions by analyzing many examples⁹⁰ and iteratively feeding data into an algorithm to improve output.⁹¹ Quite simply, accessibility of computing power, advancements in algorithms and analytics, availability of a massive amount of data, and new digitized use cases have presented a new AI technological revolution.⁹² AI presents benefits over traditional statistical inference techniques with its ability to discern useful patterns in large data sets accurately, continuously, efficiently, and predictively.⁹³ However, these advantages come with a cost of transparency, since the combination of variables and the relationships between them is extremely complex and non-intuitive.

The application of AI, such as inside of AI-based tools to produce AI-generated output, produces some unpredictable⁹⁴ results⁹⁵ but is

learning is a new type of machine learning, and it concerns the acquisition of knowledge from large amounts of data in a way that models the human brain, with machines to process information).

89. See Matt Taddy, *The Technological Elements of Artificial Intelligence*, in THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA, at 1 (Ajay Agrawal et al. eds., 2019) (explaining that AI automates tasks previously done by humans and requires a massive amount of data and continued data generation to detect patterns and make predictions).

90. See MUNRO & MIRSHARIF, *supra* note 58, at 4, 6, 10.

91. See Semmler & Rose, *supra* note 1, at 86–87.

92. See FIRTH-BUTTERFIELD & CHAE, *supra* note 1, at 5; see also Kimberly A. Houser & Anjanette H. Raymond, *It Is Time to Move Beyond the 'AI Race' Narrative: Why Investment and International Cooperation Must Win the Day*, NW. J. TECH & INTELL. PROP. (forthcoming 2021) (“The reason for the notable advances in AI stem from the recent availability of big data and cloud computing.”).

93. In simple terms, unlike traditional statistical techniques, state of the art of AI, which is known as deep learning, is unstructured—data is fed in without constraints about what the system should be looking for. The algorithms run, form connections between different layers of data, test for connections, and find relationships in the unstructured data that the machine learning programmer might not have ever thought to consider.

94. See Yanisky Ravid & Liu, *supra* note 11, at 2224–27 (suggesting that AI systems, which are based on algorithms and data, are unpredictable and evolving; specifying that AI systems produce “random mutations that result in unpredictable routes to the optimal solution” and that AI systems “continue to evolve and change according to new data” so as to “produce results that differ from the initial plan of the programmers or operators of the system”).

95. This Article clarifies the meaning of “unpredictable” identified in other non-technical AI literature to suggest that the trained weights themselves are unpredictable. Weights (in the AI context) refer to an indicator of a mathematical strength that can be explained by an example. Weights near zero mean that changing the input will not change the output, and thus, many AI algorithms set the weight to zero, so as to simplify the application of the AI. The AI-technological reasoning is: if the weights do change, then the output would be unpredictable. In general, the weights would not change, since

reproducible,⁹⁶ contrary to non-technical AI literature.⁹⁷ In other words, AI inverts the traditional computer programming paradigm, where mathematical functions take inputs to produce the desired output. Instead, AI produces a learned function—an algorithm that produces the lowest error and one that closely matches the inputs' actual outputs—when given a large enough set of inputs and outputs.⁹⁸ Once a learned function is produced, it can be used to make predictions on previously unseen data.⁹⁹ As applied to real-world examples, a learned function (that is part of an algorithm inside of an AI-based tool) can silently improve output of the quality of photos taken on a digital camera, help security screeners detect anomalies at airports, detect financial fraud, and improve online health results.¹⁰⁰

The disclosure challenge with AI-based tools stems from the inability to explain how the learned function (or algorithm) operates. The disclosure challenge with the AI-generated output stems from being unable to understand what algorithm method produced the particular output. As this Article argues, such disclosure deficits necessitate enhanced AI patent disclosure.

the goal of training is to obtain the best possible weights, so they can be reused. Since the training process is expensive, one would not want to retrain and, as a result, the AI-generated output should be predictable if the weights do not change. In other words, when one starts with the same training data, then one will get the same training out each time when the weights do not change.

96. See Eleanor Bird et al., *The Ethics of Artificial Intelligence: Issues and Initiatives*, EUROPEAN PARLIAMENTARY RESEARCH SERV., Mar. 2020, at 48; Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1210, 1213 (2017); Matthew Hutson, *Artificial Intelligence Faces Reproducibility Crisis*, 359 SCIENCE, Feb. 16, 2018, at 725–26; Joshua A. Kroll et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 659 (2017); Yew-Soon Ong & Abhishek Gupta, *Five Pillars of Artificial Intelligence Research*, IEEE TRANSACTIONS ON EMERGING TOPICS IN COMPUTATIONAL INTELLIGENCE, Oct. 2019, at 411, 413, <https://bit.ly/2EtPFKB> (explaining that AI algorithms raise questions concerning rationalizability, resilience, reproducibility, realism, and responsibility, and further clarifying that reproducibility of AI systems refers to a replication crisis).

97. This Article takes issue with the reproducibility crisis identified in other non-technical AI literature, and instead argues for reproducibility, while also acknowledging the unpredictability problem. The AI-technological reasoning is: once a system is trained, then weights are attained. If one were to restart the training, then one would be able to arrive at the same point with all of the weight values from the original training. As a result, the training process is reproducible and, thus, the AI-generated output should be reproducible as well.

98. See Michael Borella, *How to Draft Patent Claims for Machine Learning Inventions*, PATENT DOCS (Nov. 25, 2018), <https://bit.ly/36PGGh3>.

99. See *Patentability of Artificial Intelligence and Machine Learning Inventions in Europe*, WITHERS & ROGERS (Mar. 22, 2018), <https://bit.ly/2ZXDNJB>.

100. See Borella, *supra* note 98; *Patentability of Artificial Intelligence and Machine Learning Inventions in Europe*, *supra* note 99.

B. *Transparency Viewed from a Taxonomy of AI*

While AI provides new opportunities as a technological tool, it presents challenges for administrative processes, whose public policy objectives require transparency as a social contract ideal.¹⁰¹ Democratic governments should require that the public assess the efficacy and fairness of the governmental process.¹⁰² Transparency is integral to a fair society and a legitimate government, and visibility of government institutions is implicated in administrative law.¹⁰³ The USPTO is an administrative agency that issues patents to inventors, through a process called patent prosecution, by examining patent applications and determining whether the statutory requirements of patentability are met.¹⁰⁴ The prospect of inventors utilizing AI raises legal and policy transparency questions for the USPTO.¹⁰⁵

In general, AI presents a transparency paradox since requiring more information about the underlying black box¹⁰⁶ of AI creates benefits as well as risks and costs.¹⁰⁷ The USPTO has stressed that transparency is important by stating, “[a]s a federal agency, it is important for the USPTO to be able to explain all [patent] prosecution decisions made. Because of this, solution capabilities must be transparent to the USPTO and as well as to the general public. Black box solutions will not be accepted.”¹⁰⁸ From a normative standpoint, the USPTO should be able to value transparency as a mechanism for improving patent disclosure. The evaluation of this normative consideration requires assessing the use of AI by inventors, including a taxonomy of AI and unraveling the black box of AI.

A taxonomic exercise characterizes, classifies, and explains how AI generates inventive output and, in so doing, presents transparency considerations. Notably, there are different AI types, which fall on a spectrum with varying degrees of human input, from capturing human cognitive abilities to requiring some human input.¹⁰⁹ Machine learning,

101. See Finck, *supra* note 5, at 6.

102. See Brauneis & Goodman, *supra* note 5, at 109.

103. See Coglianesi & Lehr, *supra* note 96, at 18–19.

104. See Ebrahim, *supra* note 53, at 1187, 1193–95.

105. See Rai, *supra* note 7, at 2638–40; Ebrahim, *supra* note 53, at 1189, 1230.

106. See BELZER ET AL., *supra* note 63, at 35 (describing components of a learning system, or a black box); Nicholson Price, *supra* note 63, at 10; *Black-Box Medicine*, *supra* note 63.

107. See Burt, *supra* note 5.

108. OFF. OF THE CHIEF INFO. OFFICER, U.S. PATENT & TRADEMARK OFFICE, USPTO’S CHALLENGE TO IMPROVE PATENT SEARCH WITH ARTIFICIAL INTELLIGENCE: REQUEST FOR INFORMATION (RFI) 2 (2018).

109. See Rainer Mühlhoff, *Human-Aided Artificial Intelligence: Or, How to Run Large Computations in Human Brains? Toward a Media Sociology of Machine Learning*, NEW MEDIA & SOC’Y, Nov. 2019, at 1, 10 (characterizing human and AI interaction as

which is a subset and type of AI, is comprised of software programs that can learn from experience and improve their performance with time.¹¹⁰ Machine learning¹¹¹ utilizes algorithms to change its output based on new data,¹¹² and “the entire process relies on the analysis of data” with several stages.¹¹³ Deep learning, which is a subset of and the state of the art of machine learning, refers to extracting complex patterns from massive volumes of data by emulating the human brain’s ability to observe, analyze, learn, and make decisions.¹¹⁴

Whichever of these AI types is utilized, there are distinct situations with AI and patent-disclosure interplay (or lack thereof) that present transparency challenges. These situations are defined and utilized throughout this Article and fall within one of four distinct scenarios: (1) Disclosed AI-based tool; (2) Undisclosed AI-based tool; (3) Disclosed AI-generated output; and (4) Undisclosed AI-generated output. These scenarios are shown in Figures 1A & 1B, each of which is described in greater depth below. Figure 1A shows a taxonomy of AI categories, and Figure 1B provides social-policy-related considerations associated with Figure 1A’s quadrants. The taxonomy of AI categories, as shown below in Figures 1A and 1B, is helpful for ascertaining whether AI inventions may cause concern for patent law and society. Each scenario turns on the type of AI and the inherent level of disclosure by that type of AI. The type of AI shown in Figure 1A is repeated in Figure 1B, which overlays

“cybernetic AI,” and suggesting that the involvement of humans with AI technology presents a hybrid human-machine interaction).

110. See Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 89 (2014).

111. See *id.* at 88–89 (defining “machine learning techniques” as algorithms that have the ability to improve in performance over time on some task, by detecting patterns in data in order to automate complex tasks and make predictions).

112. See INFO. COMM’RS OFF., *BIG DATA, ARTIFICIAL INTELLIGENCE, MACHINE LEARNING AND DATA PROTECTION*, Sept. 4, 2017, at 7–8, <https://bit.ly/2XJ02QJ> (defining “machine learning” generally as being “the set of techniques and tools that allow computers to ‘think’ by creating mathematical algorithms based on accumulated data”; specifying that “supervised learning” involves algorithms based on labelled datasets, such that the algorithms are trained how to map form input to output with the provision of correct values assigned to them, and where the initial training phase creates models of the world on which predictions can be made in a subsequent prediction phrase; and specifying that “unsupervised learning” involves algorithms that are not trained, but are left to find regularities in input data “without instructions as to what to look for”) (first citing Deb Landau, *Artificial Intelligence and Machine Learning: How Computers Learn*, INTEL (Aug. 17, 2016); and then citing ETHEM ALPAYDIN, *INTRODUCTION TO MACHINE LEARNING* (3d ed. 2014)).

113. Josef Drexler et al., *Technical Aspects of Artificial Intelligence: An Understanding from an Intellectual Property Law Perspective*, in MAX PLANCK INSTITUTE FOR INNOVATION AND COMPETITION RESEARCH PAPER SERIES 4 (Max Planck Inst. for Innovation & Competition, Research Paper No. 19-13, 2019).

114. See Najafabadi et al., *supra* note 87, at 1–2, 4.

the concern and policy consideration associated with the particular type of AI.




Figure 1A: Taxonomy of AI categories

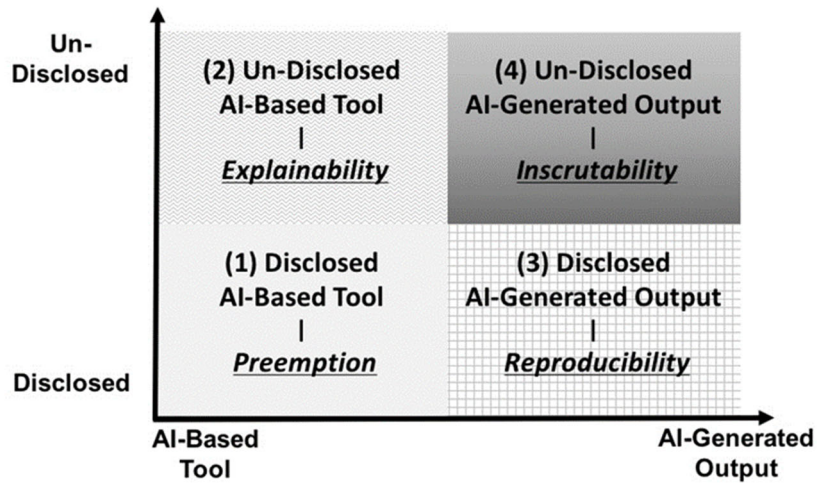


Figure 1B: Social-policy-related considerations based on a taxonomy of AI categories

An evaluation of the interplay (or lack thereof) of distinct AI categories and patent disclosure must begin with the particular characteristics of the categories. The fundamental characteristics of AI-based tools and AI-generated output, as shown in Figure 1A, provide

some guidance for the social-policy-related considerations shown in Figure 1B. The use cases described below provide new contexts and more motivation for the debate between proponents and a skeptic of patent disclosure. Unlike prior software technologies and statistical analysis techniques, AI presents new considerations through: (1) the state of the art of AI-based tools based on deep learning¹¹⁵ that emulates human cognitive abilities; and (2) an exponential increase in data (in the form of text, recordings, images, etc.) and computing power that enable AI-based tools to be applied in domains¹¹⁶ where it may be difficult to identify whether AI-based tools were used. Thus, AI presents a metaphysical shift in software technology. Recognizing explainability of an AI tool's operation and inscrutability of AI-generated output is vital for justifying the patent system¹¹⁷ and for normative accounts of how the patent system should evolve.¹¹⁸

Although an AI-based tool and AI-generated output are shown in the Figures as being distinct, it should be noted that some AI-based tools can be applied to produce AI-generated output. An AI-based tool refers to an AI-engine or AI-algorithm that is embedded into a module. AI-generated output refers to the result of applying an AI-based tool and is the outcome from analysis of data with several stages.¹¹⁹ Consequently, the AI-based tool provides mathematical infrastructure, which, when applied to training data and new data, analyzes patterns in the data to produce a result that may be a prediction or representation of the physical world. These descriptions of characteristics of AI are clarified by a description of the social-policy-related terms associated with AI, such as explainability and inscrutability.

Before delving into the details of the taxonomy, an overview is warranted: In the first scenario, Quadrant (1) provides an undisclosed AI-based tool, which refers to the claiming of an AI-based tool along with disclosure of the AI itself and how the AI operates. In the second scenario, Quadrant (2) provides an AI-based tool without disclosure of how the AI operates. In the third scenario, Quadrant (3) provides an application of an AI-based tool to generate output that is claimed in a patent application, where there is disclosure of the AI-generated output. In the fourth scenario, Quadrant (4) provides an application of an AI-

115. See Sengupta et al., *supra* note 87, at 1, 2 (suggesting that multilayer artificial neural networks comprise deep learning, which represents the complexity of the aggregation chain across many hidden layers to learn sufficiently detailed representations).

116. Eldred, *supra* note 83, at 2–3 (citing Eric Martin, *Moore's Law Is Alive and Well*, MEDIUM PREDICT (Dec. 21, 2018), <https://bit.ly/3hp13X5>).

117. See *infra* Part II.

118. See *infra* Part III.

119. See Drexler et al., *supra* note 113, at 4.

based tool to generate output that is claimed in a patent application, but there is no disclosure of the application of AI.

As shown earlier in Figures 1A and 1B, there are distinct scenarios with AI and patent-disclosure interplay (or lack thereof) that present transparency challenges. Two of the scenarios—(1) Disclosed AI-based tool; and (3) Disclosed AI-generated output—are only briefly addressed since they do not present social-policy-related considerations within the scope of this Article. However, the two other scenarios—(2) Un-disclosed AI-based tool and (4) Un-disclosed AI-generated output—are more critically analyzed, since recognition of explainability of an AI-tool’s operation and inscrutability of AI-generated output are vital for disclosure-related tensions. Descriptions and use-case examples provided as context for disclosure-related tensions are shown in Figures 2A and 2B and explained further in Part II.¹²⁰

In the first scenario, “Disclosed AI-based tool” refers to the claiming of an AI-based tool along with disclosure of the AI itself and how the AI operates. In other words, the inventor claims the AI, which refers to the AI engine or the AI techniques inside of the AI-based tool, and discloses how the AI-based tool can be applied to make AI-generated output. A patent on an AI-based tool can have a preemptive effect in the market regarding the use of data on application of the AI-based tool.¹²¹ Additionally, there are both patenting of AI-based tools¹²² and open-source modules of AI-based tools,¹²³ but an analysis of patents and open-source¹²⁴ considerations on AI-based tools is outside of the scope of this Article.

