

November 8, 2019

**BY EMAIL (AIPartnership@uspto.gov)**

U.S. Patent and Trademark Office  
P.O. Box 1450  
Alexandria VA 22313-1450

**Re: Patenting Artificial Intelligence Inventions  
PTO-C-2019-0029  
Comments of R Street Institute**

To whom it may concern:

The R Street Institute respectfully submits these comments in response to the request for comments dated August 27, 2019 (PTO-C-2019-0029), for which the deadline for response was subsequently extended to November 8, 2019. The R Street Institute thanks the USPTO for the opportunity to discuss the important issues presented in the request for comments.<sup>1</sup>

R Street is a nonprofit, nonpartisan, public policy research organization that engages in policy research and outreach to promote free markets and limited, effective government.

The Request for Comments poses several questions that could be answered either as a matter of current positive law or as a matter of normative recommendations for congressional policy change. While the Office obviously has important statutory responsibility to advise the Executive on matters of patent policy, *see* 35 U.S.C. §§ 2(b)(8)–12, the first and foremost responsibility of the agency is to apply current law in the course of patent examination. Accordingly, this comment first provides answers to the USPTO’s questions under current law, and it then offers suggestions on future policy directions.

## **I. Application of Current Law**

This section considers the application of current statutory law to the questions presented in the request for comments.

### **A. Question 2**

The USPTO asks which natural persons should be considered to have contributed to the “conception of an AI invention.” Answering this question requires only reference to the

---

<sup>1</sup> Commenter thanks Harvard Law School’s Cyberlaw Clinic for its work on this comment.

longstanding definition of conception as “formation in the mind of the inventor, of a definite and permanent idea of the complete and operative invention, as it is hereafter to be applied in practice.” *Hybritech Inc. v. Monoclonal Antibodies, Inc.*, 802 F.2d 1367, 1376 (Fed. Cir. 1986); *see also Bd. of Educ. ex rel. Bd. of Trs. of Fla. State Univ. v. Am. Bioscience, Inc.*, 333 F.3d 1330, 1338 (Fed. Cir. 2003); *REG Synthetic Fuels, LLC v. Neste Oil Oyj*, 841 F.3d 954, 962 (Fed. Cir. 2016); U.S. Patent & Trademark Office, *Manual of Patent Examining Procedure* § 2138.04 (9th ed., 08.2017 rev. Jan. 2018).

In the case where the invention is a general-purpose AI algorithm (e.g., a mechanism for training an AI more efficiently), the answer is no different than that for any other form of software: The inventor is the person who formed the idea of the algorithm sufficient to enable implementation in practice. *See* 35 U.S.C. § 112(a). This applies where the invention is a trained AI system. The Office asks whether one who sets the “weighting adaptations” (presumably referring to hyperparameters, the settings that are determined outside of the learning algorithm itself, *see* Ian Goodfellow et al., *Deep Learning* 95, 117–19 (2016), <http://www.deeplearningbook.org/>) should be considered as having conceived the invention. The answer will depend on whether the hyperparameters were within reasonable experimentation by one having ordinary skill in the art—sometimes they will be, but not always. *See id.* at 415–19.

When the invention is one produced by a generative AI system, the answer is even simpler: Conception occurs in the mind of the first person to observe the invention as output of the AI system. A person who trains a general-purpose AI that is later used to design improved toothbrushes, for example, will have no mental possession of the improved toothbrushes the system would produce and thus does not contribute to conception; only the first person to observe the toothbrush design qualifies.

Perhaps it seems strange that one who may have done no inventive work is deemed the inventor simply by virtue of being the first to observe and thus conceive of the invention. But the Patent Act makes clear that “[p]atentability shall not be negated by the manner in which an invention is made.” 35 U.S.C. § 103. Furthermore, the strangeness of giving a patent to one who contributed little militates in favor of granting no patent at all, for reasons discussed below.

It is possible, of course, that a person who trains the AI process does so with a certain result in mind, and that result contributes to the conception of the invention itself. In that case, the AI trainer should be a joint inventor under traditional rules of joint inventorship; no change to the law is necessary to achieve that ordinary result.

### **B. Question 3**

The USPTO asks whether its regulations need to be revised to account for situations where “entities other than a natural person contributed to the conception of an invention.” On the assumption that “entities other than a natural person” refers to having computers as named inventors, the answer is that no changes need to be made, since computers cannot be inventors for at least two plain reasons.

