

Intellectual Property & Technology Law Journal

Edited by the Technology and Proprietary Rights Group of Weil, Gotshal & Manges LLP

VOLUME 31 • NUMBER 3 • MARCH 2019

The Role of Explainable Artificial Intelligence in Patent Law

By Brian Higgins

Although the notion of “explainable artificial intelligence” (AI) has been suggested as a necessary component of governing AI technology, at least for the reason that transparency leads to trust and better management of AI systems in the wild, one area of U.S. law already places a burden on AI developers and producers to explain how their AI technology works: patent law.

Background

Patent law’s focus on *how* AI systems work was not borne from a Congressional mandate. Rather, the U.S. Supreme Court gets all the credit – or blame, as some might contend – for this legal development, which began with the Court’s 2014 decision in *Alice Corp. Pty Ltd. v. CLS Bank International*. *Alice* established the legal framework for assessing whether an invention fits in one of the patent law’s patent-eligible categories (i.e., any “new and useful process, machine, manufacture, or composition of matter” or improvements thereof) or is

a patent-ineligible concept (i.e., law of nature, natural phenomenon, or abstract idea).¹

AI systems based on machine learning models generally transform input data into actionable output data.

Understanding how the idea of “explaining AI” came to be following *Alice*, one must look at the very nature of AI technology. At their core, AI systems based on machine learning models generally transform input data into actionable output data, a process U.S. courts and the U.S. Patent and Trademark Office (Patent Office) have historically found to be patent-ineligible.

Consider a decision by the U.S. Court of Appeals for the Federal Circuit, whose judges are selected for their technical acumen as much as for their understanding of the nuances of patent and other areas of law, that issued around the same time as *Alice*: “a process that employs mathematical algorithms to manipulate existing information to generate additional information is not patent eligible.”² While *Alice* did not specifically address AI or mandate anything resembling explainable AI, it nevertheless spawned a progeny of Federal Circuit, district court, and Patent

Brian Higgins is a partner at Blank Rome LLP, where he applies his engineering and law knowledge to help clients maximize business valuation through the strategic use of intellectual property rights. He may be contacted at higgins@blankrome.com.

Office decisions that did just that. Notably, those decisions arose not because of notions that individuals impacted by AI algorithmic decisions ought to have the right to understand how those decisions were made or why certain AI actions were taken, but because explaining how AI systems work helps satisfy the *quid pro quo* that is fundamental to patent law: an inventor who discloses to the world details of what she has invented is entitled to a limited legal monopoly on her creation (provided, of course, the invention is patentable).

The Rise of Algorithmic Scrutiny

Alice arrived not long after Congress passed patent reform legislation called the America Invents Act (AIA) of 2011, provisions of which came into effect in 2012 and 2013. In part, the AIA targeted a decade of what many consider a time of abusive patent litigation brought against some of the largest tech companies in the world and thousands of mom-and-pop and small business owners who were sued for doing anything computer-related. This litigious period saw the term “patent troll” used more often to describe patent assertion companies that bought up dot-com-era patents covering the very basics of using the Internet and computerized business methods, and who then sued to collect royalties for alleged infringement. Not surprisingly, some of the same big tech companies that pushed for patent reform provisions now in the AIA to curb patent litigation in the field of computer technology also filed *amicus curiae* briefs in the *Alice* case to further weaken software patents. The Supreme Court’s unanimous decision in *Alice* helped curtail troll-led litigation by formalizing a procedure, one that lower court judges could readily adopt, for excluding certain software-related inventions from the list of inventions that are patentable.

Under *Alice*, a patent claim – the language used by inventors to describe what he or she claims to be his or her invention – falls outside § 101 when it is “directed to” one of the patent-ineligible concepts noted above. If so, *Alice* requires consideration of whether the particular elements of the claim, evaluated “both individually and ‘as an ordered combination’” add enough to “transform the nature of the claim” into one of the patent-eligible categories.³ While conceptually straightforward, it took years of court and Patent Office decisions to explain how

that two-part test is to be employed, and only more recently how it applies to AI technologies.⁴

Today, the Patent Office and courts across the United States routinely find that algorithms are abstract (even though algorithms, including certain mental processes embodied in algorithmic form performed by a computer, are by most measures *useful processes*). According to the Federal Circuit, algorithmic-based data collection, manipulation, and communication – functions most AI algorithms perform – are abstract.

