Innovating under Uncertainty: The Patent-Eligibility of Artificial Intelligence after Alice Corp. v. CLS Bank International

Ryan Whalen and Raphael Zingg[1]

Abstract

Artificial intelligence-related inventions raise complex questions of how to define the boundaries around patentable subject matter. In the United States, many claim that the recent doctrinal developments by the Supreme Court have led to incoherences and excessive uncertainty within the innovation community. In response, policymakers and stakeholders have suggested legislative amendments to address these concerns. We first review these developments, and subsequently use the patent examination record to empirically test the claims of increased uncertainty. We find that, although uncertainty did spike following the Supreme Court's holding in Alice, it quickly returned to levels comparable to its historic norm. This has implications both for those advocating for legislative changes to the law of eligible subject matter, as well as other jurisdictions considering adopting a test similar to that applied in Alice.

1. Introduction

As the economic and strategic importance of artificial intelligence (AI) has increased in recent years, so too has pressure on policymakers to develop policies that support and facilitate further development of the technology. Policy debate in patent law often focuses on striking a balance between providing sufficient incentive for investors and innovators, without unduly slowing innovation by blockading knowledge behind privately-owned patents. This challenge holds true across patent law, but is particularly challenging in fields like AI due to the comparatively abstract nature of the inventions in question. The patentability standards that have evolved in response to the rapid growth in patent applications claiming AI-related inventions have led to perceptions of uncertainty about patentability, and claims that recent responses by courts and policymakers have increased that uncertainty—often to the detriment of future innovation in AI.

Patent legislation and related common law have long recognized a number of exceptions to patentability. For instance, laws of nature, natural phenomena, mathematical formulae, and abstract ideas are unpatentable in most jurisdictions. The precise boundaries of these exceptions, however, are subject to heated debate and in some jurisdictions, changing policy. In the United States, the Supreme Court has restricted the scope of patentable subject matter via a handful of high-profile decisions. These developments have inspired a recent legislative attempt to rewrite Section 101 of the Patent Act, which delineates the scope of patentable subject matter. To motivated these

[1]Whalen: National University of Singapore, Faculty of Law, Singapore 119077 (e-mail: *ryanwhalen@nus.edu.sg*). Zingg: Waseda University, Institute for Advanced Study, Tokyo 169-8050 and ETH Zurich, Center for Law & Economics, Zurich 8006 (e-mail: *rzingg@ethz.ch*).

proposed revisions, Senator Coons, who co-introduced the proposed legislation declared that 'U.S. patent law discourages innovation in some of the most critical areas of technology, including artificial intelligence, medical diagnostics, and personalized medicine'[1]. In a set of hearings before the Senate Judiciary Subcommittee on Intellectual Property in June, current patent-eligibility standards were described as "truly [..] a mess"[2] and "unnecessarily ambiguous and uncertain"[3]. For software and computer-implemented inventions, the blame for undue uncertainty often goes to the Supreme Court decision Alice Corp. v. CLS Bank International (Alice)[4] and its interpretation by the Patent Office and the Federal Circuit, the appeals court charged with hearing patent disputes. Below, we present the framework developed by the Supreme Court to assess whether a patent claim is an abstract idea, and therefore unpatentable, for inventions in the field of AI. We then challenge the common wisdom that uncertainty has unduly increased in recent years due to Alice.

2. Abstract Ideas in Artificial Intelligence

Patentable subject matter has been defined broadly as including "anything under the sun that is made by man"[5]. The abstract idea doctrine is the most significant exemption to this rule, with the aim of channeling patent protection away from unnecessary upstream patents[6]. As the Supreme Court stated in Brenner v. Manson (1966), patenting an abstract idea would grant the patent holder a 'vast, unknown, and perhaps unknowable area' and the power to block off whole areas of scientific development[7]. For AI especially, where there are fears that fundamentals techniques may become enclosed by private patent rights, distinguishing abstract ideas from implementations is key. The patenting of dropout methods by Google, of approaches to design memory networks by Facebook, or active machine learning methods by Microsoft, all illustrate how firms seek to monopolize basic concepts of AI technology, reinforcing the need for objective standards for patent-eligibility[8].

