

Undressing AI: Transparency Through Patents

Henry H. Perritt, Jr.*

Abstract

Pressures for more flexible acceptance of patent applications involving computer inventions and pressures to require greater transparency of generative artificial intelligence (“AI”) products may represent a perfect storm for patent eligibility clarification and for AI transparency. The revolution in artificial-intelligence technology has led to an upsurge in applications for patents covering AI inventions. Whether these inventions are eligible for patent protection intensifies a vigorous debate that has surrounded two Supreme Court decisions in the last ten years that erect barriers to statutory eligibility for patents. These two cases, known collectively as “Alice/Mayo,” superimpose judicial exceptions disqualifying subject matter that literally qualifies under the language of the patent statutes. The controversy, even before AI’s arrival on the scene, pits innovators seeking broader intellectual property protection over computer software and biotechnology against those who claim that aggressive intellectual property rights and enforcement stifles rather than stimulates innovation.

The unusual characteristics of generative AI technology mean that applications for genuine inventions in the area can break down some barriers that have artificially circumscribed patent eligibility. Properly described, these characteristics, combined with commentator suggestions about how Alice/Mayo should be understood and applied, can permit patent law to adapt appropriately, even without legislative reform.

One of the loudest calls for regulating AI demands transparency. Transparency is a requirement for obtaining a patent. Developers have been grudging in their response because of the need to protect trade secrets. Greater reliance on patents to protect AI innovations will meet the proprietary needs of AI developers to protect

* Professor of Law Emeritus, and former dean, Chicago-Kent College of Law, Illinois Institute of Technology. Author of 25 books and more than 100 law review articles on labor and employment law, administrative law, law and technology, and international relations, including *Trade Secrets for the Practitioner* (3d ed 2024). Democratic nominee for the U.S. House of Representatives, Illinois 10th District, 2002. Former member, National Academy of Sciences Computer Science and Telecommunications Board. Member of the bar: Virginia, Pennsylvania (inactive), District of Columbia, Maryland, Illinois (retired), USPTO, and the Supreme Court of the United States. Commercial helicopter, private instrument airplane, and drone pilot. Extra-class radio amateur (K9KDF). The author wrote an expert system that evaluated the law of employment terminations during the mid-80s boom in AIA. See Henry H. Perritt, *Artificial Intelligence Techniques for Evaluating Employee Terminations on a Personal Computer* 13 RUTGERS COMPUTER & TECH. L. J. 342 (1987).

against free riding on their innovations, while also responding to regulatory and legislative movements to make AI transparent.

Legislation now pending in Congress to amend the patent act and get rid of the judicial exceptions probably is not going to get enacted, but the clamor to regulate AI may intrude into the intellectual property arena, giving new political impetus to changing perspectives on patents for twenty-first century technology. A perfect storm may allow inventors and their patent lawyers to retire from the battlefield and go back into the lab.

Table of Contents

I.	Introduction	139
II.	The basics.....	143
	A. What is AI?	143
	1. How it works.....	143
	2. Demands for regulation.....	153
	B. What is a patent?	153
III.	Is the disclosure required to obtain a patent sufficient to satisfy the demands for transparency?.....	157
	A. What is AI transparency?.....	157
	1. What needs to be transparent?	159
	2. To what end?.....	162
	3. Who is going to look at it? When?.....	162
	B. What is patent disclosure?.....	163
	C. Convergence between AI transparency and patent disclosure.....	165
IV.	Are AI inventions eligible for patent protection?	167
	A. Evolution of § 101.....	167
	1. Statutory subject matter	167
	2. Judicial exceptions	167
	3. Something more—an “inventive concept”	168
	4. Avoiding preemption	169
	B. Criticism and proposals for reform	170
	C. What the USPTO is doing.....	173
	D. Does the status quo allow patents for AI inventions?.....	180
V.	Will the creators of AI inventions rely on patent law rather than continuing to rely on trade secret law?	182
	A. Economic arguments.....	182
	1. Risk of free riding	183
	2. Patent protection and its detractors	183
	3. Trade secret alternative	185
	4. First-mover advantage	187
	5. The patent bargain.....	187
	6. Is the incentive enough?.....	188
	B. Political context: A perfect storm	189

C. Legal context: Patent protection as substitute compensation for a taking of trade secrets	190
VI. Good policy and good law	191

I. Introduction

Cavendish Molton and Parish Barrow were fraternity brothers in Phi Kappa Theta at Carnegie Mellon University. Cavendish graduated with a bachelor’s degree in electrical engineering and stuck around to complete a master’s degree in computer science. The title of his master’s thesis was “Effective use of transformers from large language models applied to music and images.”

Parish got a bachelor’s degree in electrical engineering and then moved to Washington, where he worked as a patent examiner in the U.S. Patent and Trademark Office (“USPTO”) while he went to law school on a part-time basis at George Washington University. Parish has been admitted to the District of Columbia, Maryland, and Patent Office bars.

Cavendish intensified and expanded relationships he had developed with his fellow graduate students and has been working hard on developing generative AI engines similar to ChatGPT and Google’s Gemini, but with lower cost operation and easier adaptability to specific user application environments.

The post-COVID remote work flexibility has enabled both Cavendish and Parrish to set up shop in a small crabbing and sailing village in Rock Hall, on Maryland’s Eastern Shore, opposite Baltimore. They share a 1940s-era single family house at 5878 Beach Rd, ten yards from the waters of the Chesapeake Bay. Unlike many work-from-home aficionados, who use it as an excuse to shirk from home, Cavendish and Parish work sixteen to twenty hours a day.

Their professional activities intersect only occasionally, but they spend many evenings looking out at the lights of ships coming up and down the channel in the Chesapeake Bay and the lights of Baltimore Harbor. During these sessions they talk about new developments in generative artificial intelligence (“AI”) and in patent law.

Cavendish initially was pre-occupied with the possibility of patenting his AI inventions—it would be neat to get a patent; it would legitimate him as an inventor, but his technology partners—other members of his LLC, which Parish set up for him—pooh-poohed the value of patents. They cost a fortune to obtain, take forever to get approved, and cover obsolete technologies by the time they issue, he was told. Rely on trade secret protection, they said. Make sure everyone you talk to signs a non-disclosure agreement. They have met Parish on many occasions and urge Cavendish to redirect Parish from patent prosecution to trade-secret protection.

“It may not do you any good,” Parish said when Cavendish told him about his investors’ urgings.

“What?”

“Trade secret protection.”

“Then why—”

“All the anti-technology folks are clamoring for regulation of AI, and their favorite form of regulation is a requirement for ‘transparency’: disclosure of the innards of your deep machine learning techniques, learning database, and algorithms.”

“All that stuff is exactly my invention, what I was going to seek a patent on, and now keep as a trade secret.”

“I’m not sure we could get a patent, anyway. If it’s simply computerizing an algorithm that could be worked by hand, it’s not patentable. If it’s a computer program that makes use of signals and electrical values from sensors, or which drives mechanical apparatus, it’s very likely to be patentable.”

“So it must be linked to some kind of physical machine? We have to extend our invention to the operation of a drone?”

“That was the position of the patent office for a time, but that’s not necessary anymore, if the applicant for a patent can show that the invention does something genuinely ‘more’ than merely execute an algorithm.”

“How is ‘more’ defined?”

Parish laughed so hard, he almost upset the table. “Trying to come up with a definition has kept thousands of patent examiners, patent lawyers, and judges tied up in knots for a decade—ever since *Alice*.”

“Who is *Alice*?”

“*Alice* is the name of a Supreme Court decision that circumscribes the subject matter eligible for patents, basically disqualifying algorithms. It’s been a huge problem for inventors of novel computer programs.”

“So I wouldn’t be able to get a patent, anyway.”

“You might be. And the trade secret protection that your investors are so fond of may evaporate under the pressure to regulate AI by making it more ‘transparent.’”

“In other words, the law has abandoned inventors.”

“Not necessarily. A patent has to disclose the details of an invention, but it protects the inventor from people copying his invention. The required disclosure is likely to meet the demands for AI transparency.”

“So, while my investors are urging me to dress up my inventions in elaborate clothing and armor, you are suggesting that it do precisely the opposite: that I use the patent system to undress my inventions.”

“Exactly.”

“I should have gone into banking,” Cavendish said.

Patent disclosure provides a way of responding to calls for transparency in AI

systems, as the story suggests. The most frequently advanced proposal for regulating AI involves requiring generative AI systems to be transparent.¹ A transparency requirement might mandate detailed disclosure of the machine-learning techniques used, for example generative adversarial networks (“GANs”) for images, or recurrent neural networks (“RNNs”)² or transformers for language systems; it might require disclosure of the algorithms developed from such machine learning, although it is not clear that algorithms expressed in the form of tensors with thousands of dimensions would be understandable by human observers; it might mean a post-hoc assessment of results obtained from feeding standard queries, combined with some objective standard of acceptability of results.

Developers of generative AI systems have been slow to commit to transparency, concerned that disclosure of the details of their machine learning systems and their resulting algorithms would vitiate their trade secret protection, which afford a competitive advantage.³

Transparency also is a requirement of patent law generally.⁴ Greater transparency increases the likelihood of patent eligibility for inventions, such as those involving AI, which may involve algorithms.⁵ The patentability decisions involving AI consistently reject applications that are opaque about exactly how machine learning is applied in the invention, and the granted patents consistently show a higher degree of transparency with respect to those techniques.⁶

AI developers can respond to the calls for transparency by applying for patents and affording significant transparency in their patent applications, thus accommodating the chief goal of the pro-regulation AI forces, while not giving up their intellectual property.

The pressure for a more intelligent approach to granting software patents and the pressure to regulate AI may represent a perfect storm, allowing both types of demand to be met through the patent system. For that to happen, patent law must crystallize its emerging trends to allow software patents based on detailed disclosure, and AI inventors must recognize that politics is weakening the reed of trade secrets, forcing them to seek patent protection for their innovative efforts. If the USPTO opens the door to well-crafted AI patents, innovators can be protected and AI transparency can be enhanced.

¹ See Naser Bakhshi, *A Call for Transparency and Responsibility in Artificial Intelligence*, DELOITTE (Sept. 14, 2023), <https://www2.deloitte.com/nl/nl/pages/innovatie/artikelen/a-call-for-transparency-and-responsibility-in-artificial-intelligence.html> (calling for transparency in AI).

² See *infra* Section I for an explanation of GANs, RNNs, and transformers.

³ See Clark D. Asay, *Artificial Stupidity*, 61 WM. & MARY L. REV. 1187, 1222 (2020) (noting that risks of trade secret disclosure push many AI inventors to prefer patents).

⁴ 35 U.S.C. § 112 (requiring “written description of the invention, and of the manner and process of making and using it, in . . . full, clear, concise, and exact terms.”).

⁵ See *infra* Section IV.A.2 explaining and analyzing patent-eligibility under 35 U.S.C. § 101, *Alice Corp. v. CLS Bank Int’l*, 573 U.S. 208, 216 (2014), and *Mayo Collaborative Servs. v. Prometheus Lab’ys, Inc.*, 566 U.S. 66, 77–79 (2012).

⁶ See *infra* Section III (analyzing decisions on patent applications involving AI).

Patents are not a vehicle for implementing the full range of proposals for regulating AI. But patents do provide a mechanism for realizing one of the most challenging of the regulatory proposals: requiring transparency of AI data and methods.

Assessing the validity of this idea requires answering three questions:

- Are AI inventions eligible for patent protection?
- Would the disclosure required in patent applications and patents meet the demands for transparency?
- Are the political and economic incentives such that AI inventors will seek patent protection instead of relying on trade secret protection?

Concrete proposals for requiring transparency mainly relate to transparency of the fact of use and transparency with respect to results. Few, so far, focus on transparency with respect to how the systems work. Nevertheless, vagueness in the meaning of transparency, transparency's pervasiveness as a proposal in the debate over AI regulation, and ignorance about how generative AI systems actually work may well transition into stronger focus on how AI works. It is the possibility of such a shift in regulatory attention that motivates this article.

Following this introduction, Part II provides background on patent law and on AI. It reviews statutory standards for patent eligibility, novelty, nonobviousness, and disclosure. It explains what AI is: use of huge advances in computing power and digital storage to decades old statistical analysis techniques.

Part III considers whether the disclosure required in patent applications and patents meets the demands for transparency. It analyzes calls for AI systems to be more transparent, explores different types of transparency and compares them with the types of disclosure that have been found adequate to meet the requirements of patent law.

Part IV considers whether AI inventions qualify for patent protection, summarizing the debate over Supreme Court interpretations of § 101 in the *Alice Corp. v. CLS Bank Int'l* ("Alice") and *Mayo Collaborative Servs. v. Prometheus Labs., Inc.* ("Mayo") decisions, and exploring what the patent office actually is doing with hundreds of applications for patents on AI inventions it already has received.

Part V considers whether AI inventors are likely to rely on patent protection for their inventions, as contrasted with trade-secret protection or first-mover advantage. Necessarily, this part revisits long-standing economic arguments for and against broader patent protection. Assessment of the politics of patent reform, as impacted by widespread calls to regulate AI, supplement the economic analysis. It suggests that AI innovators are likely to face erosion of trade secret protection for their inventions and should consider greater reliance on patent protection instead, resulting in greater transparency of their methods.

Part VI proposes that the USPTO and the courts guide the perfect storm and use

the upsurge in applications for AI patents to make interpretation of § 101 eligibility more coherent and predictable and the § 112 disclosure requirements more rigorous.

The subject of this article is patentability and regulation of AI technology. It does not address AI agents as inventors, or the use of AI by patent attorneys to aid patent prosecution, or the use of AI by patent examiners to evaluate applications. The three subjects not addressed by this article are the focus of an earlier article by this author.⁷

II. The basics

A. What is AI?

No one knows. “Although ‘AI’ is nearly ubiquitous, it has no agreed-upon definition.”⁸

The recent definition in Executive Order 14110⁹ would sweep under the definition Microsoft Word spellcheck, Microsoft Excel arithmetic formulas, Westlaw and Lexis searches, and Google searches.¹⁰ The EU AI Act¹¹ seeks to avoid this result: “[T]he definition should be based on key characteristics of AI systems that distinguish it from simpler traditional software systems or programming approaches and should not cover systems that are based on the rules defined solely by natural persons to automatically execute operations.”¹²

1. How it works

AI is a form of computer programming. Like other programming, it relies on abstraction,¹³ representing real-world objects and behaviors through algorithms. Its

⁷ See generally Henry H. Perritt, Jr., Robot Inventors, Robot Patents, Robot Examiners, and Robot Patent Prosecutors 1–66 (Nov. 27, 2024) (unpublished manuscript) (on file with the South Carolina Law Review).

⁸ Nikola L. Datzov, *The Role Of Patent (In)Eligibility in Promoting Artificial Intelligence Innovation*, 92 UMKC L. REV. 1, 15–16 (2023) (internal quotations omitted).

⁹ Section 3(b) of Exec. Order No. 14110 defines artificial intelligence:

The term “artificial intelligence” or “AI” has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

(c) The term “AI model” means a component of an information system that implements AI technology and uses computational, statistical, or machine-learning techniques to produce outputs from a given set of inputs.

§§ 3(b), (c), Exec. Order No. 14110, 88 Fed. Reg. 75191 (Oct. 30, 2023).

¹⁰ These popular software examples are “machine-based systems” that “make predictions” for “a given set of human-defined objectives.” The use models and computational and statistical techniques to product outputs from a given set of inputs.

¹¹ Regulation 2024/1689 of Apr. 19, 2024, Artificial Intelligence Act, O.J. (L.) [hereinafter Artificial Intelligence Act].

¹² *Id.* at ¶ 12.

¹³ See Jeff Kramer, *Is Abstraction the Key to Computing?*, 50 COMM’N.S OF THE ACM 36 (2007).

hardware “understands” information represented as data organized in structures¹⁴ that associate values with attributes of things and characteristics of behaviors. Its software makes use of iteration¹⁵ to perform repetitive calculations rapidly and infers conclusions by applying algorithms¹⁶ to appropriate data. It accomplishes these things and much more by stringing together simple operations like fetch this value from memory, compare it with that one, and if they are the same, add them and put them somewhere else in memory.¹⁷ These characteristics of computer programming had crystallized reasonably well by the late 1930s and were the principal tools of computer programmers by the time the first digital computers became useful during and after World War II.¹⁸ Computer science theory developed around these practices in the 1950s and 1960s.¹⁹

Digital computers sped up the computation of statistics according to formulas developed in the nineteenth century, and facilitated tagging, sorting, and comparing information artifacts—operations at the heart of ChatGPT.²⁰

The term “artificial intelligence” first became popular right after World War II when the popular press fueled excitement over the new “electronic brains.”²¹ British cryptographer Alan Turing popularized the still-used test for AI: whether a human being interacting through an opaque communication mechanism can tell whether what is on the other end is a human being or a machine.²²

¹⁴ See generally, NIKLAUS WIRTH, *ALGORITHMS + DATA STRUCTURES = PROGRAMS* (1976) (explaining that all computer programming comprises a combination of the two).

¹⁵ *What are the Basic Fundamental Concepts of Programming?*, EDUCATIVE, <https://www.educative.io/answers/what-are-the-basic-fundamental-concepts-of-programming> (last visited April 29, 2024) (identifying iteration as one of the fundamental concepts).

¹⁶ *What Is An Inference Engine in Machine Learning?*, RUN: AI, <https://www.run.ai/guides/machine-learning-inference/inference-engine> (last visited April 29, 2024) (“Reasoning algorithms are the logic that the inference engine uses to analyze the data and make decisions. The algorithms take the data from the knowledge base and apply logical rules to it to infer new information.”).

¹⁷ Phillip M. Sherman, *Programming and Coding the IBM 709-7090-7094 Computers 105–06* (1963).

¹⁸ See WORLD INTELLECTUAL PROPERTY ASSOCIATION (WIPO), *PATENT LANDSCAPE REPORT: GENERATIVE ARTIFICIAL INTELLIGENCE 12* (2024) [hereinafter “2024 WIPO REPORT”] (relating history of AI, beginning in 1930s); here is some programming code from that early era: “Bi x 00010,” transfer block of characters from input interface to core memory; “ck x 01011” compare contents of accumulator with memory location. *A Whirlwind Instruction Set*, O’REILLY https://www.oreilly.com/library/view/anfsq-7-the-computer/9783486990911/e9783486727661_bm02.html (last visited April 29, 2024). Whirlwind was one of the first digital computers, developed by Jay Forrester at MIT; *Whirlwind: Preparing the Way for SAGE*, COMPUTER HISTORY MUSEUM, <https://www.computerhistory.org/revolution/real-time-computing/6/123> (last visited April 29, 2024).

¹⁹ Jeffrey Shallit, *A Very Brief History of Computer Science* (1995), <https://cs.uwaterloo.ca/~shallit/Courses/134/history.html> (noting developments in computers in the 1940s to 1950s, and first expressions of computer science in the 1960s).

²⁰ See 2024 WIPO REPORT, *supra* note 18, at 13 (identifying more powerful computers, big data, and better machine-learning algorithms as drivers of 21st century rediscovery of AI).

²¹ See Hannah Gresham, *The Mechanical Monster and Discourses of Fear and Fascination in the Early History of the Computer*, 7 HUMANITIES & SOC. SCI. COMM. 1, 5–6 (2020) (describing the reaction to “electronic brain” represented by Whirlwind, Erica, and Univac).

²² See Melanie Mitchell, *The Turing Test and our Shifting Conceptions of Intelligence*, 385 SCIENCE

AI enjoyed another upsurge in popularity in the 1980s. Experts in various fields trained computers to act as *expert systems*,²³ working with knowledge engineers, specialists skilled in human processing of information, to attempt to get computers to mimic them.²⁴

AI applications can be distinguished between those that use *discriminative* models and those that use *generative* models.²⁵ Face recognition applications use discriminative models: they predict whether a new face matches another face in a database. Generative models extend learning from past data to predict new data.²⁶ A multiple regression equation used to predict next year's GDP from past economic variables is an example of a generative model. Both types of models may deal with *sequential* data, like text or music, or with *three-dimensional shapes* portraying objects, like photographs, or crime scenes.²⁷

AI models enabling systems to match or predict are developed by analyzing probabilities with respect to learning data. A natural language program takes each character and develops an estimate of the probability that it will be followed by each other character. It takes each word in the vocabulary and predicts how likely it is to be followed by each other word in the vocabulary. And then it can do the same thing at a conceptual or semantic level, predicting how likely sentence about patent law is to be followed by a sentence about toilet training for boxer dogs. It is a bottom-up learning process: first distinguishing individual words, then evaluating the frequency with which words appear together in phrases, then associating the phrases with concepts, and then building a hierarchy of concepts, creating a semantic tree,²⁸ not

(2024), <https://www.science.org/doi/10.1126/science.adq9356> (describing the Turing test).