Artificial Intelligence, Meet Alice

In a bit of ironic foreshadowing, the Supreme Court issued *Alice* in the same year that major advances in AI technologies were being announced, such as Google’s deep neural network architecture that prevailed in the 2014 ImageNet challenge (ILSVCR) and Ian Goodfellow’s generative adversarial network (GAN) model, both of which were major contributions to the field of computer vision. Even as more breakthroughs were being announced, U.S. courts and the Patent Office began issuing *Alice* decisions regarding AI technologies and explaining why it is crucial for inventors to explain how their AI inventions work to satisfy the second half of *Alice*’s two-part test.

In *Purepredictive, Inc. v. H2O.AI, Inc.*, for example, the U.S. District Court for the Northern District of California considered the claims of U.S. Patent 8,880,446, which, according to the patent’s owner, involves “AI driving machine learning ensembling.” The district court characterized the patent as being directed to a software method that performs “predictive analytics” in three steps.⁵ In the method’s first step, it receives data and generates “learned functions,” or, for example, regressions from that data. Second, it evaluates the effectiveness of those learned functions at making accurate predictions based on the test data. Finally, it selects the most effective learned functions and creates a rule set for additional data input. The court found the claims invalid on the grounds that they “are directed to the abstract concept of the manipulation of mathematical functions and make use of computers only as tools, rather than provide a specific improvement on a computer-related technology.” The claimed method, the district court said, is merely “directed to a mental process” performed by a computer, and “the abstract concept of using mathematical algorithms to perform predictive analytics” by collecting and analyzing information. The court

explained that the claims “are mathematical processes that not only could be performed by humans but also go to the general abstract concept of predictive analytics rather than any specific application.”

In a bit of ironic foreshadowing, the Supreme Court issued *Alice* in the same year that major advances in AI technologies were being announced.

In *Ex Parte Lyren*, the Patent Office’s Appeals Board, made up of three administrative law judges, rejected a claim directed to customizing video on a computer as being abstract and thus not patent-eligible. In doing so, the board disagreed with the inventor, who argued the claimed computer system, which generated and displayed a customized video by evaluating a user’s intention to purchase a product and information in the user’s profile, was an improvement in the technical field of generating videos. The claimed customized video, the Board found, could be any video modified in any way. That is, the rejected claims were not directed to the details of *how* the video was modified, but rather to the *result* of modifying the video. Citing precedent, the board reiterated that “[i]n applying the principles emerging from the developing body of law on abstract ideas under section 101, ... claims that are ‘so result-focused, so functional, as to effectively cover any solution to an identified problem’ are frequently held ineligible under section 101.”⁶

Notably, the claims in *Ex Parte Lyren* were also initially rejected as failing to satisfy a different patentability test – the written description requirement.⁷ In rejecting the claims as lacking sufficient description of the invention, the Patent Office Examiner found that the algorithmic features of the inventor’s claim were “all implemented inside a computer, and therefore all require artificial intelligence [(AI)] at some level” and thus require extensive implementation details “subject of cutting-edge research, e.g.[,] natural language processing and autonomous software agents exhibiting intelligent behavior.” The Examiner concluded that “one skilled in the art would not be persuaded that Applicant possessed the invention” because “it is not readily apparent how to make a device [to] analyze natural language.” The Appeals Board

disagreed and sided with the inventor who argued that his invention description was comprehensive and went beyond just artificial intelligence implementations. Thus, while the description of how the invention worked was sufficiently set forth, Lyren’s claims focused too much on the results of the application of the technology and thus were found to be abstract.