Alice Corporation was the assignee of four patents disclosing a scheme for intermediated financial risk settlement. The claims represented variants of a method for exchanging financial obligations, a computer system configured to carry out the method for exchanging obligations, and a computer-readable medium containing program code for performing the method of exchanging obligations. The District Court for the District of Columbia and the Court of Appeals held that all of the claims were directed to an abstract idea and therefore ineligible for patent protection. In a heavily political litigation with over 50 amicus briefs filed for either side[9], the Supreme Court was asked to set the boundaries for patentable subject matter for computer-implemented inventions under Section 101.

In Alice, the Supreme Court established a two-part test to determine patent-eligibility of software patents, building on its jurisprudence in Mayo v. Prometheus (2012)[10]. Under the test, a court must first establish whether the patent claims are directed to a patent-ineligible concept such as a law of nature, a natural phenomena, or an abstract idea. Many AI patent claims must overcome the obstacle of being drawn to an abstract idea, that is to an "idea "of itself""[11]. The basic concept of organizing information through mathematical correlations has been found to be an abstract idea by the courts[12]. Similarly, key machine learning and AI techniques such as the steps of producing new labeled observations, evaluating an observation, comparing scores, identifying a feature, and updating a model have been deemed abstract, since they are examples of organizing information[13].

When the claims at issue are directed to an abstract idea, the second step of the test requires a court to search for an "inventive concept", i.e. to examine whether the elements of the claim include additional limitations amounting to "significantly more" than the exception. In Alice, the Supreme Court found that the claims amounted to nothing more than an instruction to apply the

abstract idea of intermediated settlement using some unspecified, generic computer. A number of characteristics of "significantly more" were spelled out by the Supreme Court and have been developed in subsequent Federal Circuit case law. For instance, patent claims can purport to improve the functioning of the computer itself[14]. In this case, the claims must improve the way a computer functions as a tool; recitations of specific data structures, rules, combinations of steps or hardware configurations improving the functioning of a computer therefore cover eligible subject-matter[15]. Further, claims which effect an improvement in any technology or technical field may be patentable as well. This would be the case for claims directed towards improvements in machine vision in labeling the contents of images if they are directed to improvements in the technological process, rather than simply in determining an algorithm or classifier[16]. Alternatively, the inventor can show that a claim is a particular technical solution to a technical problem[17]. The patent office nonetheless finds that retraining or updating a machine learning model is merely an improved algorithm and that the updating process is abstract – data manipulated to create other data and thus insufficiently technical[18]. Finally, the inventor can demonstrate that the claimed improvement describe and solve a problem necessarily rooted in computer technology[19]. The Patent Trial and Appeal Board found this to be true for claims addressing the problem of monitoring the operations of machines and to predict maintenance needs and equipment failures using several artificial intelligence classification technologies[20].

Years after Alice, the definition of what constitutes an abstract idea for AI inventions remains unclear. According to some, uncertainty has now increased and upsets AI development; inventors being better off seeking patent protection in China and Europe[21]. At the recent Congressional hearings, Nokia contended that its Section 101 rejection rate for AI-related inventions increased by about 50% after Alice, while another estimate claims AI patent rejections increased overall nearly four times since the decision[22]. While the Alice two-part test has been applied by courts and the patent office, commentators find that its application is incoherent and inconsistent[23]. Critics point out that the uncertainty in administrative and judicial outcome defeats the economic incentive purpose of patent law: AI patents risk being challenged in opportunistic litigation[24], businesses cannot efficiently allocate research and development resources[25] nor diffuse associated technologies via patent licensing[26] . Proponents, however, argue that Alice does not necessarily harm innovation. The software industry was innovative in the decades prior to the expansion of patentable subject matter in the 1990s, when the viability of these patents was unclear and may well continue to be so regardless of how the law evolves[27].