²³ Intelligent Question-Answering Method and Device Based on Large Language Model, China Patent Application No. 117520491A, (unknown trans.), (filed Oct. 27, 2023) (published Feb. 6, 2024) (Background section):

The background of intelligent question-answering systems can be traced back to early expert systems and information retrieval techniques. Expert systems are artificial intelligence systems based on rules and a knowledge base that answer the user's questions by reasoning and matching rules. The information retrieval technology is to find out the answers related to the user questions by indexing and retrieving the document library. With the rapid development of natural language processing and machine learning, intelligent question-answering systems are gradually moving from rule-based and retrieval-based approaches to statistical and machine learning-based approaches. The method utilizes large-scale corpus data for training, and improves the accuracy and naturalness of the system through pattern recognition and probability calculation.

²⁴ See Henry H. Perritt, Jr., *Artificial Intelligence Techniques for Evaluating Employee Terminations on a Personal Computer*, 13 RUTGERS COMPUTER & TECH. L. J. 342 (1987) (describing development and use of expert system for employment law). See generally WIPO, TECHNOLOGY TRENDS 2019—ARTIFICIAL INTELLIGENCE 19 (identifying periods of periodic excitement about AI: 1956, 1974–1980, 1980–1987, 1993–2011, punctuated by first AI winter in 1974–1980 and second AI winter in 1987–1993).

²⁵ DAVID FOSTER, GENERATIVE DEEP LEARNING: TEACHING MACHINES TO PAINT, WRITE, COMPOSE, AND PLAY 5 (2d ed. 2023) (explaining difference between generative and discriminative modeling); 2024 WIPO REPORT, *supra* note 18, at 20 (noting difference between discriminative and generative tasks).

²⁶ FOSTER, *supra* note 25, at 11 (explaining role of data-generating distribution in generative modeling).

²⁷ 2024 WIPO REPORT, *supra* note 18, at 25 (summarizing music applications).

²⁸ See Jesús Giménez & Lluís Márquez, *Linguistic Measures for Automatic Machine Translation*

unlike a conventional thesaurus.²⁹

A deep learning system does this again and again, using the result of the model at each stage to predict the next token, measuring the error in an error function, and then doing everything all over again to find the set of weights in the neural network to minimize the error function, sometimes called a loss function. The error function reflects how well the network does in predicting. Sophisticated networks feed back their results to compare them with actual data, permitting error functions to be computed and then improved.³⁰

This is the ultimate in brute force analysis. The combinatorial possibilities are enormous: twenty-six squared for character analysis in English, 470,000 squared, or 220 billion, for word analysis in English³¹ and at least the square of the number of concepts (11,000 squared, or 121 million) in a thesaurus for higher level modeling.³² The computational demands are even greater than these numbers suggest, considering the iterative nature of the necessary process, working upward in levels of abstraction to refine next-word and next-character predictions and then moving to the next character and the next word and doing it all over again.

Most of the development of new statistical models and analytical techniques in the machine learning field have been devoted to identifying new types of representation—new models—that can simplify what otherwise would exceed computational capability.

Representation is a core concept in AI, as it is in human knowledge and communication. A human being can be represented by a face, and the person's face can be represented by a collection of pixels comprising a digital image. A physical phenomenon such as bright sunlight can be represented by a sentence such as, "The sun is shining."

Ever since the Greek mathematician Euclid first wrote about geometric algebra,³³ real-world phenomena have been represented by numeric values, and their

Evaluation, 24 MACH. TRANS. 209, 232, 236 (2010) (noting use of semantic trees in machine translation of languages); Hyejin Youn, *On the Universal Structure of Human Lexical Semantics*, 113 PROC. NAT'L ACAD. SCI. U.S. AM. 1766, 1769 (2016); Jean-Pierre Koenig & Anthony R. Davis, *The KEY to Lexical Semantic Representations*, 42 J. LINGUISTICS 71, 73 (2006).

²⁹ Warren Sack, *Conversation Map: An Interface for Very Large-Scale Conversations*, 17 J. MGT INFO. SYS. 73, 79–80 (Winter, 2000/2001) (explaining analysis of email messages and usenet posts with reference to semantic trees and thesauri).

³⁰ See Jason Brownlee, *Loss and Loss Functions for Training Deep Learning Neural Networks*, MACH. LEARNING MASTERY (Oct. 23, 2019), <https://machinelearningmastery.com/loss-and-loss-functions-for-training-deep-learning-neural-networks/>.

³¹ *How Many Words are there in English?*, MERRIAM-WEBSTER, <https://www.merriam-webster.com/help/faq-how-many-english-words> (last visited July 21, 2024) (estimating one million words in English; noting 470,000 in dictionary).

³² AMAZON, <https://www.amazon.com/Rogets-Desk-Thesaurus-Random-House/dp/0375425691> (last visited April 29, 2024) (describing Roget's thesaurus).

³³ J.J. O'Connor & E.F. Robertson, *Euclid of Alexandria*, MAC TUTOR (last updated Jan. 1999), <https://mathshistory.st-andrews.ac.uk/Biographies/Euclid/> (last visited July 21, 2024) (reviewing Euclid's work).

relationships by mathematical connections among those values. With the advent of digital computers during World War II, representation extended into identifying almost anything by a set of binary values—a collection of bits in millions of computer memory units. With the introduction of word processing in the 1970s,³⁴ binary digitization combined with natural language tokens such as characters, words, sentences, and concepts.

Modern AI is defined by its use of *machine learning* to infer means of representation rather than relying on human tagging of characteristics—learning from examples rather than rules.³⁵ Machine learning can determine, for example, what pattern of pixels represents a human face, as opposed to a cement block. There is nothing magical about the analysis: it is factor analysis,³⁶ which has been used as a social science methodology for more than 100 years.³⁷

Most sophisticated machine learning systems use binary neural networks to represent data and to learn relationships. In binary neural networks, each node is connected to every other node. Links between each pair of nodes have weights from zero to one. For example, a binary neural network for language modeling at the character level would have a node for the character, “A,” and links to every other letter in the alphabet. It would have a node for the character “B” and similar links to every other character. In English, “A” is more like to be followed by “P” than by “O,” and so the link weight between the node “A” and the node for “P” would be greater than the weight for the link connecting the node for “A” to the node for “O.”

Various levels of abstraction can be used to facilitate statistical analysis and to increase the power of models.³⁸ A face also may be represented by a simple characterization such as black, white, wrinkled, or cherubic. A sunny day also can be represented by a checkmark in a table of meteorological conditions. Data capture and communications can be made more efficient by increasing the level of abstraction of representation, for example, by sending a message in a criminal alert to look for a “white wrinkled face” rather than a “black cherubic” one, instead of sending all of

³⁴ Microcomputer word processing applications such as WordStar, Perfect Writer and XyWrite were preceded by applications on stand-alone typewriters and mini-computers. See Brian Kunde, A Brief History of Word Processing (Through 1986) (December 1986) (Final Paper, Foothill College), <https://web.stanford.edu/~bkunde/fb-press/articles/wdprhist.html> (recounting IBM magnetic tape and magnetic card selectric typewriters, Lexitron, and Vydec, from 1969 to 1973); *Vydec 1800*, R.I. COMPUT. MUSEUM, <https://www.ricomputermuseum.org/collections-gallery/equipment/vydec-1800> (last visited on July 21, 2024) (reporting on Exxon’s Vydec and Wang); *Wang Laboratories: From Success to Success to . . .*, COMPUT. HIST. MUSEUM, <https://www.computerhistory.org/revolution/minicomputers/11/364> (last visited on July 21, 2024).

³⁵ See 2024 WIPO REPORT, *supra* note 18, at 14 (heart of machine learning is learning with examples rather than with rules).

³⁶ See *Factor Analysis*, STATISTICS SOLUTIONS, <https://www.statisticssolutions.com/factor-analysis-sem-factor-analysis/> (last visited on July 21, 2024).

³⁷ Charles Spearman, Demonstration of Formulæ for True Measurement of Correlation, 18 AM. J. PSYCH. 161, 161–69 (1907).

³⁸ FOSTER, *supra* note 25, at 12–13 (role of representation-abstractions to reduce number of data points).

the pixels necessary for a complete image.

Different statistical equations and procedures for applying them in neural networks are defined by *models*. David Foster identifies six basic types of models for machine learning: *variational encoders*, useful for photographic image modeling; *energy-based models*; *diffusion models*; *normalizing models*; *generative adversarial networks*; and *auto regressive models*, including *Recursive Neural Networks* (useful for modeling text).³⁹ These models are not mutually exclusive, and they are usually combined to some degree in practical generative AI applications.⁴⁰ The predicting is somewhat different depending on whether elements in a sequence, such as a sequence of words or a sequence of musical notes, is being modeled and/or whether a three dimensional image or pattern is being modeled or generated.

Transformers, often considered to be an extension of Recursive Neural Network models, vastly improve efficiency by simplifying representation and by facilitating parallel processing so that the computer can think about many different things at the same time as it learns and develops link weights and refines error functions.⁴¹

Before transformers, the state of the art in natural language processing was a *recurrent* neural network (“RNN”), which processed each word in a sentence sequentially.⁴² Transformers take entire sentences and “remember” the relationship of each word with all the others. Transformers allow models to take context beyond sentence boundaries.⁴³ Transformers are associated with developing models for

³⁹ *Id.* at 19 (identifying six basic types of model); See 2024 WIPO Report, *supra* note 18, at 10 (reporting that the most important generative AI models are generative adversarial networks (“GANs”) for images, variational autoencoders (“VAEs”), and decoder-based large language models (“LLM’s”) for text, speech, and music, with the strongest growth in GAN and LLMs). See *id.* at 22–23 (providing brief explanation of GAN, LLM and VAEs).

⁴⁰ FOSTER, *supra* note 25, at 19.

⁴¹ *Id.* at 235 (describing transformers); The most sophisticated machine learning systems employ transformers—the “T” in ChatGPT. Transformers are a type of neural network architecture that “remembers” what came before, maybe long before, in a sequence of words, ideas, sounds, or images. They enable a better and more efficient prediction of “what comes next” than the common alternative, recurrent neural networks. Transformers use convolutional neural networks combined with attention models. See Giuliano Giacaglia, *How Transformers Work*, TOWARDS DATA SCIENCE (Mar. 10, 2019), <https://towardsdatascience.com/transformers-141e32e69591> (explaining transformers and comparing them with recursive neural networks; giving examples of language translation systems); Rick Merritt, *What Is a Transformer Model?*, NVIDIA (Mar. 22, 2022), <https://blogs.nvidia.com/blog/2022/03/25/what-is-a-transformer-model/> (explaining how transformers “pay attention”) (“A transformer model is a neural network that learns context and thus meaning by tracking relationships in sequential data like the words in this sentence.”).

⁴² See Jakob Uszkoreit, *Transformer: A Novel Neural Network Architecture for Language Understanding*, GOOGLE (Aug. 31, 2027), <https://research.google/blog/transformer-a-novel-neural-network-architecture-for-language-understanding/> (explaining how transformers work, compared to RNNs; summarizing work of eight Google AI researchers); see Giacaglia, *supra* note 41; Merritt, *supra* note 41.

⁴³ Madhumita Murgia, *Generative AI Exists Because of the Transformer*, FINANCIAL TIMES (Sept. 12, 2023), <https://ig.ft.com/generative-ai/> (step-by-step, animated, explanation and demonstration of how transformers work in large-language machine processing).

sequential expression, like text and music.⁴⁴ They are, however, also useful for three-dimensional image generation as well.

These techniques, though only iterative improvements in centuries-old mathematical concepts and decades-old digital computation methods,⁴⁵ have engendered great public excitement after an extremely successful public relations and marketing campaign by Sam Altman,⁴⁶ the CEO of OpenAI, an enterprise funded largely by Microsoft.⁴⁷ The natural language processing capability of ChatGPT,⁴⁸ a publicly released, free generative AI system, astounded reporters and others with its fluency in understanding and responding to natural language queries and its ability to generate relevant responses from machine learning applied to the enormous store of information available through the Internet.⁴⁹

The most visible of the recent advances involve pattern matching applied by cheap and powerful new semiconductor chips to the huge stock of information available through the Internet. *Machine learning in large language models*⁵⁰ enables the systems to develop algorithms that predict what comes next in information artifacts. Generative image systems recreate movement by predicting where an object will be from one frame of a video recording to the next.⁵¹

⁴⁴ FOSTER, *supra* note 25, at 297 (music is sequential, like text; both use transformers).

⁴⁵ *Ex parte* Daniel J. Ferranti, No. 2022-002794, 2023 WL 9061302, at *10 (P.T.A.B. Dec. 28, 2023) (rejecting a claim limitation of “natural language processing (NLP) algorithms” because it “well-known, conventional, and routine,” going back into the 1950s).

⁴⁶ Samantha Kelly, *Sam Altman Warns AI Could Kill us All. But He Still Wants the World to Use it*, CNN (Oct. 31, 2023), <https://www.cnn.com/2023/10/31/tech/sam-altman-ai-risk-taker/index.html>.

⁴⁷ Jordan Novet, *Microsoft’s \$13 billion Bet on OpenAI Carries Huge Potential Along with Plenty of Uncertainty*, CNBC (Apr. 9, 2023), <https://www.cnbc.com/2023/04/08/microsofts-complex-bet-on-openai-brings-potential-and-uncertainty.html>.

⁴⁸ CHATGPT, <https://chat.openai.com/> (last visited Oct. 9, 2024).

⁴⁹ See Kevin Roose, *The Brilliance and Weirdness of ChatGPT*, N.Y. TIMES, Dec. 9, 2022, at B1.

Perritt: “Write one paragraph on whether patent transparency can meet the demands for ai transparency.”

ChatGPT:

Patent disclosure can enhance AI transparency by providing detailed and publicly accessible information about the technical specifics of AI systems. However, it falls short of fully meeting the demands for AI transparency due to several limitations. The complexity of AI technologies can make patent documents difficult for non-experts to understand, and companies may withhold critical information as trade secrets. Additionally, the rapid evolution of AI can render patents outdated quickly, and strategic patenting practices can obscure rather than clarify the technology. Therefore, while helpful, patent disclosure alone cannot fully satisfy the comprehensive transparency needed for AI systems.

ChatGPT 4 session with the author on 22 July 2024 at 9:16 PM.

⁵⁰ *What are Large Language Models (LLM)?*, AMAZON AWS, <https://aws.amazon.com/what-is/large-language-model/> (last visited Apr. 29, 2024) (explaining how neural networks and other machine learning techniques are used to build LLMs).

⁵¹ The disaggregation of movement was pioneered by Walt Disney’s animators in the 1930s. See *Walt Disney: Animation Pioneer*, NATIONAL INVENTORS HALL OF FAME (Oct. 21, 2022), <https://www.invent.org/blog/inventors/walt-disney-multiplane-camera> (describing early animation techniques, involving inking each animation frame by hand). Generative AI has just digitized it and sped it up.

The same predictive techniques can be applied to music: in Western 12-tone tonal music is a G more likely to be followed by a C or a C-sharp?⁵² Is an imperfect cadence or deceptive cadence likely to be followed by a perfect cadence?⁵³ Is a rhythmic pattern likely to be repeated in the next measure?⁵⁴

Similarly, images can be processed by predicting whether the next pixel to the right, scanning horizontally, is likely to be lighter or darker, the same or a different color, and if a different color, which? Or at a higher level, is a lip more likely to be followed by a cheek or an ear?⁵⁵ In image recognition systems, a very large number of image samples are processed by machine learning.⁵⁶ Some of the samples contain the target image, and others contain something else. Thus, a robocowboy might be trained to recognize cattle by presenting hundreds of thousands of images of different kinds of animal, tagging only those that represent cows, bulls, steers, and calves.⁵⁷ Machine-learning techniques can be used to accommodate the challenges associated with recognizing the target image despite different orientations, different lighting conditions, and different backgrounds.⁵⁸

The techniques work by scanning the lines of an image, much as a laser printer or office scanner does, and looking for discontinuities in brightness and color. A model of an image then is constructed identifying the locations of those discontinuities.⁵⁹ Then, statistical techniques, enabled by a complex hierarchy of neural networks⁶⁰ implementing statistical algorithms, compare the location of

⁵² A G is more likely to be followed by a C, because that movement is part of an authentic cadence: V-I. See William S. Rockstro et al., *Cadence*, OXFORD MUSIC ONLINE (2001), <https://www.oxfordmusiconline.com/grovemusic/display/10.1093/gmo/9781561592630.001.0001/omo-9781561592630-e-0000004523>.

⁵³ A cadence in music is a two-chord progression. If a cadence ends with a V in the scale, it is a half cadence, and sounds incomplete, begging for resolution. If a cadence ends with a VII going to the I, it is an authentic cadence, giving a sense of final resolution. *Phrases and Cadences*, MUSICTHEORY.NET, <https://www.musictheory.net/lessons/55> (last visited Oct. 9, 2024).

⁵⁴ See *Songwriting Tutorial: Part Four—Rhythmic Considerations*, MUSICTECH (Jan. 30, 2015), <https://musictech.com/tutorials/songwriting-4-rhythm/> (discussing considerations in writing good song rhythms, including repetition).

⁵⁵ See *generally* Ger. Patent No. 11,2020,001,253 (filed Mar. 9, 2023) (explaining how machine learning techniques are applied to develop image matching systems).

⁵⁶ See U.S. Patent Publication No. 2014/0105467 ¶¶ [0017]–[0026] (filed Dec. 2, 2013) (identifying preceding face recognition patents); Henry H. Perritt, Jr., *Defending Face-Recognition Technology (And Defending Against It)*, 25 FLA. J. TECH. LAW & POL'Y 42, 46 (2020) (explaining how image recognition systems use machine learning).

⁵⁷ See Henry H. Perritt, Jr., *The 21st Century Cowboy: Robots on the Range*, 43 UNIV. ARK. LITTLE ROCK L. REV. 149, 161–69 (2020) (exploring feasibility of robot cowboy who herds cattle; describing machine learned aimed at recognition of cattle).

⁵⁸ This involves the second step in most typologies: *alignment*.

⁵⁹ A line connecting the discontinuities represents an “edge” in the image, such as the edge of a cheek in a human face.

⁶⁰ See Divyanish Dwivedi, *Face Recognition for Beginners*, TOWARDS DATA SCIENCE (Apr. 28, 2018), <https://towardsdatascience.com/face-recognition-for-beginners-a7a9bd5eb5c2> (explaining how neural networks can facilitate use statistical techniques such as Principal Component Analysis, Linear Discriminant Analysis, Independent Component Analysis, Discrete Cosine Transforms, Gabor Filters, and Markov Models for face recognition); Arun Alvappillai & Peter Neal Barrina,

different types of discontinuities between images, and thus identify images that are most similar. The indicia of similarity are the particular facial features that discriminate a cow from a wolf—or one face from another.⁶¹

The foundational layer examines a large inventory of images of faces (the *training database*) to learn what a face is.⁶² This process is much like the one described for the robocowboy learning how to recognize a cow. The statistical model of a face is tagged to identify those features that vary from one face to another, for example: spacing of the eyes, height of the forehead, thickness of the lips, width of the nose, coloration, and so on.

At the conclusion of this layer of processing, the program has a template for a face. It now can take the image of a new face (the *probe face*) and determine the values for each of the facial features that have differentiating effect.⁶³ Most programs work with about 100 reference points that comprise individual features.⁶⁴ Then the computer software compares the faceprint of the probe face to the faceprints in the enrolled database.⁶⁵

Models of semantic relationships in language are developed by interactive comparison of sentences separately deconstructed by neural networks and computing a loss function that represents their similarity or lack thereof.⁶⁶

Face Recognition using Machine Learning (June 16, 2017) (class project, University of California, San Diego), <http://noiselab.ucsd.edu/ECE285/FinalProjects/Group7.pdf> (describing the architecture of face recognition algorithms).

⁶¹ This is the *feature extraction* step. Some commentators call this identifying “landmarks.” See Adam Geitgey, *Machine Learning is Fun! Part 4: Modern Face Recognition with Deep Learning*, MEDIUM (July 24, 2016), <https://medium.com/@ageitgey/machine-learning-is-fun-part-4-modern-face-recognition-with-deep-learning-c3cffc121d78>.

⁶² This is the “detection” step. See Oleksii Kharkovyna, *An Intro to Deep Learning for Face Recognition*, MEDIUM (June 26, 2019), <https://perma.cc/KZ2P-78VK> (identifying detection, alignment, feature extraction, and feature matching to database as the steps in face recognition); See U.S. Patent Publication No. 2014/0105467 ¶¶ [0009]–[0015] (filed Dec. 2, 2013) (describing Detection, Alignment, Normalization, Representation, and Matching steps in pattern matching); Geitgey, *supra* note 61 (describing and illustrating the steps with examples of actual faces of celebrities; providing links to Python programs that execute the algorithms).