120. See *infra* Part II.

121. See Simon & Sichelman, *supra* note 50, at 382.

122. See Haney, *supra* note 88.

123. See generally ANTONIO GULLI & SUJIT PAL, DEEP LEARNING WITH KERAS: IMPLEMENTING DEEP LEARNING MODELS AND NEURAL NETWORKS WITH THE POWER OF PYTHON (2017) (explaining use of Keras for deep learning applications); ELI STEVENS ET AL., DEEP LEARNING WITH PYTORCH (2020) (explaining use of PyTorch for deep learning applications); GIANCARLO ZACCONE & MD. REZAUL KARIM, DEEP LEARNING WITH TENSORFLOW: EXPLORE NEURAL NETWORKS AND BUILD INTELLIGENT SYSTEMS WITH PYTHON (2d ed. 2018) (explaining use of Python for deep learning applications); KERAS, <https://bit.ly/2Am6Qf8> (last visited July 12, 2020) (showing an open-source library designed to enable fast experimentation with AI, specifically deep neural networks); PYTORCH, <https://bit.ly/2ySt9Pu> (last visited July 13, 2020) (offering an open-source AI library used for applications such as computer vision); TENSORFLOW, <https://bit.ly/2XEuCLa> (last visited July 14, 2020) (providing a free and open-source software library for dataflow and differentiable programming for use in a range of AI applications).

124. See Greg R. Vetter, *Claiming Copyleft in Open Source Software: What if the Free Software Foundation’s General Public License (GPL) Had Been Patented?*, 2008 MICH. ST. L. REV. 279, 281 (2008) (noting that important strands in the open-source movement include available public source code disclosure); Greg R. Vetter, *The Collaborative Integrity of Open-Source Software*, 2004 UTAH L. REV. 563, 594–96

In the second scenario, “Un-disclosed AI-based tool” refers to the claiming of an AI-based tool without disclosure of how the AI operates. In other words, the inventor claims the AI, which refers to the AI engine or the AI techniques inside of the AI-based tool, but does not disclose how an invention is made by the AI-based tool. Thus, the inventor does not explain how or why an AI-based tool would generate certain AI-generated output. Assuming that the inventor can overcome a patent eligibility rejection for abstractness, the inventor can attain a patent grant on the AI-based tool by meeting patentability. However, even with a lack of an explanation of how and why the AI-based tool produces AI-generated output, there could still be patent protection of an after-arising technology¹²⁵ under the doctrine of equivalents.¹²⁶ In effect, an AI-based tool can be a generator for various types of AI-generated output since different data fed into its AI algorithms would produce different output. Accordingly, the lack of explainability of the AI-based tool suggests that patent protection of AI-based tools would extend to unforeseeable uses. For example, in the area of computational chemistry, the inventor can claim an AI-based tool (in the form of an engine, mechanism, or technique) without explaining how it can predict the optimal function, structure, and properties of a compound for a particular application.¹²⁷ Thus, the mere fact that an inventor has claimed an AI-based tool may be an issue for explainability, which is a major point of analysis of this Article. This scenario is related to the fourth scenario (“Un-disclosed AI-generated output”) since the inner workings of an AI-algorithm may not comprehensibly relate the AI-based tool to the AI-generated output.

(2004) (explaining the open-source approach as having a kinship and freedom, and finding software has more value with source code); Greg R. Vetter, *Commercial Free and Open Source Software: Knowledge Production, Hybrid Appropriability, and Patents*, 77 *FORDHAM L. REV.* 2087, 2091 (2009) (explaining the start and development of open-source projects); Greg R. Vetter, *Exit and Voice in Free and Open Source Software Licensing: Moderating the Rein over Software Users*, 85 *OR. L. REV.* 183, 183 (2006); Greg R. Vetter, “*Infectious*” *Open Source Software: Spreading Incentives or Promoting Resistance?*, 36 *RUTGERS L.J.* 53, 57 (2004); Greg R. Vetter, *Open Source Licensing and Scattering Opportunism in Software Standards*, 48 *B.C. L. REV.* 225, 225 (2007); Greg R. Vetter, *Opportunistic Free and Open Source Software Development Pathways*, 30 *HARV. J.L. & TECH.* 167, 167–68 (2017).

125. See Christopher A. Cotropia, “*After-Arising*” *Technologies and Tailoring Patent Scope*, 61 *N.Y.U. ANN. SURV. AM. L.* 151, 151 (2005) (“An ‘after-arising’ technology is a technology that ‘come[s] into existence after the filing date of a[] [patent] application.’”) (alteration in original) (quoting *In re* Application of Hogan, 559 F.2d 595, 605 (C.C.P.A. 1977)).

126. See Michael J. Meurer & Craig Allen Nard, *Invention, Refinement and Patent Claim Scope: A New Perspective on the Doctrine of Equivalents*, 93 *GEO. L.J.* 1947, 1948 (2005) (explaining that the Doctrine of Equivalents “allows a court to expand patent scope beyond the rights literally claimed in the patent”).

127. See Ebrahim, *supra* note 19, at 595–96.

In the third scenario, “Disclosed AI-generated output” refers to the application of an AI-based tool to generate output that is claimed in a patent application, where there is disclosure of the AI-generated output. By comparison, the first scenario discloses how the AI-based tool can be applied to make AI-generated output, whereas this third scenario discloses how the AI-generated output is made by an AI-based tool (without necessarily claiming the AI-based tool). An AI-based tool that utilizes the same AI-technique may fail to replicate the AI-generated output since the inventor may not know how or why the AI-based tool produced the AI-generated output. Thus, the mere fact that an inventor has claimed AI-generated output may be an issue for reproducibility, which is a technological concern outside the scope of this Article.

In the fourth scenario, “Un-disclosed AI-generated output” refers to the application of an AI-based tool to generate output that is claimed in a patent application, but where there is no disclosure of the application of AI. In other words, the generated output appears to be a product of human invention but, in fact, it was created through the use of an AI-based tool. A patent examiner would be unable to distinguish whether the invention was based on an AI-based tool or based on human ingenuity. For example, an AI-based tool could be utilized in computational chemistry for predicting the optimal function, structure, and properties of a compound for a particular application.¹²⁸ But the inventors do not disclose that they used an AI-based tool for AI-generated output. In such a scenario, the disclosure in the patent application would reveal how to make or use the compound, but it would not disclose that an AI-based tool has been utilized for providing AI-generated output in the form of the compound’s composition of matter and method of manufacturing.¹²⁹ This scenario presents a problem with inscrutability. In many such situations, the patent system is presented with a new challenge, since the disclosure would not enable a *human* having ordinary skill in the art to make or use the invention without undue experimentation, but would require *either AI or a person using AI* to make or use the invention without undue experimentation. In this scenario, even if a *human* having ordinary skill in the art could make or use the invention, it may require undue experimentation. Thus, the mere fact that an inventor has claimed AI-generated output may be an issue for inscrutability, which is a major point of analysis of this Article.

128. *See id.*

129. In patent law practice in the unpredictable arts (such as chemistry-based technologies), enablement and written description rejections are routine, particularly since patent examiners often want to know from the inventor specifics of how chemical compounds are made or used in greater detail.

Explainability of an AI-based tool's operation refers to the black-box nature of AI, which is unexplainable.¹³⁰ For example, an artificial neural network is provided where the complex relationship between the inputs and outputs, the ability to find patterns in data, and the capture of the statistical structure in a set of data representing observed variables are unexplainable¹³¹ since the network is self-learning.¹³² AI

130. See Finck, *supra* note 5, at 11, 15, 17.

131. In simple terms, being unexplainable means being indiscernible to the relevant audience. The reason is that, as AI-based tools become non-linear, a human cannot explain how one particular variable relates to another variable. In contrast, with simple linear equations, humans may be able to explain how variables relate to one another. In the context of complex AI-based tools, such as those utilizing deep learning, a human would not be able to explain what inspired the underlying algorithm to develop the weights associated with the deep learning layers. As a result, a human cannot aggregate and diagnose the details between the deep learning layers, which contributes to a lack of explainability.

132. See Michael D. Stein, *Patenting Inventions in Machine Learning: Part 2*, LAW360 (Dec. 22, 2016), <https://bit.ly/3do7h7s> (providing as an example a patent claim to a computer-implemented classification system comprising an artificial neural network (ANN), which serves as an AI-based tool where the underlying algorithm and its operation is unexplainable). Consider, as an example, U.S. Patent No. 6,792,412 (filed Feb. 2, 1999) owned by plaintiff Hyper Search, LLC, and asserted against defendant Facebook in *Hyper Search, LLC v. Facebook, Inc.*, where the first patent claim states:

1. A system for controlling *information output* based on user feedback about the information comprising:

A plurality of information sources providing information;
At least one neural network module that selects one or more of a plurality of objects to receive information from the plurality of information sources based at least in part on a plurality of inputs and a plurality of weight values;
At least one server, associated with the neural network module, that provides one or more of the objects to one or more recipients;
The recipients enabling for one or more users to generate feedback about the information; and
Wherein the *neural network module* generates a rating value for a plurality of the objects at the end of an epoch, redetermines the weight values using the rating values, and selects which objects to receive information during a subsequent epoch using the redetermined weight values and the inputs for the subsequent epoch.

Hyper Search, LLC v. Facebook, Inc., No. 17-1387-CFC-SRF, 2018 WL 6617143, at *10 (D. Del. Dec. 17, 2018) (emphasis added) (quoting '412 Patent, col. 19 l. 49–67).

Here, the lack of explainability refers to the neural network module (an AI-based tool) identified above, which utilizes an AI-method (a neural network) and which the district court found did not show how the system was an improvement in computer capability. See *id.* As the court's reasoning suggests, the neural network's operation would be indiscernible to the relevant audience (even after reading the detailed description of the patent), since the neural network periodically re-evaluated the weights based on ratings from the objects but provided no detail on how the model's input or output was processed. See *id.* Here, inscrutability refers to the inability to understand how (or by which method) information output (AI-generated output), as identified above, was produced or how the model operated or was trained. See *id.*

explainability refers to the difficulty of making plain or comprehensible the inner workings of an algorithm to humans.¹³³ Thus, explainability of AI-based tools can be achieved only when one opens the black box and analyzes the underlying computational logic.¹³⁴ However, even data scientists, developers, and software engineers find it difficult to interpret such AI-based tools, and there is much ongoing research into explaining the underlying technological methods.¹³⁵

Inscrutability refers to defying practical inspection and resisting comprehension due to the underlying complex and numerous calculations in the AI context.¹³⁶ Thus, AI-generated output is inscrutable, since one cannot understand how, or from which method the output was produced.¹³⁷ The application of AI-based tools applies

133. See Valérie Beaudouin et al., Flexible and Context-Specific AI Explainability: A Multidisciplinary Approach 7–9, 12–13 (Mar. 23, 2020) (unpublished manuscript), <https://bit.ly/3eGgFDD> (suggesting that explainability refers to providing an explanation of the functioning of an algorithm in its entirety or leaving uncertainty out by relegating a particular algorithmic decision to a human user, and in so doing, suggesting that explainability is a part of transparency, which in turns provides traceability, auditability, and accountability to the system).

134. See Finck, *supra* note 5, at 15.

135. See Mühlhoff, *supra* note 109, at 2–3 (stating that deep learning “is a form of distributed orchestration of human cognition through networked media technology” such that “human cognitive skills . . . [are] embedded in machine networks,” and “harvesting human cognition in computing networks through new forms of labor and machinized power relations”).

136. See Selbst & Barocas, *supra* note 8, at 1094.

137. See Guided Image Composition on Mobile Devices, U.S. Patent No. 10,516,830 B2 (issued Dec. 24, 2019). Consider as an example the Guided Image Composition on Mobile Devices patent, which is assigned to Adobe and which relates to a system that automatically crops a picture with a good composition, for which patent claim 14 states:

A system comprising:
 a camera;
 a display;
 a non-transitory computer-readable medium storing computer-executable instructions of an image processing application; and
 a processing device communicatively coupled to the non-transitory computer-readable medium for executing the computer-executable instructions, wherein executing the computer-executable instructions configures the system to perform operations comprising:
 receiving image data corresponding to a field of view of the camera, wherein the image data is received in a preview mode of the image processing application;
 rendering, in the preview mode, a major view of the display, the major view presenting a previewed image based on the image data;
 receiving, from *a deep-learning system*, a composition score of a cropped image generated by cropping the previewed image;
 rendering, in the preview mode, a sub-view on the display, the sub-view presenting the cropped image; and

complex computations from feature inputs to outcome predictions that are incomprehensible to humans¹³⁸—the result is AI-generated inventions. AI-based tools can be applied toward developing AI-generated output, which can allow for: (1) claiming variations beyond what the inventor actually invented to broaden claim scope; (2) disclosing (without claiming) variations on existing patent claims and preventing future improvements to destroy novelty; (3) claiming next incremental steps and adding claims on demand in response to innovative developments in certain fields; (4) claiming, by utilizing analytical techniques and correlations to assist inventors in predicting performance and properties of inventions in the physical world; and (5) claiming prophecies with high accuracy through computational experimentation to allow for an early patent grant.

In summary, the second scenario, “Un-disclosed AI-based tool,” and the fourth scenario, “Un-disclosed AI-generated output,” are the central concerns of this Article. These particular scenarios are the context for the disclosure-related tensions shown in Figure 2B and explained

rendering, in the preview mode, *the cropped image* in the major view based on a user interaction with the sub-view, wherein the user interaction is (i) a swipe left or right to update the previewed image with a cropped image displayed in a sub-view, (ii) a swipe up or down to perform a zoom operation, or (iii) a tap to display a composition score.

Id. (emphasis added).

Here, inscrutability refers to AI-generated output in the form of a cropped image, which is the data that is produced by the use of an AI-based tool in the form of a deep-learning system. *See id.* The cropped image would not be understandable if details of how it was produced from the deep-learning system was not provided in the detailed description of the patent application (which although is not shown here, specifies, “A computer-implemented method for real-time image cropping”). *Id.* To be more precise, the composition score could be the number of pixels of a particular color region, and different AI-based tools can give different composition scores, which would result in different cropped images; as a result, the removal of the phrase “by a deep-learning system” (assuming it would still be awarded a patent), would result in an unidentified invention of AI-generated output in the form of a cropped image. *Id.*

While this analysis recognized that this particular patent provided the deep learning system and a first claim specifying the method to provide the cropped image, many AI-invention patents where an AI-based tool is applied to produce AI-generated output do not specify how the AI-generated output was made (while meeting the enablement and written description requirements). *See id.* Thus, had claim 1 not been included by the inventor and had claim 14 been allowed by the USPTO, claim 14’s cropped image would be an unidentified invention since it would constitute inscrutable AI-generated output that was not understandable as to how (or which method) it was produced from. *See id.* While this reasoning recognizes claim 1 is provided to explain the method of providing the cropped image, here an example of inscrutability of claim 14 is shown for demonstration purposes. *See id.* Moreover, a closer inspection of this patent reveals insufficient explanation of the inner workings of the deep learning system, further demonstrating that the AI-generated output of in the form of the cropped image is not understandable. *See id.*

138. *See* Selbst & Barocas, *supra* note 8, at 1094.

further in Part II.¹³⁹ The second scenario, “Un-disclosed AI-based tool,” concerns identified AI inventions, for which this Article questions the adequacy of patent disclosure of the inner workings of AI-based tools and the inventive method of AI-generated output. The fourth scenario, “Un-disclosed AI-generated output,” concerns unidentified inventions, for which this Article cautions against a non-detection problem, an overall lack of disclosure of unidentified inventions. Even if AI-generated output is descriptively just like human-generated output in appearance, there may be normative justifications against recognizing it for patentability.

Indeed, since it is foreseeable that AI-generated output may be examined to meet patentability by the USPTO, scholars and society may expect that AI-generated output is worthy of patent protection. One policy lever, which this Article assesses, is strengthening patent disclosure. Prior to revisiting the explainability and inscrutability concepts that this Article discussed in the Introduction¹⁴⁰ and Part I, Part II introduces theories and normative assessments¹⁴¹ that provide a foundation for revisiting explainability and inscrutability in a prescriptive sense in Part III.¹⁴² Therefore, it is necessary to consider the normative accounts of promulgating disclosure with patents and preventing protection through trade secrets to determine where the justification for stronger patent disclosure stands concerning AI-based tools and AI-generated output.

II. THEORIES & NORMATIVE ASSESSMENTS OF PATENT DISCLOSURE WITH AI

Part I concluded that patent protection may conceivably inhere in AI-generated output. However, this conclusion still begs the question of whether patent protection should conceivably inhere in AI-generated output or, alternatively, if trade secrecy should come into play. Why should society care whether there is patent protection of AI-generated output? Is this question not as silly as being concerned about patent protection of output generated through the use of any other tool or software?

One crucial difference mentioned in Part I¹⁴³ is that contemporary AI increasingly blurs the lines between real and unreal.¹⁴⁴ One might be

139. *See infra* Part II.

140. *See supra* INTRODUCTION.

141. *See infra* Part II.

142. *See infra* Part III.

143. *See supra* Part I.

144. *See* Ryan Abbott & Alex Sarch, *Punishing Artificial Intelligence: Legal Fiction or Science Fiction*, 53 U.C. DAVIS L. REV. 323, 331 (2019).

surprised to learn that AI-based tools are increasingly applied to produce AI-generated output¹⁴⁵ and that the USPTO has recognized unidentified inventions as being increasingly prevalent for patent examination consideration.¹⁴⁶ As a result, the emergence of AI presents strains on conceptual foundations, gaps, and unnecessary overlap between patent and trade secret protection. At first glance, patents and trade secrets appear mutually exclusive; however, they can be utilized in complementary ways.¹⁴⁷ Inventors may choose not to pursue patents on AI due to AI's lack of detectability and the inventors' inability to discern whether a competitor is using a disclosed AI technique. While there are several challenges to attaining patent protection of AI inventions that may prompt seeking trade secrecy protection,¹⁴⁸ this Article considers an enhanced AI patent disclosure as a policy lever to equilibrate this legal protection decision.