3. Patenting under Uncertainty

To assess whether claims about AI patenting uncertainty are supported by the empirical record, we identify AI-related patents by relying on the classification developed by the United States Patent and Trademark Office (USPTO), which classifies AI patents in its USPC class 706 titled 'Data Processing – Artificial Intelligence'. This generic class includes inventions for artificial intelligence type computers and digital data processing systems and corresponding data processing methods and products for emulation of intelligence. We extract data from the USPTO Office Action Research Dataset for Patents, a curated dataset on actions by patent examiners, and from the Patent Examination Data System (PEDS) and the data it makes available on all published US patent applications. Our search identified 20,866 AI patent applications (filed 1985 to 2019), and detailed rejection data for 5,806 of them (filed 2008 to 2016). The patent examination process may be simplified as follows. When an application is filed, it is assigned to an examiner. The patent examiner conducts a prior art search identifying relevant prior art that supports a decision of either claim allowance or non-final rejection. In the instance of a rejection, the applicant may enter in one or more rounds of amendments and negotiations. Even after receiving a final rejection, applicants may still submit amended applications or lodge an appeal to the Board of Patent Appeals and Interferences. Examples of recently granted patents in the area of AI range from methods to train neural networks, learning with deep methods, classification of data objects, and intelligent medical image detection[28].

To assess the direct and immediate impact of Alice, a decision rendered in June 2014, we study the likelihood of AI patent applications to be granted over time. Because Alice focused on patent subject matter eligibility, we distinguish between the applications in which patent examiners raised a non-final Section 101 rejection – with an examiner refusing to allow claims based on subject matter – and those for which no such rejection occurred. Figure 1a shows that in both scenarios, the grant rate dramatically dropped for patent applications filed in 2014. Strikingly, only 40% of the AI patent applications for which a non-final Section 101 rejection was raised ended up granted, compared to about 80% in other years. This trend holds true across subfields – such as machine learning and neural networks patents (see Figure A2a to A2d in the Appendix). The decreased grant rate following the Alice decision perhaps contributed to the perceived uncertainty about the patent-eligibility of AI inventions that current proposed reforms seek to mitigate.

However, it appears that for patent applications filed after the decision, the grant rate quickly returns to levels similar to those seen before Alice. In the same vein, the likelihood applicants would receive a Section 101 rejection did not vary significantly over time (Figure 1b). We conduct a range of additional tests regarding a potential substitution towards Section 102 rejections (novelty) and Section 103 rejections (inventive step) and a comparison to another technological class, which all tend to support the finding that Alice's impact on the certainty of a patent grant was relatively short-lived (Figure A3, A4, A7, and A8 in the Appendix). These findings are in line with a new study by the USPTO according to which Alice increased subject-matter eligibility rejections in the short-term[29].