First, inventors must be human under the Patent Act. The term “inventor” is defined as the “*individual* or, if a joint invention, the *individuals* collectively who invented or discovered the subject matter of the invention.” § 100(f) (emphasis added). “Individual” refers only to natural persons. *See Beech Aircraft Corp. v. Edo Corp.*, 990 F.2d 1237, 1247–48 (Fed. Cir. 1993); *MBO Labs., Inc. v. Becton, Dickinson & Co.*, 602 F.3d 1306, 1309 n.1 (Fed. Cir. 2010); *Karrer v. United States*, 138 Ct. Cl. 385, 390 (1957); *see also Mohamad v. Palestinian Auth.*, 566 U.S. 449, 454 (2012) (Supreme Court “routinely uses ‘individual’ to denote a natural person”) (citing *Goodyear Dunlop Tires Operations, SA v. Brown*, 564 U.S. 915, 924 (2011)); 1 U.S.C. § 1 (distinguishing artificial persons from “individuals”); *cf. Naruto v. Slater*, 888 F.3d 418, 425–26 (9th Cir. 2018) (holding that non-human species lack standing unless a statute explicitly provides otherwise). The Patent Act also states that a patent may be obtained by “*whoever* invents or discovers any new and useful process, machine, manufacture, or composition of matter . . . .” 35 U.S.C. § 101 (emphasis added). Finally, it instructs inventors to apply for a patent jointly if the invention was “made by two or more persons jointly.” § 116(a). Since AI systems are not considered individuals or people, they do not meet the statutory requirements for inventorship.

Second, computers do not have minds. Even artificial neural networks, despite the name, “are generally not designed to be realistic models of biological function.” Goodfellow et al., *supra*, at 13. Since, as explained above, conception requires “formation in the mind” of an invention, an AI system cannot satisfy the legal requirements of inventorship.<sup>2</sup>

### C. Question 4

The USPTO asks whether entities “other than a natural person, or company to which a natural person assigns an invention,” should be able to “own a patent on the AI invention.” The answer is no. Under current law, only the inventors or assignees thereof may be issued a patent. *See* 35 U.S.C. § 118 (providing that others may only apply for patent “on behalf of and as agent for the inventor”); § 152. Those persons are free to assign the patent, but a patent cannot issue to a legal entity except under specific circumstances where the entity can show an obligation to assign the patent right.

It is unclear why the answer should be different for AI-generated inventions, as opposed to any other type of invention. The Office suggests that one who “trains the artificial intelligence process” might merit ownership of a patent arising out of the AI’s produced inventions (presumably in situations where, as discussed above, the trainer did not contribute to conception of the invention). But this is analogous to a company that develops complex biochemical

---

<sup>2</sup> The commenter is not unsympathetic to the possibility that a substantially advanced AI system might be considered to have consciousness and free will to the point that it merits treatment on par with humans under the law. No other aspect of the law is known to act upon that possibility, and for patent law to take it into account would require a bevy of other legal changes as well, in areas such as property law, criminal and tort liability, and constitutional rights.

reagents, such as Cas9 enzymes for CRISPR gene editing,<sup>3</sup> calling for patent rights over patentable biochemicals produced using those reagents. Neither the AI trainer nor the enzyme manufacturer would have any claim to inventorship or ownership of downstream patents.

#### **D. Question 5**

The USPTO asks whether, under § 101, there are “patent eligibility considerations unique to AI inventions.” In addition to the ordinary issues of patentability of software (which are not unique to AI inventions), there is an additional issue: whether inventions produced by AI systems are products of nature.

Consider a generative AI system that designs new inventions completely automatically, where that system is fully in the public domain and operated in many different instances by many different people. Those inventions could be considered analogous to products of nature, where a mere discoverer or even isolator of the result would not receive a patent. To be sure, the generative AI system itself is not “natural” by any means, but that is not a requirement of the eligibility doctrine. In *Mayo Collaborative Services v. Prometheus Laboratories, Inc.*, the Supreme Court held a natural correlation between administration of a drug and a blood chemical marker to be ineligible for patenting, despite the fact that the administered drug was a non-natural manufacture. *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66, 73–75 (2012). By the same token, an invention arising out of the natural operation of an AI system could be treated as an ineligible product of nature.