In *Ex Parte Homere*, claims directed to “a computer-implemented method” involving “establishing a communication session between a user of a computer-implemented marketplace and a computer-implemented conversational agent associated with the market-place that is designed to simulate a conversation with the user to gather listing information, the Appeals Board affirmed an Examiner’s rejection of the claims as being abstract.⁸ In doing so, the Appeals Board noted that the inventor had not identified anything in the claim or in the written description that would suggest the computer-related elements of the claimed invention represent anything more than “routine and conventional” technologies. The most advanced technologies alluded to, the Board found, seemed to be embodiments in which “a program implementing a conversational agent may use other principles, including complex trained Artificial Intelligence (AI) algorithms.” However, the *claimed* conversational agent was not so limited. Instead, the Board concluded that the claims were directed to merely using recited computer-related elements to implement the underlying abstract idea, rather than being limited to any particular advances in the computer-related elements.

In *Ex Parte Hamilton*, a rejection of a claim directed to “a method of planning and paying for advertisements in a virtual universe (VU), comprising...determining, via the analysis module, a set of agents controlled by an Artificial Intelligence . . .,” was affirmed as being patent ineligible.⁹ The Appeals Board found that the “determining” step was insufficient to transform the abstract idea of planning and paying for advertisements into patent-eligible subject matter because the step represented an insignificant data-gathering step and thus added nothing of practical significance to the underlying abstract idea.

In *Ex Parte Pizzorno*, the Appeals Board affirmed a rejection of a claim directed to “a computer implemented method useful for improving artificial

intelligence technology” as abstract.¹⁰ In doing so, the Board determined that the claim was directed to the concept of using stored health care information for a user to generate personalized health care recommendations based on Bayesian probabilities, which the Board said involved “organizing human activities and an idea in itself, and is an abstract idea beyond the scope of § 101.” Considering each of the claim elements in turn, the Board also found that the function performed by the computer system at each step of the process was purely conventional in that each step did nothing more than require a generic computer to perform a generic computer function.

In *Ex Parte Pizzorno*, the Appeals Board affirmed a rejection of a claim directed to “a computer implemented method useful for improving artificial intelligence technology” as abstract.

Finally, in *Ex Parte McAfee*, the Appeals Board affirmed a rejection of a claim on the basis that it was “directed to the abstract idea of receiving, analyzing, and transmitting data.”¹¹ At issue was a method that included “estimating, by the ad service circuitry, a probability of a desired user event from the received user information, and the estimate of the probability of the desired user event incorporating artificial intelligence configured to learn from historical browsing information in the received user information, the desired user event including at least one of a conversion or a click-through, and the artificial intelligence including regression modeling.” In affirming the rejection, the Board found that the functions performed by the computer at each step of the claimed process was purely conventional and did not transform the abstract method into a patent-eligible one. In particular, the step of estimating the probability of the desired user event incorporating artificial intelligence was found to be merely “a recitation of factors to be somehow incorporated, which is aspirational rather than functional and does not narrow the manner of incorporation, so it may include no more than incorporating results from some artificial intelligence outside the scope of the recited steps.”

The above and other *Alice* decisions have led to a few general legal axioms, such as:

- A claim for a *new* abstract idea is still an abstract idea;
- A claim for a *beneficial* abstract idea is still an abstract idea;
- Abstract ideas do not become patent-eligible because they are new ideas, are not previously well known, and are not routine activity; and
- The “mere automation of manual processes using generic computers does not constitute a patentable improvement in computer technology.”¹²

The Focus on How, Not the Results

Following *Alice*, patent claims directed to an AI technology must recite features of the algorithm-based system that represent how the algorithm improves a computer-related technology and is not previously well-understood, routine, and conventional. In *PurePredictive*, for example, the Northern California district court, which sees many software-related cases due to its proximity to the Bay Area and Silicon Valley, found that the claims of a machine learning ensemble invention were not directed to an invention that “provide[s] a specific improvement on a computer-related technology.”¹³ Satisfying *Alice*, thus, requires claims focusing on a somewhat narrow application of how an AI algorithmic model works, rather than the broader and result-oriented nature of what the model is used for. This is necessary where the idea behind the algorithm itself could be used to achieve many different results. For example, a claim directed to a mathematical process (even one that is said to be “computer-implemented”), and that could be performed by humans (even if it takes a long time), and that is directed to a result achieved instead of a specific application, will seemingly be patent-ineligible under today’s *Alice* legal framework.