Figure 1: Patenting AI

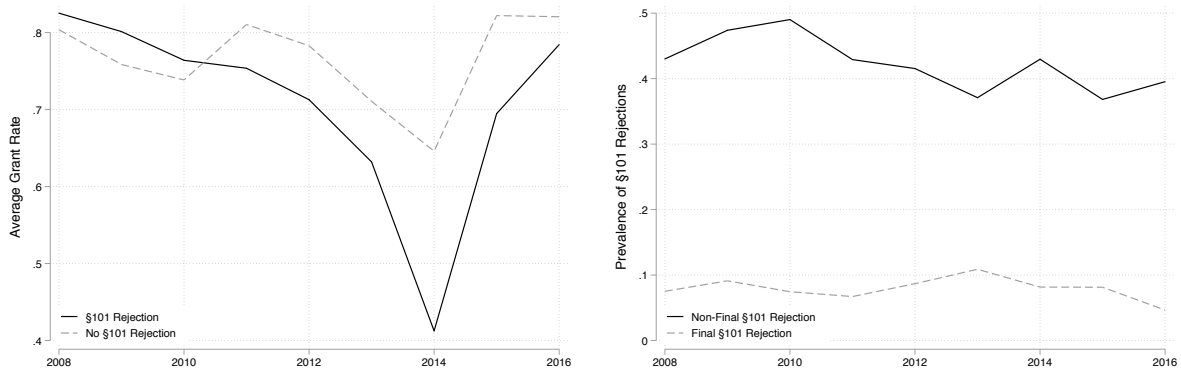

(a) Grant Rate over Time.

(b) Rejections over Time.

Figure 1a depicts the average grant rate for AI patent applications, differentiating them based on whether or not they received a non-final Section 101 rejection; Figure 1b depicts the prevalence of a non-final and final Section 101 rejection for AI patent applications over time (n=5,806). Pending applications were excluded.

The likelihood that an application receives a Section 101 rejection or is granted provides some insight into the uncertainty that patent applicants may be faced with. However, the application process is lengthy, with multiple stages at which an applicant's claimed intellectual property can be curtailed, altered, or granted. To further explore the degree of patenting uncertainty, and the extent to which it may have changed over time, we construct an uncertainty index. This index is based on a number of potential outcomes for applicants during prosecution: receiving a non-final rejection, receiving a final rejection, abandoning the application and having claims allowed. The index is built by averaging the standardized outcome variables, with a score normalized to [0,1] (details in the Appendix). Figure 2 shows that, when simultaneously assessing these various outcomes that each reflect a degree of uncertainty in the application process, uncertainty indeed increased after Alice in 2014. Yet, those uncertainty levels remained lower than in the early 2000s, and, additionally, decreased in 2016.

Figure 2: Uncertainty in Patenting AI

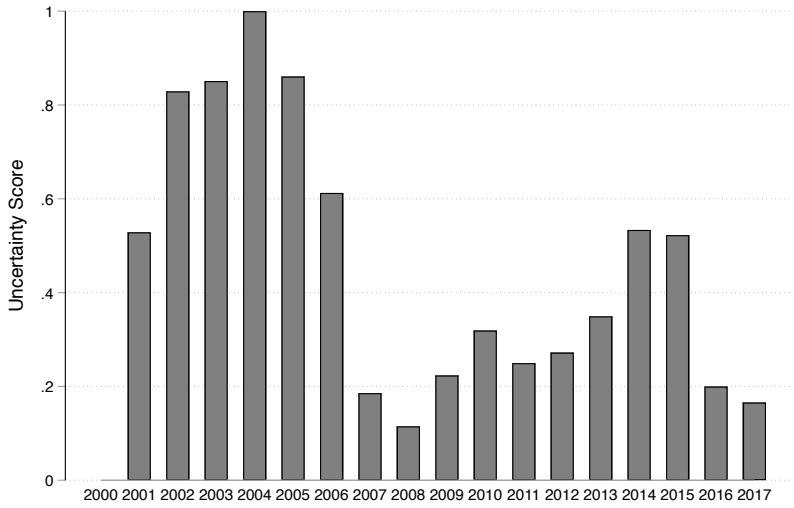

Figure 2 represents the normalized score of uncertainty in patenting in AI as a result of averaging the yearly likelihood of patent applicants to receive a non-final rejection, to receive a final-rejection, to abandon the application, and to have claims not allowed (n= 12,831). Pending applications were excluded.