A. *Theories of the Patent and Trade Secret
Trade-off Decision with AI*

The traditional view of the dichotomy between patents and trade secrets emphasizes that patents are premised on disclosure, and trade secrets are premised on a lack of disclosure.¹⁴⁹ Whereas patent law (as well as copyright law)¹⁵⁰ is meant to promulgate disclosure, trade secret law has the opposite function of preventing disclosure by maintaining secrecy. Although patents and trade secrets may seem to be mutually exclusive, they can be used in complementary ways.¹⁵¹ Patents and trade secrets are not necessarily economic substitutes, and various forms of AI may not be protected in a complementary fashion, which could yield

145. See *supra* Part I.

146. See *supra* Section I.B.

147. See ASHISH ARORA, ANDREA FOSFURI & ALFONSO GAMBARDILLA, *MARKETS FOR TECHNOLOGY: THE ECONOMICS OF INNOVATION AND CORPORATE STRATEGY* 262 (2001); Jessica M. Meyers, *Artificial Intelligence and Trade Secrets*, 11 *LANDSLIDE*, No. 3, 2019, <https://bit.ly/2ZMB2ea>; Michael Risch, *Trade Secret Law and Information Development Incentives*, in *THE LAW AND THEORY OF TRADE SECRECY: A HANDBOOK OF CONTEMPORARY RESEARCH* 152 (Rochelle C. Dreyfuss & Katherine J. Strandburg eds., 2010).

148. See Asay, *supra* note 1, at 1200–19 (2020) (suggesting difficulties of AI patenting with patentable subject matter, disclosure requirements, and novelty and non-obviousness).

149. See generally Mark A. Lemley, *The Surprising Virtues of Treating Trade Secrets as IP Rights*, 61 *STAN. L. REV.* 311 (justifying trade secrets as a form of intellectual property, and while providing the traditional viewpoint that trade secrets suppress disclosure, asserting paradoxically that trade secret law encourages disclosures based on investments into secrets by certain industries); Meyers, *supra* note 147.

150. See generally Jeanne C. Fromer, *An Information Theory of Copyright Law*, 64 *EMORY L. J.* 71 (2014) (suggesting that expressive works are valuable to society by using that expression to communicate knowledge).

151. See Meyers, *supra* note 147.

socially harmful results.¹⁵² This Article argues, for AI inventions as a particular type of invention, that societal benefits from greater disclosure with patent law outweigh limiting disclosure with trade-secret law.

AI in the form of AI-based tools and AI-generated output can be naturally hidden from human view. Against this backdrop, trade secrets may seem like a natural choice of intellectual property protection. Yet the interface between AI and patents is interesting on many levels and presents normative implications for what should be protected with patents rather than with trade secrets. The unique technological features of AI, compared to other forms of software and statistical techniques in general, present considerations for greater patent disclosure. Unlike other software-based technologies, which, regardless of complexity, are expensive to create but relatively easy to reproduce,¹⁵³ AI technologies present new social policy considerations.¹⁵⁴ As this Article showed in Part I, AI-based tools present challenges with explainability, and AI-generated output presents challenges with inscrutability. As a result, one policy consideration is greater patent disclosure of the inner workings of AI-based tools and the inventive method of AI-generated output.

A restriction on patenting of AI-based tools would simultaneously reduce incentives for research and development of AI algorithms, deprive the public of knowledge and information of such tools, lessen knowledge spillovers,¹⁵⁵ and increase incentives to protect them with trade secrecy.¹⁵⁶ Scholars have noticed that weak patent disclosure has

152. See Brenda M. Simon & Ted Sichelman, *Data-Generating Patents*, 111 NW. U. L. REV. 377, 377, 430 (2017).

153. See Steven J. Frank, *What AI Practitioners Should Know About the Law Part I*, AI MAG., Spring 1988, at 63.

154. See Philippe Lorenz & Kate Saslow, *Demystifying AI & AI Companies*, STIFTUNG NEUE VERANTWORTUNG, July 2019, at 1, 2, 32 (noting that “AI [is] a hot topic for social and economic policy,” including foreign policy, and furthermore that “[p]olicy makers around the world have begun to understand the importance of AI [which] is reflected in numerous international initiatives and fora devoted to AI [g]overnance”).

155. See Adam B. Jaffe et al., *The Meaning of Patent Citations: Report on the NBER/Case-Western Reserve Survey of Patentees*, in PATENTS, CITATIONS, AND INNOVATIONS: A WINDOW ON THE KNOWLEDGE ECONOMY 379, 379–80, 388 (2002); Wesley M. Cohen et al., *R&D Spillovers, Patents and the Incentives to Innovate in Japan and the United States*, 31 RES. POL’Y 1349, 1349–50 (2002); Ouellette, *supra* note 39, at 548, 550.

156. While restricting patenting of AI-based tools, or weakening patents would reduce incentives for innovation, there are beneficial effects with promoting equitable access and promoting distributive values. See Peter Lee, *Towards a Distributive Agenda for U.S. Patent Law*, HOU. L. REV. 55, 321, 323–24 (2017) (suggesting that enforcing exclusive rights with patents is consonant with access, equity, and distributive justice issues, but stressing that patent law scholarship should consider the effect on incentives on technological development on marginalized communities); see also Elif Kavusturan, *Reforming U.S. Patent Law to Enable Access to Essential Medicines in the Era of Artificial Intelligence*, NW. J. TECH. & INTELL. PROP. (forthcoming), <https://bit.ly/321UIBa> (recognizing that while incentives for patents are important, there

led innovators to avoid the competition-promoting function of patent expiration and switch to trade secrecy.¹⁵⁷ As a result of these effects, this Article's normative position is that there should be enhanced AI patent disclosure. The ensuing discussion presents the theoretical trade-off of patents and trade secrets; after which, there is a normative assessment of enhanced AI patent disclosure. To understand the theory and normative assessment, some preliminary discussion on the theory of self-disclosing and non-self-disclosing inventions is necessary.

The terms "self-disclosing" and "disclosing" refer to the level of ease or difficulty of reverse-engineering the invention¹⁵⁸—a self-disclosing invention is considered easy to reverse engineer; in contrast, a non-self-disclosing invention is expensive to reverse engineer.¹⁵⁹ Reverse engineering refers to the process by which an object is deconstructed through a backward-looking inspection to deduce its architecture, design, or features or the methodology for its development or manufacturing.¹⁶⁰ Software reverse engineering includes identifying the blueprints of source code and object code as well as their interrelationships in a high level of abstraction.¹⁶¹ AI inventions fall within a specialized technological area where reverse engineering would be more challenging than other forms of software. In general, software patents often cover the functionality embedded in source code but do not reveal all of the technical details that make reverse engineering easy. AI inventions are even more difficult to reverse engineer because they are neither explainable nor scrutable.¹⁶²

The distinction between self-disclosing and non-self-disclosing inventions is helpful in understanding societal trade-offs between patents and trade secret protection. Under the theory of "self-disclosing" and "disclosing" inventions, self-disclosing inventions are easily copied from the commercial embodiment, whereas non-self-disclosing inventions are

are tradeoffs with access, such as, for health care and pharmaceuticals, including effects concerning pricing, health risks, and essential medicines); W. Keith Robinson, *Access to the Patent System*, NEV. L.J. (forthcoming), <https://bit.ly/3fhfon8> (noting that patenting of artificial intelligence affects underrepresented innovators, including small businesses, women, and underrepresented minorities, as well as socially immobile communities).

157. See Anton et al., *supra* note 63, at 1 (noting that weak patents cause firms to rely more heavily on secrecy and potentially reduce the incentives to innovate).

158. See Julie E. Cohen & Mark A. Lemley, *Patent Scope and Innovation in the Software Industry*, 89 CAL. L. REV. 1, 14, 17 (2001) (noting that "intellectual property regimes that have traditionally protected software permit reverse engineering").

159. See Ouellette, *supra* note 39, at 588.

160. See Tonya M. Evans, *Reverse Engineering IP*, 17 MARQ. INTEL. PROP. L. REV. 61, 88 (2013).

161. See Dennis S. Karjala, *Copyright Protection of Computer Documents, Reverse Engineering, and Professor Miller*, 19 U. DAYTON L. REV. 975, 991–95 (1994).

162. See *supra* Part I.

not easily copied.¹⁶³ Typically, patents are sought on inventions that are self-disclosing, or relatively cheap to reverse engineer, whereas an inventor will only patent a non-self-disclosing invention if it seems likely that others will recreate the invention before the patent expires.¹⁶⁴ In other words, patenting of non-self-disclosing inventions would not occur if the invention would reach the public inevitably.

The self-disclosing/non-self-disclosing distinction requires understanding more about whether sufficient information is contained in a patent application and about the ease of reverse engineering the invention to gauge whether the invention would be fruitful to the public. The relevant question is whether society would want to enforce more robust patent disclosure of non-self-disclosing inventions. In other words, the normative inquiry is whether enhanced patent disclosure will serve the public good by enhancing the potential for follow-on inventions and steering others away from duplicative efforts to reinvent the patented invention. Figures 2A and 2B reflect this tension; then, a normative assessment of enhanced AI patent disclosure is provided. Figures 2A and 2B provide a graphical representation of the available legal protections of the taxonomy of AI categories shown earlier,¹⁶⁵ with the central analysis of this Article depicted as a dotted line below.

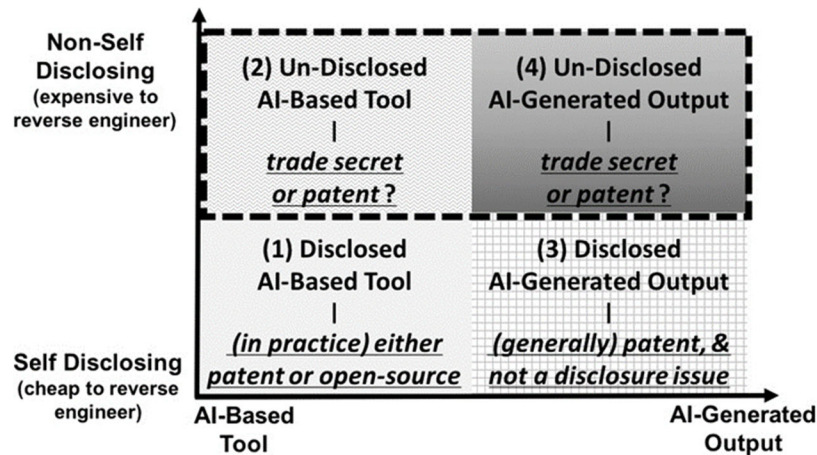


Figure 2A: Relating Degree of Reverse Engineering with AI Taxonomy and IP Protections

163. See Katherine J. Standburg, *What Does the Public Get? Experimental Use and the Patent Bargain*, WIS. L. REV., 2004, at 81, 112.

164. See Oullette, *supra* note 39, at 546.

165. See *supra* Part I.

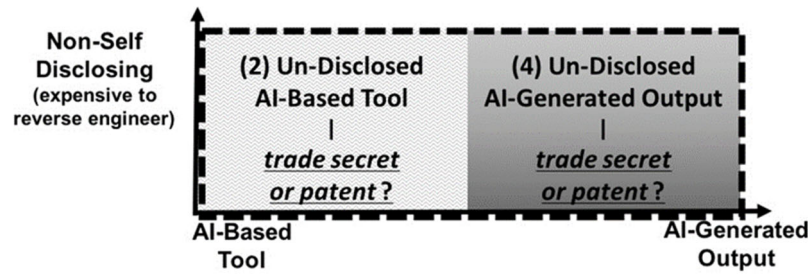


Figure 2B: Patent and Trade Secret for Non-Self Disclosing and Un-Disclosed AI

Patent disclosure is particularly relevant with AI-generated-output inventions, which are non-self-disclosing (that is, expensive to reverse engineer)¹⁶⁶—this is counter to the theory of self-disclosing/non-self-disclosing inventions,¹⁶⁷ which provides that inventors seek patents for non-self-disclosing inventions if it seems likely that others will recreate the invention before the patent expires. By contrast, and as central AI-specific themes in this Article: (1) AI inventors seek patents for AI-generated output inventions, although it seems likely that others will *not* recreate the invention before the patent expires,¹⁶⁸ and (2) AI is not only included in *identified* inventions, but could also be in *unidentified* inventions.¹⁶⁹ The first theme relates to inventors’ rush to seek patents on AI-generated output for defensive patenting.¹⁷⁰ The second theme relates to patents on AI-generated output where a patent examiner does not understand whether an AI-based tool was put into operation. These themes are interrelated and stem from the challenge of AI-transparency—including explainability and inscrutability—in patent disclosure and AI inventors’ strategic decision between patent protection and trade secret protection, which this Part has asserted are not necessarily substitutes. Subsequently, this Article provides theories held by a skeptic of and proponents of the disclosure function of patents and

166. See Ouellette, *supra* note 39, at 546.

167. See Standburg, *supra* note 163, at 104–06; Ouellette, *supra* note 39, at 546.

168. See *infra* Section II.B.1 (suggesting that this phenomenon could be due to risk aversion, a response to weak patentability criterion for AI inventions, or an early gold rush for AI patenting, or some combination thereof).

169. See *infra* Section II.B.2.

170. The mere fact that there is a lot of AI patenting from observation might indicate an arms race in which firms are patenting defensively. Perhaps all firms would be better off if they could agree to a patenting moratorium. In relation to this point, it would be helpful to understand whether AI patents might facilitate licensing-based business models that are inherently difficult to execute based on a trade secrecy business model.

an assessment of why there is a need for enhanced AI patent disclosure, the normative position of this Article.

B. Theories of Patent Disclosure

Section 112 of the Patent Act requires inventors to explain their invention—known as claim definiteness—as well as how to make or use the invention—the disclosure requirements of enablement and written description.¹⁷¹ Disclosure is an important feature of patents since it sets the boundaries of patents and represents what the inventor actually invented.¹⁷² Before addressing the need for enhanced AI patent disclosure and mechanisms for enhancing patent disclosure with AI, some preliminary introduction about the theories of patent disclosure is warranted.

1. Theories by a Skeptic of Patent Disclosure

Inventors must satisfy the legal standard of disclosure as required under 35 U.S.C. § 112.¹⁷³ Patent disclosure is carried out by satisfying two separate requirements—written description and enablement—and serves to adequately disclose the invention to the public.¹⁷⁴ The enablement requirement refers to providing an explanation of the invention to enable others to recreate it,¹⁷⁵ and the written description requirement describes the invention in a way that one skilled in the art could clearly conclude that the invention was that of the inventor.¹⁷⁶ Furthermore, the patent claims, which demarcate the boundaries of the invention, are part of the disclosure. The disclosure in a patent application must allow a person having ordinary skill in the art (“PHOSITA”) to practice the invention without undue

171. See 35 U.S.C. § 112 (2018).

172. See John R. Allison & Lisa Larrimore Ouellette, *How Courts Adjudicate Patent Definiteness and Disclosure*, 65 DUKE L.J. 609, 611 (2016).

173. See 35 U.S.C. § 112 (2018).

174. See Guang Ming Whitley, Comment, *A Patent Doctrine Without Bounds: The “Extended” Written Description Requirement*, 71 U. CHI. L. REV. 617, 617–18 (2004) (suggesting that written description has traditionally served as a policing doctrine to prevent patent applicants from improperly amending patent claims after submitting a patent application, whereas the enablement requirement ensures that public knowledge is enriched by the patent specification commensurate with the scope of the patent claims).

175. See *In re Wands*, 858 F.2d 731, 737 (Fed. Cir. 1988); ROBERT P. MERGES & JOHN F. DUFFY, *PATENT LAW AND POLICY: CASES AND MATERIALS* 247–323 (7th ed. 2017) (discussing disclosure and enablement); Sean B. Seymore, *Uninformative Patents*, 55 HOUS. L. REV. 377, 384–85 (2017); Seymore, *The Enablement Pendulum Swings Back*, *supra* note 24, at 279; Seymore, *Heightened Enablement in the Unpredictable Arts*, *supra* note 24, at 127.

176. See *Regents of the Univ. of Cal. v. Eli Lilly & Co.*, 119 F.3d 1559, 1566 (Fed. Cir. 1997) (quoting *Lockwood v. Am. Airlines, Inc.* 107 F.3d 1565, 1572 (Fed. Cir. 1997)); Holbrook, *supra* note 25, at 127.

experimentation.¹⁷⁷ In effect, patent law's disclosure function is a way for the inventor to reveal the invention to society. The legal and policy debates about patent disclosure are those of costs versus benefits for society. The underlying issues are whether stronger patent disclosure should be a central goal for quid pro quo of the patent system rather than keeping an invention secret.