To illustrate, consider an image classification system, one that is based on a convolutional neural network. Such a system may be patentable if the claimed system improves the field of computer vision technology. Claiming the invention in terms of how the elements of the computer are technically improved by its deep learning architecture and algorithm, rather than simply claiming a deep learning model using results-oriented language, may survive an *Alice* challenge, provided the claim

does not merely cover an automated process that human used to do. Moreover, claims directed to the multiple hidden layers, convolutions, recurrent connections, hyperparameters, and weights could also be claimed if those elements help explain how the claimed computer vision system works.

By way of another example, a claim reciting “a computer-implemented process using artificial intelligence to generate an image of a person,” is likely abstract if it does not explain how the image is generated but instead merely claims a computerized process a human could perform. But a claim that describes a unique AI system that specifies how it generates the image, including the details of a generative adversarial network process and its various inputs provided by physical devices (not routine data collection), its connections and hyperparameters, has a better chance of passing muster (keeping in mind, this only addresses the question of whether the claimed invention is *eligible* to be patented, not whether it is, in fact, patentable, which is an altogether different analysis and requires comparing the claim to prior art).

Uncertainty Remains

Although the issue of explaining how an AI system works in the context of patent law is still in flux, the number of U.S. patents issued by the Patent Office mentioning “machine learning,” or the broader term “artificial intelligence,” has jumped in recent years. Just this year alone, U.S. machine learning patents are up 27 percent compared to the same year-to-date period in 2017 (thru the end of November), according to available Patent Office records.¹⁴ Even if machine learning is not the focus of many of them, the annual upward trend in patenting AI over the last several years appears unmistakable.

The number of U.S. patents issued by the Patent Office mentioning “machine learning” or the broader term “artificial intelligence,” has jumped in recent years.

But with so many patents invoking AI concepts being issued, questions about their validity may arise. As the Federal Circuit has stated, “great uncertainty yet remains” when it comes to the test

for deciding whether an invention like AI is patent-eligible under *Alice*, this despite the large number of cases that have “attempted to provide practical guidance.”¹⁵ Calling the uncertainty “dangerous” for some of today’s “most important inventions in computing,” specifically mentioning AI, the Federal Circuit expressed concern that perhaps the application of the *Alice* test has gone too far, a concern mirrored in testimony by Andrei Iancu, Director of the Patent Office, before Congress in April 2018 (stating, in response to Judiciary Committee questions, that *Alice* and its progeny have introduced a degree of uncertainty into the area of subject matter eligibility, particularly as it relates to medical diagnostics and software-related inventions, and that *Alice* could be having a negative impact on innovation).

If explaining how AI works is to have a future meaningful role in patent law, the courts or Congress will need to provide clarity.

Absent legislative changes abolishing or altering *Alice*, a solution to the uncertainty problem, at least in the context of AI technologies, lies in clarifying existing decisions issued by the Patent Office and courts, including the decisions summarized above. While it can be challenging to explain why an AI algorithm made a particular decision or took a specific action (due to the black box nature of such algorithms once they are fully trained), it is generally not difficult to describe the structure of a deep learning or machine learning algorithm or how a system embodying it works. Even so, it remains unclear whether and to what extent fully describing how one’s AI technology and including “how” features in patent claims will ever be sufficient to “add[] enough to transform the nature of an abstract algorithm into a patent-eligible [useful process].” If explaining how AI works is to have a future meaningful role in patent law, the courts or Congress will need to provide clarity.

Notes

1. *Alice Corp. Pty Ltd. v. CLS Bank International*, 134 S. Ct. 2347, 2354–55 (2014); 35 USC § 101.
2. *Digitech Image Techs, LLC v. Elecs. v. Imaging, Inc.*, 758 F.3d 1344, 1351 (Fed. Cir. 2014).