Although this empirical approach provides insight into the uncertainty facing patent applicants, it does have a number of limitations. Applicants might well be drafting different type of patent applications, for instance by including narrower claims, or they may attempt to counter the uncertainty by putting more time and resources into the grant procedure[30]. Studies by Chien et al. (2020)[31] and Kesan and Wang (2020)[32] find that software patent holder adjusted their application strategies by filing amended patent claims and fewer patent applications post-Alice. These alternate avenues through which uncertainty might manifest deserve further research. Our data is also limited in time; the impact of new examiner guidelines and court decisions cannot be examined. In its investigation of Alice, the USPTO found that since the issuance of the USPTO's memorandum "Change in Examination Procedure Pertaining to Subject Matter Eligibility, Recent Subject Matter Eligibility Decision" in April 2018 and the subsequent Revised Patent Subject Matter Eligibility Guidance in January 2019, uncertainty in Alice-affected technologies has declined[33].

4. Conclusion

The empirical assessment of AI patent prosecution suggests Alice's impact on uncertainty has perhaps been exaggerated in the current policy discourse. Patent applicants appear able to internalize new legal standards rapidly, even in the face of a highly debated and politically charged court decision. The levels of uncertainty, may in fact be linked to the inherent difficulty in delineating abstract AI claims from technical implementations, or simply to difficulties in patenting new, emerging technologies. The cost of the existing uncertainty is notable; patent applicants must invest resources to overcome the mismatch of their perception of patent-eligibility and that of patent examiners. As a result, smaller firms and start-ups might well be the most severely affected.

The temporary uncertainty that followed Alice highlights the implications of relying on courts to deal with complex patent policy issues. This approach may harm applicants when standards change abruptly as a result of court decisions. Firms with pending AI patent applications in the US in 2014 exemplify this – only a fraction of their patents ended up being granted. With countries worldwide discussing the introduction and merits of an Alice test, this study suggests that those jurisdictions considering changes to their patent eligibility policy should do so carefully, so as to minimize harm to inventors and firms. The current proposed legislative changes to Section 101 in the US represent an alternate path to patent policy development. Whereas Alice injected relatively sudden—and apparently temporary—uncertainty into the patent system, legislative changes by their nature provide greater notice to participants in the patenting process and are likely to result in lower degrees of uncertainty shock.

Notes

1. Press release under https://www.tillis.senate.gov/public/index.cfm/press-releases?ID=B521846C-594A-46BE-B17A-0E113
2. Oral testimony, Kappos D. J., https://www.judiciary.senate.gov/imo/media/doc/Kappos%20Testimony.pdf.
3. Oral testimony, Dickinson T. Q., https://www.judiciary.senate.gov/imo/media/doc/Dickinson%20Testimony.pdf.
4. Alice Corp. v. CLS Bank International, 573 U.S. 208, 134 S. Ct. 2347 (2014.).
5. Diamond v. Chakrabarty, 447 U.S. 303 (1980).
6. Burk, Dan L. and Lemley, Mark A, 'Policy Levers in Patent Law' (2003) 89(7) Virginia Law Review 1642.
7. Brenner v. Manson, 383 U.S. 519 (1966).
8. US9406017 'System and Method for Addressing Overfitting in a Neural Network' granted in 2016; US10489701 'Generating Responses Using Memory Networks' granted in 2019 and US20160162802A1 'Active Machine Learning ' granted in 2019, see Zingg, Foundational Patents in AI, in: Hilty/Lee/Liu, Artificial Intelligence and Intellectual Property, Oxford University Press, 2021, [page].
9. Available under https://www.scotusblog.com/case-files/cases/alice-corporation-pty-ltd-v-cls-bank-international.
10. Mayo Collaborative Servs. v. Prometheus Labs., Inc. - 566 U.S. 66, 132 S. Ct. 1289 (2012).
11. Gottschalk v. Benson, 409 U.S. 63 (1972).
12. Digitech Image Technologies, LLC v. Electronics For Imaging, Inc., Case No. 13-1600 (Fed. Cir. July 14, 2014).
13. Zingg, ibid, [page].
14. Enfish, LLC v. Microsoft Corp., 822 F.3d 1327 (Fed. Cir. 2016).
15. Bockman, J., Kim, R. Y. and Yuan, A. Patenting Artificial Intelligence in the U.S. – Considerations for AI Companies, Morrison & Foerster LLP, November 8, 2018 https://www.mofo.com/resources/publications/181108-patenting-artificial-i html.
16. Non-Final Rejection by the United States Patent and Trademark Office, Application 14/562,747, August 30, 2018 'Active Machine Learning'.
17. Bascom Global Internet Servs., Inc. v. AT&T Mobility LLC, 827 F.3d 1341, 1350 (Fed. Cir. 2016).
18. Non-Final Rejection by the United States Patent and Trademark Office, Application 14/562,747, August 30, 2018 'Active Machine Learning'.
19. This strategy has been followed by a number of high-profile and prolific applicants. Google states that: 'we have found that when we draft applications and claims to clearly explain how the invention provides a technical solution to a technical problem, we draft higher-quality applications that meet with more success both at the U.S. PTO and in foreign patent offices, see Comments of Google Inc. before the USPTO on Legal Contours of Patent Subject Matter Eligibility, January 18, 2017.
20. Patent Trial and Appeal Board, October 10, 2019, Application 14/815,940 "Method for calling for preemptive maintenance and for equipment failure prevention".
21. The Impact of Alice: A Swinging Pendulum, World Intellectual Property Review, May 10, 2017, https://www.worldipreview.com/article/the-impact-of-alice-a-swinging-pendulum.
22. Oral testimony, Kappos D. J., https://www.judiciary.senate.gov/imo/media/doc/Kappos%20Testimony.pdf.
23. Saltiel, Joseph, In the courts: five years after Alice – five lessons learned from the treatment of software patents in litigation, WIPO Magazine, August 2019.