⁶³ See, e.g., U.S. Patent No. 5,859,921 (Feb. 28, 1996) (issued Jan. 12, 1999) (describing apparatus that takes a facial image, corrects for lighting differences, and detects eyes).

⁶⁴ There are about 80 nodal points on a human face. A few of the nodal points that are measured by the FACEIT software are: distance between eyes; width of nose; depth of eye sockets; cheekbones; jaw line; and chin. These nodal points are measured to create a numerical code that represents the face in a database. This code is referred to as a faceprint and only fourteen to twenty-two nodal points are necessary for the FACEIT software to complete the recognition process.” U.S. Patent Publication No. 2014/0105467 ¶ [0009] (filed Dec. 2, 2013) (issued Apr. 17, 2014).

⁶⁵ See Scott Jeffrey Klum, *FaceSketchID: A System for Facial Sketch to Mugshot Matching* 14–20 (2014) (Thesis, Michigan State University), http://biometrics.cse.msu.edu/Publications/Thesis/ScottKlum_FaceSketchID_MS14.pdf (describing algorithms and specifications for face-matching program); U.S. Patent Publication No. 2014/0105467 (filed Dec. 2, 2013) (issued Dec. 29, 2015) (disclosing a method and system for matching unknown facial image to image of celebrity by finding, comparing, contrasting, and identifying similar facial characteristics).

⁶⁶ WIPO Patent No. 2018/212584 Claim 1 (filed May 16, 2018).

Natural language user queries can be preprocessed to chunk a set of sentences of the natural language user query into a set of smaller sentences and retaining the references between chunks of the set of sentences. For each chunk of preprocessed user query a Name Entity Recognition (NER) ensemble extracts a domain specific name entity from the chunked preprocessed user query.⁶⁷

Classification techniques model the topics of each chunk. Then, in appropriate applications, sentiment analysis⁶⁸ can determine the sentiment of each chunk of the user query. User queries are converted to system queries by a combination of different natural-language-processing functions, such as NER, sentiment analysis, part of speech tagging, canonicalization,⁶⁹ classification, and translation.

Then, pattern matching techniques are used to compare the system query to the knowledge model to determine the closest state in the knowledge model. The system returns a set of decisions scored according to their degree of match with the system query. A winner state in the knowledge model is the knowledge state with the highest score.⁷⁰

Systems can be fine-tuned by taking large language models trained on the full array of data usually used for training such models and then connecting them with more specialized learning databases. The machine learning system constructs new concept vectors enabling a branching by subject matter before more finely grained responses are constructed.⁷¹

Like all technological innovations, AI builds on foundations established long ago: mathematical and statistical theories, computational methods, computer architectures. Generative AI is a young art, and the press, media, and government excitement over it, however overblown, indicates its potential as further innovation occurs. The enormous amounts of capital invested in it⁷² ensure that lots of smart people will be inventing new processes and systems involving AI technology.

⁶⁷ Named Entity Recognition (“NER”) detects semantic entities such as names, locations, business enterprises, events, products, themes, topics, monetary values, and percentages. Once they are detected, the entities are tagged and linked to an overarching neural network. Nick Barney, *Named entity recognition (NER)*, TECHTARGET (last updated Oct. 2024), <https://www.techtarget.com/whatis/definition/named-entity-recognition-NER>.

⁶⁸ *What is Sentiment Analysis*, AMAZON AWS, <https://aws.amazon.com/what-is/sentiment-analysis/> (last visited April 29, 2024).

⁶⁹ Converting stochastic expression into standardized forms.

⁷⁰ Convolutional State Modeling for Planning Natural Language Conversations, U.S. Patent Publication No. 2020/0387672 ¶¶ [0005]–[0006] (filed Aug. 25, 2020) (published Dec. 10, 2020) (abandoned Sep. 13, 2022).

⁷¹ See Intelligent Question-Answering Method and Device Based on Large Language Model, China Patent Publication No. 117520491A ¶¶ [0006]–[0027], (unknown trans.), (filed Oct. 27, 2023) (published Feb. 6, 2024) (Disclosure of Invention section).

⁷² *AI Investment Forecast to Approach \$200 billion Globally by 2025*, GOLDMAN SACHS (Aug. 1, 2023), <https://www.goldmansachs.com/insights/articles/ai-investment-forecast-to-approach-200-billion-globally-by-2025> (“Innovations in electricity and personal computers unleashed investment booms of as much as 2% of U.S. GDP as the technologies were adopted into the broader economy. Now, investment in artificial intelligence is ramping up quickly and could eventually have an even bigger impact on GDP.”).

2. *Demands for regulation*

Claims of doom because computers have suddenly become intelligent and willful⁷³ are based on ignorance of how generative AI works, based on marketing efforts by vendors, or reflect efforts by policy makers and commentators to get attention.⁷⁴ Nevertheless, the verbal alarm bells are widespread and having an effect on policymakers.⁷⁵

The proposals to regulate AI are wide-ranging. They include outright bans on deployment or use of AI, licensing requirements, disclosure, prohibition of discriminatory results, mandatory training, censorship, robot taxes, physical inspection, stakeholder involvement in entrepreneurial decision making, and enforcement of private guidelines.⁷⁶

One commentator suggests three elements to AI regulation. “First, regulation should enforce more transparency on the part of developers. That would mean transparency about when AI is being used, as in the case of companies using AI algorithms to sort through job or rental applications.”⁷⁷ Second, developers should “make available to researchers the data on which chat bots are trained.”⁷⁸ Third, “AI should declare that it is AI.”⁷⁹

A frequent demand is that AI systems be “transparent,” a subject considered in § A.

B. What is a patent?

A patent is a twenty-year⁸⁰ monopoly granted by an agency of the U.S. government called the United States Patent and Trademark Office (USPTO).⁸¹

The Patents and Copyrights clause of the United States Constitution authorizes the Congress “[t]o promote the Progress of Science and the Useful Arts. . . by securing for limited Times to . . . Inventors the exclusive Right to their respective . . .

⁷³ See Matt Egan, *AI Could Pose ‘Extinction-Level’ Threat to Humans and the US Must Intervene, State Dept.-Commissioned Report Warns*, CNN (Mar. 12, 2024), <https://www.cnn.com/2024/03/12/business/artificial-intelligence-ai-report-extinction/index.html>.

⁷⁴ The press is an echo chamber, managed by lemmings. It has consistently stoked public hysteria over new technologies, starting with the power loom and the spinning Jenny. It would have done the same with the printing press, except that the press did not exist before that innovation.

⁷⁵ See NAT’L CONF. OF STATE LEGISLATURES, A.I. 2024 LEGIS. (updated Sept. 9, 2024); 2024 WIPO REPORT, *supra* note 18, at 68–69 (reporting on different approaches in different countries to regulate AI).

⁷⁶ Henry H. Perritt, Jr., *Robot Regulations*, 75 S.C. L. REV. 219, 241–268 (2023).

⁷⁷ Daniel Van Boom, *Elon Musk Is Right: We Need to Regulate AI Now*, CNET (May 17, 2023), <https://www.cnet.com/tech/elon-musk-is-right-we-need-to-regulate-ai-now/>.

⁷⁸ *Id.*

⁷⁹ *Id.* See also Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. REV. 1347, 1380-1386 (2023) (identifying impact assessments, audits, testing, precautionary bans, and post-market mechanisms as regulatory tools).

⁸⁰ Extendable to compensate for delays caused by the Patent Office. 35 U.S.C. § 156.

⁸¹ 35 U.S.C. § 2(a)(1) (granting authority to USPTO); 35 U.S.C. § 154(a)(2) (defining term of patent as twenty years).

[d]iscoveries.”⁸² The very first Congress exercised that authority by enacting the Patent Act of 1790 (“the Patent Act”),⁸³ which authorized a committee of the cabinet made up of the Secretary of State, the Secretary of War, and the Attorney General to grant patents “if they shall deem the invention or discovery sufficiently useful and important.”⁸⁴

The prerequisites for patents have become more detailed since 1790 and the machinery for granting them more complex, but the basic concepts have remained the same. Most recently, the America Invents Act of 2011⁸⁵ replaced the traditional first-to-invent system of priority with a new first-inventor-to-apply system similar to that in use by the rest of the world.

To qualify for a patent, inventions must involve eligible subject matter, defined by § 101 as involving a process, machine, manufacture, or composition of matter.⁸⁶ Even when an invention falls into one of the categories enumerated in § 101, it nevertheless is non-patentable if it falls within one of the judicial exceptions to § 101: if it comprises a law of nature, a natural phenomenon, or an abstract idea.⁸⁷ Section II.B considers statutory eligibility and the judicial exceptions in more detail.

To qualify for a patent, inventions also must be novel, under § 102.⁸⁸ This means that they must not be anticipated by any prior art reference. A prior art reference defeating novelty is said to *anticipate* the invention. The statute defines prior art as “patented, described in a printed publication, or in public use, on sale, or otherwise available to the public.”⁸⁹ Novelty is defeated when a single prior art reference exists that contains every element of the claimed invention.⁹⁰

In addition to being novel, an invention must be nonobvious under § 103.⁹¹ Obviousness under § 103 results when one or more prior art references collectively contain all of the elements of the claimed invention and something motivates a person skilled in the relevant art to combine those references to come up with the invention. The Supreme Court, in *Graham v. John Deere Co.*,⁹² offered what the USPTO

⁸² U.S. CONST. art. I, § 8, cl. 8.

⁸³ Patent Act of 1790, ch. 7, 1 Stat. 109–12 (April 10, 1790).

⁸⁴ *Id.* at § 1.

⁸⁵ Leahy-Smith America Invents Act, Pub. L. 112–29, 125 Stat. 284 (Sep. 16, 2011).

⁸⁶ 35 U.S.C. § 101.

⁸⁷ *See Alice Corp. v. CLS Bank Int’l*, 573 U.S. 208, 216 (2014) (holding that claims were drawn from the abstract idea of financial settlement and generic computer implementation did not sufficiently transform them into a patent-eligible invention).

⁸⁸ 35 U.S.C. § 102.

⁸⁹ 35 U.S.C. § 102(a)(1). In *Helsinn*, the Supreme Court made it clear that secret sales may nevertheless constitute prior art. *Helsinn Healthcare S.A. v. Teva Pharms. USA, Inc.*, 586 U.S. 123, 125 (2019) (holding that sale of an invention to third party who is contractually obligated to keep invention secret places invention “on sale” within meaning of section 102(a)).

⁹⁰ MPEP § 2120(III) (9th ed. Rev. 10, 2019) (using the term “anticipate”); MPEP § 2131.01 (9th ed. Rev. 8, 2017) (asserting that 102 anticipation usually involves only one reference).

⁹¹ 35 U.S.C. § 103.

⁹² *Graham v. John Deere Co. of Kan. City*, 383 U.S. 1, 17 (1966) (rejecting argument that new section 103 was meant to change the law; instead “the section was intended merely as a codification of

understood to be a relatively simple formula; in addition to the elements of the claimed invention, prior art had to contain some teaching, suggestion, or motivation (“TSM”) to combine them.⁹³ The *John Deere* Court also accepted the idea that “secondary considerations” could play a role in assessing obviousness:

Such secondary considerations as commercial success, long felt but unsolved needs, failure of others, etc., might be utilized to give light to the circumstances surrounding the origin of the subject matter sought to be patented. As indicia of obviousness or nonobviousness, these inquiries may have relevancy.⁹⁴

That TSM test was applied as the core analytical principle of obviousness analysis until the Supreme Court decided *KSR International Co. v. Teleflex Inc.* (“*KSR*”)⁹⁵ in 2007. In *KSR*, the Supreme Court rebuked the United States Court of Appeals for the Federal Circuit, the appeals court with exclusive jurisdiction over patent appeals, for applying too simplistic and rigid a test and held that a more flexible assessment of obviousness was required.⁹⁶

The USPTO synthesized *KSR* for its examiners, giving a number of examples.⁹⁷ The consistent theme in all of the examples is whether a person of ordinary skill in the art (“POSITA”)⁹⁸ would have a “reasonable expectation of success,” in combining elements from the prior art references.⁹⁹ If a POSITA would not have had a reasonable expectation of success in combining elements of different sources, the invention is not obvious.¹⁰⁰ This reasonable-expectation-of-success standard is similar in many respects to the without-undue-experimentation test for enablement under § 112.¹⁰¹ The specification must be sufficiently clear and definite that one can know what the inventor is claiming as his invention, and it must be sufficiently detailed and concrete that someone skilled in the art would be enabled to make and use the invention with little more than reading the specification.¹⁰²

An application for a patent, even for a novel and nonobvious invention, must be

judicial precedents embracing the Hotchkiss [Hotchkiss v. Greenwood, 11 How. 248 (1851)], condition.”).

⁹³ *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 407 (explaining USPTO’s “‘teaching, suggestion, or motivation’ test (TSM test), under which a patent claim is only proved obvious if ‘some motivation or suggestion to combine the prior art teachings’ can be found in the prior art, the nature of the problem, or the knowledge of a person having ordinary skill in the art.”).

⁹⁴ *Graham*, 383 U.S. at 17–18.

⁹⁵ *See KSR Int’l Co.*, 550 U.S. 398.

⁹⁶ *Id.* at 415.

⁹⁷ *See* MPEP § 2141(III) (9th ed. Rev. 7, 2022); *See generally* MPEP § 2143 (9th ed. Rev. 7, 2022).

⁹⁸ *See* MPEP § 2141.03 (9th ed. Rev. 7, 2022) (defining a person of ordinary skill in the art).

⁹⁹ *See generally* MPEP § 2143 (9th ed. Rev. 7, 2022).

¹⁰⁰ *Salix Pharms., Ltd. v. Norwich Pharms. Inc.*, 98 F.4th 1056, 1065–67 (Fed. Cir. 2024) (applying reasonable expectation of success standard in affirming finding of obviousness).

¹⁰¹ 35 U.S.C. § 112; MPEP § 2164 (9th ed. Rev. 11, 2013) (discussing section 112’s enablement requirement).

¹⁰² Henry H. Perritt, Jr., *Literary Fantasies as Prior Art: Eclipsing True Invention*, 104 J. PAT. TRADEMARK OFF. SOC’Y 453, 465–69, 496–504 (2024) (criticizing presumption that prior art that anticipates is enabling).

accompanied by a specification published as part of the patent. It must be sufficiently clear and definite such that one can know what the inventor is claiming as his invention, and it must be sufficiently detailed and concrete such that someone skilled in the art would be enabled to make and use the invention with little more than reading the specification.¹⁰³

Section 112 imposes two requirements: first, that the specification particularly point out and distinctly claim the subject matter which the inventor or a joint inventor regards as the invention,¹⁰⁴ and that it describe the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable one skilled in the art to make and use the invention.¹⁰⁵ The first function is largely performed by the claims section of the specification, required to be separated and labeled as such.¹⁰⁶

Second, § 112 requires that the applicant make the metes and bounds of a claimed invention clear.¹⁰⁷ That requirement is performed by the claims but is not limited to the claims; the more discursive description is intended to flesh out an understanding of the claims of the patent. The enablement requirement of § 112 is exported into § 102 in § 103 analysis: a prior art reference is not anticipatory under § 102 unless it is enabling,¹⁰⁸ and a combination of prior art references do not make an invention obvious unless they are collectively enabling.¹⁰⁹

Patents are issued to “inventors.”¹¹⁰ The inventorship concept was refined in U.S. patent law during its first 223 years, when priority was determined based on when invention occurred. Priority disputes turned on who was the first to invent, and to resolve those disputes necessitated rules for deciding when an invention was complete. Was conception enough, or must there be reduction to practice? And what constitutes reduction to practice? Under *Perry v. Cornell*, “[t]here is no law requiring the applicant to reduce his invention to actual use before he can obtain a patent.”¹¹¹

The statute defines how the USPTO reviews applications and interacts with applicants,¹¹² and the USPTO has, by rule, filled in the details.¹¹³ Decisions by examiners can be reviewed in an administrative appeal by the Patent Trial and Appeals Board (“PTAB”).¹¹⁴

After patents are issued, they may be revised by being withdrawn and re-issued,¹¹⁵ or in ex parte re-examinations triggered by the patent owner or a third

¹⁰³ 35 U.S.C. § 112.

¹⁰⁴ 35 U.S.C. § 112(b).

¹⁰⁵ 35 U.S.C. § 112(a).

¹⁰⁶ MPEP § 608.01(m) (9th ed. Rev. 10, 2019) (describing required form of claims).

¹⁰⁷ MPEP § 2174 (9th ed. Rev. 11, 2013) (explaining purpose of definiteness requirement for claims).

¹⁰⁸ MPEP § 2121(I) (9th ed. 8, 2017) (enablement is presumed).

¹⁰⁹ *Raytheon Techs. Corp. v. Gen. Elec. Co.*, 993 F.3d 1374, 1376 (Fed. Cir. 2021).

¹¹⁰ 35 U.S.C. § 101.

¹¹¹ *Perry v. Cornell*, 19 F. Cas. 267, 271. (C.C. D.C. 1847).

¹¹² 35 U.S.C. §§ 131–33.

¹¹³ See generally MPEP chaps. 600, 700, 2100.

¹¹⁴ 35 U.S.C. § 134.

¹¹⁵ 35 U.S.C. § 251; MPEP chap. 1400 (reissue of patent).

party.¹¹⁶ Since 2011, they also can be challenged in administrative litigation before the PTAB in Post Grant Review on almost any grounds for nine months after issuance.¹¹⁷ After nine months, a patent may be challenged through an Inter Partes Review,¹¹⁸ where a challenger can marshal patents or printed publications¹¹⁹ that establish a reasonable likelihood that the petitioner would prevail with respect to at least one of the claims challenged in the petition.¹²⁰ The validity of issued patents also can be reviewed by U.S. district courts adjudicating actions brought by patent holders for infringement¹²¹ or in declaratory judgment actions brought by potential infringers.¹²²

The system in the U.S. is integrated with other systems around the world through the Patent Cooperation Treaty, providing for international filings and recognition of filings in one country by the patent offices of other countries.¹²³

III. Is the disclosure required to obtain a patent sufficient to satisfy the demands for transparency?

The thesis of this article is that a properly applied patent disclosure requirement can be a vehicle for achieving transparency of AI inventions.¹²⁴ The discussion that follows considers primarily AI systems developed by deep machine learning. Traditional systems that accept computer coded conditions and produce structured computer coded results may fall within an imprecise definition of AI, but they pose no new regulatory challenges.

A. What is AI transparency?

No one knows.¹²⁵ Transparency may apply to the fact that AI technology was used, to the way it works, or to its results. How it works is the type of transparency most relevant to this article.

Most of the proposals for regulation of generative AI have called for increased transparency of the development of models through machine learning, of the databases used, and of precautionary measures implemented by developers and users.¹²⁶ Notably, however, Sam Altman, in his May 2023 testimony before the Senate Judiciary Committee, did not embrace transparency as a regulatory goal, even

¹¹⁶ 35 U.S.C. §§ 302–307; MPEP chap. 2200 (ex-parte re-examination).

¹¹⁷ 35 U.S.C. §§ 321–29.

¹¹⁸ 35 U.S.C. §§ 311–19.

¹¹⁹ 35 U.S.C. § 311(b).

¹²⁰ 35 U.S.C. § 314(a).

¹²¹ 35 U.S.C. § 282(b)(2).

¹²² See *Medimmune, Inc. v. Genentech, Inc.*, 549 U.S. 118, 137 (holding that party seeking declaratory judgment of patent invalidity need not breach license agreement first).

¹²³ 35 U.S.C. §§ 361–76; MPEP § chap. 1800 (implementing Patent Cooperation Treaty).

¹²⁴ See Charlotte A. Tschider, *Beyond the “Black Box”*, 98 DENVER L. REV. 683, 715 (noting suggestions by Roger Ford and Nicholson Price to leverage patent system to improve AI transparency and reduce reliance on trade secrets).

¹²⁵ Datzov, *supra* note 8, at 16 (identifying absence of definition of AI as problem for policy evaluation and legal critique).

¹²⁶ 2024 WIPO REPORT, *supra* note 18, at 69 (listing transparency at top of list of best practices).

though he advanced pretty far out in front of other industry spokesmen in embracing the idea of regulation in general, through a new agency.¹²⁷

Much of the debate over regulation of AI seeks disclosure by system designers.¹²⁸ This is a particularly popular suggestion in the debate over generative AI regulation, usually expressed as assuring *transparency*. The simplest form of disclosure is labelling, similar to that required for foods by the Federal Food, Drug, and Cosmetic Act (“FDA”).¹²⁹ Generative AI systems can be required to disclose that they are such rather than a real human being. This would follow the example of FDA food labeling and consumer product labeling requirements.