No doubt the scope of this doctrine would be complex and difficult to discern, but it would appear to be supported under current law and thus ought to be considered by the Office.

#### **E. Questions 6–7**

The USPTO asks whether unique “disclosure-related considerations” or enablement considerations apply to AI inventions. The answer is no: The same rules apply to those inventions as others.

That being said, the USPTO must ensure that it applies the written description and enablement requirements fully to AI inventions. The training of AI systems is a highly sensitive and unpredictable process, with small changes to hyperparameters or training data having potentially major effects on the quality of the trained system. See *Goodfellow et al.*, *supra*, at 409–10. The assumption often made with ordinary computer software, that a simple flowchart or description suffices to enable the person having ordinary skill in the art to implement the patented invention, will often be incorrect in the AI context.

Disclosure of exact neural network models, hyperparameter and training settings, and training data sources will in many cases be necessary to fulfill written description and enablement.

---

<sup>3</sup> For factual background, see generally *Regents of the Univ. of Cal. v. Broad Inst., Inc.*, 903 F.3d 1286, 1289–90 (Fed. Cir. 2018).

Alternatively, patent applicants may choose to disclose complete weight matrices for patent applications on specific neural networks, which would undoubtedly satisfy the § 112 requirements with respect to AI inventions. The USPTO ought to be equipped to receive these potentially massive disclosures, likely by building infrastructures analogous to those for receiving disclosures of gene and amino acid sequences.

#### F. Questions 8–9

The USPTO asks about the effects of AI on the obviousness doctrine, and particularly whether AI would “impact the level of a person of ordinary skill in the art.” The answer is yes: The capabilities of extant AI technology increase the capabilities of one of ordinary skill in the art, thereby rendering many more technologies obvious and thus unpatentable. In particular, the advancement of AI may significantly impact the standard of obviousness under the “obvious to try” doctrine, which the Supreme Court articulated in *KSR International Co. v. Teleflex Inc.*:

When there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely the product not of innovation but of ordinary skill and common sense. In that instance the fact that a combination was obvious to try might show that it was obvious under § 103.

550 U.S. 398, 421 (2007).

The Federal Circuit has interpreted the “finite” number of solutions to be a number “small in the context of the art,” or an “easily traversed, small and finite number of alternatives.” *Ortho-McNeil Pharm., Inc. v. Mylan Labs., Inc.*, 520 F.3d 1358, 1364 (Fed. Cir. 2008). Thus, where a person having ordinary skill in the art “merely pursues ‘known options’ from ‘a finite number of identified, predictable solutions,’ the resulting invention is obvious under Section 103.” *Eurand, Inc. v. Mylan Pharm., Inc. (In re Cyclobenzaprine Hydrochloride Extended-Release Capsule Patent Litig.)*, 676 F.3d 1063, 1070 (Fed. Cir. 2012) (quoting *KSR*, 550 U.S. at 421). By contrast, an invention is not obvious where the inventor arrives at it “by ‘merely throw[ing] metaphorical darts at a board’ in hopes of arriving at a successful result, but ‘the prior art gave either no indication of which parameters were critical or no direction as to which of many possible choices is likely to be successful.’” *Id.* at 1070–71 (citing *In re Kubin*, 561 F.3d 1351, 1359 (Fed. Cir. 2009)). Thus, it is likely inappropriate to read “finite number” literally as any number that is not infinite when determining obviousness under the “obvious to try” rationale.

AI almost certainly affects what is considered a “small” number of alternatives “in the context of the art.” What is a large number of alternatives for a human may be “small” and “easily traversed” by an AI system, perhaps guided by a human. This difference in the number of alternatives that can be “easily traversed” may broaden the standard of obviousness once human inventors can take advantage of AI assistive technologies, and it certainly augments the scope of obviousness once AI technologies can produce inventions themselves. *See* Ryan Abbott,

*Everything Is Obvious*, 66 UCLA L. Rev. 2, 33–37 (2019) [hereinafter Abbott 2019]. As such, the availability of AI for the development of an invention must be considered when determining whether the invention was “obvious to try.”

It should be irrelevant whether the patent applicant actually used AI in the development of the invention. Obviousness is an objective standard based on the prior art, not the manner in which the invention was in fact made. *See* 35 U.S.C. § 103. The question for obviousness should be whether, at the time of filing a patent application on an invention, an AI that could have assisted in developing the invention was in the prior art and thus available to a person having ordinary skill in the art.