24. Opinion by Judge Newman, CLS Bank International v. Alice Corporation Pty. LTD, 768 F. Supp. 2d 221 (D.D.C. 2011).

25. Brief of Amici Curiae Trading Technologies International Inc. et al. in Alice Corp. v. CLS Bank (available under http://sblog.s3.amazonaws.com/wp-content/uploads/2013/11/13-298-Trading-Technologies-International-Inc.-et-al.pdf).

26. Brief of Amicus Curiae Gibbons Institute of Law, Science and Technology in Alice Corp. v. CLS Bank (Available under http://sblog.s3.amazonaws.com/wp-content/uploads/2013/11/Alice_brief.pdf).

27. Brief of Amici Curiae Law, Business, and Economics Scholars in Alice Corp. v. CLS Bank (available under digitalcommons.law.scu.edu/facpubs/853).

28. (1) US9390373 'Neural network and method of neural network training' granted in 2016; (2) US9679258 'Methods and Apparatus for Reinforcement Learning' granted in 2017; (3) US9946783 'Methods and systems for classifying data using a hierarchical taxonomy' granted in 2018; (4) US10096107 ' Intelligent medical image landmark detection' granted in 2018.

29. USPTO, Office of the Chief Economist, Adjusting to Alice, USPTO patent examination outcomes after Alice Corp. v. CLS Bank International, IP Data Highlights 3, April 2020.

30. Aboy, M., Crespo, C., Liddell, K., Minssen, T., and Liddicoat, J. Mayo's impact on patent applications related to biotechnology, diagnostics and personalized medicine, Nature Biotechnology 37, 513–518 (2019).

31. Chien, Colleen V., Halkowski, Nicholas, He, Maria, Swartz, Rodney, The Impact of 101 on Patent Prosecution – Post Guidance Updates, Patently-O Patent Law Journal 20 (2020).

32. Kesan, Jay P., and Wang, Runhua, Eligible Subject Matter at the Patent Office: An Empirical Study of the Influence of Alice on Patent Examiners and Patent Applicants, Minnesota Law Review 2020, 527.

33. USPTO, Office of the Chief Economist, Adjusting to Alice, USPTO patent examination outcomes after Alice Corp. v. CLS Bank International, IP Data Highlights 3, April 2020.