The National Institute of Standards and Technology (“NIST”) has published a risk management framework for generative AI.¹³⁰ This framework enumerates specific risks associated with the use of AI¹³¹ and also identifies more than 500 mitigation measures that can reduce those risks.¹³² Some of them involve additional transparency.¹³³

Colorado has enacted a general AI statute¹³⁴ that requires deployers of high risk AI systems to perform risk assessments and to disclose to consumers: the reasons for

¹²⁷ See generally Written Testimony of Sam Altman Chief Executive Officer OpenAI Before the U.S. Sen. Comm. On the Judiciary & Sub Comm. On Privacy, Tech., & the Law, S.Hrg. 118-37, No. J-118-16 (1st Sess. 2023).

¹²⁸ Perritt, *supra* note 76, at 251–53.

¹²⁹ See 21 U.S.C. § 343 (detailing requirements for food labeling).

¹³⁰ U.S. Dept. of Commerce: National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1 April 2024) (withdrawn draft), <https://airc.nist.gov/docs/NIST.AI.600-1.GenAI-Profile.ipd.pdf>.

¹³¹ *Id.* at § 2.

¹³² *Id.* at § 3.

¹³³ See *id.* GV-5.1-002 at 20 (“Disclose interactions with GAI systems to users prior to interactive activities.”); *Id.* GV-5.1-003 at 20 (“Engage independent experts to audit models, data sources, licenses, algorithms, and other system components, Consider sponsoring or engaging in community-based exercises (e.g., bug bounties, hackathons, competitions) where AI actors assess and benchmark the performance of AI systems, including the robustness of content provenance management under various conditions.”); *Id.* MP-1.1-005 at 25 (“Document GAI system ownership, intended use, direct organizational value, and assumptions and limitations.”); *Id.* MP-2.1-001 at 26 (“Define GAI system’s task(s) that relate to content provenance, such as original content creation, media synthesis, or data augmentation while incorporating tracking measures.”); *Id.* MP-2.1-004 at 27 (“Institute audit trails for data and content flows within the system, including but not limited to, original data sources, data transformations, and decision-making criteria.”); *Id.* MS-2.5-004 at 39 (“Document the construct validity of methodologies employed in GAI systems relative to their context of use.”); *Id.* MS-2.5-007 at 39 (“Measure, monitor, and document prevalence of erroneous GAI output content, system availability, and reproducibility of outcomes via field testing or other randomized controlled experiments.”); *Id.* MS-2.8-010 at 44 (“Use interpretable machine learning techniques to make AI processes and outcomes more transparent, and easier to understand how decisions are made.”); *Id.* MS-2.9-003 at 45 (“Document GAI model details including: Proposed use and organizational value; Assumptions and limitations, Data collection methodologies; Data provenance; Data quality; Model architecture (e.g., convolutional neural network, transformers, etc.); Optimization objectives; Training algorithms; RLHF approaches; Fine-tuning approaches; Evaluation data; Ethical considerations; Legal and regulatory requirements.”).

¹³⁴ Colo. Rev. Stat. Ann. § 6-1-1703.

“consequential decisions” made with the aid of high risk AI systems; the “degree to which and manner in which” the AI system contributed to the decision; the type of data that was process in making the consequential decision; and the sources of the data. It also requires deployers to give an opportunity for the consumer to correct any incorrect personal data used, an opportunity to appeal any adverse decision, and an opportunity to insist on human involvement in the review if that is feasible.¹³⁵

The statute does not make it clear whether the data disclosure requirement concerns only data identified with the consumer or whether it extends more broadly to data used to train, validate, or censor the AI system.

1. *What needs to be transparent?*

The calls for transparency of AI are meaningful only when one answers the question, “What needs to be transparent?” Datzov’s three layers—data, software, and systems¹³⁶—are useful ways of attempting to answer this question. The World Intellectual Property Organization (“WIPO”) report offers another: computer programs, type of input or output, and type of application.¹³⁷

Advocates of AI transparency have not articulated what transparency means very concretely. Some lists of the hallmarks of transparency are fairly broad, including things like the name and contact information for the vendor, version numbers for software, and disclosure of intended uses.

Article 13 of the European Union Artificial Intelligence Act (“EU AI Act”), adopted in July 2024, is the most formal and concrete list of transparency requirements. The EU AI Act has no direct effect in the U.S., and U.S. patent law has no direct effect in Europe, but the EU AI Act is held out as a model by many U.S. advocates of AI regulation.¹³⁸

The most detailed U.S. regulation of AI systems is contained in the January 2024 transparency rule of the U.S. Department of Health and Human Services.¹³⁹ The rule sets standards and requirements for certification of information technology used to facilitate patient care. In particular it regulates *Predictive Decision Support Intervention* (“Predictive DSI”), which is defined as:

[T]echnology that supports decision-making based on algorithms or models that derive relationships from training data and then produces an output that results in prediction, classification, recommendation, evaluation, or analysis.¹⁴⁰

¹³⁵ *Id.*

¹³⁶ See Datzov, *supra* note 8.

¹³⁷ 2024 WIPO REPORT, *supra* note 18, at 16 (reporting that AI patents can be addressed to (1) computer programs used, or models; (2) type of input of output, or modes; or (3) applications of generative AI).

¹³⁸ See John Hillman, *Smart Regulation: Lessons from the Artificial Intelligence Act*, 37 EMORY INT’L L. REV. 775, 822 (2023) (calling for U.S. to catch up to Europe and to use EU Act as a model for regulating AI).

¹³⁹ See Health Data, Technology, and Interoperability, 89 Fed. Reg. 1192 (Jan. 9, 2024).

¹⁴⁰ 45 C.F.R. § 170.102.

The rule emphasizes that its transparency requirements do not require disclosing or sharing intellectual property.¹⁴¹ It notes, with respect to its final rule:

The information required in source attributes is not intended to include detailed information on model parameters, hyperparameters, detailed specifics around how input or output variables are defined, transformed, or otherwise operationalized. We do not believe that information at that level of detail is necessary for source attributes in § 170.315(b)(11)(iv)(B) or necessary for users of a Predictive DSI to determine whether it is fair, appropriate, valid, effective, and safe.¹⁴²

The requirements reflect a preoccupation with bias:

[B]ias associated with AI and machine learning algorithms could create substantial risks if they are presented to the end user without information to understand how they were constructed, evaluated, and should be maintained. . . . We have only amended existing requirements for evidence-based DSIs by asking for specific data elements to be identified when used by the DSI, including race, ethnicity, language, sexual orientation, gender identity, sex, date of birth, SDOH, sexual orientation, and health assessments data elements (*e.g.*, disability status).¹⁴³

It is thus much more limited than the EU requirement.

The justification for mandating disclosure usually focuses on the content of the databases from which generative AI systems learn.¹⁴⁴ If researchers, observers, and critics know what databases have been used in the learning process, they are better positioned to pick apart the learning that occurs and to isolate instances of infringement of intellectual property or undesirable viewpoints or sources that the machine gives prominence to.

Transparency with respect to AI data means that the contents of training databases and validation databases should be disclosed so that the public can evaluate the representativeness of those databases. If an AI system meant to generate text learned from a database comprised mainly of images, the result will not be satisfactory. Exposure of the data used to train and validate also may be helpful discerning imperfections in the way the data is represented. Tags of data always manifest judgments about what characteristics of entries are relevant to the purposes of the system. Databases tagging individuals by race, age, national origin and sexual orientation will be highly suspect. And other, more subtle, shortcomings may be revealed. The size of the learning database may be too small to produce meaningful results, given the statistical techniques used; it may be skewed and underrepresent identifiable parts of the relevant population.

¹⁴¹ Health Data, Technology, and Interoperability, 89 Fed. Reg. 1192, 1259, 1261.

¹⁴² *Id.*

¹⁴³ *Id.* at 1295.

¹⁴⁴ See Alex Engler, *The Case for AI Transparency Requirements*, BROOKINGS (Jan. 22, 2020), <https://www.brookings.edu/articles/the-case-for-ai-transparency-requirements/> (Advocating for AI disclosure policies in the name of transparency).

Full data dumps are not likely to be very useful, however. State-of-the-art deep learning requires billions of items in the training database, and making those individual items available not only would be expensive and require huge amounts of memory for anyone who wants to examine them, but they would also be meaningless without some means of summarizing and measuring their characteristics. Moreover, privacy limitations would be implicated by finely grained disclosure.¹⁴⁵ Transparency of data, to be meaningful, must refer to meta-data, information about how the data was selected, from what sources, and how it was represented.

Transparency of AI software refers to disclosure of the particular statistical methods used to learn from the data, the algorithms that result from that learning, the error functions, and the protocols used to minimize errors. If the system uses validation beyond minimization of error functions, disclosure of validation methods also is relevant.

Knowledge about these aspects of the software layer can expose methodological weaknesses, such as the application of techniques more suited for image processing to large language databases, or the use of transformers, well suited for large language models, to image processing. Exposure of these aspects of the software also can enable debate and criticism about the power of the methods used and their propensity to produce desirable hallucinations,¹⁴⁶ or to match images poorly.

Disclosure of the design of the neural network transformers and the algorithms they use permit scrutiny of how loss functions are managed and optimized and how censorship layers are implemented, if at all.

On the other hand, algorithms expressed in the form of tensors with thousands of dimensions are not likely to be understandable by human observers.¹⁴⁷

Transparency of the software layer also means that developers and users should disclose any blacklists or censor layers they use.

Transparency of the hardware layer would expose the particular processor chips, data links, memory devices, centralized architectures, and electric power consumption efficiencies.

Going beyond Datzov's three layers, transparency also can be applied to results. Indeed, transparency of results may be more useful than transparency of Datzov's three layers. Mandates to disclose that AI technology was used can require a simple statement to that effect without more, or it can require disclosure of the brands of AI products and/or the underlying models such as ChatGPT, Google Gemini, or DALL-E 3. Such simple use disclosure can be coupled with requirements like those in the

¹⁴⁵ DELOITTE, *supra* note 1 (noting that publishing AI training data is problematic because of privacy regulations).

¹⁴⁶ Henry H. Perritt, Jr., *Robot Slanderer*, 46 U. ARK. LITTLE ROCK L. REV. 169, 172 (2023) (noting that generative AI systems are designed to hallucinate; that is the essence of generating new imagery, text, music, or computer code).

¹⁴⁷ DELOITTE, *supra* note 1 (noting that publishing algorithms is problematic because of intellectual property issues and because most people do not know how to make sense of AI models).

Colorado statute, summarized in § III.A, for the availability of appeals or human intervention, or they may be coupled with some form of consent provision, allowing those potentially affected to withhold their consent to the use of the technology.

Mandating transparency of results is not particularly meaningful unless it is coupled with some standard by which those results can be evaluated and the designation of an agency to perform the evaluation. It certainly is conceivable that standard test protocols could be developed and deployed, perhaps by NIST, and the systems performance in conducting those protocols evaluated, either by the marketplace in the case of a mere disclosure requirement, or by some regulatory agency.

Third parties are unlikely to know about all of the uses to which an AI system is put, and to the extent that they have standing to opine about the uses, some new means of opening a window into results is necessary. That might mean a requirement for vendors to sample results periodically and to make them available to the public or some third-party auditor. If such a requirement were imposed, user confidentiality would have to be respected in some fashion.¹⁴⁸

2. *To what end?*

Requiring transparency of a particular aspect of AI systems is unjustifiable unless the proponent of such transparency can offer a justifiable reason for the requirement. To what end must an element be disclosed? The preceding section, enumerating types of transparency discusses some of the motivations for requiring transparency.

The debate over AI has been dominated by expressions of concern about bias and hallucination. Transparency of both data and software is likely to fuel arguments and claims of bias.

Transparency of data and software also can enable a more informed debate about hallucination and how it can be managed, although argument on that subject is likely to outrun rigorous knowledge and experience.

3. *Who is going to look at it? When?*

Requirements for transparency usually contemplate some mechanism for scrutinizing whatever is disclosed. Formal scrutiny takes the form of *audits*. The National Telecommunications and Information Administration (“NTIA”) issued a request for comment on April 13, 2023.¹⁴⁹ The request:

[F]ocuse[d] on self-regulatory, regulatory, and other measures and

¹⁴⁸ Users are entitled to trade secret protection for their queries and the results of those queries, because each may confer a competitive advantage by virtue of not being generally known. *See* HENRY H. PERRITT, JR., *TRADE SECRETS FOR THE PRACTITIONER* chap. 1 (3d ed. 2024) (defining trade secret); Trade secrets are property, protected by the Fifth Amendment from compelled disclosure. *Ruckelshaus v. Monsanto*, 467 U.S. 986, 1003–1004 (1984) (holding that trade secrets are property protected by the Takings Clause of the Fifth Amendment).

¹⁴⁹ AI Accountability Policy Request for Comment, 88 Fed. Reg. 22433, 22433 (Apr. 13, 2023).

policies that are designed to provide reliable evidence to external stakeholders—that is, to provide assurance—that AI systems are legal, effective, ethical, safe, and otherwise trustworthy. NTIA will rely on these comments, along with other public engagements on this topic, to draft and issue a report on AI accountability policy development, focusing especially on the AI assurance ecosystem.¹⁵⁰

Transparency can be a regulatory requirement without more; it need not be accompanied by any sort of licensing, approval, or certification requirement. Developers and users of generative AI systems would be obligated to disclose specific characteristics and features of their systems, but any scrutiny of that information would take place only through public discourse, and not formal legal review or action.

Compelling transparency of AI systems confronts substantial practical impediments.¹⁵¹ The iterative interactions in a machine learning system among data, statistical models, and the resulting algorithms (constantly being redefined as further learning takes place) are difficult to explain. The learning databases frequently comprise billions of entries. The derived algorithms are vectors with thousands of dimensions.

B. What is patent disclosure?

No one knows.

Some critics claim that the disclosure required by patent law is inadequate to inform technologists.¹⁵² Information and software-technology inventors are criticized for their use of vague “patent-speak” jargon.¹⁵³ One professor, W. Keith Robinson, concludes that patent disclosure is irrelevant.¹⁵⁴ According to his view, patent disclosure would be insufficient to satisfy the demands for AI transparency. But

¹⁵⁰ *Id.*

¹⁵¹ The reality of modern AI is that neural networks and deep-learning applications obscure details of the AI system to such an extent that transparency is nearly impossible to achieve—even an AI’s creators do not fully understand how it works. AI algorithms, in their most effective state, also continue to learn and change automatically to become more effective, so even if an algorithm could be explained, it would only be explainable temporarily before the algorithm changed again. Furthermore, the focus on the *algorithm* is misplaced: the methods, system, and processes that create and sustain the algorithm and its functionality actually control what the algorithm is, what it does, how reliable or fair it is, and ultimately how safe it will be.

Tschider, *supra* note 124, at 689 (internal footnotes omitted); *Id.* at 700 (explaining why the calls for AI transparency may be undesirable and nearly impossible to achieve).

¹⁵² Tabriz Y. Ebrahim, *Artificial Intelligence Inventions and Patent Disclosure*, 125 PENN STATE L. REV. 147, 188 (2020).

¹⁵³ *Id.* at 159 (internal footnotes omitted) (citing 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”)); Allan Devlin, *The Misunderstood Function of Disclosure in Patent Law*, 23 HARV. J.L. & TECH. 401, 403 (2010) (characterizing patents in the information technology industry as being notorious for their vague language).

¹⁵⁴ W. Keith Robinson, *Enabling Artificial Intelligence*, 60 HOUS. L. REV. 331, 356–58 (2022) (arguing that enablement and patent disclosure in general is irrelevant to other innovators; assessing whether enablement requirement is a barrier to getting AI patents).

“many scholars have provided empirical evidence that patent disclosure has benefits for the use of technical information and provides knowledge spillovers. 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.”¹⁵⁵ Tabriz Y. Ebrahim thinks current disclosure requirements are inadequate for AI inventions. He proposes intensifying the disclosure requirements when AI inventions are sought to be patented.¹⁵⁶

Ebrahim claims that “AI fundamentally challenges disclosure in patent law (“patent disclosure”), which has not kept up with rapid advancements in AI technology.”¹⁵⁷ He argues that “nonintuitive data relationships between AI-based tools and AI-generated output should compel AI inventors to strive for greater patent disclosure.”¹⁵⁸ His focus is on AI as an inventor or joint inventor, rather than on AI as the subject of inventive effort. AI as an inventor presents greater transparency problems than AI inventions.¹⁵⁹ Generative AI is “inscrutable,” to use Ebrahim’s term.¹⁶⁰ Applying any form of transparency or disclosure requirement to it is challenging.¹⁶¹ The same difficulties do not exist with AI inventions. To understand what an inventor has invented, he must be able to explain it. He knows what comprised the learning database, the statistical techniques used to develop algorithms from it, and the mathematical function used to represent and to minimize error functions.

In *Ex parte Rubendran Amarasingham et al.*,¹⁶² the PTAB affirmed an examiner’s § 112 rejection:

Paragraph 64 discloses “artificial intelligence model tuning process utilizes adaptive self-learning capabilities using machine learning technologies,” but does not disclose an algorithm for implementing either the process or adaptive self-learning. Likewise, paragraph 65 discloses capabilities including periodically reassessing weightings, examining thresholds, and determining whether to update but fails to disclose how these capabilities are to be implemented. Figure 5 depicts AI Model Tuning as a single block, i.e., a generic black box without any details of the device itself and devoid of the requisite corresponding structure.¹⁶³

In *Ex Parte Guy Lev, et al.*,¹⁶⁴ the applicant disclosed quite a lot about its use of an AI model but still suffered a § 112 rejection, which the PTAB affirmed. The

¹⁵⁵ Ebrahim, *supra* note 152, at 189 (internal footnotes omitted).

¹⁵⁶ See generally *id.* (arguing for enhanced disclosure requirements for AI patents and patent applications).

¹⁵⁷ *Id.* at 153.

¹⁵⁸ *Id.* at 158 (internal footnotes omitted).

¹⁵⁹ *Id.* at 171–81.

¹⁶⁰ See, e.g., *Id.* at 161.

¹⁶¹ *Id.* at 171–73 (identifying challenges in making AI processes more transparent).

¹⁶² See *Ex parte Rubendran Amarasingham et al.*, No. 2023-001276, 2024 WL 2179415 (P.T.A.B. May 14, 2024).

¹⁶³ *Id.* at *7.

¹⁶⁴ See *Ex parte Guy Lev et al.*, No. 2023-001664, 2024 WL 3220638 (P.T.A.B. June 28, 2024).

applicant claimed use of GANs to deter anomalies in computer system abuse. The examiner found a lack of detail resulting in an inadequate disclosure.¹⁶⁵

The USPTO decisions show that patent disclosure requirements under § 112 considerably exceed any of the calls for AI transparency.

The 2024 WIPO Report¹⁶⁶ makes it clear that national patent systems all over the world have embraced disclosures with respect to AI inventions that expose the details of data, models, and algorithms.

C. Convergence between AI transparency and patent disclosure

The patent system has no obvious role to play in implementing some of the proposals for regulation, such as a prohibition on deployment absent regulatory approval, limitations on the purposes for which AI systems may be used, and auditing and reporting requirements. But a central purpose of the patent system is congruent with one of the loudest calls for AI regulation: transparency.

Charlotte Tschider considered the relationship among demands for AI transparency, trade secret law, and the patent disclosure requirement.¹⁶⁷ She asserts that patent law's disclosure requirements may enhance AI transparency, but does not really develop the idea,¹⁶⁸ apparently because she thinks that patent-eligibility restrictions are likely to make AI algorithms unpatentable,¹⁶⁹ and because "patents are uninformative for purposes of use or duplication."¹⁷⁰

¹⁶⁵ [T]he application provides no direction as to the following:

- 1) No objective function or loss function;
- 2) No code or pseudocode;
- 3) Not a single parameter/weight/coefficient;
- 4) No benchmark results or known dataset comparison, e.g., MNIST or CIFAR;
- 5) No architecture layout illustrating the relation between generator and discriminator;
- 6) No unlabeled data to convey unsupervised learning;
- 7) No latent representation of variables or discussion of feature space;
- 8) No vectors or matrices and no operations for transformation, e.g., element-wise;
- 9) No reconstruction error or regularization/normalization for training;
- 10) No validation or testing of data, e.g., synthetic dataset
- 11) No sampling of group data;
- 12) No attention mechanism, i.e., context vector suggesting transformer;
- 13) No back-propagation or discussion of directionality;
- 14) No discussion of dimensionality or singular value decomposition;
- 15) No calculation specified for anomaly probability, non-integer or otherwise;
- 16) No definition of terms: context or conditional as used throughout;
- 17) No form of solution suggested, e.g., sub-quadratic or exponential; and
- 18) No stability or convergence discussion to stave mode collapse or even tractability.

Id. at *6.

¹⁶⁶ 2024 WIPO REPORT, *supra* note 18, at 16.

¹⁶⁷ Tschider, *supra* note 124, at 683 (evaluating relationship between demands for AI transparency, trade-secret protection, and patent-law's disclosure requirement).

¹⁶⁸ *Id.* at 715–17.