One commentator has criticized disclosure theory as a justification for the patent system for several reasons, including: (1) patents do not disclose much useful information to technologists and contain inadequate specifications; (2) disclosure is inconsistent with the patent system's normative foundations; and (3) strengthening disclosure is unnecessary since patents are distinct from the underlying technology.¹⁷⁸ As to this first reason, this commentator is not alone. Numerous other skeptics point out that the inadequacies and inefficiencies of patent disclosure are particularly problematic in information technology and software, for which patents do not convey meaningful information, and their sheer volume makes distinguishing from prior art exhausting.¹⁷⁹ In this realm, one commentator points out that the patent system encourages disclosure of useful information in only a narrow class of inventions.¹⁸⁰ Second, this skeptic of patent disclosure argues that the benefits of disclosure are secondary to economic incentives for invention and commercialization,¹⁸¹ and that these goals often come into conflict.¹⁸² This viewpoint suggests that the patent system is a means to induce innovation and remedy a public-goods problem, and that disclosure is a disincentive and cost that comes hand-in-hand with the patent grant.¹⁸³ Third, this skeptic distinguishes the disclosure function from incentives

177. See *In re Wands*, 858 F.2d at 737; Seymore, *The Enablement Pendulum Swings Back*, *supra* note 24, at 279.

178. See Devlin, *supra* note 42, at 404.

179. See Mark A. Lemley & Bhaven Sampat, *Is the Patent Office a Rubber Stamp?*, 58 EMORY L.J. 181, 202 (2008); Arti K. Rai, John R. Allison & Bhaven N. Sampat, *University Software Ownership and Litigation: A First Examination*, 87 N.C. L. REV. 1519, 1531 (2009); Jerome H. Reichman, *Intellectual Property in the Twenty-First Century: Will the Developing Countries Lead or Follow?*, 46 HOUS. L. REV. 1115, 1135–36 (2009).

180. See Benjamin N. Roin, Note, *The Disclosure Function of the Patent System (Or Lack Thereof)*, 118 HARV. L. REV. 2007, 2016 (2005).

181. See Adam B. Jaffe, *The U.S. Patent System in Transition: Policy Innovation and the Innovation Process* 41 (Nat'l Bureau of Econ. Research, Working Paper No. 7280, 1999) (summarizing arguments for why software patents are undesirable from a commercial and market perspective, including: the untenable need to secure many different licenses in order to market any given software-based product, the result of driving smaller software firms out of business since the need for multiple licenses favors large firms that can bargain for cross-licensing, and the requirement that software interfaces with other pieces of software necessitating interoperability).

182. See Devlin, *supra* note 42, at 404, 406.

183. See *id.* at 412, 416, 418, 419.

and points out that the invention itself is self-revealing, even if the patent document does not provide sufficient clarity.¹⁸⁴ In sum, skepticism of the normative goals of patent disclosure suggests that greater disclosure is ineffective and detrimental to the patent system's social goals.

2. Theories by Proponents of Patent Disclosure

To be fair, the patent-disclosure skeptic's views are not entirely without merit in a limited circumstance, in the sense that patent disclosure may not be good enough in practice. However, the majority of the scholarly patent community agrees that patents should provide disclosure and that patent disclosure has significant societal benefits. Proponents of patent disclosure provide insights into and distinctions of AI that are worth exploring in-depth. Proponents point out that patent disclosure provides one source of useful technical information to researchers,¹⁸⁵ demarcates the contours of property rights,¹⁸⁶ achieves the goal of predictable patent claim boundaries,¹⁸⁷ and prevents over-claiming.¹⁸⁸ Contrary to critics of disclosure's benefits, many scholars have provided empirical evidence that patent disclosure has benefits for the use of technical information and provides knowledge spillovers.¹⁸⁹ The U.S. Supreme Court has stated that patent disclosure advances the storehouse of knowledge that leads to further advancements and improvements in the technology area.¹⁹⁰

In addition to the legal and technological benefits of patent disclosure, scholars have noted that there are also strategic business benefits of patent disclosure that yield a competitive advantage. One scholar has shown that a firm leading a patent race has an incentive to disclose in order to reduce its rival's expected payoff and encourage the rival to quit the race.¹⁹¹ This perspective suggests that patent races prompt greater disclosure of new information to the public to preempt the issuance of a patent to a firm that is ahead of its competitor in the patent race. Other scholars have determined that a firm trailing in a patent race has an incentive to disclose to extend a patent race since the

184. *See id.* at 405, 406, 411, 418, 427.

185. *See Ouellette, supra* note 39, at 533.

186. *See Devlin, supra* note 42, at 409.

187. *See Miller, supra* note 40, at 180.

188. *See Rantanen, supra* note 41, at 370.

189. *See supra* note 155 and accompanying text.

190. *See Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470, 481 (1974).

191. *See generally* Douglas Lichtman, Scott Baker & Kate Kraus, *Strategic Disclosure in the Patent System*, 53 VAND. L. REV. 2175 (2000) (explaining the strategic use of patent disclosure in the realm of patent races, and in so doing, analyzing the incentive to disclose for preempting a rival's patent when the laggard lacks the ability to leapfrog the leader, and concluding that a leader in a patent race has an incentive to disclose to cause the rival to quit the race).

trailing firm needs more time to catch up in the race.¹⁹² This perspective, for which the disclosing firm still plans to pursue patents, suggests that patent disclosure serves as a strategy to enable continuation of the race. Another view suggests that firms that are not planning to pursue patents see patent disclosure as a defensive business strategy to thwart a rival firm's patents or keep the rival firm from obtaining patents.¹⁹³

C. Normative Assessment with AI Disclosure

A relevant question concerning AI inventions and patent disclosure is: are the benefits of stronger disclosure underappreciated? As patent law and society begin to grapple with AI-generated inventions, there are new interactions between the unique technological features of AI¹⁹⁴ and patent law. There is much about AI that is decidedly unfamiliar to disclosure for patentability and detection of disclosure for patent examination. For instance, by disclosure for patentability, an inventor might not disclose that they used an AI-based tool to create AI-generated output.¹⁹⁵ Patent law generally encourages inventors to use tools in the invention process and to invest in research and development of tools for which they seek patent protection (or alternatively trade secrecy).¹⁹⁶

Harder to explain, though, are some other aspects of AI in patent examination. For instance, no capability exists at the USPTO to determine whether an invention is based on AI-generated output from the use of an AI-based tool. Even if patent examination could ascertain whether an invention was AI-generated output, there is, strangely, nothing in patent law that prohibits such a practice. Recently, the USPTO and the EPO have indicated that AI cannot be an inventor,¹⁹⁷ and the Patent Trial and Appeal Board (PTAB) has held patent claims involving AI to be patent-eligible under 35 U.S.C. §101;¹⁹⁸ additionally, and most

192. See Scott Baker & Claudio Mezzetti, *Disclosure as a Strategy in the Patent Race*, 48 J.L. & ECON. 173, 173–94 (2005) (describing that since patents are evaluated in light of the prior art, patent disclosures make it more difficult for another firm to make a related patent claim, and therefore, patent disclosure extends a patent race since the subsequent invention must have a sufficient advance over an expanded prior art).

193. See Justin P. Johnson, *Defensive Publishing by a Leading Firm*, 28 INFO. ECON. & POL'Y 15, 15 (2014).

194. See *supra* Section I.A.

195. See *supra* Section I.B.

196. See *supra* Section II.A.

197. See *In re* Application of Application No.: 16/524,350, 2020 Dec. Comm'r Pat.; *EPO Publishes Grounds for Its Decision to Refuse Two Patent Applications Naming a Machine as Inventor*, EUROPEAN PAT. OFF. (Jan. 28, 2020), <https://bit.ly/3d5TGBu>.

198. See *ex parte* Hannun, No. 2018-003323 (P.T.A.B. Dec. 11, 2019) (describing that the claims involved a method for speech recognition using a trained neural network and that the patent examiner asserted that the claims recited a mathematical relationship/formula and certain methods of organizing human activity and a mental process; noting that the PTAB found the claimed steps could not “practically be

importantly for patent practitioners, there are no court decisions addressing whether the issue involved AI-generated output and no mechanism for detecting AI-generated inventions.¹⁹⁹ In fact, AI-generated inventions do not even have a visible distinction from human-generated inventions—yet not one fully AI-generated invention has been acknowledged by the USPTO. When asked about AI-generated inventions, the USPTO acknowledged that AI is included in more than one-third of all identified inventions.²⁰⁰ However, the USPTO has yet to acknowledge what may be the alarming rise of “un-identified” (that is, AI-generated) inventions.

In the world of AI inventions, one may ask whether the distinction between AI-generated and human-generated inventions is truly necessary. By studying AI, this Part argues that inventions should be addressed by patent law, and that we can gain insight into how norms and laws of inventing become even more “artificial.”

1. Disclosure for Patentability

AI-based tools, which were introduced earlier in this Article, result in “artificial” inventions,²⁰¹ but they are not “magic.”²⁰² Using less lofty language, one may call this “the use of a super-human tool to assist a human inventor.” AI-based tools are just one type of tool in the larger environment of tools that inventors utilize in the invention process. Other tools include rulers and slide rules, calculators, microscopes, mass spectrometers, chromatographs, debugging tools, compilers, computer-aided design, and even advanced software-based tools, such as finite-element analysis and computational fluid dynamics. Unlike other tools, AI’s state-of-the-art machine learning,²⁰³ deep learning, and

performed mentally,” so as to determine that the claim was not a mental process and found the claim to be patent eligible); Braden M. Katterheinrich et al., *Artificial Intelligence at the Patent Trial and Appeal Board*, FAEGRE DRINKER (Feb. 11, 2020), <https://bit.ly/30BX0B2> (explaining that *ex parte Hannun* and other PTAB decisions have found that “AI-related inventions are more likely to be found patent-eligible at the USPTO when the claims do not explicitly recite mathematical formulas and instead recite AI-related features that are technologically-specific and that cannot practically be replicated in one’s mind”).

199. See *supra* Section I.B.

200. See Iancu AI Policy Update, *supra* note 48.

201. See *supra* INTRODUCTION.

202. See *supra* Section I.A.

203. See, e.g., ALEX SMOLA & S. V. N. VISHWANATHAN, INTRODUCTION TO MACHINE LEARNING 3–6 (2008) (describing a variety of machine learning applications, where there exists a nontrivial dependence between some observations for which a simple set of deterministic rules is not known, such as: (1) *web page ranking*, which is a process of submitting a query to a search engine to find webpages relevant to the query and returning them in an order of relevance; (2) *collaborative filtering*, where Internet bookstores utilize users’ “past purchase and viewing decisions” information to predict

reinforcement learning capabilities can be used to conceive of new inventions and to create a new class of patent applications without much, if any, human input.²⁰⁴ Such inventions would otherwise be too complex for a human's limited reasoning capabilities.

This "artificial world" of AI inventions encompasses AI-based tools that can be applied toward producing AI-generated output. AI-based tools allow a human inventor to simulate complexity and scale of analysis beyond human capabilities, and in so doing, allow the human inventor to creatively meet the patentability criteria. AI-based tools can be taken a step further by inventors using them to provide AI-generated output.²⁰⁵ AI-based tools challenge, among other things, patentability by introducing a non-human²⁰⁶ element in the conception process of invention.²⁰⁷ Obviously, this is a new concept for inventions, if it is even properly characterized as an invention at all now. For instance, in the invention process, AI-based tools allow for supplementing human input,²⁰⁸ predicting outcomes before they may occur in the physical

future viewing and purchase habits of similar users; (3) *speech recognition*, where an audio sequence is annotated with text or where handwriting is annotated with a sequence of strokes; and (4) *classification*, where a spam filtering program can identify whether an email contains relevant information or not, such as a frequent traveler email from an airline, based on the type of user). See generally GIANLUCA BONTEMPI, HANDBOOK: STATISTICAL FOUNDATIONS OF MACHINE LEARNING (2017), <https://bit.ly/3gvO2uA> (explaining how to extract knowledge from large amounts of data, including procedures such as knowledge extraction, pattern analysis, and data processing, and how to model such processes using computational techniques).

204. See also Saptarshi Sengupta et. al., *A Review of Deep Learning with Special Emphasis on Architectures, Applications and Recent Trends*, KNOWLEDGE-BASED SYSTEMS, Apr. 2020, at 1.; Zahangir Alom et. al., *A State-of-the-Art Survey on Deep Learning Theory and Architectures*, ELECTRONICS, Mar. 2019, at 1; Hironobu Fujiyoshi et. al., *Deep Learning-Based Image Recognition for Autonomous Driving*, 43 IATSS RES. 244, 244 (2019); Najafabadi et al., *supra* note 87, at 4. See generally CHARU C. AGGARWAL, NEURAL NETWORKS AND DEEP LEARNING (2018) (explaining the theory and algorithms of deep learning, with a focus on concepts concerning neural network architectures).

205. See *supra* Section II.B.

206. See Tabrez Y. Ebrahim, *Data-Centric Technologies: Patent & Copyright Doctrinal Disruptions*, 43 NOVA L. REV. 287, 317–20 (2019).

207. Robin C. Feldman & Nick Thieme, *Competition at the Dawn of Artificial Intelligence*, J. ANTITRUST ENFORCEMENT (forthcoming) (manuscript at 6–8, <https://bit.ly/36zdPgQ>) (explaining that conception requires "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," and that patent law requires some degree of consciousness by its inventors).

208. See generally GIUSEPPE SANSEVERINO, THE ABILITY TO CHASE DOWN OUR DREAMS: INVENTIVE STEP AND ARTIFICIAL INTELLIGENCE (2018), <https://bit.ly/2zv9Gyt> (suggesting some human involvement with AI in the current state of artificial narrow intelligence technology, including for a preliminary phase, a subsequent stage of processing the results, and examining the innovative results, and asserting that tensions in patent law arise only when these phases are subject to human intervention); Mühlhoff, *supra* note 109 (suggesting that today's artificial intelligence technology is dependent on

world,²⁰⁹ and doing hundreds to millions of human calculations and decisions in a fraction of the time of an ordinary human.²¹⁰

Perhaps because AI-based tools support this kind of rich and “artificial” interaction with inventors in the invention process, many scholars who have studied AI have identified doctrinal tensions with patent law, including with inventorship,²¹¹ patent eligibility,²¹² non-obviousness,²¹³ and enablement.²¹⁴ For instance, with AI-based tools, some scholars have suggested that computers should be considered inventors rather than humans.²¹⁵ As an additional example, scholars have argued that the PHOSITA standard should be modified to consider “thinking machines”²¹⁶ since “everything is obvious”²¹⁷ with the use of AI-based tools. AI-based tools pressure the patent system since they reduce the time for humans to conceive of an invention and constructively reduce it to practice, simplify the complexity of the inventive task, and produce an earlier patent grant. But one may reply—so what? And how should patent law respond?²¹⁸

2. Detection of Disclosure for Patent Examination

AI-generated output, which this Article introduced earlier, results in “artificial” inventions.²¹⁹ “Artificial” means the inventions are fictitious and imaginary but appear as if they had been created in the physical world. These inventions may be considered as being either imaginary, never-achieved, or unworkable to the inventor, but may appear as if they were created, tested, or made workable to reasonable onlookers of the patent document. In other words, inventors can utilize an AI-based tool to provide AI-generated output, from which they develop a patent

human participation, but that advanced and developing artificial intelligence technology such as deep learning will provide a socio-technological transformation between humans and machines that will entail new models of subjectivation with less human involvement and digital labor).

209. See Ebrahim, *supra* note 19, at 596 (explaining the use of AI simulations to describe hypothetical chemical structures and their properties in the specification of a patent application, such as with prophetic patent claims and broad genus patent claims).

210. See Drexel et al., *supra* note 113, at 4 (explaining machine learning, including its subcategories and related fields, training processes, black-box nature, outputs, and evolutionary algorithms).

211. See Abbott, *supra* note 1, at 1079; Feldman & Thieme, *supra* note 207, at 9–11; Ravid & Liu, *supra* note 11, at 2215; Schuster, *supra* note 11, at 1959.

212. See Mizuki Hashiguchi, *The Global Artificial Intelligence Revolution Challenges Patent Eligibility Law*, 13 J. BUS. & TECH. L. 1, 5 (2017).

213. See Vertinsky, *supra* note 12; Ryan Abbott, *supra* note 12, at 2.

214. See Früh, *supra* note 5, at 6, 13.

215. See Abbott, *supra* note 1, at 1079.

216. Vertinsky, *supra* note 12.

217. Abbott, *supra* note 12, at 2.

218. See *infra* Part III.

219. See *supra* INTRODUCTION.

application that includes legitimate-sounding diagrams and descriptions of how an invention might work without ever having taken, checked, and tested that it actually works in the physical world. One might reasonably ask how an inventor that utilized an AI-based tool could obtain hundreds to thousands to millions of patents protecting AI-generated output that either was hardly (or never) developed or effectively was concealed through an unexplainable algorithmic inventive process.

As a result of the unexplainable (that is, indiscernible to the relevant audience) AI-based tool's operation, such AI-generated output is inscrutable (that is, not understandable as to how, or by which method it was produced). In a legal sense, patent law scholars may deem such AI-generated output as being prophetic²²⁰ or resulting in too early of a patent grant.²²¹ Yet AI-generated output claimed in a patent application appears “real,” even though it is “unreal” in terms of actual reduction to practice.²²² If the patent system accepted AI-generated output as real, then we may need to consider banishing from patent examination all manner of inventions, including purely human-generated inventions, machine-assisted human inventions, and—most importantly for society's purposes—AI-generated inventions, since it is impossible to distinguish between them.