-
3. Elec. Power Grp., LLC v. Alstom S.A., 830 F.3d 1350, 1353 (Fed. Cir. 2016) (quoting *Alice*, 134 S. Ct. at 2355).
 4. See 84 Fed. Reg. 50-56, Request for Comment, 2019 Revised Patent Subject Matter Eligibility Guidance (Jan. 7, 2019)
 5. Purepredictive, Inc. v. H2O.AI, Inc., slip op., No. 17-cv-03049-WHO (N.D. Cal. Aug. 29, 2017).
 6. Ex Parte Lyren, No. 2016-008571 (PTAB, June 25, 2018) (citing Affinity Labs of Texas, LLC v. DirecTV, LLC, 838 F.3d 1253, 1265 (Fed. Cir. 2016) (quoting *Elec. Power Grp., LLC v. Alstom S.A.*, 830 F.3d 1350, 1356; see also Ex parte Colcernian et al., No. 2018-002705 (PTAB, Oct. 1, 2018) (rejecting claims that use result-oriented language as not reciting the specificity necessary to show how the claimed computer processor's operations differ from prior human methods, and thus are not directed to a technological improvement but rather are directed to an abstract idea).
 7. 35 USC § 112.
 8. Ex Parte Homere, Appeal No. 2016-003447 (PTAB Mar. 29, 2018).
 9. Ex Parte Hamilton et al., Appeal No. 2017-008577 (PTAB Nov. 20, 2018).
 10. Ex Parte Pizzorno, Appeal No. 2017-002355 (PTAB Sep. 21, 2018).
 11. Ex Parte McAfee, Appeal No. 2016-006896 (PTAB May 22, 2018).
 12. Synopsys, Inc. v. Mentor Graphics Corp., 839 F.3d 1138, 1151 (Fed. Cir. 2016); Ariosa Diagnostics, Inc. v. Sequenom, Inc., 788 F.3d 1371, 1379-80 (Fed. Cir. 2015); Ultramercial, Inc. v. Hulu, LLC, 772 F.3d. 709, 715-16 (Fed. Cir. 2014); Credit Acceptance Corp. v. Westlake Servs., 859 F.3d 1044, 1055 (Fed. Cir. 2017); see also SAP Am., Inc. v. Investpic, LLC, slip op. No. 2017-2081, 2018 WL2207254, at *2, 4-5 (Fed. Cir. May 15, 2018) (finding financial software patent claims abstract because they were directed to “nothing but a series of mathematical calculations based on selected information and the presentation of the results of those calculations (in the plot of a probability distribution function)”; but see *Apple, Inc. v. Ameranth, Inc.*, 842 F.3d 1229, 1241 (Fed. Cir. 2016) (noting that “[t]he Supreme Court has recognized that all inventions embody, use, reflect, rest upon, or apply laws of nature, natural phenomena, or abstract ideas[] but not all claims are directed to an abstract idea.”).
 13. See also *Neochloris, Inc. v. Emerson Process Mgmt LLLP*, 140 F. Supp. 3d 763, 773 (N.D. Ill. 2015) (explaining that patent claims including “an artificial neural network module” were invalid under § 101 because neural network modules were described as no more than “a central processing unit – a basic computer’s brain”).
 14. See also WIPO Technology Trends 2019 - Artificial Intelligence (Jan. 2019) (“Filings of machine learning-related patent have grown annually on annual average by 28 percent, with 20,195 patent applications filed in 2016 (compared with 9,567 in 2013).”) (*available at* https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf).
 15. *Smart Systems Innovations, LLC v. Chicago Transit Authority*, slip. op. No. 2016-1233 (Fed. Cir. Oct. 18, 2017).

Copyright © 2019 CCH Incorporated. All Rights Reserved.
Reprinted from *Intellectual Property & Technology Law Journal*, March 2019, Volume 31,
Number 3, pages 3–8, with permission from Wolters Kluwer, New York, NY,
1-800-638-8437, www.WoltersKluwerLR.com