¹⁶⁹ *Id.* at 717.

¹⁷⁰ *Id.* at 718.

A useful exercise is to take the transparency language of the EU AI Act and consider whether a disclosure satisfying its requirements would be likely to satisfy § 112, and, conversely, whether a disclosure satisfying § 112 would satisfy the language of the EU AI Act.

Some requirements in the EU AI Act, such as the requirement that developers disclose “any known or foreseeable circumstance . . . which may lead to risks to the health and safety or fundamental rights,”¹⁷¹ or the requirement that developers disclose “the name of the provider and the version of the system reflecting its relation to previous versions,”¹⁷² are outside the scope of patent transparency requirements. But others, such as the requirement that “[h]igh-risk AI systems shall be designed and developed in such a way as to ensure that their operation is sufficiently transparent to enable deployers to interpret a system’s output and use it appropriately,”¹⁷³ and the requirement that they “be accompanied by instructions for use in an appropriate digital format or otherwise that include concise, complete, correct and clear information that is relevant, accessible and comprehensible to deployers,”¹⁷⁴ reinforce patent law’s requirement that the disclosure be sufficiently detailed as to enable one skilled in the art to make and use the invention.¹⁷⁵

Likewise, the following requirements from the EU AI Act completely overlap with patent law’s requirement for a “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 make and use [it]”:¹⁷⁶

A detailed description of the elements of the AI system and of the process for its development, including:

- (a) the methods and steps performed for the development of the AI system, including, where relevant, recourse to pre-trained systems or tools provided by third parties and how those were used, integrated or modified by the provider;
- (b) the design specifications of the system, namely the general logic of the AI system and of the algorithms; the key design choices including the rationale and assumptions made, including with regard to persons or groups of persons in respect of who, the system is intended to be used; the main classification choices; what the system is designed to optimise for, and the relevance of the different parameters; the description of the expected output and output quality of the system; the decisions about any possible trade-off made regarding the technical solutions . . .;
- (c) the description of the system architecture explaining how software components build on or feed into each other and integrate into the overall processing; the computational resources used to

¹⁷¹ Artificial Intelligence Act, *supra* note 11, at Art. 13(3)(b)(iii).

¹⁷² *Id.* at Annex IV(1)(a).

¹⁷³ *Id.* at Art. 13(1).

¹⁷⁴ *Id.* at Art. 13(2).

¹⁷⁵ 35 U.S.C. § 112(a).

¹⁷⁶ *Id.*

develop, train, test and validate the AI system.¹⁷⁷

The PTAB decisions insisting on more disclosure of AI inventions to satisfy the requirements of § 112¹⁷⁸ make it clear that patent disclosure is more than enough to satisfy any requirements for transparency.

IV. Are AI inventions eligible for patent protection?

Patent disclosure will satisfy demands for AI transparency only if AI is patentable.¹⁷⁹ Whether AI inventions are eligible for patent protection requires application of the United States Supreme Court decisions in *Alice* and *Mayo*,¹⁸⁰ interpreting 35 U.S.C. § 101.

A. Evolution of § 101

1. *Statutory subject matter*

The eligibility requirements have not changed much over the life of the U.S. patent system. Under the Patent Act, eligible subject matter was defined as “any useful art, manufacture, engine, machine, or device.”¹⁸¹ Under current law, the subject matter of an invention must involve a process,¹⁸² machine,¹⁸³ manufacture, or composition of matter.¹⁸⁴ In *Diamond v. Diehr* (“*Diehr*”),¹⁸⁵ the Supreme Court quoted legislative history to indicate that “[c]ongress intended statutory subject matter to include anything under the sun that is made by man.”¹⁸⁶ It cautioned that “courts should not read into the patent laws limitations and conditions which the legislature has not expressed.”¹⁸⁷ The Court held that a process for curing synthetic rubber which includes in several of its steps the use of a mathematical formula and a programmed digital computer was patentable subject matter under 35 U.S.C. § 101.¹⁸⁸

2. *Judicial exceptions*

Thirty-five years later, however, the Court held that even when an invention falls into one of the categories enumerated in § 101, it nevertheless is non-patentable if it comprises a law of nature, a natural phenomenon, or an abstract idea.¹⁸⁹ In *Alice*, the

¹⁷⁷ Artificial Intelligence Act, *supra* note 11, at Annex IV(2).

¹⁷⁸ See discussion *supra* Section B.

¹⁷⁹ One can apply for a patent, even if a patent eventually is denied. Patent applications, regardless of whether they result in patents are published eighteen months after filing, with some exceptions. 35 U.S.C. § 122(b). So even unsuccessful AI applications could result in disclosure and transparency. But AI inventors will not apply for patents if they are unlikely to get them.

¹⁸⁰ *Datzov*, *supra* note 8, at 12 (reviewing the continuing debate over *Alice/Mayo*).

¹⁸¹ Patent Act of 1790, ch. 7, 1 Stat. 109–12 (April 10, 1790).

¹⁸² AI inventions qualify as method (process) claims because they involve sequences of steps.

¹⁸³ AI inventions qualify as apparatus (machine) claims because they all run on computer hardware.

¹⁸⁴ 35 U.S.C. § 101.

¹⁸⁵ See *Diamond v. Diehr*, 450 U.S. 175 (1981).

¹⁸⁶ *Id.* at 182.

¹⁸⁷ *Id.*

¹⁸⁸ *Id.* at 177.

¹⁸⁹ *Alice Corp. v. CLS Bank Int’l*, 573 U.S. 208, 216 (2014) (holding that claims were drawn to abstract idea of financial settlement and that requiring generic computer implementation did not transform

Supreme Court reiterated these three “judicial exceptions,” noting that they have applied for more than 150 years, and applied the framework articulated two years earlier in *Mayo*.¹⁹⁰

Applying the judicial exceptions is known as step one¹⁹¹ of the *Alice*¹⁹²/*Mayo*¹⁹³ analysis. The concern that animates the judicial exceptions is preemption; upholding a patent covering these exceptions:

would pre-empt use of this approach in all fields, and would effectively grant a monopoly over an abstract idea. Laws of nature, natural phenomena, and abstract ideas are the basic tools of scientific and technological work. Monopolization of those tools through the grant of a patent might tend to impede innovation more than it would tend to promote it, thereby thwarting the primary object of the patent laws.¹⁹⁴

3. *Something more—an “inventive concept”*

Even if an invention involves a judicial exception, it nevertheless may be patentable if it utilizes an inventive concept—something more.

First, we determine whether the claims at issue are directed to one of those patent-ineligible concepts. If so, we then ask, “[w]hat else is there in the claims before us?” To answer that question, we consider the elements of each claim both individually and ‘as an ordered combination’ to determine whether the additional elements ‘transform the nature of the claim’ into a patent-eligible application. We have described step two of this analysis as a search for an ‘inventive concept’ — *i.e.*, an element or combination of elements that is sufficient to ensure that the patent in practice amounts to significantly more than a patent upon the ineligible concept] itself.¹⁹⁵

So, under *Alice/Mayo*, step one¹⁹⁶ determines whether the invention falls within

them into patent-eligible invention).

¹⁹⁰ *Id.* at 216, 217–218 (quoting *Mayo Collaborative Servs. v. Prometheus Lab’ys, Inc.*, 566 U.S. 66 (2012)).

¹⁹¹ USPTO numbers the steps somewhat differently. In its system, step 1 involves determining if an invention involves one of the four statutory categories. MPEP § 2106.03 (9th ed. Rev. 7, 2022). Step 2A asks if a claim is directed to a judicial exception. MPEP § 2106.04 (9th ed. Rev. 7, 2022). Step 2B asked if a claim amounts to significantly more. MPEP § 2106.05 (9th ed. Rev. 7, 2022).

¹⁹² *See Alice*, 573 U.S. at 226–27 (holding that claims were drawn to abstract idea of financial settlement and that requiring generic computer implementation did not transform them into patent-eligible invention).

¹⁹³ *See Mayo*, 566 U.S. 66.

¹⁹⁴ *Alice*, 573 U.S. at 216.

¹⁹⁵ *Id.* at 217–18 (internal quotations and citations omitted); MPEP § 2106.05 (9th ed. Rev. 7, 2022) (noting terminology).

¹⁹⁶ The numbering of the steps is confusing because the patent office starts with the section 101 textual assessment and assigns the number one to it, while labeling the *Alice/Mayo* inquiries as steps two-A and two-B. Compare MPEP § 2106.03 “Eligibility Step 1: The Four Categories of Statutory Subject Matter” (9th ed. Rev. 7, 2022), MPEP § 2106.04 “Eligibility Step 2A: Whether a Claim is Directed to a Judicial Exception” (9th ed. Rev. 7, 2022), and MPEP § 2106.05 “Eligibility Step 2B: Whether a Claim Amounts to Significantly More” (9th ed. Rev. 7, 2022) with *Alice Corp. Pty., Ltd. v. CLS Bank Int’l*, 573 U.S. 208, 217 (2014) (“We have described step two of this analysis as a search for

one of the statutory categories: process, machine, article of manufacture, or composition of matter.¹⁹⁷ Step two¹⁹⁸ asks whether the invention also falls within one of the judicial exceptions: an algorithm, a law of nature, a natural phenomenon, or a product of nature.¹⁹⁹ Step three²⁰⁰ asks whether the invention nevertheless is patent eligible because it includes something more—an inventive step—beyond the judicial exceptions.²⁰¹

4. *Avoiding preemption*

The concern of *Alice/Mayo* is to avoid the preemption of knowledge already in the public domain.²⁰² Nature supplies much of what is in the public domain, so laws of nature are not eligible for patent. Mathematical operations are akin to laws of nature, although they are human conceptions. They are sufficiently necessary for large areas of human activity, and they are very old. So, they qualify as part of the public domain as well.

AI inventions are problematic under *Alice/Mayo*, although not under the text of § 101. A generative AI models is undeniably a “system,” and it runs on a “machine.” But AI systems are collections of data and algorithms, likely with a judicial exception under *Alice/Mayo*. Not only that, but their techniques have also been in the public domain for quite some time.²⁰³ So to be eligible for a patent, an AI invention must involve “something more,” an “inventive concept.”

Diehr recognizes that the preemptive effect of a patent can be circumscribed by tying the use of something in the public domain to a machine.²⁰⁴ Monopolizing the application of a concept on a particular machine with unique characteristics does not preempt use of the concept on other types of machines.

But the utility of the *Diehr* approach depends on how narrowly the machine is specified. If one were to obtain a patent for applying Newtonian equations for forces and moments to a simple lever, the inclusion of the lever should not result in patent eligibility, because a lever is so general.

Conversely if one ties the moment and force equations to a truss design, the invention might be eligible for a patent if the inventor can show that the truss design

an “inventive concept,”) and *In re Killian*, 45 F. 4th 1373, 1379 (Fed. Cir. 2022) (“We have held that mental processes are abstract ideas under *Alice/Mayo* step one.”). Whatever numbers are assigned, assessing patent eligibility under 35 U.S.C. § 101 is a three-step process under *Alice/Mayo*, and this article characterizes the steps as 1, 2, and 3.

¹⁹⁷ 35 U.S.C. § 101.

¹⁹⁸ This is step one in the Supreme Court's numbering system. See *Alice*, 573 U.S. at 217; See also *In re Killian*, 45 F. 4th 1373 at 1379.

¹⁹⁹ MPEP § 2106.04 (9th ed. Rev. 7, 2022) (enumerating judicial exceptions).

²⁰⁰ This is step two in the Supreme Court's numbering system, and step 2B in the USPTO numbering. *Alice*, 573 U.S. at 217; MPEP § 2106.05 (9th ed. Rev. 7, 2022).

²⁰¹ MPEP § 2106.05 (9th ed. Rev. 7, 2022).

²⁰² *Alice*, 573 U.S. at 216 (observing that concern about pre-emption drives limitations on scope of patent eligibility).

²⁰³ See *supra* Section 1 (describing history of AI's basis techniques).

²⁰⁴ *Diamond v. Diehr*, 450 U.S. 175, 187 (1981).

is novel and nonobvious.²⁰⁵

B. Criticism and proposals for reform

Critics attack the *Alice/Mayo* approach as vague and subjective; producing unpredictable results; legally flawed; adding requirements beyond the statute; conflating § 101, § 102, and § 103 analysis; putting the U.S. at a competitive disadvantage; and disincentivizing innovation by narrowing the realm of patents.²⁰⁶ *Alice/Mayo* makes patents harder to get.²⁰⁷ The reform proposals and congressional testimony about them make it clear that advocacy reform is motivated by a desire to make certain types of patents more available. Section V.A reviews the economic argument for and against patents.

The leading proposal for reform is the Patent Eligibility Restoration Act (“PERA”),²⁰⁸ introduced in June 2023 by North Carolina Republican Representative Thomas Tillis and cosponsored by Delaware Democratic Senator Christopher A. Koons.²⁰⁹ The bill includes a finding that existing law has made patents less available. “Judicial exceptions to § 101 have made an increasing number of inventions ineligible for patent protection.”²¹⁰ It eliminates the judicial exceptions.²¹¹

In hearings held in January 2024 by a subcommittee of the Senate Judiciary Committee, seven of eight witnesses, including two former directors of the USPTO, testified in favor of the bill. Andrei Iancu, former director of the USPTO and now a partner at Sullivan and Cromwell testified in favor of the bill,²¹² lamenting:

patchwork of decisions over time, struggling to keep up with fast-changing technologies, has created significant confusion and uncertainty as to what is in and what is outside the bounds of the statute. These court decisions also have resulted in certain de facto rules—such that diagnostic techniques, for example, are generally not eligible for a patent.²¹³

²⁰⁵ Note the use, in this hypothetical example, of concepts of novelty, defined by § 102, and nonobviousness, defined by § 103, into the § 101 eligibility analysis: a conflation of distinct requirements much criticized, but probably inevitable. See discussion *infra* Section IV.B; see discussion *infra* Section VI.

²⁰⁶ Kennedy Stanley, *The Plot Thickens in the Convoluted Saga of Section 101 Eligibility: Where Do We Go From Here?*, 23 TUL. J. TECH. & INTELL. PROP. 137, 147–49 (2021) (summarizing criticisms of *Alice/Mayo*); Asay, *supra* note 3, at 1195, 1201–07 (concluding that patent law is limited in its utility as an incentive for AI innovation because of *Alice/Mayo*).

²⁰⁷ Lidiya Mishchenko, *Alice: Through the Formalist Looking Glass*, 97 J. PAT. & TRADEMARK. OFF. SOC’Y 214, 215 (2015) (observing that the *Alice* and *Mayo* decisions may have been motivated to reduce the number of “inferior patents,” being abused by “patent trolls.”).

²⁰⁸ Patent Eligibility Restoration Act of 2023, S. 2140, 118th Cong. (2023).

²⁰⁹ *Id.*

²¹⁰ *Id.* § 2(2).

²¹¹ *Id.* § 2(5)(A).

²¹² The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of Andrei Iancu, Partner, Sullivan & Cromwell LLP).

²¹³ *Id.* at 1–2.

David J. Kappos, former director of the USPTO and partner at Cravath, Swaine & Moore,²¹⁴ characterized PERA as “sensible and practical legislation that fixes the problems with 101 in a calibrated way, moving away from the unduly constrictive, un-administrable test we have been struggling with based on highly subjective concepts like ‘abstractness,’ to much more definable, specific and objective tests.”²¹⁵

Courtney C. Brinckerhoff, a patent attorney at Foley & Lardner who represents chemical, biotech, and pharmaceutical clients before the USPTO,²¹⁶ testified that “PERA would go a long way towards restoring not only patent-eligibility, but also predictability and confidence in patent rights that can better foster investment in innovations in life sciences technologies.”²¹⁷

Phillip S. Johnson testified on behalf of the Steering Committee of the Coalition for 21st Century Patent Reform, commonly referred to as “21C”, which is an organization that includes manufacturers such as: 3M, Bristol Myers Squibb, Boeing, Caterpillar, Eli Lilly, General Electric Aerospace, Johnson & Johnson, and RTX Corporation (the company formed by the merger of United Technologies and Raytheon).²¹⁸ “[T]he law of patent eligibility in this country is now a mess,” he said. “Congress should pass PERA to restore patent eligibility to its traditional scope by closing these judicially-created loopholes.”²¹⁹

Adam Mossoff, a professor of law at George Mason University, attacked *Alice/Mayo*²²⁰: “The Alice-Mayo inquiry has created a tremendous amount of uncertainty for innovators and severely restricted the patent eligibility of high-tech and biopharmaceutical innovations.”²²¹

He urged enactment of PERA.²²²

Mark Deem, Operating Partner at Lightstone Ventures, a firm that invests in innovative biotech and medtech companies, testified in favor of²²³ PERA as critical and much needed patent protection necessary as an incentive for continued innovation in his industries.²²⁴

²¹⁴ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of David J. Kappos).

²¹⁵ *Id.* at 2.

²¹⁶ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of Courtney C. Brinckerhoff, Partner, Foley & Lardner LLP).

²¹⁷ *Id.* at 3.

²¹⁸ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of Phillip S. Johnson, Esq.).

²¹⁹ *Id.* at 1.

²²⁰ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of Adam Mossoff).

²²¹ *Id.* at 2.

²²² *Id.* at 3.

²²³ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of Mark Deem, Operating Partner, Lightstone Ventures).

²²⁴ *Id.* at 2.

Two witnesses opposed the bill, one from the biotech industry, and one from the microprocessor industry. David W. Jones testified on behalf of the High Tech Inventors Alliance, which represents companies that manufacture microprocessors and telecommunications equipment.²²⁵ He explained:

why HTIA is skeptical that legislative changes are necessary to reform current patent eligibility jurisprudence. While advocates of PERA contend that current law is uncertain, the evidence tells a different story. Empirical studies and the Patent Office's own data show that the predictability of patent eligibility is on par with other major patent law doctrines.²²⁶

He opposed PERA:

PERA would create greater uncertainty while inappropriately expanding patent eligibility to non-technological innovation. PERA would, for the first time in our nation's history, amend the Patent Act to broadly allow the patenting of nontechnological innovation. The literal text of PERA would expand patenting and patent infringement liability to any area of human activity that uses or relies on any type of machine or manufacture—from computers and routers to pencils and paper—or to any subject matter that the patentee has simply chosen to claim in conjunction with using technology. The hard experience of the pre-Alice era shows that this would not only lead to absurd patents but would do real damage to the U.S. economy.²²⁷

The Electronic Frontier Foundation, though it did not testify in the hearings, opposes PERA:

The majority of high-tech patent lawsuits are brought by patent trolls—companies that exist not to provide products or services, but primarily have a business using patents to threaten others' work.
* * *

The Patent Eligibility Restoration Act, S. 2140, . . . would be a huge gift to patent trolls, a few tech firms that aggressively license patents, and patent lawyers. For everyone else, it will be a huge loss.²²⁸

A coalition of left-of-center organizations²²⁹ expressed support for existing patent eligibility law and opposition to PERA.²³⁰

²²⁵ The Patent Eligibility Restoration Act: Hearing Before the Subcomm. on Intellectual Property of the S.Comm. of the Judiciary, 118th Cong. 1 (2024) (Statement of David W. Jones, High Tech Inventors Alliance).

²²⁶ *Id.* at 2.

²²⁷ *Id.* at 6.

²²⁸ Joe Mullin, *This Bill Would Revive the Worst Patents on Software—And Human Genes*, ELECTRONIC FRONTIER FOUNDATION (EFF) (Sep. 15, 2023), <https://www.eff.org/deeplinks/2023/09/bill-would-boost-worst-patents-software-and-human-genes>.

²²⁹ Alex Moss, *Public Interest Organizations to Senate: Patent Eligibility Limits are Vital to Innovation, Prosperity, and Public Health*, PUBLIC INTEREST PATENT LAW INSTITUTE (Jan. 30, 2024), <https://www.piplus.org/news/public-interest-organizations-to-senate-patent-eligibility-limits-are-vital-to-innovation-prosperity-and-public-health>.

²³⁰ Wayne Brough, *Coalition Letter in Opposition to S. 2140 – Patent Eligibility Restoration Act of 2023 (PERA)*, R STREET (Jan. 30, 2024) <https://www.rstreet.org/outreach/coalition-letter-in-opposition-to-s-2140-patent-eligibility-restoration-act-of-2023-pera/>.

The fate of S.2140 is uncertain.