Of course, this assessment reveals that there are many senses of the words “artificial” and “real.” Ontologically speaking, “artificial inventions” in the form of AI-generated output claimed in a patent application have much in common with virtual reality.²²³ Augmented

220. See Janet Freilich, *Prophetic Patents*, 53 U.C. DAVIS L. REV. 663, 666–68 (2019) (referring to prophetic patent claims as being permissible, common, and comprising “made-up experiments and fictional data in patents”) (first citing *Atlas Powder Co. v. E.I. du Pont De Nemours & Co.*, 750 F.2d 1569, 1577 (Fed. Cir. 1984); and then citing MPEP (9th ed. 2015) § 608.01(p)); Christopher M. Holman, *Nuvo v. Dr. Reddy and the Patentability of Prophetic Pharmaceutical Inventions Based on Unexplained Inventive Insight*, 38 BIOETICHNOLOGY L. REP. 207, 207 (2019). Holman defines a “prophetic claim” as a:

[P]atent claim with respect to which, as of the filing date, no embodiment falling within the scope of the claim has actually been reduced to practice, either by it having been made (in the case of a product) or performed (in the case of process), or if it has been made or performed it has been physically demonstrated that the claimed product or process provides the practical utility that the invention is purported to provide.

Id. See also Ebrahim, *supra* note 19 at 596–98, 606–09, 618, 621, 636, 640–41, 649.

221. See Christopher A. Cotropia, *The Folly of Early Filing in Patent Law*, 61 HASTINGS L.J. 65, 68–69 (2009) (discussing how the early-filing nature of a patent application well before a commercial application results in patents before the commercial viability of the invention, consequently creating an ever-rising number of underdeveloped patents).

222. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2138.05.

223. See Mark A. Lemley & Eugene Volokh, *The Real Law of Virtual Reality*, 51 U.C. DAVIS L. REV. 51, 51 (2017) (describing virtual reality as being about virtual

reality generates realistic images and other sensations that simulate physical presence in a virtual environment.²²⁴ Virtual reality masterfully blends the distinctions between unreal and real.²²⁵ Similarly, AI-generated output seems physically real, but that physical reality is largely a faux representation of what an invention might be if there were any semblance of reduction to practice by the inventor.²²⁶ For instance, biological and chemical compounds, drugs and pharmaceuticals, and materials and nanotechnology—or what patent law scholars use as examples of the “unpredictable arts”²²⁷—can appear in a patent application as AI-generated output representations of what may appear in the physical world, but in fact are prophecies.²²⁸ Conventional composition-of-matter patent claims and product patent claims are not limited by a reduction to practice²²⁹ that occurs in the physical world but can be computationally experimented upon²³⁰ with the use of AI-based tools to produce AI-generated output on the computer. Should patent law embrace the unreal AI-generated output, and if so, how can the unreal be detected in patent examination with a fictitious disclosure? Should society expect patent examiners to reasonably detect such imaginary, never-achieved, or unworkable disclosure, and if not, then what are the patent administration possibilities?

Indeed, many inventors are embracing the unreality of AI-generated output by the application of AI-based tools. Despite probable skepticism,

interactions and communications with someone in another part of the world, and providing gaming as an application example).

224. See Mark A. Lemley & Eugene Volokh, *Law, Virtual Reality, and Augmented Reality*, 166 U. PA. L. REV. 1051, 1054 (2018) (explaining that augmented reality “allows digital content to be layered over the real world,” and that, “[u]sing special glasses or . . . a smartphone, [augmented reality] users can see the real world as it actually exists, but with digital images superimposed on the world so that they seem to exist as part of the world”).

225. See W. Keith Robinson & Joshua T. Smith, *Emerging Technologies Challenging Current Legal Paradigms*, 19 MINN. J.L. SCI. & TECH. 355, 360–61 (2018).

226. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2138.05; Mark A. Lemley, *Ready for Patenting*, 96 B.U. L. REV. 1171, 1172, 1177 (2016).

227. See Seymore, *Heightened Enablement in the Unpredictable Arts*, *supra* note 24, at 137–39.

228. See Freilich, *supra* note 220, at 715 (explaining that prophecies in patent law may be enough to get a patent granted but may not be enough to convince investors as being valuable, and also that prophecies in patents may not represent something that is functional or workable).

229. Reduction to practice has not historically required any physical embodiment, just an adequate description of how to make and use. Thus, drawings and diagrams of a machine, for example, are sufficient with an explanation, and an inventor need not have built the physical embodiment.

230. See generally Ebrahim, *supra* note 19 (exploring the use of AI computational technologies to experiment upon the structure, function, and properties of chemical compounds on the computer).

AI-generated output is inscrutable and significant in patent examination for three primary reasons.

First, from the standpoint of anticipating future patenting capabilities, contemporary AI-based tools, as the precursors to state-of-the-art deep learning,²³¹ will present challenges for patent examination in the future. As AI usage becomes increasingly important in the future of invention activity, there will be a related problem of detection of the use of AI at the USPTO for patent examination.²³² Contemporary AI-based tools do not yet emulate the human brain, but AI usage in patent prosecution can allow a patent practitioner to more easily overcome rejections in office actions by analyzing past patterns of the corresponding patent examiner and the particular art unit. In so doing, the inventor can produce AI-generated output to assist in overcoming patentability rejections provided in office actions. The concern from a patent examination standpoint is whether a patent examiner will detect that such AI capabilities are being utilized by the inventor.

Second, the use of AI-based tools to meet patentability is related to the issue of AI-generated output being indistinguishable from human-generated output in patent applications. This phenomenon of the use of AI for meeting patentability also illustrates a second reason why AI-generated output is worthy of consideration, namely the institutional mechanisms for detecting, monitoring, and distinguishing between the “real” patent claims and unreal patent claims based on AI-generated output. The USPTO will need to assess the right amount and depth of disclosure in patent applications, and algorithmic governance²³³ at the USPTO will require proactive design and maintenance of systems capable of detecting AI-generated output that are transparent, unbiased among inventor types, and accountable as to why patentability decisions are made by AI-aided patent examiners. It will be a challenge to craft AI-based tools that detect AI-generated output and that meaningfully fulfill transparency to ensure fidelity to the USPTO’s public accountability mandate. With an AI-aided patent examiner, an interrelated challenge is the development of evaluation metrics of patentability, such as why an inventor got a rejection or why there was a notice of allowance.

A third reason for exploring AI-generated output is that it represents an amazing experiment in internal AI-capacity building for USPTO

231. See Zahangir Alom et al., *A State-of-the-Art Survey on Deep Learning Theory and Architectures*, ELECTRONICS, Mar. 2019, at 1, 44, <https://bit.ly/2Yrv7st>.

232. See Ebrahim, *supra* note 53, at 1240–44 (suggesting that a counteracting artificial intelligence institution will need to be created in order to aid in the detection of the use of AI at the USPTO during patent examination in order to “serve as a guarantee of patent applications derived from artificial intelligence”).

233. See Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, EMORY L.J. (forthcoming 2020).

patent examination.²³⁴ The USPTO will need to invest in their technical and data infrastructure,²³⁵ which will require addressing the interrelated challenges of deploying AI-based tools that require building internal staff capacity.²³⁶ The inevitable implications are that: (1) the USPTO will need to grapple with budgeting and other human-resource constraints,²³⁷ and (2) the USPTO will need to develop standardization across its technology centers and its art units.

D. Theories for Justification of Patent Disclosure of AI-Generated Output

Given the previous descriptive account in Section I.B²³⁸ and the problems with AI disclosure described in Part II,²³⁹ there is substance to the claim that there might be patent claiming of AI-generated output,²⁴⁰ much of which may not be identified in patent examination. AI-generated output is no less a reasonable concept of patent protection than human-generated output. Many legal scholars have explored the justifications for the existence of various economic incentives to the patent system.²⁴¹ Therefore, it is necessary to consider normative accounts of AI-generated output to determine where the justification for AI-generated output stands in relation to established forms of human-generated output for patent law.

There are normative accounts of patents, and these accounts differ in significant ways to yield different conclusions about whether a given output is appropriately characterized as being worthy of patent protection. Strangely, the two normative theories conclude—though with qualifications—that AI-generated output, as claimed in patents, is, in

234. See Ebrahim, *supra* note 53, at 1241.

235. See generally Prithwiraj Choudhury et al., *The Future of Patent Examination at the USPTO*, HARV. BUS. SCH. PUB., Apr. 2017 (suggesting that the USPTO will need to invest in new tools as AI and machine learning uses have emerged and applied in the context of patent applications that are filed with the USPTO).

236. See Rai, *supra* note 7, at 2636–38.

237. See generally Naira Rezende Simmons, *Putting Yourself in the Shoes of a Patent Examiner: Overview of the United States Patent and Trademark Office (USPTO) Patent Examiner Production (Count) System*, 17 J. MARSHALL REV. INTELL. PROP. L. 32 (2017) (describing the USPTO's Production Units and the “count” system, which evaluate the performance of a patent examiner, and as a result of such metrics, the USPTO has budget and human-resource constraints).

238. See *supra* Section I.B.

239. See *supra* Part II.

240. See Iancu AI Policy Update, *supra* note 48.

241. See Edmund W. Kitch, *The Nature and Function of the Patent System*, 20 J.L. & ECON. 265, 285 (1977); Kesan, *supra* note 28, at 898 (examining several theories that explain and justify the role of patents in the modern economy, including traditional *ex ante* justifications and *ex post* justifications, as well as how these economic rationales may differ across industries).

reality, considered human-created when the “reality” is that a human applied AI-based tools (or at the very least was significantly assisted by AI²⁴²). From this discussion of two main normative theories of patents, there is a surprising conclusion: there are strong normative grounds for finding that patent protection should inhere in AI-generated output, whether the output is generated by AI-based tools or the combination of human-AI interaction, even when there is significant assistance by AI.

Further, based on the earlier discussion in Part I,²⁴³ the conclusion is that there is no descriptive disconnection between human-generated output and AI-generated output. From both descriptive and normative positions,²⁴⁴ it seems that AI-generated output does, or should, attract patent protection. Assuming that there should be patent protection in AI-generated output, Part III considers how disclosure should be heightened in response and the related effect on trade secrecy.

Before delving into the prescriptions for enhanced AI patent disclosure in Part III,²⁴⁵ a discussion of the theories justifying²⁴⁶ greater AI patent disclosure, including utilitarian and Lockean theories, is warranted. Scholars have introduced other justifications and economic reasoning for the patent system, and each has shed a new light on the consequences for incentives and societal influence of patents. Edmund Kitch’s 1977 article *The Nature and Function of the Patent System* theorized that the patent system offers broad prospects at early stages of development that encourage investment in that prospect *ex post* of the patent grant.²⁴⁷ Clarissa Long’s 2002 article *Patent Signals* suggested that patents reduce information asymmetries between inventors and observers and could serve as signals to convey information and firm attributes.²⁴⁸ Ted Sichelman’s 2010 article *Commercializing Patents* stressed the role of patents for commercialization and proposed decoupling the invention and commercialization functions of patents into dual rights.²⁴⁹ This Part analyzes utilitarian economic justifications and Lockean labor theory justifications for patents and is based on incentives

242. See McLaughlin, *supra* note 80, at 240–43 (establishing a spectrum of human intervention in computer-assisted AI-generated inventions).

243. See *supra* Part I.

244. See *supra* Section II.C.

245. See *infra* Part III.

246. See Hilty et al., *supra* note 27 (addressing justifications for AI under traditional IP theories).

247. See Kitch, *supra* note 241, at 267–71.

248. See Clarissa Long, *Patent Signals*, 69 U. CHI. L. REV. 625, 627 (2002).

249. See generally Ted Sichelman, *Commercializing Patents*, 62 STAN. L. REV. 341 (2010) (proposing “a new ‘commercialization patent,’ granted in exchange for the commitment to make and sell a substantially novel product,” and in doing so, decoupling patent law into an invention function and a commercialization function that would improve commercialization).

such that innovation would be at suboptimal levels absent the incentives.²⁵⁰

1. Utilitarian Theories for Adequate AI-Generated Output

Utilitarianism is the dominant justification for U.S. patent law. The U.S. Constitution's grant for creating patent (and copyright) protection is utilitarian and states that Congress may "promote the Progress of Science and useful Arts."²⁵¹ Much of the emphasis in U.S. patent law and many of the normative views of patent scholars are utilitarian in nature.²⁵² As an example, the U.S. Supreme Court has invoked utilitarianism in deciding patent-law cases.²⁵³ The utilitarian principle of seeking the greatest good for the greatest number and its conceptions of utility, rather than deontological perceptions of good, provide that we should grant private property rights if doing so would increase overall social welfare.²⁵⁴

Utilitarianism suggests that invention will only occur if there is some sort of grant of exclusive rights, and it assumes that the patent grant increases the production of new inventions. Although some scholarship has noted that patents do not promote innovation,²⁵⁵ utilitarian principles are invoked for the proposition that the grant of patents in an object will increase the production of such objects.²⁵⁶ Thus, if a type of invented object is assumed to be socially desirable, then an increase in these

250. See Kesan, *supra* note 28, at 898–99.

251. U.S. CONST., art. I, § 8, cl. 8.

252. See Lee, *supra* note 29, at 325 (providing a counterargument and exception to the generally held view that utilitarianism is dominant in U.S. patent law scholarship).

253. See Precision Instrument Mfg. Co. v. Auto Maint. Mach. Co., 324 U.S. 806, 816 (1945).

254. See generally RICHARD A. POSNER, *FRONTIERS OF LEGAL THEORY* (Harv. Univ. Press 2004) (examining the interdisciplinary study of economics, history, psychology, epistemology, and empiricism to provide a unified framework for disparate phenomena, and in so doing, contributing to utilitarianism and applying cost-benefit analysis to parties and judicial economy).

255. See generally Soma Dey, *Are Patents Discouraging Innovation?* (Nat'l Univ. of Sing.), <https://bit.ly/3fjfKcm> (suggesting that the strategic complementarity between patenting and R&D is relatively weaker in the presence of licensing and proposing that a patent regime's effect on R&D depends on the licensing environment); Bernard Girard, *Does 'Strategic Patenting' Threaten Innovation? And What Could Happen If It Did?* (Univ. of Que. at Montreal), <https://bit.ly/2UCzujf> (suggesting that overprotecting intellectual property in general, and patents specifically, is not the only solution to innovation; proposing that stronger intellectual property laws in the U.S. may not be good for other nations; and noting that patent law may cause speculative bubbles that could harm certain industries).

256. See Justin Hughes, *The Philosophy of Intellectual Property*, 77 GEO. L.J. 287, 303 (1998) (suggesting that the utilitarian foundation promotes labor, which in turn promotes public good).

objects is also socially desirable. What, then, is the utilitarian justification for patent protection of AI-generated output?

For most of the types of AI-generated output in question,²⁵⁷ their production is presumably undetectable. However, to the inventor, the creation of AI-generated output is of some significance. AI-assisted inventors are spending money and time on research and development for the application of AI-based tools and are placing a high value on the AI-generated output that they create. However, the amount of time and money that AI-assisted inventors spend on the AI-generated output that they create is significantly less than it would be without the use of AI-based tools. From the utilitarian perspective, society should consider that the grant of patents on adequate AI-generated output is justified based on its value to the inventor. Even with a narrow view of the social utility of AI-generated output, utilitarianism provides adequate justification for considering such output as beneficial to society.

There are some objections to the grant of patents in AI-generated output based on utilitarianism, the first of which stems from the application of intellectual property law theories. While utilitarianism may warrant some provision of exclusive rights to inventors, it is not an unfettered warrant but, instead, has limitations on the granting of such rights. For example, there may be limitations on the granting of patent rights, such as the exclusivity period, for certain patentable subject matter, even upon meeting the other patentability criteria. At this stage in the development of AI, it may be too soon to ascertain the appropriate balance of utilitarian interests and limitations. The second objection is that granting patent protection to AI-generated output would reduce the welfare and utility of human inventors and also reduce general societal wellbeing.

This Article responds to such objections and concludes that patent protection of AI-generated output can be recognized through engagement in a social balancing act with enhanced patent disclosure.²⁵⁸ In doing so, it argues that utilitarianism justifies patent protection of adequate AI-generated output but not for the allocation of those interests. Thus, the

257. *See supra* Section I.B (summarizing the types of AI-generated output in question as being potentially claimed in patent claims by: (1) claiming variations beyond what the inventor actually invented to broaden claim scope; (2) disclosing (without claiming) variations on existing patent claims and preventing future improvements to destroy novelty; (3) claiming next incremental steps and adding claims in response on demand to innovation developments in a certain field; and (4) claiming, by utilizing analytical techniques and correlations, to assist inventor in predicting performance and properties of inventions in the physical world; and (5) claiming prophecies with high accuracy through computational experimentation to allow for an early patent grant).

258. *See infra* Part III.

allocation issue is addressed through greater efforts required by the inventor through enhanced patent disclosure.²⁵⁹

2. Lockean Theories for Adequate AI-Generated Output

It is fitting to consider how Locke's theory of property might be applied to AI-generated output. Locke's conception of property stemmed from his view of America as a boundless land.²⁶⁰ Despite the era in which Locke wrote, one can make a comparison of Locke's view on property to the seemingly boundless modern-day environment of AI, which is open-ended in its possible applications. At the same time, given the scholarly debate on whether AI inventions can be attributed to the human or to the AI,²⁶¹ it is perhaps amusing that Lockean theory is considered in an area like AI, for which an invention may be artificial and arguably not human.