C. What the USPTO is doing

The USPTO published guidance on subject matter eligibility under § 101 in January 2019.²³¹ The USPTO published examples to be used with the 2019 guidance.²³² Example 39 pertains to an AI invention for a method for training a neural network for facial detection.²³³

The example concludes that the claim “recites a series of steps and, therefore, is a process,” which is a statutory category.²³⁴ It concludes that no judicial exception is recited because: “the claim does not recite any mathematical relationships, formulas, or calculations.”²³⁵

Efforts to use Example 39 to overturn examiner findings of non-eligibility have not fared well. “Claim 8 here is readily distinguishable from Example 39 both in terms of its data inputs and how the analysis is performed.” In Example 39, the data used to train the model are “digital facial images” to which transformations are applied including “mirroring, rotating, smoothing, or contrast reduction to create a modified set of digital facial images.” On the other hand, as discussed above, the Appellant does not specify how its algorithm is automatically adjusted or trained or is in any way similar to the claims of Example 39. “As best we can tell, the adjusting and applying comprise analyses that are abstract ideas.”²³⁶

In *Ex parte Aman Grover*,²³⁷ the PTAB distinguished Example 39 because the application on appeal recited use of a machine learning algorithm at a high level of generality to perform steps that could be performed by a human.²³⁸

In *Ex parte Laura Arroyo*,²³⁹ the PTAB likewise distinguished Example 39.

[T]he claimed invention here, including its adjustment of configurable parameters, can be performed entirely mentally or using pen and paper. . . . Further, claims reciting *training* a machine have been held ineligible under § 101 where, as here, the claimed invention did not improve existing technological or computer functionality, but simply used computers to automate work otherwise performed through human labor.²⁴⁰

Example 39 does not offer any help for applying step two of *Alice*, where most of the recent AI applications have run aground.

²³¹ See 2019 Revised Patent Subject Matter Eligibility Guidance, 84 Fed. Reg. 50 (Jan. 7, 2019).

²³² See generally *Subject Matter Eligibility Examples: Abstract Ideas*, USPTO (Jan. 7, 2019), https://www.uspto.gov/sites/default/files/documents/101_examples_37to42_20190107.pdf.

²³³ *Id.* at 8, ex. 39.

²³⁴ *Id.* at 9, ex. 39.

²³⁵ *Id.*

²³⁶ *Ex parte Arten Petakov*, No. 2022-004039, 2023 WL 4842002, at *14 (P.T.A.B. July 20, 2023).

²³⁷ *Ex parte Aman Grover et al.*, No. 2020-003327, 2020 WL 6055291 (P.T.A.B. Oct. 8, 2020).

²³⁸ *Id.* at *7–8.

²³⁹ See *Ex parte Laura Arroyo et al.*, No. 2023-001360, 2024 WL 2292735 (P.T.A.B. May 20, 2024).

²⁴⁰ *Id.* at *6.

Despite the modest effect of Example 39, the USPTO regularly grants patents on AI inventions and is receiving more and more applications. The USPTO has granted seventy-seven patents using the term “generative AI” between January 1, 2024, and the time of this writing (June 4, 2024).²⁴¹

An April 29, 2024, advanced search in USPTO’s Patent Public Search database²⁴² with the term “artificial intelligence” resulted in 106,061 records, representing issued patents. The same search run against the published applications database resulted in 154,378 records. An advanced search of the issued patents database with the term “generative AI” produced 33 results, and the same search against the published applications database produced 203 results.²⁴³

Internationally, the upsurge in AI patents is even more pronounced.²⁴⁴

New generative AI inventions qualify for patents if they meet the tests for patentability under § 101,²⁴⁵ as interpreted in *Alice/Mayo*.²⁴⁶ AI applications are computer programs, embedded in powerful computing hardware. The complexity of the software and the power (and power consumption) of the hardware is singular and largely unprecedented.²⁴⁷ But the legal tests for patent eligibility are more than 200 years old. The generative AI invention must involve a machine, process, item of manufacture or compound; it must be more than a mere algorithm or theory or phenomena of nature.²⁴⁸ It helps assure patentability when an idea is expressed in detailed pseudocode or actual programming code for computers, and it helps when the invention involves implementation of software on a physical piece of hardware. But the innovative content of most AI inventions is in the software; not the hardware. That forces the USPTO and the courts away from the anachronistic machine or transformation test²⁴⁹ and into an exploration of other indicia of originality.

It is clear that prompts for an AI system can satisfy the conception requirement.²⁵⁰

²⁴¹ Google Patents Search for “Generative AI”, GOOGLE, (search done on June 4, 2024) (screenshot on file with author).

²⁴² *Patent Public Search*, (April 17, 2024), <https://ppubs.uspto.gov/pubwebapp/>.

²⁴³ See Datzov, *supra* note 8 (urging temperance in calls for legislative reform to make AI patents easier to get; reviewing upsurge in patent applications mentioning AI).

²⁴⁴ See 2024 WIPO REPORT, *supra* note 18, at 7 (reporting that number of patent families involving generative AI has grown from 733 to 14,000 in the ten years preceding 2023 and that they have increased more than 800% since the introduction of transformers in 2017).

²⁴⁵ 35 U.S.C. § 101.

²⁴⁶ See *supra* Section IV, discussing *Alice/Mayo* test for eligibility of subject matter under 35 U.S.C. § 101.

²⁴⁷ See modeling programs for meteorology and nuclear weapons.

²⁴⁸ The “more than” here represents the “inventive step” of *Alice/Mayo*. See *supra* Section IV.A.3.

²⁴⁹ Mark A. Lemley, et. al., *Life After Bilski*, 63 STAN. L. REV. 1351, 1338–39 (2011) (criticizing machine-or-transformation test as a “failure.”).

²⁵⁰ Inventorship Guidance for AI-Assisted Inventions, 89 Fed. Reg. 10043, 10048 (Feb. 13, 2024) (Guiding Principle No. 2) (“A significant contribution could be shown by the way the person constructs the prompt in view of a specific problem to elicit a particular solution from the AI system.”).

U.S. Patent No. 11,941,378 B1 *System and Method for Utilizing Production Insights in Generative AI Models* covers a method and system for evaluating runtime computer code, compiled from base code (the source code).²⁵¹ The technical field of the patent is identified as “generative AI in general and . . . utilizing production insights by a generative AI model in particular.”²⁵²

The invention develops human readable comments (“production insights”) pertaining to code elements such as functions or code blocks and records them in appropriate places in the base code.²⁵³ The insights may report data such as the error rate of the code element when executed, the execution time of the code element, the indication frequency of the code element, and/or a resource utilization metric.²⁵⁴ One embodiment utilizes generative AI to debug the code base, to code a new feature into the code base, to revise the code base to improve performance of the runtime code, and to provide suggestions regarding source code optimization.²⁵⁵

Neither the claims nor the specification refers to any particular hardware for running the invention. Indeed, one paragraph of the specification notes that the invention may run on a general-purpose computer.²⁵⁶ The intricate detail with respect to how the production insight is transferred between the source code and the run-time file is the only apparent stab at an inventive concept. Little detail is provided on how the productive insight is developed.

U.S. Patent No. 11,908,449 B2 *Audio and Video Translator* uses AI to translate audio and video that accounts for variations such as a motion, pacing, idioms, sarcasm, jokes, tone, and phonemes.²⁵⁷ The invention represents a relatively modest improvement over existing language translation applications by using nonverbal data obtained from video and audio files to improve the translation of transcripts also provided as inputs.²⁵⁸ Nowhere in the claims or specification is any reference made to any specific hardware configuration for execution of the invention. The only aspect of the specification that appears to satisfy the inventive concept in the third step of § 101 analysis is the level of detail provided with respect to the processing of video images and sounds.

U.S. Patent No. 11,958,410 B2 *Artificially Intelligent Mobility Safety System* uses camera imagery, a radar sensor, a Lidar sensor, a proximity sensor, an inertial measurement unit, GPS data, and at least one trained model to warn bicycle riders of hazards pertinent to vehicle operation in a traffic environment.²⁵⁹

The specification contains several paragraphs describing devices that constitute

²⁵¹ U.S. Pat. No. 11,941,378 (filed Oct. 19, 2023) (issued Mar. 26, 2024).

²⁵² *Id.* at col. 1 l. 7–9.

²⁵³ *Id.* at col. 5 l. 32–55.

²⁵⁴ *Id.* at col. 23 l. 8–20.

²⁵⁵ *Id.* at col. 23 l. 55–62.

²⁵⁶ *Id.* at col. 21 l. 15–22.

²⁵⁷ U.S. Pat. No. 11,908,449 abstract (filed Nov. 29, 2022) (issued Feb. 20, 2024).

²⁵⁸ *Id.* at col. 25 l. 55–col. 26 l. 20 (noting audio files and transcripts as inputs).

²⁵⁹ U.S. Pat. No. 11,958,410, at [57] (issued 16 Apr. 2024).

the system such as sensors, computers, batteries, and physical enclosures. It provides a fair amount of detail about the capabilities of each of these,²⁶⁰ identifying horns and sirens as alerting mechanisms, three-axis accelerometers, three-axis gyroscopes, three-axis magnetometers, CMOS cameras, and onboard computers with “microprocessors, memory, input output interfaces, and other components required for operation.”²⁶¹

It refers to machine learning as a way of developing risk estimates from sensor data.²⁶² It explicitly refers to convolutional neural networks, RNNs, and transformers, as elements of the desirable machine learning system.²⁶³ The specification describes how the system would alert a bicycle rider to various scenarios including an approaching vehicle from behind, the bicyclist’s approaching pedestrians, the bicyclist pursued by a canine, lane deviation, intersection collision avoidance, and emergency stop.²⁶⁴ The method claims are relatively indifferent to the hardware on which they are implemented,²⁶⁵ but the apparatus claims are tied directly to features of the hardware.²⁶⁶

This patent illustrates the strongest inventive concept of the five reviewed.

U.S. Patent No. 11,947,902 B1 *Efficient Multi-Turn Generative AI Model Suggested Message Generation* covers systems and methods for using generative AI to construct suggested replies to messages such as text messages or email messages.²⁶⁷ Figures 4 through 7 illustrate the logical sequence of steps in deciding how to respond to a message. The specification explains how generative AI models work and how they might be applied to draft appropriate responses to messages.²⁶⁸ The claims cover the system described in terms of its functions and purposes rather than its technical characteristics.²⁶⁹ The method claims likewise are indifferent to hardware specifics.²⁷⁰ This patent seems to have dodged a bullet because of the weakness of its disclosure or claim of any inventive concept.

U.S. Patent No. 11,875,240 B1 *Tuning Generative Artificial Intelligence Model* covers a system and methods for automating the iterative process of submitting successive prompts to a generative AI model.²⁷¹

The technical field description in the specification places the invention in the category of “management of generative artificial intelligence models.”²⁷² The

²⁶⁰ *Id.* at col. 4–9.

²⁶¹ *Id.* at col. 7 l. 36–38.

²⁶² *Id.* at col. 22.

²⁶³ *Id.* at col. 15 l. 21–37, col. 16 l. 11–17.

²⁶⁴ *Id.* at col. 24–26.

²⁶⁵ *See id.* at col. 38–39.

²⁶⁶ *See id.* at col. 39–40.

²⁶⁷ U.S. Pat. No. 11,947,902 fig. 4–7 (filed Mar. 3, 2023) (issued 2 Apr. 2024).

²⁶⁸ *Id.* at col. 4.

²⁶⁹ *Id.* at col. 29–30.

²⁷⁰ *Id.* at col. 30–32.

²⁷¹ U.S. Pat. No. 11,875,240 (filed July 25, 2023) (issued 16 Jan. 2024).

²⁷² *Id.* at col. 1 l. 6–7.

specification explains how the natural language processing model works, including the use of transformers and other accepted techniques.²⁷³ It does not, however, refer to any particular hardware implementation. The system claims provide no concrete details on hardware.²⁷⁴ The details on the working of the natural language model present a strong inventive concept.

This sample of five recent AI patents shows that the machine or transformation test embraced in *Bilski v. Kappos*²⁷⁵ is a weak reed when it comes to AI patents. That is as it should be. The machine or transformation test was a bridge, an anachronistic one, to a world in which innovation involved physical machines. Such a test is irrelevant in a world where innovation mostly involves new ways to use computers—and new ways to understand the elements of life forms.

These five patents and the others involving AI inventions show that *Alice/Mayo* is not making it impossible—or even particularly difficult to get patents on this new subject matter.

The sample of recent AI patents also illustrates the application of the traditional criteria for patentability to new technologies. Examination of AI inventions does not only not present new problems, but it also provides an opportunity to cut through the chaff of debate over *Alice/Mayo* and rediscover basic patentability principles.

The PTAB has considered appeals of rejections of inventions involving assertions of AI in two-dozen cases reported since the beginning of 2024.²⁷⁶ Several of the cases involved rejections under §§ 102, 103, and 112. In only one, did the applicant survive explicit § 101 scrutiny. In *Ex parte Kristin E. McNeil et al.*, the PTAB, reversing the examiner, determined that claims were not directed to a judicial exception under Step 2A, Prong Two:

The Examiner's determination that the testing and determining steps (steps [v] and [vi]) of claim 1 recite mental processes is insufficient to support the rejection.

* * *

We determine the testing and determining steps (steps [v] and [vi]) of claim 1 are not merely part of the abstract idea; rather, they are additional steps performed *in the context of a trained cognitive computing system* that cannot be performed by human thought alone or by a human using pen and paper.²⁷⁷

²⁷³ *Id.* at col. 13–14.

²⁷⁴ *See id.* at col. 22–24.

²⁷⁵ *See Bilski v. Kappos*, 561 U.S. 594, 604 (2010) (disagreeing with Federal Circuit and holding that machine-or-transformation text, though a useful and important clue, is not the only test for patent eligibility of a process claim).

²⁷⁶ *See, e.g., Ex parte Tim Frandsen et al.*, No. 2023-000756, 2024 WL 959418, at *11 (P.T.A.B. Mar. 5, 2024) (affirming rejection under section 101; “the elements that Appellant identifies are, in combination, the abstract idea itself; they are not additional elements to be considered in determining whether claim 4 includes an additional element or combination of elements that integrates a recited abstract idea into a practical application.”).

²⁷⁷ *Ex parte Kristin E. McNeil et al.*, No. 2023-000011, 2024 WL 197214, at *15 (P.T.A.B. Jan. 17, 2024).

The other cases involved a judicial exception and nothing more. *Ex parte Melinda Min Xiao-Devins* involved a § 101 judicial-exception rejection.²⁷⁸ The PTAB agreed with the examiner that an invention of a computerized personal assistant merely constructed a use task list, which amounts to “managing personal behavior.”²⁷⁹ It rejected the applicant’s argument that the invention’s improvements in electronic communications infrastructure made it patent eligible under Step 2 of *Alice/Mayo*.²⁸⁰

In *Ex parte Corville O. Allen*, the PTAB applied step two, saying that:

[T]he claim recites no technical or technological details on how the ingesting is performed such that not re-ingesting is a technical improvement or an improvement to the functioning of the computer. And to the extent the “expenditure of computer resources” is reduced, this reduction is a business improvement in terms of cost and use, and not a technical or technological improvement to the computer elements themselves.²⁸¹

In *Ex parte Kevin Lyman et al.*, the PTAB similarly rejected a claim at step two:

The claimed invention does not provide a technical improvement as the computing devices are used as a tool to implement the abstract idea.

* * *

Merely requiring generic computer implementation does not move into § 101 eligibility territory.²⁸²

In *Ex parte Eduardo Cardoso et al.*, the PTAB said more about the flaw in all of these rejected applications:

Taking the claim elements separately, the operation performed by the computer at each step of the process is expressed purely in terms of results, devoid of implementation details. Steps 1-4 recite generic computer processing expressed in terms of results desired by any and all possible means and so present no more than conceptual advice.

Steps 1-3 recite using an artificial intelligence component, but do not recite any technological implementation details as to how this is done or how the artificial intelligence component is constructed or operates. All recitations are no more than an instruction to apply a black box to labelled data inputs.

Limitation 5 recites that the artificial intelligence component is

²⁷⁸ *Ex parte Melinda Min Xiao-Devins*, No. 2023-003736, 2024 WL 2151074, at *1 (P.T.A.B. May 13, 2024).

²⁷⁹ *Id.* at *5.

²⁸⁰ *Id.* at *6 (internal quotations omitted).

²⁸¹ *Ex parte Corville O. Allen*, No. 2022-003311, 2024 WL 339904, at *13 (P.T.A.B. Jan 26, 2024).

²⁸² *Ex parte Kevin Lyman et al.*, No. 2023-000805, 2024 WL 621157, at *8 (P.T.A.B. Feb. 12, 2024) (internal quotations and citations omitted).

trained using historical employee off-boarding data to create a specialized machine learning model that determines the employee retention insights and the action steps for respective employers for particular employee types. This again simply recites applying some labelled data as input to a black box with no technological implementation details. The recitation of a specialized machine learning model that determines the employee retention insights and the action steps for respective employers for particular employee types is a conclusory description as a “specialized learning model” with no recitation of how it is so specialized or how the model is constructed or operates, followed by the aspirational objective for the model with no recitation as to how that objective is achieved.

Limitation 6 recites that the specialized machine learning model increases performance and accuracy regarding analytical and predictive capabilities of the artificial intelligence component thereby increasing performance of the computer, itself. This is likewise a conclusory recitation with no recitation of how the performance and accuracy are improved or how computer performance is improved.

All purported inventive aspects reside in how the data are interpreted and the results desired, and not in how the process physically enforces such a data interpretation or in how the processing technologically achieves those results.²⁸³

At least one of the PTAB cases rejected the application, not only under § 101, but also under § 112. This shows how § 112 can properly be applied when a patent applicant makes a generic claim to AI or to machine learning. In *Ex parte David Gustafson et al.*, the PTAB explained that a “black box” limitation in a specification pertaining to machine learning does not satisfy the requirements of § 112 because it does not show that the inventors possessed the invention.²⁸⁴

Section 101 eligibility analysis overlaps considerably with § 112 definiteness analysis. In *Berkheimer v. HP Inc.* (“*Berkheimer*”),²⁸⁵ the Federal Circuit explained that “[t]he second step of the *Alice* test is satisfied when the claim limitations ‘involve more than performance of well-understood, routine, [and] conventional activities previously known to the industry.’”²⁸⁶ It explained application of that test by analogy to application of § 112.²⁸⁷

²⁸³ *Ex parte* Eduardo Cardoso et al., No. 2022-003546, 2023 WL 7647584, at *7–8 (P.T.A.B. Nov. 14, 2023).

²⁸⁴ *Ex parte* David Gustafson et al., No. 2020-000900, 2020 WL 4673713, at *3 (P.T.A.B. Aug. 3, 2020) (“[T]he specification itself must explain how the claimed function is achieved to demonstrate that the applicant had possession of it.”).

²⁸⁵ *Berkheimer v. HP Inc.*, 881 F.3d 1360, 1370–71 (Fed. Cir. 2018) (vacating invalidation of certain claims as ineligible under section 101).

²⁸⁶ *Id.* at 1367 (quoting *Content Extraction and Transmission LLC v. Wells Fargo Bank, Nat. Ass’n*, 776 F.3d 1343 (Fed. Cir. 2014)).

²⁸⁷ *Id.* at 1368. This is yet another example of the criticized but inevitable conflation of § 101 eligibility

In April 2018, the USPTO issued a memorandum explaining how the USPTO will apply *Berkheimer*.²⁸⁸ It said that the analysis of whether an element is widely prevalent or in common use for purposes of applying *Alice/Mayo* step 2 is the same as the analysis under § 112(a) as to “whether an element is so well-known that it need not be described in detail in the patent specification.”²⁸⁹

The PTAB elaborated on this step for a mundane computer implementation:

[T]he relevant question under Step 2A, Prong Two²⁹⁰ is not whether the claimed invention itself (i.e., extracting information from a communication channel of a UCaaS platform and processing the information using machine learning models to develop a user task list) is a practical application. Instead, the question is whether the claim includes additional elements beyond the judicial exception that integrate the judicial exception into a practical application. Here, the only additional elements recited in claim 1 beyond the abstract idea are “a first machine learning model” and “a second machine learning model,” i.e., generic computer components.²⁹¹

D. Does the status quo allow patents for AI inventions?

Analysis of recently granted patents involving AI, and recent PTAB cases suggest that current law is entirely workable as applied to AI inventions.²⁹² Patent examiners apparently are allowing patents for AI inventions that particularize how they advance the art, and the PTAB opinions make coherent distinctions between patentable AI and purported AI that represents little more than a collection of known algorithms and techniques.

Nikola Datzov²⁹³ suggests considering patent eligibility with respect to three layers:²⁹⁴

- The data layer, including training databases, testing and validation data, and production input and output data.²⁹⁵
- The application (software) layer, including deep-learning algorithms, he says,²⁹⁶ but his explanation relates more to traditional if-then procedural programming than to machine learning.

analysis with other aspects of patentability. Compare *supra*, Section II.B, with *infra*, Section VI.

²⁸⁸ Memorandum from Robert W. Bahr, Deputy Commissioner For Patent Examination Policy, to Patent Examining Corps. (Apr. 19, 2018) (available at <https://www.uspto.gov/sites/default/files/documents/memo-berkheimer-20180419.PDF>).

²⁸⁹ *Id.* at 3.

²⁹⁰ See MPEP § 2106.05(b) (9th ed. Rev. 7, 2022) (Particular Machine).

²⁹¹ *Ex parte* Melinda Min Xiao-Devins, No. 2023-003736, 2024 WL 2151074, at *6 (P.T.A.B. May 13, 2024).

²⁹² See Datzov, *supra* note 8, at 41–42 (summarizing numerous studies as supporting conclusion that AI patenting activity is flourishing both in the U.S. and globally).

²⁹³ *Id.* at 3 (urging caution in amending patent law to relax the *Alice/Mayo* requirements in order to promote more innovation in AI).

²⁹⁴ *Id.* at 24.

²⁹⁵ *Id.* at 25.

²⁹⁶ *Id.* at 25–26.

- The system (hardware) layer, comprising computer processors, memories, input-output devices, voice-to-text translators, fingerprint scanners, and humanoid robots.²⁹⁷

“The argument for patent ineligibility of data layer AI inventions is likely the strongest of all three layers.”²⁹⁸ Nevertheless, methods for collecting, organizing, storing, manipulating, and presenting data may be patent eligible.²⁹⁹

For his application layer, the biggest barrier to patent eligibility is that:

Much of AI software is focused on emulating human behavior, which is precisely within the sights of the abstract idea exception. This far from dooms all AI software claims into ineligibility, but it requires focusing the patent claims on *how* to mimic human behavior rather than attempting to broadly capture mimicking some human behavior through software in functional terms (i.e., broad functional claiming). For example, inventors have found success in patenting claims for mimicking human behavior by writing claim language with sufficient specificity. Another popular category of AI software inventions has focused on image recognition AI software to increase efficiency and accuracy for quickly making determinations through visual input. The key to the eligibility of such inventions is to focus on the claim language and how the invention is performing the activity or making the decisions, rather than focusing on what function or objective the software is fulfilling.³⁰⁰

The third category is presumptively eligible: “AI hardware systems are squarely within the scope of patent eligible subject matter.”³⁰¹ But, he says, “[u]nlike AI inventions that invent a new sensor, a new camera, or a new microphone, AI inventions that rely on existing hardware can often be characterized as utilizing well-understood and conventional technology, and as such, the hardware components of such claims fade into the background.”³⁰²

USPTO guidance, combined with its jurisprudence in deciding patentability of AI inventions, has created an interesting dilemma for drafting claims. On the one hand, the claims must be detailed enough about concrete techniques for using and deploying AI to satisfy the requirements for an inventive concept under the third step of § 101 analysis. On the other hand, the claims must not use poisonous language such as reciting mathematical steps, mental processes, or organization human activity, which would trigger a judicial exception, as counseled by the 2019 Example 39. Jon Grossman, of Blank Rome LLP, offers suggestions on how to deal with the

²⁹⁷ *Id.* at 26; see also 2024 WIPO REPORT, *supra* note 18, at 16 (reporting that AI patents can be addressed to (1) computer programs used, or models; (2) type of input of output, or modes; or (3) applications of generative AI).

²⁹⁸ Datzov, *supra* note 8, at 27.

²⁹⁹ *Id.*

³⁰⁰ *Id.* at 28–29.

³⁰¹ *Id.* at 29.

³⁰² *Id.* at 30.

dilemma.³⁰³

V. Will the creators of AI inventions rely on patent law rather than continuing to rely on trade secret law?

Even if a patent disclosure is perfectly congruent with the type of transparency desired for AI patent disclosure, it will not meet the demand for transparency if AI developers do not apply for patents. Public policy or economic reality must cause them to shift their preferences from trade protection or first-mover advantage to patent law to protect their innovations.

Innovators have generally been wary of transparency requirements because they want to protect their trade secrets in methods and strategies that give them a competitive advantage in the marketplace. Trade secret protection is only available so long as they keep these methods secret.³⁰⁴ They may be entitled to copyright protection in their specific computer code, but copyright does not extend to ideas or algorithms.³⁰⁵ Many of the systems may be entitled to patent protection, and that would indeed protect the generative AI entrepreneurs against misappropriation of their methods and ideas. But obtaining a patent costs tens of thousands of dollars, and it typically takes two to three years before one is granted.

AI inventors will shift their preference for intellectual property protection to patent law away from trade secret law only if they have greater confidence that they can obtain patents and only if the demands for regulation of AI—in particular the demands for transparency of AI innovations—convince them that they will lose their trade secret protection and need patent protection to protect themselves against piracy.

A. Economic arguments

The debate over the efficacy of patent protection historically has involved arguments about whether patents serve as incentives for innovation that would not

³⁰³ Jon Grossman makes four concrete suggestions:

- A. Draft a claim that recites the specific function(s) or the improvement(s) explicitly tied to the AI features.
- B. Draft a claim explicitly reciting the AI technology.
- C. Draft a claim and a specification that do not merely improve the abstract idea of the claimed AI technology, but can directly associate the AI technique with improved hardware performance.
- D. Draft a specification that discloses and supports hardware tied to the claimed AI technology not in terms of listing that hardware as generic components but as an improvement or a solution for a problem tied to the performance of such components.

Jon Grossman, *AI Inventions and Subject Matter Eligibility*, 35 INTELL. PROP. & TECH. L. J. 1, 8 (2023).

³⁰⁴ See generally HENRY H. PERRITT, JR., TRADE SECRETS FOR THE PRACTITIONER § 6:2 (2d ed. 2005 & Supp. 2022) (analyzing secrecy requirement for trade secret protection).

³⁰⁵ 17 U.S.C. § 102(b); See generally Henry H. Perritt, Jr., *Copyright for Robots?*, 57 IND. L. REV. 139 (2023) (analyzing possibility of copyright protection for generative AI).

occur otherwise.³⁰⁶ Those who claim positive incentive effect marshal the traditional argument that protection is necessary to prevent free riding and to reduce the inconvenience and uncertainty of trade secret protection. Those who claim negative effects argue that the risks of patent infringement or accusations of patent infringement deter innovation and that the cost, delays, and uncertainties of patent protection undermine its theoretical positive incentives.

1. *Risk of free riding*

An innovator always is free to put his invention in the public domain. If he does so, however, he is leaving himself open to free riding: the phenomenon in which a competitor can copy his invention and sell it at a cost lower than the inventor's total cost, taking into account the inventor's cost of developing and refining the invention. Economic theory says that inventor and the free rider have comparable marginal costs, but the free rider has no fixed costs associated with development, while that category of costs may be substantial for the inventor.

William Landes and Richard A. Posner explain, in more formal microeconomic terms, why producers have difficulty recovering their fixed costs of research and development when their products embody inventions that are easy to copy.³⁰⁷ Pirates copy the products and appropriate the inventions without incurring the development costs and thus deprive the producers of a price that fully covers their investments. An invention shifts the inventor's cost curve downward, assuming that the effect of the invention is to reduce the cost of production or distribution. Absent patent protection, the inventor enjoys that amount of cost advantage in the marketplace until his competitors—existing or new—appropriate the invention. Once they appropriate it, they have cost curves even lower than that of the inventor because they need not recover the development cost of the invention.³⁰⁸

In 1962, Kenneth Arrow examined the need for government intervention to facilitate the production of innovation by easing transfers of knowledge.³⁰⁹

2. *Patent protection and its detractors*

Patents protect inventors against free riding and thus are an incentive to invent. Simple microeconomic supply and demand models show that the availability of patent protection will increase the supply of inventions.³¹⁰ Many commentators

³⁰⁶ “The age-old debate on whether a narrower scope of patent eligibility serves to promote or stifle innovation has now carried into the AI Revolution.” Datzov, *supra* note 8, at 35 (summarizing arguments for and against patent protection as an incentive for innovation).

³⁰⁷ See William Landes & Richard A. Posner, *The Economic Structure of Intellectual Property Law* 294 (2003).

³⁰⁸ *Id.* at 294–97.

³⁰⁹ Kenneth J. Arrow, *Economic Welfare and the Allocation of Resources for Invention*, in *The Rate and Direction of Inventive Activity*, NAT'L BUREAU OF ECON. RESEARCH 609, 623 (1962) (out of print) (available at <https://www.nber.org/system/files/chapters/c2144/c2144.pdf>).

³¹⁰ Landes and Posner offer a simple supply and demand model and graph showing that patent produces no significant change in industry output or price but allows the patent holder to maximize its profits either by capturing the entire industry output or by earning royalties by licensing the patent to

defend patents on the traditional ground that they incentivize innovation.³¹¹

Factors other than direct incentives for inventors matter. Industry structure and custom in the industry matter. Inventors wanting to commercialize in an industry where patent protection is pervasive, as in the biosciences or semiconductor fields are likely to find investors unwilling to invest unless they see patent protection.³¹²

On the other side, a number of empirical studies seem to show that assertion of patent protection discourages innovation. Certainly this appears to have been the case with respect to aeronautical technology in the decades after the Wright brothers flight³¹³ and with respect to sewing machine technology almost 100 years before.³¹⁴ More recently, concerns with patent thickets and the transaction costs of patent pools as a solution have given rise to a phenomenon of defensive patenting which has little to do with the original concepts of patents as incentives for innovation.³¹⁵ Many serious students of intellectual property in both law and economics question whether the patent system is fulfilling its goals.³¹⁶ Other, less serious, attacks on the patent

existing firms. LANDES AND POSNER, *supra* note 307, at 297–298 and fig. 11.1.

³¹¹ See Maureen K. Ohlhausen, *Patent Rights in a Climate of Intellectual Property Rights Skepticism*, 30 HARV. J. LAW & TECH. 1, 4 (2016) (finding “ample evidence that patents serve a materially valuable role in promoting innovation in at least some settings.”). In the absence of patent protection, an inventor will protect herself against this free riding by keeping her invention secret, thus depriving society of its benefits. LANDES AND POSNER, *supra* note 307, at 294; Ebrahim, *supra* note 152, at 188 (“[P]atent law’s disclosure function is a way for the inventor to reveal the invention to society.”). Thus, patent law grants a monopoly to protect the inventor against free riding, to induce her to give up her trade secret, but only a temporary one, accompanied by a requirement that the details of the invention be fully disclosed to the whole world immediately and fully in the public domain after the patent expires. LANDES AND POSNER, *supra* note 307, at 294–95.

³¹² Landes and Posner argue that the justification for patents varies by industry. LANDES & POSNER, *supra* note 307, at 326–27. “Whether a given degree of patent protection is socially desirable depends on the patentee’s fixed costs, the inherent difficulty in inventing around the patent . . . , and the extra profits that the patentee can expect to receive from greater protection.” *Id.* at 300. Inventors and venture capitalists in the Internet communication sector may find venture capitalists and potential commercial partners uninterested in patents. The evidence is mixed on whether patent protection attracts investors, notwithstanding the belief in some quarters that start up enterprises will not be able to obtain capital unless they can demonstrate intellectual property protection for their products and services. Compare Datzov, *supra* note 8, at 46 (questioning assertion that investor will withhold capital from those without patent rights, with some analysis), with Tschider, *supra* note 124, at 717, 721 (asserting that patents valuable to investors; no analysis).

³¹³ See Herbert A. Johnson, *The Wright Patent Wars and Early American Aviation*, 69 J. AIR L. & COM. 21, 43 (2004) (describing Orville Wright’s buyout of investors, sale of company, and merger with Martin); *The Wright/Smithsonian Controversy: The Patent Pool*, THE WRIGHT BROTHERS AEROPLANE COMPANY: A VIRTUAL MUSEUM OF PIONEER AVIATION, https://www.wright-brothers.org/History_Wing/History_of_the_Airplane/Doers_and_Dreamers/Wright_Smithsonian_Controversy/08_The_Patent_Pool.htm (last visited Apr. 13, 2024) (discussing NACA’s encouragement of patent pool, Wright-Martin Company’s reluctance to enter pool, and postwar criticism of pool as barrier to innovation and competition).

³¹⁴ See Adam Mossoff, *The Rise and Fall of the First American Patent Thicket: the Sewing Machine War of the 1850s*, 53 ARIZ. L. REV. 165, 167 (2011).

³¹⁵ See Henry H. Perritt, Jr., *Drowning in Patent Pools: Is Statutory Invention Registration a Lifeguard?*, 127 W.VA. L. REV. (2024).

³¹⁶ See Ohlhausen, *supra* note 311 (identifying major critics of patent system and summarizing their views). “Historical evidence suggests that in countries with patent laws, the majority of innovations

system are shrill.³¹⁷

Without a doubt, the economic evidence is mixed on whether patent protection promotes innovation or restrains it.³¹⁸ One reason the verdict on this inquiry is so indefinite is that innovation is hard to measure.³¹⁹ The truth probably is that some inventors use patent protection to block competitive innovation by others; that patent protection in some fast-evolving industries is irrelevant because it takes too long to obtain, is too expensive, and is too uncertain; that some patent holders stop innovating and concentrate on suing innovators for infringement;³²⁰ and that patents in some industries facilitate innovation because they accelerate the dissemination of knowledge about innovation and they are licensed on sophisticated terms that encourage licensees to practice the patented inventions.

Analysis also is complicated by the multiplicity of firms involved in making patents into economically successful products. Most inventors must seek support from other entities to commercialize their patents.³²¹

Despite all the controversy, patent protection is a reality, and any substantial modification of the current patent statutes in the U.S. to weaken their effect is unlikely politically.

The debate over § 101 eligibility arguments is fundamentally a debate over whether more patents are a good idea. The advent of AI and widespread concerns about it may provide a new justification for the existence of patent law as an engine of transparency.

3. *Trade secret alternative*

A patent monopoly, however, is not the only way to protect innovators from free riding and thus to induce them to invent. Trade secret protection is an alternative. To protect himself against free riding, an inventor may keep the details of his invention

occur outside of the patent system. . . . In industries where secrecy was effect, inventors were less likely to use patents.” Petra Moser, *Patents and Innovation: Evidence from Economic History*, 27 J. ECON. PERSPECTIVES 23, 40 (2013).

³¹⁷ See generally MICHELLE BOLDRIN & DAVID K LEVINE, *AGAINST INTELLECTUAL MONOPOLY* (2008). See also Toshiko Takenaka, *Patents for Sharing*, 26 MICH. TECH. L. REV. 93, 115 (2019) (arguing that traditional utilitarian theory of patents to allow inventor to recover the cost of invention and keep free-riders at bay no longer applies, given changes in technology and market structures).

³¹⁸ See JACOB SCHMOOKLER, *PATENTS, INVENTION, AND ECONOMIC CHANGE* 72–77 (1972) (evaluating data and inferring that the main stimulus for invention is the level of expenditure on capital goods in an industry, leading invention rather than lagging it).

³¹⁹ *The Trouble with Measuring Innovation*, WALL STREET JOURNAL, <https://deloitte.wsj.com/cio/the-trouble-with-measuring-innovation-01671231403> (last visited Apr. 22, 2024); Jennie Bjork, *Measuring Innovation Effectively—Nine Critical Lessons*, 66 RSCH. TECH. MGMT. 17, 18 (2023) (noting difficulty of measurement).

³²⁰ Significant evidence suggests that the decline of the Wright Company resulted from the brothers devoting their attention to patent infringement litigation instead of more technological innovation. “Orville Wright’s preoccupation with patent litigation doomed the Wright’s fate in the field of aircraft construction.” Johnson, *supra* note 313, at 48.

³²¹ *Id.* (using Wright Brothers’ early lack of success to illustrate the inability of most inventors to finance large scale production of pioneering inventions).

secret, while still being able to sell it. The law has developed the doctrine of trade secret protection to allow him to do this.³²²

Innovators choose between protecting their inventions by patent or trade secret according to a number of different considerations.³²³ A number of observers believe that innovators in the information technology field prefer trade-secret protection.³²⁴ Clark Asay has explored the relative effectiveness of patent protection and trade secret protection as incentives for innovation in AI.³²⁵

If the rate of innovation is high, patent protection may not be satisfactory because it takes too long to get a patent. By the time a patent issues, five or six years after it is applied for, the technology claimed has been superseded by better technology.

If trade secrets are difficult to reverse engineer from products in public markets, trade secret protection may be satisfactory;³²⁶ patent protection against innocent imitation is not needed. Conversely if reverse engineering is easy, trade secret protection does not offer much after products are introduced to the public.

Trade secret protection is weaker, the larger the number of different people who must know the trade secret. Not everyone honors commitments made under nondisclosure agreements and the more people who should be bound by them, the more the risk of breach and disclosure. Conversely if only a handful of core venturers who trust each other need to know the trade secret it offers good protection.

Whether the trade secret is likely to qualify for a patent makes a difference, as well. If the technology and the product anticipated to come from it are close to the boundary of eligible subject matter under § 101 of the Patent Act,³²⁷ an inventor should be reluctant to incur the cost of seeking a patent and the certainty of disclosure and extension of the trade secret when the patent application is published. Likewise, if prior art is robust, the likelihood is high that the patent office will find the invention anticipated³²⁸ or obvious.³²⁹ Small but remunerative refinements in prior art may best be commercialized and protected by being kept secret.³³⁰

³²² See generally Henry H. Perritt, Jr., *Trade Secrets for the Practitioner* (3d ed. 2024).

³²³ See generally *id.* at ch. 2 (analyzing trade secret protection in comparison to alternatives such as patents).

³²⁴ Ebrahim, *supra* note 152, at 183–94. (Internal footnotes omitted) (nothing tendency for innovators to shift to trade-secret, rather than patent, protection); Tschider, *supra* note 124, at 710–713 (noting that trade secrecy frustrates transparency goals, but that patent ineligibility may encourage reliance on trade secret protection).

³²⁵ See Asay, *supra* note 3, at 1188 (evaluating power of trade-secret and patent protection as incentives for innovation in the AI field; concluding that neither is a good fit; preferring government subsidies).

³²⁶ The Coca Cola trade secret is a favorite example. It has remained secret for 120 years.

³²⁷ 35 U.S.C. § 101.

³²⁸ Not novel as required by 35 U.S.C. § 102.

³²⁹ 35 U.S.C. § 103 (prohibiting patents for obvious inventions).

³³⁰ See Asay, *supra* note 3, at 1219 (noting that trade secret law is largely displacing patent law as the preferred form of protection for AI and explaining why).

4. *First-mover advantage*

Innovators may not need any form of legal protection against free riding, because they enjoy a first-mover advantage in the marketplace.³³¹ Harvard economist Frank Scherer identifies four types of first-mover advantage, which often are cumulative: delays in competitor recognition of advantages conferred by invention and getting up to speed to imitate it; consumer perceptions that the first-mover is a superior supplier; declining variable costs with experience; and economies of scale.³³² He also identifies ability to protect trade secrets as an advantage.

Landes and Posner note that some highly innovative research-intensive industries such as the computer software industry do not rely predominantly on patents as a method of preventing free riding.³³³ They note that the pace of innovation often means that an imitator is at a disadvantage in the marketplace because he is relying on old technology.³³⁴

First-movers have more time to accumulate and master technical knowledge. They can preempt competitor access to scarce assets such as talented employees and key suppliers. They can build a critical mass of customers so as to make it inconvenient or costly for those customers to switch to competitors.³³⁵

Fernando F. Suarez and Gianvito Lanzolla conclude that “gradual evolution in both technology and markets provides first-movers with the best conditions.”³³⁶ This is not necessarily so. A nimble first-mover will have developed a second invention using new technology by the time a competitor copies the first invention with older technology. A strategy premised on constant innovation—staying one step ahead of competitors—may require neither trade secret nor patent protection. Anyway, a patent will be obsolete by the time it issues.

5. *The patent bargain*

Patent law exists, in major part, to induce inventors to give up their trade secrets and to make public their inventions. In other words, to become more transparent. The often-reiterated patent bargain is this:

If inventors rely on trade secrets, the public never gets the benefit of the enhanced knowledge that gave rise to and is embedded in the invention.³³⁷ The patent

³³¹ See F. M. Scherer, *First Mover Advantages and Optimal Patent Protection* (2015) (working paper, Harvard University) (available at https://www.hks.harvard.edu/sites/default/files/centers/mrcbg/files/RPP_2015_05_Scherer.pdf) (quantifying first-mover advantage as alternative to patent).

³³² *Id.* at 2.

³³³ LANDES & POSNER, *supra* note 307, at 312.

³³⁴ *Id.* at 313.

³³⁵ Fernando F. Suarez & Gianvito Lanzolla, *The Half-Truth of First-Mover Advantage*, HARVARD BUS. REV. (Apr. 2005), <https://hbr.org/2005/04/the-half-truth-of-first-mover-advantage> (identifying two factors that influence a first-mover advantage: pace at which technology is evolving and pace at which market is expanding).