Locke's central thesis was that if a person removes a resource from nature and applies labor to it, that resource would become that person's property.²⁶² In effect, Locke's theory asserted that a person who expended labor converting something in nature into something valuable deserved some benefit for the effort.²⁶³

AI-generated output may have a property claim based on Lockean theory. Like any other asset that emerges from an inventor's time and effort, AI-generated output is the product of an inventor's labor. Some might claim that the application of an AI-based tool to produce AI-generated output is not labor. The Lockean perspective would hold that there is no discernable distinction between AI-generated output and human-generated output, especially in a world where AI usage is both prevalent and undetectable. Anyone who has invented with the assistance of AI-based tools can verify, "inventing" in the AI world involves less labor and time than in the pre-AI world.

Some would object to granting patents for adequate AI-generated output based on Locke's labor theory; the first objection deals with the degree of labor. The standard objection to Locke's labor theory uses as an example an action that provides little change, such as affecting only a few molecules.²⁶⁴ However, there are defenses to this objection. For one, Locke's labor theory only grants property to labor that makes the greatest part of the value, even in a small portion of the entire interest. Thus, for

259. See *infra* Part III.

260. See JOHN LOCKE, TWO TREATISES OF GOVERNMENT 290–91 (Peter Laslett ed., Cambridge Univ. Press 1988) (1690).

261. See Abbott, *supra* note 1, at 1080; Abbott, *supra* note 80.

262. See LOCKE, *supra* note 260, at 290–91.

263. See *id.*

264. See Ofer Tur-Sinai, *Beyond Incentives: Expanding the Theoretical Framework for Patent Law Analysis*, 45 AKRON L. REV. 243, 258 (2012).

example, with AI-generated output, an inventor may not be able to claim patent protection over its entirety, but only where the inventor's labor makes the greatest part of the value. Of course, this is exactly what makes AI challenging, since it blurs the distinctions between the real and artificial worlds.²⁶⁵ Additionally, another objection stems from the Lockean proviso that the property is granted only where there is "enough and as good left in common for others."²⁶⁶ There is some merit to an objection based on the Lockean proviso, since unlike the physical world, the provision of property interests in adequate AI-generated output may seem unlimited in a sense.

This Article responds to such objections and concludes in favor of recognizing patent protection of AI-generated output, which, through enhanced disclosure, enables humans to provide more labor.²⁶⁷ In so doing, it argues that Lockean theory provides a justification for patent protection of adequate AI-generated output, which is limited through acts demonstrating labor and acts that inventors would deem worthy of spending time and effort when they are of more value. Thus, the labor issue is addressed through greater efforts required by the inventor through enhanced patent disclosure.²⁶⁸

E. Toward an Enhanced Patent Disclosure & Responding to Objections

One policy response to the challenge of AI-generated output and unidentified inventions is to enhance the disclosure requirements for obtaining a patent. For example, Congress might amend the Patent Act to require greater disclosure on the part of the patent applicants, or courts might interpret existing doctrines as mandating more stringent disclosure. A requirement for greater disclosure would address any utilitarian and Lockean objections²⁶⁹ to AI-generated output and would, in theory, mitigate concerns with explainability and inscrutability of AI.²⁷⁰ Moreover, enhanced disclosure would be consistent with legal principles suggesting that patentees should not enjoy exclusive rights while concealing useful information about their inventions.²⁷¹ Indeed, it would be very difficult for the USPTO to ascertain whether an invention was produced as AI-generated output from the application of AI-based tools. For example, a patent examiner cannot know the subjective

265. See *infra* Part III.

266. Hughes, *supra* note 256, at 324.

267. See *infra* Part III.

268. See *infra* Part III.

269. See *supra* Sections II.D.1–II.D.2.

270. See *supra* Section I.B.

271. See *In re Gay*, 309 F.2d 769, 772 (C.C.P.A. 1962).

contents of an inventor's mind and, therefore, would be unaware of whether the inventor utilized AI-based tools to produce AI-generated output or produced human-generated output.²⁷²

While there are trade-offs with enhanced disclosure, this Article proposes several prescriptions²⁷³ to maximize efficiency and minimize drawbacks. In terms of maximizing efficiency, a trade-off arises regarding enhancing patent disclosure. A heightened AI disclosure may certainly impose high *ex ante* costs or time on inventors.²⁷⁴ Additionally, an AI-technology-specific intervention has some perils that scholars have noted, prompting them to advocate against abandoning uniformity. As it may relate to both of these objections, a move toward a specialized response to AI would entail higher administrative costs for the USPTO and legislative reform. While these objections present challenges for political economy and are not without merit, recognizing that AI presents unique technological features²⁷⁵ would benefit patentability and patent examination.

One of the central insights of the economics of patent law is that context matters,²⁷⁶ and the key is to strike the right balance for any technology-specific solution. In the context of software, patents have become more difficult to identify and assess for patent examination at the USPTO. As software patents have become increasingly obscure due to abstruse writing and incomprehensible scope, software patents have become increasingly difficult to search among and differentiate at the

272. However, the USPTO might have some options. Patent practitioners have a duty of candor to the office, and USPTO regulations qualify that the duty of candor requires disclosure of AI means used in the patent. The duty of candor requires that the inventor and all individuals associated with the filing of the patent application disclose to the USPTO all information that is known to be material to patentability, and as such, the inventor should disclose when an AI-based tool has been used for AI-generated output. A violation of the duty of candor can raise inequitable conduct issues during patent litigation that could invalidate the entire patent—this should motivate the inventor to specify the use of an AI-based tool. Nonetheless, without strong mechanisms (that are easy to implement and cost-effective) for ascertaining a breach of the duty of candor, an inventor may not be motivated to disclose the use of an AI-based tool to produce AI-generated output.

Also, a patent examiner may consider written description and enablement to be insufficient and require inventors to provide declarations supplementing the disclosure or explaining how a person having ordinary skill in the art would understand how to make and use the invention. Despite these options that are available to the patent examiner, the lack of detectability when applying an AI-based tool necessitates greater patent disclosure from societal and theoretical standpoints.

273. See *infra* Part III.

274. See *infra* Section III.A.1.

275. See *supra* Part I.

276. See Robert Merges & Richard R. Nelson, *On the Complex Economics of Patent Scope*, 90 COLUM. L. REV. 839, 839 (1990) (explaining that the economic significance of a patent depends on its scope and, as a result, “the broader the scope, the larger the number of competing products and processes that will infringe the patent”).

USPTO.²⁷⁷ In many ways, this behavior reveals that software patents demonstrate a related narrative of disclosure in patent application and patent examination. AI magnifies this problem with software and necessitates technology-specific intervention.

Additionally, because the current patent examination process is incapable of distinguishing between AI-generated output and human-generated output, an enhanced patent disclosure would promote the incentives and justifications of the patent system. While there are trade-offs with economic efficiency due to higher *ex ante* costs or time for inventors, as seen in Part III, there are wider welfare-enhancing benefits.²⁷⁸ Somewhat analogously, scholars have recommended other policy levers to affect patenting behaviors with preferential treatment of other technological areas.²⁷⁹ Ultimately, these technology-specific developments help define a narrative of the policy response to technological developments in patent law.

U.S. patent law should embrace the reality that AI-generated output is already part of the invention process. The use of AI-based tools in the invention process to produce AI-generated output is wholly consistent with and advances the aims of the patent system, but only when there is enhanced patent disclosure. Rather than viewing such AI inventions as peripheral or exceptional, Congress, courts, and the USPTO should promote a more robust vision of modern-day inventing. Accordingly, a sketch of the contours of an enhanced AI patent disclosure—as provided next in Part III²⁸⁰—would cultivate a more nuanced understanding of what the patent system should be in an AI world.

III. NORMATIVE IMPLICATIONS, PRESCRIPTIONS, & FUTURE DIRECTIONS

The technological distinctions of AI explored in Part I and the normative assessments explored in Part II raise several important implications, normative considerations, and additional questions for further study. Having examined the theories, justifications, and

277. Stephen Lindholm, *Marking the Software Patent Beast*, STAN. J.L. BUS. & FIN., Spring 2005, at 82, 82 (“[S]oftware patents are, practically speaking, hidden away in the recesses of the patent office and practically impossible to find.”).

278. See *infra* Part III.

279. See Lucas S. Osborn et al., *A Case for Weakening Patent Rights*, 89 SAINT JOHN’S L. REV. 1185, 1189, 1191 (2015) (assessing reducing the patent term as a policy lever, and in so doing, arguing that since emerging technologies are altering relative costs and benefits of the patent system, lawmakers should weaken patent rights by 25–50%); Anton et al., *supra* note 63, at 2 (suggesting that the level of patent strength is a policy lever that will affect innovation, organizational incentives, and entrepreneurial activity tied to patents).

280. See *infra* Part III.

normative assessment of enhanced patent disclosure for AI, an important question arises: how should the law respond? Accordingly, this Part prescribes more robust patent disclosure of AI. Although this Part's primary aim is to describe various proposals for enhancing patent disclosure of AI-generated output, some preliminary normative considerations and normative implications are in order.

A. Normative Implications of Enhanced AI Disclosure

Having examined the technological distinction of AI and the application of AI-based tools to produce AI-generated output, an important question arises: what are the necessary considerations and implications of any proposed prescriptions? Accordingly, it is important to engage in a normative evaluation of any prescription. Indeed, much is at stake here, for the use of AI represents a great portion of inventing and results in unidentified inventions.²⁸¹ Society should care about patents associated with imaginary, never-achieved, or unworkable disclosure but that appear to reasonable onlookers of the patent document as though they were created, tested, or made workable. Congress should enact reforms to the patent system to require greater disclosure of AI-generated output of inventions that were hardly (or never) developed or were effectively concealed through an unexplainable algorithmic inventive process. This Article argues that while AI-generated output is worthy of patent protection, there should be enhanced patent disclosure to balance societal welfare.²⁸² The prescriptions presented here raise two primary concerns: (1) greater *ex ante* inventor costs and efforts for AI inventors, and (2) a possible move toward trade secrecy for AI inventions.

1. Greater *ex ante* Inventor Costs and Efforts for AI Inventors

There is an optimal patent policy balance between encouraging the creation and disclosure of advances in AI technology by providing incentives for innovation while weighing the costs and efforts for the inventor.²⁸³ That being said, while enhanced AI disclosure may require greater incentives for the AI invention, such concerns raise special considerations in the context of more effort toward reduction to practice for an AI inventor. This analysis provides a normative evaluation of *ex ante* inventor costs and efforts required by enhanced AI patent disclosure. In particular, proposals to enhance AI patent disclosure would

281. See Iancu AI Policy Update, *supra* note 48.

282. See Miller, *supra* note 40, at 196.

283. See *id.* at 196–97 (recognizing that enhanced patent disclosure would result in an increased cost to patent applicants in initial filing of the patent application but would result in a large social benefit).

have to counter the historical departure from greater disclosure in patent law's history.

Early in patent law's history, the requirement for a close connection to the physical "working" of the claimed invention in the form of physical embodiments was high.²⁸⁴ Since then, there has been a shift away from actual reduction to practice of the invention, which has brought patent disclosure to the forefront. Due to a host of developments, interpretation of federal patent law and statutes, and interpretations by courts and the USPTO, the physical-based requirements gradually shifted from submission of a physical model of the invention to a commercial working of the invention to constructive reduction to practice.²⁸⁵

Contemporary patent law views the filing of the patent application as a representation that the invention is capable of actually working if it were built, but the inventor is not required to build it, practice it, or make it.²⁸⁶ The shifting normative status of the type of disclosure has reduced the amount of evidence required to show an invention would actually work. This shifting disclosure is part of a broader structural trend in patent law of not accounting for the evolution of the PHOSITA with the progression of computing and software technology.²⁸⁷ In many ways, the developments of AI technology described here reflect increased tensions between software technologies and patent disclosure.

Such a decreased emphasis on disclosure in patent law raises a special concern in the context of AI, for its technological distinction with explainability and inscrutability further dampens the benefits of the teaching function and knowledge spillovers of patent disclosure.²⁸⁸ Furthermore, the high degree of trifling inventions that are reaching the USPTO,²⁸⁹ especially in light of the potential flood of unidentified inventions with AI-generated output at the USPTO, further justifies enhancing modern-day patent disclosure. The USPTO has provided guidance for AI inventors to better demonstrate greater AI patent disclosure.²⁹⁰ The USPTO's guidance documents aim to improve the

284. See Cotropia, *supra* note 32, at 1545.

285. See *id.* at 1545, 1549–51, 1554, 1563; John F. Duffy, *Reviving the Paper Patent Doctrine*, 98 CORNELL L. REV. 1359, 1368–71 (2013).

286. See Duffy, *supra* note 285, at 1366.

287. See Samuel Adams, *Law and Economics of Software Patent Disclosure* 38–39 (Nov. 10, 2008) (unpublished LL.M. IP Thesis, Munich Intell. Prop. L. Ctr.), [://bit.ly/2ZG49PZ](https://bit.ly/2ZG49PZ).

288. See Deepak Hegde et al., *Patent Publication and Innovation* 2–5 (Jan. 20, 2019) (unpublished manuscript), <https://bit.ly/3efTKi6>.

289. See Jeanne C. Fromer, *The Layers of Obviousness in Patent Law*, 22 HARV. J.L. & TECH. 75, 75 (2008).

290. See generally 2019 Revised Patent Subject Matter Eligibility Guide, 84 Fed. Reg. 50 (Jan. 7, 2019) (providing a guide, as well as a flowchart and examples, concerning assessment of the threshold question of whether a patent's claimed subject

clarity, consistency, and predictability of patent examination of AI inventions, particularly by defining what is and is not patent-eligible as applied to inventions associated with the USPTO's Class 706.²⁹¹ Turning to the prescriptive, Section III.B²⁹² returns to the question of how the patent system should promote greater AI disclosure,²⁹³ primarily through the lens of other possible USPTO actions.

2. Possible Move Toward Trade Secrecy for AI Inventions

Patents and trade secrets occupy opposite ends of the disclosure spectrum and achieve their goals through different mechanisms. The patent bargain requires a disclosure of the invention for limited-term monopoly protection until the patent's expiration, whereas a trade secret lacks an expiration date and loses protection with reverse engineering or independent invention. Patents have strength in enforcement over trade secrets since patents can be enforced against anyone making, using, or selling the invention. A recent AI patent infringement²⁹⁴ case,

matter is within the realm of topics that are even eligible for patenting and, in so doing, demonstrating what can be disclosed and written to meet patent eligibility and indirectly provide greater disclosure).

291. See *Classification Resources: Data Processing – Artificial Intelligence*, U.S. PAT. & TRADEMARK OFF., <https://bit.ly/2B7YGYd> (last visited May 27, 2020). The subject matter of this technological class is described as being:

[A]rtificial intelligence type computers and digital data processing systems and corresponding data processing methods and products for emulation of intelligence (i.e., knowledge based systems, reasoning systems, and knowledge acquisition systems); and including systems for reasoning with uncertainty (e.g., fuzzy logic systems), adaptive systems, machine learning systems, and artificial neural networks.

Id.

292. See *infra* Section III.B.

293. See *supra* Section II.E.

294. At first blush, AI patent infringement might seem preposterous in multiple ways.

One scenario relates to when an AI-based tool operates in a way that unexpectedly infringes a patentee's patent. In such a scenario, when an AI-based tool invents independently of a human, a question arises as to whether the AI-based tool can be found to have infringed on a patented invention when the human that uses the AI-based tool is unaware of the AI-based tool's potential to infringe another's patent. Even though the unpredictable and challenging-to-reproduce nature of an AI-based tool may cause some pause on the basic presumptions of the patent infringement statute, it should be noted that the AI-based tool is like any other tool, and the human as the user of the tool would be found liable for patent infringement (not the tool itself). The AI-based tool should not be found to be the infringer. After all, because there is a human that is behind the development of the AI-based tool in some way, the AI-based tool cannot be seen as being purely autonomous. Even though an AI-based tool can produce some AI-generated output independent of the human (or as this Article assesses, inscrutable AI-generated output from an AI-based tool lacking explainability), the AI-based tool is a product of human ingenuity, and a human can be considered a patent infringer even without knowing that the AI-based tool was producing AI-generated output that infringes on another's patent. Moreover, while a patentee may be concerned when another human uses an AI-based tool

PurePredictive, Inc. v. H2O.AI, Inc.,²⁹⁵ other AI patent litigation cases,²⁹⁶ and PTAB proceedings²⁹⁷ represent signs of more AI enforcement

that produces unexpected AI-generated output that infringes on the patentee's patent, it may never be noticed by the patentee.

Another scenario concerning the possibility of AI patent infringement relates to when the patentee of the AI-based tool sues another party for patent infringement. The other party may seek to invalidate the patentee's patent, since patents directed to AI and based on algorithms and statistics may be invalidated in district courts or in post-issuance administrative proceedings at the USPTO.