³³⁶ *Id.*

³³⁷ See Ebrahim, *supra* note 152, at 182 (“The traditional view of the dichotomy between patents and

system is intended to provide an incentive for disclosure rather than maintaining trade secrecy.³³⁸ In exchange for disclosure, the inventor gets a monopoly for a limited time—presently twenty years from the date of application for a patent—and at the end of that term, the public gets the benefit of the free use of the knowledge associated with the invention. In the meantime, other inventors, while they are not free to make or use the invention without permission of the patent holder, have access to the knowledge behind the invention and may use it to develop their own innovations.

For the system to work, an application for a patent must disclose the requisite knowledge. In particular, § 112 imposes an obligation to describe “the invention, and . . . 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 make and use [it].”³³⁹

The patent laws impose upon the inventor a requirement of disclosure. To insure adequate and full disclosure so that upon the expiration of the 17-year period “the knowledge of the invention inures to the people, who are thus enabled without restriction to practice it and profit by its use,” the patent laws require that the patent application shall include a full and clear description of the invention and “of the manner and process of making and using it” so that any person skilled in the art may make and use the invention. 35 U.S.C. § 112.³⁴⁰

Section III explores the congruence between the § 112 requirement and the desired transparency of AI innovations.

6. *Is the incentive enough?*

AI developers strike their own balance and make their own choices about how much they want to depend on patent law relative to trade secret law and first-mover advantage. Demands for transparency may mature into requirements that compromise trade secret protection and cause AI inventors to see other ways of protecting against free riding.

The first-mover advantage may justify an indifference to any form of intellectual property protection, in which case, weakening trade secret protection to promote AI transparency will have no effect on the propensity to seek patents.

On the face of it, nothing about the controversy over generative AI and demands to regulate it alters the balance between patent law and other measures. One aspect of the clamor for regulation does have potential impact, however. That is the demand for transparency. Full transparency would compromise trade secrets, which no longer

trade secrets emphasizes that patents are premised on disclosure, and trade secrets are premised on a lack of disclosure.”).

³³⁸ *Id.* at 158 (“Disclosure refers to patents being awarded as the quid pro quo for disclosing the invention through enablement and written description. 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.”).

³³⁹ 35 U.S.C. § 112(a).

³⁴⁰ *Kewanee Oil Co., v. Bicron Corp.*, 416 U.S. 470, 480–81 (1974) (internal citations and footnotes omitted).

qualify as such once they become publicly known. It is entirely possible that the policies of such a regulation will compel AI developers to disclose and thereby sacrifice their trade secrets. If it does, or if the developers fear that it is likely to do so, they can abandon trade secret protection voluntarily and place greater reliance on patent protection, which affords legal protection against piracy even though the subject matter and details of the invention have been and must be made public.

Some commentators argue that inventors may be deterred from seeking patents by the disclosure requirement,³⁴¹ but this argument has force only if trade secret comparison is available as an alternative.

B. Political context: A perfect storm

The push for AI transparency and the *Alice/Mayo* reform movement may reinforce each other, tilting AI business strategies toward patent protection even as patents on AI become easier to get. The legislation to reform § 101 and to repudiate *Alice/Mayo*, considered in § IV, has stalled.³⁴² Hearings have been held in the Senate Judiciary committee,³⁴³ but nothing has happened since then to move the bills forward in the Senate.³⁴⁴ The only patent-reform bill in the House is the Restoring America's Leadership in Innovation Act of 2024, introduced by Representative Thomas Massie of Kentucky. It is not likely to go anywhere. It would return the U.S. to a first-to-invent system.³⁴⁵ At the same time that patent reform legislation has run aground, voices are baying all over the place for aggressive regulation of AI. The advocates of such regulation are unable to explain what form it should take among traditional regulatory approaches, and, indeed, they have been unable to come up with a workable definition of what they want to regulate.³⁴⁶ But the pressure for legislative and regulatory action is real and has infected both political parties.³⁴⁷

The momentum to legislate with respect to AI might be just the boost that patent reformers need. If they recharacterize their legislation as necessary to address patents on generative AI, they may have discovered the ticket for building momentum to get something enacted. If this occurs, *sui generis* patent law amendments might be enacted for AI, perhaps conditioning eligibility for a patent on particular characteristics of an invention that woo the most powerful interest groups, say a standard that disqualifies any invention that does not have mechanisms to correct for disparate impact on certain identity groups.

³⁴¹ See Asay, *supra* note 3, at 1209–10 (arguing that patent disclosure requirements are a disincentive to seek patent protection because of limited value of patent if one is available; and because of greater power of trade-secret protection).

³⁴² *Id.* at 1247–48 (noting that enactment of proposals to expand patent eligibility would add to incentives for AI innovation, but noting that risks of overbroad patents and patent trolls suggest caution).

³⁴³ See discussion *supra* Section IV.B.

³⁴⁴ Patent Eligibility Restoration Act of 2023, S. 2140, 118th Cong. (2023).

³⁴⁵ Restoring America's Leadership in Innovation Act of 2024, H.R.8134, 118th Cong. (2024).

³⁴⁶ Exec. Order No. 14110 (giving an overbroad definition of AI); Artificial Intelligence Act, *supra* note 11.

³⁴⁷ See discussion of S.2140, *supra* Section IV.B.

Far more likely, however, is a recognition that the legal status quo, including *Alice/Mayo* and the USPTO's application of it, permits plenty of AI inventions to be patented and that a tightening of the § 112 requirements for detailed disclosure of AI inventions does not require legislation.

C. Legal context: Patent protection as substitute compensation for a taking of trade secrets

The threat of AI transparency requirements' extinguishing trade secrets in AI inventions may be more apparent than real. The Supreme Court of the United States held in *Ruckelshaus v. Monsanto*³⁴⁸ that trade secrets qualify as constitutionally protected property.³⁴⁹ To compel the disclosure of a trade secret is to destroy it and therefore to take the property. The Fifth Amendment prohibits such a taking unless accomplished for a public purpose and accompanied by just compensation. It is not unlikely that legislators or regulators can offer a plausible public purpose for compelling AI transparency. But that is not enough, they also must offer just compensation. The interesting question is whether the compensation could come in the form of existing or enhanced patent protection.

The Fifth Amendment takings doctrine is clear that the compensation requirements of the Fifth Amendment can come in forms other than immediate cash payment, particularly in the form of access to the Court of Claims. But a claim in that court only potentially results in cash compensation, while resorting to the patent system could result in a valuable alternative form of property.

In *O'Neill v. City of East Providence*, the Rhode Island Supreme Court held that forcing a property owner to take another piece of property of equal value does not satisfy the just compensation requirement of the state constitution.³⁵⁰ Under the substitute facilities doctrine, however, Fifth Amendment just-compensation can be provided to a municipality by providing substitute facilities for condemned streets and roads.³⁵¹

The context for compelling disclosure of trade secrets in an AI system is quite different from the context of taking public streets and roads from a municipality in order to build a dam. Nevertheless, the federal substitute facilities doctrine provides some support, however weak, for the proposition that substitute compensation can meet the requirements of the Fifth Amendment when it is compelled by practical necessity, in the form of difficulty in valuation.³⁵²

³⁴⁸ See *Ruckelshaus v. Monsanto*, 467 U.S. 986 (1984).

³⁴⁹ *Id.* at 1003–04.

³⁵⁰ *O'Neill v. East Providence*, 480 A.2d 1375, 1383 (R.I. 1984).

³⁵¹ See *U.S. v. 50 Acres of Land*, 469 U.S. 24, 30 (1984) (noting that doctrine is premised on absence on market value for public infrastructure). The substitute facilities doctrine may refer to paying compensation for a taking in the form of providing substitute facilities, or it may refer to using the cost of constructing a substitute facility as a way of quantifying monetary compensation.

³⁵² See Andrew J. Miller, *Transferable Development Rights in the Constitutional Landscape: Has Penn Central Failed to Weather the Storm?*, 39 NAT RES. J. 459, 506–08 (1999) (noting substitute facilities doctrine and suggesting that transferable development rights might be a constitutionally adequate

The concept of patent rights as substitute compensation is viable, if at all, in the context of a federal termination of trade secrets, and an explicit grant of some kind of patent entitlement. The Fifth Amendment taking requirement for just compensation is surely not satisfied by the government's saying: "We are taking your property, and you are free to troll for compensation."

A state may not offer something similar to a patent. Such an offering would be preempted.³⁵³

VI. Good policy and good law

Continuing to broaden eligibility for patents of AI inventions, combined with rigorous enforcement of the § 112 disclosure requirement is a good way to enhance transparency of AI. Achieving transparency through the patent system will be enhanced by recognition of the possibility and by transparency demands that undermine trade-secret protection.

Making limited monopolies available in the form of patent protection to induce inventors to make their trade secrets public is one of the oldest justifications for the patent system. It can come into play now as a way of resolving tensions between AI innovators and advocates of regulation. Patents are intended to make new technologies transparent.³⁵⁴ Encouraging generative AI innovators to place greater reliance on the patent system represents a readily available and feasible means of realizing the transparency goal.

The explosion of interest in AI stimulated by ChatGPT and other generative AI products in 2023³⁵⁵ means that the flow of applications for patents on AI inventions, already considerable, will increase. Evaluating these applications for patentability provides an opportunity to tame the controversy that has raged over the judicial exceptions to patent eligibility created by the twin Supreme Court decisions of *Alice*³⁵⁶ and *Mayo*³⁵⁷.

Bona-fide AI inventions almost entirely involve software rather than hardware. They bring arguments over whether and the degree to which computer software should be eligible for patents. At the same time, few of the bona fide advances in AI technology can be characterized as humdrum applications of known algorithms. They thus mitigate concerns over making patents too easily available for mere algorithms.

form of just compensation).

³⁵³ *Bonito Boats, Inc. v. Thunder Craft Boats, Inc.*, 489 U.S. 141, 168 (1989) (holding that Florida statute affording patent-like protection was preempted).

³⁵⁴ *Amgen Inc. v. Sanofi*, 598 U.S. 594, 604 (2023) (referring to "patent bargain" and citing Federal Papers on goal of "new designs and technologies into the public domain through disclosure.").

³⁵⁵ See 2024 WIPO REPORT, *supra* note 18, at 15 (reporting that release of ChatGPT stimulated public enthusiasm for AI).

³⁵⁶ See *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66 (2012).

³⁵⁷ *Alice Corp. v. CLS Bank Int'l*, 573 U.S. 208, 226–27 (2014) (holding that claims were drawn to abstract idea of financial settlement and that requiring generic computer implementation did not transform them into patent-eligible invention).

The USPTO has made a good start in determining patent eligibility for AI inventions. As it continues to refine its analytical approach for this new category of patents, it has an opportunity to rationalize and make more predictable § 101 standards, integrated with the other statutory tests for patentability.

Enactment of PERA or similar legislation might make patents on AI more easily available, but amending patent law is not necessary to allow AI innovation to flourish.³⁵⁸ Moreover, new legislation inevitably provokes some years of heightened uncertainty until its new terminology is interpreted by agencies and courts. So the goal of PERA supporters to reduce uncertainty by legislating § 101 reform is likely to have the opposite effect, at least for some period of time. The language of PERA is not much less indefinite than the language of the *Alice/Mayo* tests. It may be desirable to get rid of *Alice/Mayo* as an independent engine of scrutiny but it is not necessary to do so have a workable system—and the risks of attempted reform are great any initiative is likely to be captured by “stakeholder” advocates who want to use the patent systems to pursue goals other than technological innovation.³⁵⁹

Alice/Mayo’s “something more” and “inventive concept” requirement allow considerable flexibility in interpretation and application. Indeed, that flexibility is the source of the complaint that *Alice/Mayo* is unpredictable.³⁶⁰

The policy behind *Alice/Mayo*’s narrowing of the § 101 window is the concern about preemption. One way to limit preemption is to close the § 101 window more

³⁵⁸ Some commentators argue that broadening patent eligibility might not be enough to make patent protection more attractive to AI inventors. *See, e.g.,* Asay, *supra* note 3, at 1211–13 (noting that novelty is a particular barrier for AI inventions because AI’s elements have been around for a long time).

³⁵⁹ The most important objective for stewards of sound patent law is to resist “stakeholder” capture. Proposals proliferate for regulatory approaches that bring stakeholders to the table. Writ large, bringing stakeholders to the table means giving them a veto over decisions. Politicizing economic decision-making rarely produces good results. *See, e.g.,* Roger C. Cramton, et al., *A Critique of the Agencies as Presently Constituted and of the Council’s Recommendations for Realignment and Reorganization*, 23 ADMIN. L. REV. 423, 426–27 (1971) (observing that ICC set prices to prevent low cost industries or firms from capturing business from high-cost competitors); *Seaboard Coast Line R.R. Co. v. United States*, 599 F.2d 650 (5th Cir. 1979) (upholding ICC decision barring proposed rate decrease on petition by competitor); *See generally* ROBERT C. FELLMETH ET AL., *THE INTERSTATE COMMERCE OMISSION: THE PUBLIC INTEREST AND THE ICC* (1970) (noting tendency of transportation regulators to oppose “technological breakthroughs which topple a well-nurtured, highly-valued regulatory system,” citing container ships, “Big John” boxcars, “piggyback” operations in mail transport, “foreign attachments” to the AT&T telephone system, and (to a certain extent) cable television.); Mark Green & Ralph Nader, *Economic Regulation vs. Competition: Uncle Sam the Monopoly Man*, 82 YALE L. J. 871, 882 (1973) (estimating \$4 to \$8.7 billion cost of ICC increased costs and shifting of traffic from low cost to high cost modes, involving railroads, trucks, water).

Stakeholder involvement is even more harmful when the considerations that are introduced into economic decision-making have nothing to do with economics. It was one thing for railroad regulation to give farmers shipping grain and ranchers shipping cattle a position at the table; it’s quite another to give advocates of an entirely different social and political order an opportunity to bend economic decisions in their favor although they manifestly lack the votes to do so through the ordinary democratic process.

³⁶⁰ *See supra* Section B, regarding criticisms of *Alice/Mayo*.

tightly; another way is to apply the § 112 disclosure requirement more rigorously, and *Amgen Inc. v. Sanofi* (“*Amgen*”) gives some momentum to taking § 112 more seriously as a way of staving off overbroad preemption.³⁶¹ The cases analyzed in § IV.A.2 are useful guidance on rigorous disclosure under § 112, even when their focus was on § 101 eligibility, as the *Berkheimer* case, discussed in § IV.C, suggests.

The USPTO can patent applications for AI inventions with the transparency goal in mind, explicitly recognizing the congruence between transparency demands and the historic purpose of patent system. The Director of the USPTO can issue guidance to examiners encouraging rigorous application of § 112 to such applications. Unless the Patent Act is amended, neither the USPTO nor the courts would be justified in compelling additional disclosure to serve the interests advanced by calls for AI regulation. On the other hand, additional disclosure would aid a person skilled in the art in building on the knowledge underlying a particular invention, and the USPTO and the courts can use that long-standing purpose of the patent system and couple it with the inherent opacity of generative AI to justify requiring additional and more detailed disclosure. The USPTO has statutory authority to “establish regulations, not inconsistent with law.”³⁶² This authority extends only to procedural, not substantive rules.³⁶³

It also may authorize special priority for applications for “products, processes, or technologies that are important to the national economy or national competitiveness.”³⁶⁴ This could authorize special priority for examination of application involving detailed disclosure of AI inventions.

Such an approach is broadly consistent with a pre-*Alice/Mayo* observation by Dean Michael Risch of Villanova University Charles Widger School of Law: “Virtually all of the important historical patentable subject matter cases may be explained by applying each of the other requirements for patentability. When viewed through this lens, subject matter concerns are at bottom patentability concerns.”³⁶⁵ He argues that the Court’s patentable subject matter doctrine would be more consistent

³⁶¹ *Amgen Inc. v. Sanofi*, 598 U.S. 594, 610 (2023) (holding that “The more one claims, the more one must enable” under section 112).

³⁶² 35 U.S.C. § 2(b)(2).

³⁶³ In *Tafas v. Dudas*, 541 F. Supp. 2d 805, 811–12 (E.D. Va. 2008). The district court held that section 2 does not confer authority to make substantive, as opposed, to procedural rules. It invalidates the rule limiting the number of continuing applications, RCE, and claims as ultra vires. *Id.* at 817. The court of appeals affirmed in material part in *Tafas v. Doll*. *Tafas v. Doll*, 559 F.3d 1345, 1364 (Fed. Cir. 2009) (affirming conclusion that only procedural rules are authorized, reversing finding that most of challenged rules were substantive, and finding that one rule was invalid because it conflicted with section 120). Rehearing en banc was ordered, but the court subsequently dismissed the appeal as moot. *Tafas v. Doll*, 328 Fed. Appx. 658, 658 (Fed. Cir. 2009); *Tafas v. Kappos*, 586 F.3d 1369, 1371 (Fed. Cir. 2009) (dismissing appeal as moot after USPTO withdrew regulation). Notably, however, the Federal Circuit denied a motion to vacate the district court’s opinion and order. *Id.*

³⁶⁴ 35 U.S.C. § 2(b)(2)(G).

³⁶⁵ Michael Risch, *Everything Is Patentable*, 75 TENN. L. REV. 591, 598 (2008) (arguing that patentability criteria such as obviousness, novelty, utility, inventorship, written description, and enablement motivate the Supreme Court’s patentable subject matter decisions).

and rigorous if replaced with a strict application of other patentability doctrines such as obviousness, novelty, utility, inventorship, written description, and enablement,³⁶⁶ what he calls a “rigorous patentability” approach.³⁶⁷ Risch’s suggestion did not go unnoticed and unadmired.³⁶⁸ Mark Lemley, a professor of law at Stanford Law School, was hospitable to Risch’s idea: “Claims are proper when the scope of the patentee’s claims is commensurate with a practical, real world contribution the patentee has made.”³⁶⁹ But Lemley and his co-authors worried about conflating §§ 101, 102, 103, and 112.³⁷⁰

Risch’s persuasive rigorous patentability approach lost momentum. Although there was judicial commentary on it, the Supreme Court expressly rejected it: “This [Risch’s] approach, however, would make the ‘law of nature’ exception to § 101 patentability a dead letter.”³⁷¹

In an email exchange with the author, Risch lamented that his proposal to deemphasize § 101 analysis in favor of a more rigorous application of §§ 102, 103, and 112 was frustrated because the “train has left the station,” referring to this passage in *Mayo*.³⁷² Risch’s focus was on rigorous application of the novelty and nonobvious requirements. But the train has not left the station with respect to the rigorous application of the disclosure requirement under § 112. Indeed, the Supreme Court’s *Amgen* ruling suggests that the § 112 train may still be sitting in the station waiting for new passengers and freight.³⁷³ AI transparency might be the new passenger. The context of the decision involved the permissible scope of generic claims in the pharmaceutical field, but its logic reinforces the requirement to insist on meaningful

³⁶⁶ *Id.* at 594.

³⁶⁷ *Id.* at 607–09 (proposing “rigorous patentability” approach to eligibility determinations, comprising sequential application of sections 101, 102, 103, and 112).

³⁶⁸ *See, e.g.,* Jorge A. Goldstein et al., *The Time Has Come to Amend 35 U.S.C. Sec 101*, 44 AIPLA Q. 171, 195 n.84 (2016) (citing Risch’s suggestion and comparing it to Judge Newman’s dissent in *CLS Bank Int’l v. Alice Corp.*, 717 F.3d 1269, 1321 (Fed. Cir. 2013)).

³⁶⁹ Mark A. Lemley et al., *Life After Bilski*, 63 STAN. L. REV. 1315, 1339–41 (2011) (identifying factors such as generative nature of invention, particular industry reliance on cumulative invention, speed at which innovation occurs in the particular industry, claims to broad principle based on few specific applications, and importance of inventor’s contribution).

³⁷⁰ *Id.* at 1327 (criticizing proposal for conflating section 101 inquiry with section 112 requirement); *Id.* at 1329–32 (explaining why section 112 concerns are different from section 101 concerns about breadth of claims). “We recognize that, in evaluating the significance of additional steps, the § 101 patent-eligibility inquiry and, say, the § 102 novelty inquiry might sometimes overlap. But that need not always be so. And to shift the patent-eligibility inquiry entirely to these later sections risks creating significantly greater legal uncertainty, while assuming that those sections can do work that they are not equipped to do.” *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66, 90 (2012).

³⁷¹ *Mayo*, 566 U.S. at 89, 91 (citing Risch’s article and rejecting the approach).

³⁷² E-mail from Dean Michael Risch, Vice Dean and Professor of L., Villanova University; Charles Widger School of Law, and Henry H. Perritt, Jr., Professor of L., Chicago-Kent Coll. of L. (June 6, 2024, 12:59 PM) (on file with author).

³⁷³ *See Amgen Inc. v. Sanofi*, 598 U.S. 594 (2023) (affirming invalidation of biological patents for failure to disclose classes of species with sufficient specificity to be enabling).

disclosure of everything claimed.³⁷⁴

³⁷⁴ *Id.* at 606–12 (discussing 19th century brick making, glue, and telegraph patents).