The aforementioned scenarios provide a preliminary analysis. However, the prospect of AI patent infringement requires a deeper analysis of the following questions: Can AI (or some aspect of AI) infringe on patent rights? Who or what entities should be liable for actions taken by AI, if it could infringe on a patent? Should these questions differ in the case of AI that infringes on a patent whose inventor is a human versus the case of AI that infringes on a patent whose conception was performed by AI? While a preliminary analysis was performed here, a deeper assessment of these questions is outside of the scope of this Article and would require grappling with the fundamental notion of who or what type of entity could be an "infringer" and exploring the way the patent system can respond to protect patent owners against appropriations of their inventions by use of AI. A detailed exploration of the various applications and uses of AI under 35 U.S.C. § 271 that constitute patent infringement would require an assessment of the theories and causes of action of patent infringement, including the nuances of indirect infringement, divided infringement, and infringing uses with the use of AI.

295. See *PurePredictive, Inc. v. H2O.AI, Inc.*, No. 17-cv-03049-WHO, 2017 WL 3721480, at *1 (N.D. Cal. Aug. 29, 2017), *aff'd*, 741 F. App'x 802 (Fed. Cir. 2018) (addressing U.S. patent 8,880,446, which concerned automatically generating an ensemble of machine learning models, such that the collection of such models cooperated to generate accurate predictions or a classifications by repeated exposures to training data with little to no input from users through evaluating a number of different learned functions).

296. See *Elec. Power Grp., LLC v. Alstom S.A.*, 830 F.3d 1350 (Fed. Cir. 2016); *Vehicle Intelligence & Safety LLC v. Mercedes-Benz USA, LLC*, 78 F. Supp. 3d 884, 885–86 (N.D. Ill. 2015), *aff'd*, 635 F. App'x 914 (Fed. Cir. 2015) (holding claims directed to an expert system for detection of impairment of human operators were ineligible); *Hyper Search, LLC v. Facebook, Inc.*, No. 17-1387-CFC-SRF, 2018 WL 6617143, at *10 (D. Del. Dec. 17, 2018) (finding the claims to be invalid under 35 U.S.C. § 101 for reciting a neural-network module that was limited to an abstract idea); *Neochloris, Inc. v. Emerson Process Mgmt. LLP*, 140 F. Supp. 3d 763, 773 (N.D. Ill. 2015) (finding the claims to be ineligible for reciting an artificial network module akin to a black box without any limitations on the system).

297. See *generally Ex parte Hamilton*, No. 2017-008577, 2018 WL 6428478 (P.T.A.B. Nov. 16, 2018) (rejecting a claim directed to a method of planning and paying for advertisements in a virtual universe via a set of virtual agents controlled by artificial intelligence); *Ex parte Pizzorno*, No. 2017-002355, 2018 WL 4846938 (P.T.A.B. Sept. 19, 2018) (affirming a rejection of a patent claim directed to "a computer implemented method useful for improving artificial intelligence technology" as abstract); *Ex parte Lyren*, No. 2016-008571, 2018 WL 3391361 (P.T.A.B. June 25, 2018) (rejecting a claim directed to customizing video on a computer as being abstract and thus not patent eligible); *Ex parte Mitzlaff*, No. 2016-003447, 2018 WL 1737978 (P.T.A.B. Mar. 27, 2018) (rejecting claims directed to a computer-implemented method designed to simulate a conversation of a user of a computer-implemented conversational agent as being too abstract, and finding that embodiments involving complex trained artificial intelligence algorithms did not identify anything more than routine and conventional technologies).

actions,²⁹⁸ including divided infringement possibilities²⁹⁹ that suggest some strategic patent drafting considerations.³⁰⁰ The objective construction of patent liability makes patent-based protection more attractive, as it is a form of strict liability that requires no knowledge or intention on the part of the alleged infringer.³⁰¹ Moreover, patent protection can capture small changes under the doctrine of equivalents,³⁰² and secondary liability for indirect patent infringement extends to a broader class of actors. By contrast, trade-secret protection is essentially a form of protection that extends liability for unlawful acquisition without the consent of the trade-secret holder.

An enhanced AI patent disclosure would push some AI inventors toward trade secrecy. In addition to increased *ex ante* costs and efforts with providing more disclosure in the patent application, AI inventors may find trade secrecy to be beneficial in other respects.³⁰³ AI's

298. See WORLD INTELLECTUAL PROP. ORG., WIPO TECHNOLOGY TRENDS 2019: ARTIFICIAL INTELLIGENCE 111 (2019) (noting that, while not comprehensive, initial worldwide data reveals 1,264 artificial intelligence patent families mentioned in litigation cases, 4,231 mentioned in opposition cases, and 492 mentioned in both types of case; of these cases, 73% of the identified litigation cases involving artificial intelligence patents were filed in the U.S.).

299. See Nathaniel Grow, *Resolving the Divided Patent Infringement Dilemma*, 50 U. MICH. J.L. REFORM 1, 3–4 (2016); W. Keith Robinson, *Using Interactive Inventions*, 69 DEPAUL L. REV. 95, 96, 100 (2019) (introducing interactive inventions as being systems and processes that can be used by multiple actors at the same time, and in so doing, assessing liability associated with patent infringement, including the issue of divided infringement where the performance of the interactive method patent claim is split among multiple parties, actors, or devices); Christopher J. White & Hamid R. Piroozi, *Drafting Patent Applications Covering Artificial Intelligence Systems*, LANDSLIDE, Jan./Feb. 2019, at 12, 16, <https://bit.ly/3fbzcYf> (explaining that a method claim to an AI model training technique is infringed when the patentee asserts a patent infringement claim against the producer of the training software, and that divided infringement may occur with the use of AI when a party does not train models itself in an attempt to divide infringement between itself and the end users of its software).

300. Divided infringement occurs when the actions of multiple entities are combined to perform every step of a claimed method, but no single party acting alone has completed the entire patented method. Furthermore, the multiple parties can be located in different jurisdictions. In order to avoid divided infringement, there are some strategic patent-claims-drafting considerations that AI inventors should keep in mind: (1) A patent claim reciting the use of a trained model without requiring any particular training process would likely be infringed by only one party; (2) A patent claim reciting the use of a training process without requiring any particular trained model would likely be infringed by only one party; (3) Since training data may come from numerous sources, a claim reciting “receiving training data” may implicate multiple parties.

301. See Anton et al., *supra* note 63, at 5–6 (suggesting that strong patents give the patent holder stronger protection against patent infringement).

302. See Meurer & Nard, *supra* note 126, at 1948 and accompanying text.

303. Additionally, the uncertain scope of AI protection under patent law, such as with patent eligibility and non-obviousness, may make trade secrecy a more realistic alternative. A lack of clear patentability for AI makes patent protection less favorable since enforcement against patent infringement would be more problematic. While

explainability and inscrutability make reverse engineering more difficult.³⁰⁴ An enhanced AI patent disclosure would enable the inventor to capture small changes under the doctrine of equivalents³⁰⁵ and protect against something similar. Turning to the prescriptive, Section III.B³⁰⁶ returns to the question of how the patent system should promote greater AI disclosure³⁰⁷ to provide such benefits and incentives for AI inventors.

B. Assessments & Prescriptions for Enhanced AI Disclosure

The concern with inscrutability and reproducibility³⁰⁸ that results from the use of AI-based tools sheds new light on the longstanding debate over patent disclosure and incentives for inventors. The patent system is meant to incentivize inventors, but increasingly patent examiners are presented with imaginary, never-achieved, or unworkable inventions that appear to reasonable onlookers of a patent application as though they were created, tested, or made workable. Patent examiners must take the patent application at face value, yet legitimate-sounding diagrams and descriptions of the invention might in fact be illegitimate, as a result of AI-based tools; the patent system should evolve in response. The USPTO cannot reasonably expect patent examiners to confirm whether the patent application is for an invention that is fictitious or unexplainable in an era of increasing usage of AI-based tools, and heightened disclosure provides a better verification mechanism for society.

Recall that scholars have long argued over the costs and benefits of patent disclosure.³⁰⁹ Drawing on the foregoing accounts,³¹⁰ this Section argues that greater AI disclosure by the inventors has benefits that outweigh the costs and provides proposals that would counter patent protection for AI-generated output. Given the social welfare benefits yielded by greater AI disclosure, such calibration is worth pursuing.³¹¹

misappropriation and independent invention create a loss in trade secrecy protection and are concerns for an AI inventor considering trade secrecy, patent protection may not be desirable to pursue unless patent disclosure is enhanced for AI inventions.

304. It should be noted that reverse engineering barriers may not be that high in some cases, in which trade secrecy would not block involuntary dissemination. Moreover, AI patent owners may license a method to others, in which case patenting similarly does not block voluntary dissemination.

305. See Meurer & Nard, *supra* note 126, at 1948 and accompanying text.

306. See *infra* Section III.B.

307. See *supra* Section II.E.

308. See *supra* Section I.B.

309. See *supra* Section II.B.

310. See *supra* Section II.B.

311. See generally Adam Mossoff, *Testimony on the STRONGER Patents Act Before the Senate Judiciary Committee, Intellectual Property Subcommittee* (Geo. Mason L. & Econ. Research Nos. 19-32, 2019), <https://bit.ly/2UGLgtc> (proposing that Congress restore reliable and effective patent rights, in part to foster the next generation of

Enhanced patent disclosure for AI has an important role to play in equilibrating an appropriate level of quid pro quo. Along these lines, this Section takes the view that the USPTO can and should enact proposals to enhance AI patent disclosure that fit within its agency authority,³¹² without creating rulemaking in patent law.³¹³

1. Establishing Disclosure Incentives for the Inventor

Enhanced patent disclosure for AI plays an important role in incentivizing disclosure of unidentified inventions in the form of AI-generated output in patent applications examined at the USPTO. Along these lines, a prudent response is not to ban the AI-generated output, but to develop an enhanced disclosure mechanism to incentivize AI inventors to explain how the applied AI-based tools develop AI-generated output.

However, there are trade-offs to enhanced AI patent disclosure. The more incentives that are provided by the USPTO to disclose, the more difficult it will be to implement. Moreover, greater AI patent disclosure incentives will require greater *ex ante* costs and time for inventors. The added disclosure will require the inventor to spend more money on attorney time as the cost of writing an AI patent application will increase. Another cost borne by the AI inventor would be the increased risk of inequitable conduct.³¹⁴ On the one hand, the potential for added inequitable conduct risk may lead some AI inventors to refrain from providing enhanced AI patent disclosure. On the other hand, a more enticing incentive for inventors may cause them to over-disclose to avoid inequitable conduct and place a greater administrative burden on patent examiner. The risk that patent litigators may introduce inequitable conduct by AI inventors may be high enough to prevent enhanced AI patent disclosure measures.³¹⁵ Additionally, there will be increased

innovations in AI since extensive changes to the patent system have weakened patents and sowed legal uncertainty, which requires a recalibration).

312. See Robert P. Merges, *The Hamiltonian Origins of the U.S. Patent System, and Why They Matter Today*, 104 IOWA L. REV. 2559, 2588 (2019) (explaining that the USPTO is an expert agency that has the authority to make rules that govern its internal proceedings and make rules on procedural matters, but that it is not an agency that possesses the power to issue binding rules with the force of law, such as on patentability).

313. See 35 U.S.C. § 2(b)(2)(A) (2018).

314. See generally Jason Rantanen & Lee Petherbridge, *Therasense v. Becton Dickinson: A First Impression*, 14 YALE J.L. & TECH. 226 (2012) (specifying that inequitable conduct is a judicially-created doctrine developed to punish inventors who behave inappropriately during patent prosecution); see also MPEP (9th ed. Rev. 8, Jan. 2018) § 1448.

315. See Lee Petherbridge et al., *The Federal Circuit and Inequitable Conduct: An Empirical Assessment*, 84 S. CAL. L. REV. 1293, 1295, 1299 (2010) (describing that inequitable conduct is a “plague” that permeates patent litigation and strikes fear in the hearts of inventors).

administrative costs for the USPTO based on the costs of incentives for enhanced AI patent disclosure. Given the overburdened USPTO, additional fees may be necessary to offset the cost of any enhanced AI patent disclosure program. The additional costs for the inventor and the USPTO will depend on the level of enhanced disclosure required.

This Article proposes a few incentive options for enhanced AI patent disclosure. The most highly recommended proposals—optional AI patent disclosure, prioritized AI examination, and reduced maintenance fees—would necessitate a few USPTO procedural changes. Additional proposals with more substantial incentive levels—a greater patent term for enhanced AI patent disclosure and a requirement for complete AI patent disclosure—would necessitate some legislative action. In sum, the prescriptions provided represent a spectrum of disclosure solutions for AI's explainability and inscrutability challenges with patent law and address adequate disclosure to identify the AI inventive method applied for producing AI-generated output.

Foremost, optional AI patent disclosure and prioritized AI examination would necessitate a few USPTO procedural changes but not a legislative overhaul. This first prescription of disclosure incentives concerns mostly relatively easy-to-implement procedural reforms at the USPTO that fit within its agency authority. There could be an increasing range of greater incentives, which could be analogized as being a larger set of carrots on a stick.

The level of incentive could be optional AI patent disclosure that would allow an AI inventor to voluntarily identify the AI inventive method utilized to generate AI output. The USPTO could provide a survey mechanism to enable an AI inventor to identify the AI inventive method applied toward their AI-generated output. The USPTO survey could be provided in the patent application filing process and enable the AI inventor to identify their use of a particular AI-based tool, including an open-source library, training data,³¹⁶ or output data. Inventors would have minimal to no incentive to disclose because the optional AI disclosure via a survey or a form would be left up to the inventors' choice. Still, while such an optional AI patent disclosure would not provide much incentive to AI inventors, it would serve a notice-serving function and could act as defensive patenting strategy to block follow-on inventors from applying the same technique.

The second level of incentive on this spectrum of disclosure solutions is priority examination, which would enable patent applications that provide disclosure of the AI inventive process to attain a higher

316. See Taddy, *supra* note 89, at 1; see also MUNRO & MIRSHARIF, *supra* note 58, at 10.

priority for patent examination. This incentive refers to providing an inventor an earlier place in line to enable a quicker review of the patent application by the USPTO when there is greater disclosure provided by the inventor. The USPTO has already had priority examination programs,³¹⁷ and an AI disclosure program could provide a petition for specially expedited examination. The prospect of an expedited patent application could provide a generous disclosure incentive.

A third level of incentive for disclosure on this spectrum is reduced maintenance fees, which would encourage AI inventors to disclose the AI inventive process by making the cost of patent maintenance more affordable.³¹⁸ This incentive refers to charging a lower cost to maintain the patent *ex post*, or after issuance, but would require the inventor to adequately disclose the AI *ex ante*, or before issuance. If the USPTO lowered the renewal-maintenance fees of AI patents, then AI patents would become stronger.³¹⁹ Similar to reduced maintenance fees for small and micro entities, lower maintenance fees for AI patent applications would provide greater incentives for the patenting of such inventions. In sum, optional AI patent disclosure, prioritized AI examination, and reduced maintenance fees would provide enhanced disclosure incentives to AI inventors with minimal procedural changes at the USPTO.

The three incentive levels represent a spectrum of disclosure incentives, are greater incentives for AI than in present day, and would not require legislative action. Other proposals with more substantial incentive levels—a greater patent term for enhanced AI patent disclosure and a requirement for complete AI patent disclosure—would necessitate rulemaking changes. The USPTO is charged with establishing regulations that are consistent with the patent laws that govern its proceedings,³²⁰ but it does not have rulemaking authority.³²¹ Therefore, adding a greater patent term, or adding additional years of patent protection for disclosing an AI inventive method, would require an act of Congress. The Patent Act would need to be amended to identify what

317. See Press Release, U.S. Patent & Trademark Office, The U.S. Commerce Department's Patent & Trademark Office (USPTO) Will Pilot a Program to Accelerate the Examination of Certain Green Technology Patent Applications (Dec. 7, 2009) (on file with the U.S. Patent & Trademark Office).

318. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2506.

319. See Colleen V. Chien, *Reforming Software Patents*, 50 HOUS. L. REV. 325, 360–63 (2012) (discussing an increase in maintenance fees as a deterrent to non-practicing entity patent litigation); Brian J. Love, *An Empirical Study of Patent Litigation Timing: Could a Patent Term Reduction Decimate Trolls Without Harming Innovators?*, 161 U. PA. L. REV. 1309, 1357 (2013); Lucas Osborn et al., *supra* note 279, at 1191 (noting that raising maintenance and renewal fees would weaken patents, and thus, lowering maintenance and renewal fees would strengthen patents).

320. See MPEP (9th ed. Rev. 8, Jan. 2018) § 1001.

321. See 35 U.S.C. § 2(b)(2)(A) (2018).

would serve as an AI inventive disclosure and provide the extra term for the added disclosure. Subsequent to Congress's passing of this law, the USPTO would need to develop its procedure and examination capabilities for its execution. This process would take many years but, upon its conclusion, would provide a strong incentive for an AI inventor to provide enhanced AI patent disclosure. An even stronger incentive for an AI inventor would be to require complete AI patent disclosure simply to qualify for a complete patent application. In other words, this prescription would do more than merely encourage disclosure; disclosure would be a patent-application requirement to qualify for examination of a submitted patent application. Similar to other requirements, such as payment of a fee and submission of an oath under 35 U.S.C. § 111, a complete AI patent disclosure would be necessary for the AI inventor's patent application to be complete. Thus, unlike the statutory requirement in 35 U.S.C. § 112(a),³²² this proposal would address a patent application's completeness or incompleteness. Similar to the proposed greater patent term, an AI inventor would have a strong incentive to provide such disclosure or else would be unable to proceed in the patent prosecution process and would receive a Notice of Missing Parts from the USPTO. This strong incentive would require revisions to the Patent Act to specify requirements for AI inventors to satisfactorily complete the patent application under 35 U.S.C. § 111.

In sum, some proposals would require procedural changes, and others would require legislative actions. While the proposals for procedural changes would be easier to implement, the proposals for legislative actions would take time and receive resistance from various industry groups and lobbyists. In particular, organizations that would oppose greater *ex ante* inventor costs and time in the invention process could thwart any legislative change. Alternatively, a data deposit requirement would be easier to implement because the USPTO could draw on prior comparable circumstances to aid in its implementation.

2. Establishing a Data Deposit as an Alternative Disclosure

While this Article takes a general salutary view of enhanced AI patent disclosure of AI-generated output, in some cases, AI-generated output is not solely produced from the application of AI-based tools. A

322. 35 U.S.C. § 112(a) (2018). This section states that:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same

Id.

combination of human-machine interaction,³²³ or the use of AI together with human input,³²⁴ can produce AI-generated output. Thus, a potential challenge in advocating for enhanced AI patent disclosure is that AI-generated output may be the result of minimal AI application. Along these lines, some consider that a stronger AI disclosure may be unwarranted if AI technology has not advanced to the point of having the sole or even substantial ability to contribute to the conception of an invention.³²⁵ Nevertheless, when AI-based tools confer meaningful autonomy to inventions, there may be societal benefits for similar incentives and justifications³²⁶ for some form of enhanced AI patent disclosure. In such cases, enhanced AI patent disclosure may be warranted for certain AI-based tools.

A potential—though ultimately a time-consuming and difficult—legislative reform would be to enact a statute for patent protection with a data deposit requirement akin to that provided for seeds of plants, which provides biological sequence listings (such as for DNA) to the USPTO. Similar to the biological context, where the inner processes are not fully understood and a deposit of certain materials could satisfy teaching the public in lieu of a textual explanation in a patent application, a deposit of certain data in the AI context could serve an analogous role. For instance, a deposit requirement or some form of a repository could be accessible to the public and satisfy enablement of the patentability statute similar to the biological context and the Plant Patents Act (PPA).³²⁷ The objective of enacting such a statutory reform would be to provide an alternative form of patent protection and, in so doing, provide enhanced AI patent disclosure, when utility patent protection is unachievable for AI-generated output. This Article has argued that the explainability and inscrutability considerations of applying AI-based tools to produce AI-generated output warrant an enhanced AI patent disclosure system. In theory, the USPTO could heighten the disclosure requirement for AI inventions through USPTO procedural reforms,³²⁸ thus strengthening the ability of AI patents to promote knowledge spillovers and, ultimately, AI

323. See GUY A. BOY, *THE HANDBOOK OF HUMAN-MACHINE INTERACTION* 3–4 (Guy A. Boy ed., 2011) (explaining human-machine interaction from analysis, design, and evaluation perspectives, including principles, methods, techniques, and tools).

324. See Mühlhoff, *supra* note 109, at 6–8 (providing and explaining as an example the use of a CAPTCHA, which requires the response of a human user as training data for an AI application).

325. See Memorandum from R Street on Patenting of Artificial Intelligence Inventions (Nov. 13, 2019) (on file with the U.S. Patent & Trademark Office).

326. See *supra* Section II.D.

327. See 35 U.S.C. §§ 161–64 (2018).

328. See *supra* Section III.B.1.

innovation.³²⁹ However, such an approach could be unavailing for AI inventions that do not meet the definiteness requirements of 35 U.S.C. § 112 and, hence, do not warrant patent protection.

AI patent protection similar to the PPA could take the form of a data deposit requirement. The key here is to produce patent protection and greater disclosure of the Article's second scenario, "Un-disclosed AI-based tool,"³³⁰ akin to that of plant patent protection³³¹ and the biological deposits and sequence listings.³³² Just as deposits in the biological and plant contexts represent starting materials, a data deposit in the AI context could provide the starting data set for subsequent AI-generated output. Similarly, Congress should enact a statute to allow for a public deposit of data in the AI context.

Similar to the PPA, which requires an inventor to submit both a complete description of the plant variety and an explanation of how a particular type of plant includes the right to exclude others from asexually reproducing the plant,³³³ a data deposit could require similar disclosure of AI-based tools to provide patent protection. Similar to information embedded in seeds being amenable to various forms of proprietary protection,³³⁴ the information in data that is trained with AI-based tools to produce AI-generated output should have similar patent protection upon deposit of the AI training data. The AI inventor could be

329. See Exec. Order No. 13859, 84 Fed. Reg. 3967 (Feb. 11, 2019); NAT'L SCI. & TECH. COUNCIL, SELECT COMMITTEE ON ARTIFICIAL INTELLIGENCE: THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN: 2019 UPDATE, at iii (2019); CENTER FOR THE FOURTH INDUSTRIAL REVOLUTION, WORLD ECONOMIC FORUM, ARTIFICIAL INTELLIGENCE COLLIDES WITH PATENT LAW 4 (2018); Asay, *supra* note 1, at 1187; Bessen et. al., *supra* note 2; Furman & Seamans, *supra* note 2.

330. See *supra* Section I.B (referring to the claiming of an AI-based tool, without disclosure of how the AI operates, such as when an inventor claims the AI, which refers to the AI engine or the AI techniques inside of the AI-based tool, but does not disclose how an invention is made by the AI-based tool).

331. See Michael Blakeney, Patenting of Plant Varieties and Plant Breeding Methods, 63 J. OF EXPERIMENTAL BOTANY 3 (2012); Mark D. Janis & Jay P. Kesan, *Intellectual Property Protection for Plant Innovation: Unresolved Issues After J.E.M. v. Pioneer*, 20 NATURE BIOTECHNOLOGY 1161, 1161–62 (2002) (noting that plants are eligible for patent protection under the utility patent regime, the Plant Patent Act, and the Plant Variety Protection Act); Elisa Rives, *Mother Nature and the Courts: Are Sexually Reproducing Plants and their Progeny Patentable Under the Utility Patent Act of 1952?*, 32 CUMB. L. REV. 187, 187 (2002) (explaining that the case of *Pioneer Hi-Bred Int'l, Inc. v. J.E.M. AG Supply, Inc.* helped to "resolve the issue of whether the Plant Patent Act of 1930 (PPA) and the Plant Variety Protection Act of 1970 (PVPA) "[were] the exclusive forms of protection for plant life").

332. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2403.

333. See 35 U.S.C. §§ 161–64 (2018) (requiring that a complete description of the plant variety be submitted before a patent is used).

334. See Jim Chen, *The Parable of the Seeds: Interpreting the Plant Variety Protection Act in Furtherance of Innovation Policy*, 81 NOTRE DAME L. REV. 105, 108 (2005).

required to submit AI training data or an AI-based tool to the USPTO, similar to the requirement of submitting a deposit seed or other propagation material capable of producing claimed plants, wherein the deposits are treated as part of the applicant's disclosure.³³⁵ Just as a deposit of active biological material, whose inner-workings are difficult to replicate or understand, provides the public with access to how to make and use it, AI training data deposits can serve the same purpose in the AI context.

Biological deposits and sequence listings make another useful analogy to AI-based tools for patent disclosure. Biological deposits, which have been used to satisfy the enablement and written description requirement for inventions in the U.S. since 1949, provide a means for making biological material publicly available as a clearinghouse for reference samples of biological materials.³³⁶ Similar to the biological deposit system that serves to remedy the fact that newly isolated or created biological material cannot always be reproduced from a written description, AI-generated output produced by AI-based tools may not be reproducible. For example, a deposit of AI training data or an AI-based tool could be provided in written form and furnished to the USPTO³³⁷ as a deposit akin to that of nucleotide or amino acid sequences.³³⁸ Just as the biological deposit requirement can provide a solution to the enablement problem of biological deposits and sequence listings,³³⁹ a deposit of AI training data or an AI-based tool to the USPTO could more appropriately fulfill patent disclosure.

Such disclosure via a deposit would entail a description of how or why an AI-based tool would generate certain AI-generated output. Similar to the multi-tiered system to protect plant varieties³⁴⁰ and biological deposits and sequence listings,³⁴¹ data deposits for AI would provide an alternative form of protection that should be enacted into a statute, so as to enable the disclosure of a different variety of AI that may

335. See generally Daniel J. Knauss et al., *Protecting Plants Inventions*, 11 LANDSLIDE 6, July/Aug. 2019 (pointing out that it is “possible for innovators to obtain multiple formal protections for their new varieties, provided they meet the statutory requirements for each form,” when the applicant “deposit[s] seed or other propagation material capable of producing the claimed plants . . . as part of the applicant’s disclosure”).

336. See David J. Weitz, *The Biological Deposit Requirement: A Means of Assuring Adequate Disclosure*, 8 HIGH TECH L.J. 275, 280, 281 (1993).

337. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2431.

338. See *id.* § 2424.

339. See Eileen M. Woo, *Enabling Life*, 91 N.Y.U. L. REV. 1060, 1076 (2016).

340. See generally CHRISTIE M. HAYES & GREGORY RILEY, A GUIDE TO PLANT PATENTS FOR PENNSYLVANIA’S GROWERS, PENN. ST. DICK. AGRIC. L. RESOURCE & REFERENCE CTR. (2002) (describing a three-tiered U.S. system to protect plant varieties via the Plant Protection Act, the Plant Variety Protection Act, and utility patents).

341. See MPEP (9th ed. Rev. 8, Jan. 2018) § 2403.

not be promoted by disclosure incentives. A trade-off would be an increased cost of obtaining patents, which would shift the cost-benefit balance of obtaining a patent. There could be greater *ex ante* efforts by the inventor and greater *ex post* administrative examination burdens for the USPTO to carry out these disclosure mechanisms. This increased cost would also decrease the cost borne by other inventors researching a similar application or field by minimizing duplicated efforts. However, a societal loss exists when the public cannot practice the AI invention absent a deposit.

Policymakers should consider incentivizing the greater disclosure of AI-based tools and AI-generated output as a means of preventing the socially harmful tragedy of the anticommons.³⁴² Society would be negatively impacted by underuse and waste caused by unclear rights in upstream training data³⁴³ and algorithms from insufficient patent disclosure. As a result of not disclosing and demarcating rights to use of training data and algorithms when producing AI-generated output, there could be negative impact to downstream innovation.³⁴⁴

C. Future Directions

This Article broadens the perspective of contemporary inventions and observes that greater attention should be given to the role of AI in shaping patentability and patent examination through enhanced disclosure. As noted, the traditional view in patent law has focused on the role of the human in the invention process to provide a claimed invention based on human-generated output. This Article has introduced a valuable and challenging line of inquiry, which has explored the effects of both AI-based tools and AI-generated output on patentability and patent examination. As this Article has shown, AI-generated output can challenge the adequacy of patent disclosure and can have a significant impact on USPTO patent examination by making it more difficult to distinguish AI-generated output from claimed human-generated output. While some scholars have richly pursued the doctrinal patentability challenges of AI in the invention process,³⁴⁵ more attention to the structural implications of disclosure by inventors and detection of disclosure by the USPTO was warranted.

342. See Heller & Eisenberg, *supra* note 57, at 698.

343. See Hacker, *supra* note 58 (stating that training data is of fundamental importance for the development of AI applications).

344. See Jorge L. Contreras, *The Anticommons at Twenty: Concerns for Research Continue*, 361 SCIENCE, July 2018, at 335, 335–37 (suggesting that underutilization of resources would add to the cost and slowing of the pace of downstream innovation); see also *supra* note 59 and accompanying text.

345. See Ramalho, *supra* note 10, at 2.

Along these lines, this Article's theoretical contributions define a framework for verifying underpinnings concerning AI patent disclosure and provide further examination of related topics. This Article has relied on technological taxonomy,³⁴⁶ theory, and normative assessments³⁴⁷ to provide prescriptions for enhanced AI patent disclosure.³⁴⁸ While there are benefits of enhanced AI patent disclosure to an inventor, such as protection against patent infringement,³⁴⁹ a natural inquiry is whether an AI inventor will opt for trade secrecy or patent protection for AI inventions and, if so, for which aspects. In other words, this Article calls for further empirical examination to elucidate and quantify the AI intellectual-property-protection strategic decision via patents or trade secrets.

Scholars have determined that the strategic decision to pursue patents, trade secrets, or a combination thereof during the innovative process is highly contextual.³⁵⁰ Empirical studies that compare the trade-offs and use of patents versus trade secrets are scant.³⁵¹ Subsequent research in AI invention can supplement the existing literature with examples and conclusions drawn from a series of evidence-based insights drawn from semi-structured interviews with professionals (technologists, managers, business strategy leaders of multinationals, investors, technology transfer leaders, licensing professionals, attorneys at law firms, and in-house counsel involved with emerging informational technologies) in AI-dominant industries. Notably, such future research can provide a more thorough account of the debate among theorists and conflicting results in a few empirical studies as to whether patents and trade secrets act as economic substitutes or complements, specifically for AI technologies.

By understanding what aspects of AI technologies inventors consider unavailable for patent protection (aspects such as being too abstract or obvious, having a perceived lack of patentability, or being too weak or difficult to enforce), the USPTO, courts, scholars, and society

346. *See supra* Part I.

347. *See supra* Section II.C.

348. *See supra* Section III.B.

349. *See supra* Section III.A.2.

350. *See* W. Nicholson Price II, *Regulating Secrecy*, 91 WASH. L. REV. 1769, 1769, 1771, 1777–79 (2016); Michael Risch, *Why Do We Have Trade Secrets?*, 11 MARQ. INTELL. PROP. L. REV. 1, 12–13 (2007).

351. *See* David S. Levine & Ted Sichelman, *Why Do Startups Use Trade Secrets?*, 94 NOTRE DAME L. REV. 751, 754–55 (2018) (explaining that much of the data and questions concerning trade secrets has been underexplored, including “the economic relationship between patents and trade secrets, the role of trade secrets in promoting first-mover advantage, and the use of trade secrets as strategic intellectual property (IP) assets,” so much so, that a leading trade secrets researcher has called for scholars to pursue a new research agenda concerning empirical studies of trade secrets law).

can have a barometer of how well the patent system is functioning, and firms can better understand the level of necessary collaboration in AI-dominant industries. Such a future empirical project could explore narratives drawn from interviews conducted with AI inventors to conclude consistencies and inconsistencies with theoretical understandings of the patent-versus-trade-secret decision, in relation to firm-level strategy about controlling information leakage, engaging in research and development, and promoting alliances.³⁵² By identifying the perspective of AI inventors, the focused and qualitative interviews would uncover rich, descriptive views to provide broader insights into patent policy, firm-level views of trade secrets in relation to the boundaries of the firm, and the economic relationship between patent and trade secrets for AI technologies. While it is important to understand the complex ways in which patents and trade secrets impact business, it is also important to contextualize these effects within the broader ways that the USPTO and society can shape AI innovation.³⁵³

CONCLUSION

This Article has shed new light on the longstanding debate on the effects of enhanced patent disclosure, specific to AI inventions. A skeptical viewpoint toward disclosure considers patent disclosure as not disclosing useful information, as being inconsistent with the patent system's normative goals, and as being distinct from the underlying technology. Conversely, an influential body of scholars has advocated for the positive benefits of patent disclosure, including technical

352. See Jonathan Barnett & Ted M. Sichelman, *The Case for Noncompetes*, 87 U. CHI. L. REV. 953, 986–87 (2020) (noting that “patents . . . protect against knowledge leakage resulting from employee movement . . . [and] after consummation of [an] acquisition”); Orly Lobel, *Non-Competes, Human Capital Policy & Regional Competition*, J. CORP. L. (forthcoming), <https://bit.ly/2ZW2Afw> (arguing that a policy that protects the ability of employees to move freely between competitors supports the goals of economic development, and in so doing, noting patenting rates in relation to high-mobility regions and correlations with innovation).

353. See Remarks by Director Iancu at the AI Event, *supra* note 3. The USPTO Director stated:

America's national security and economic prosperity depend on the United States' ability to maintain a leadership role in AI and other emerging technologies. Additionally, the National Security Strategy calls for the U.S. to prioritize critical emerging technologies and particularly emphasizes AI as a field that is advancing rapidly. On a global stage, the U.S. is working with the other G5 members to “advance a shared understanding of how to best seize the opportunities presented by AI.” From countries as small as Singapore to ones as large as China, nations around the world have become extremely competitive in the innovation ecosphere and especially assertive in emerging technologies like AI. Only by innovating faster and in key areas will the U.S. be able to remain competitive.

Id.

information and knowledge spillovers, demarcation of rights and prevention of overclaiming, and strategic business benefits. This Article provided a taxonomy of AI and presented challenges with explainability and inscrutability to argue that the proponents of enhanced AI patent disclosure are correct. It reasons that the application of AI-based tools produces AI-generated output in the form of unidentified inventions. Indeed, it observed that because AI-generated output may be examined to meet patentability by the USPTO, expectations that AI-generated output is worthy of patent protection are entirely foreseeable. This finding provides law and policymakers with a more robust understanding of the nuanced ways that AI strains conceptual foundations and gaps and the unnecessary overlap between patent and trade secret protection. While it is important to recognize the *ex ante* costs and efforts of enhanced AI patent disclosure, the increasing blurring of artificial and real inventions suggests that disclosure incentives for the inventor and patent protection from data deposits for AI will provide positive welfare gains for society.