## **II. Policy for Future Legislation**

This section considers the presented questions in the context of possible future legislation.

### **A. Question 3**

The USPTO asks whether “current patent laws . . . regarding inventorship” ought to account for “entities other than a natural person” as inventors, again likely regarding granting patents to AI systems. To the extent that the USPTO is interested in whether it should advocate for policy change to enable machines to receive patents, the answer is no.

First, AI systems currently cannot respond to the incentives of the patent system at all. An AI system cannot receive money, hold bank accounts, enjoy luxury goods, or do any other thing that would enable it to enjoy the fruits of profit from its inventions. While the developers of AI systems obviously can respond to such incentives, as discussed in the section below, the computers themselves are impervious to the value of patent protection.

Furthermore, there is no determinate way to issue patents to nonhuman machines. AI systems, at least today, are software: computer programs that can be replicated and run simultaneously on multiple machines, being turned on and off at will. Issuing a patent to as ephemeral a thing as running computer code would create numerous novel issues of ownership, legal liability, inheritance, and so on, that the law is fundamentally not equipped today to handle. Patent law attempts to conclusively determine the inventor of a patent, but AI technologies are “autonomous, creative, unpredictable, rational, and evolving systems,” such that “one cannot conclusively determine an owner for these rights within the scope of patent law.” Shlomit Yanisky-Ravid & Xiaoqiong (Jackie) Liu, *When Artificial Intelligence Systems Produce Inventions: The 3A Era and an Alternative Model for Patent Law*, 39 Cardozo L. Rev. 2215, 2221–22 (2018).

Perhaps more practically, the USPTO suggests that the person who sets the “weighting adaptations,” or hyperparameters of an AI system, should merit a right to patents on inventions resulting from the AI system, even where that person did not contribute to conception of the invention. This again is fraught with indeterminacy. The particular hyperparameters that work effectively for any given AI application are often composed by a “host of parties” often working at completely independent firms. W. Michael Schuster, *Artificial Intelligence and Patent*

*Ownership*, 75 Wash. & Lee L. Rev. 1945, 1988 (2018). The chain of potential contributors could thus become an endless source of disputes and litigation.

### **B. Questions 2 and 4**

Questions 2 and 4 suggest a possibility that inventorship or ownership of patents ought to accrue to some outside the traditional boundaries of conception. In particular, the USPTO appears interested in the possibility that one who sets hyperparameters and trains an AI system that produces new inventions might receive an interest in patents on those downstream inventions. Giving an upstream inventor a stake in downstream patents would represent a radical change to the structure of the patent system, and it would be inadvisable as a matter of policy. The inventor of an AI system that generates new inventions already receives a handsome reward: a United States patent on the AI system itself. While it may be debated whether the current level of patent protection suffices, the solution for any error lies in Congress adjusting the patent right, not in granting the inventor more patents downstream. Indeed, there is at least some evidence that no additional incentive is necessary in high-tech industries such as AI, where patents tend to “have transitory value given the pace of technological change” and thus “account for only a small proportion of pioneers’ perceived quality advantages.” Yanisky-Ravid & Liu, *supra*, at 2252–53.

The USPTO’s suggestion may be read as a proposal for a mandatory grant-back clause on AI-developed inventions.<sup>4</sup> That insight reveals two reasons why such legislation would be undesirable. First, AI technologists already can obtain patent rights on downstream inventions, through private-ordering license arrangements, so legislation is unnecessary. Second, courts have occasionally held grant-back arrangements anticompetitive. *See* Richard Schmalbeck, *The Validity of Grant-Back Clauses in Patent Licensing Agreements*, 42 U. Chi. L. Rev. 733 (1975). It would be remarkable for patent law to mandate effective grant-backs on all generative AI technologies, regardless of the competitive consequences.

As noted above, the inventor of a technology produced by a generative AI system could be the human who processes the computer’s output, and thus the first individual who “discovers” the invention. In this case, the inventions generated by an AI system would only be patentable when a human being recognizes and evaluates the significance of the AI system’s results. This would also be an efficient allocation of patent rights, as one scholar finds, because “aggregate welfare is maximized by allocating the rights to them.” Schuster, *supra*, at 1950. To be sure, this would not recognize the contributions of those who designed and trained the algorithm, but as noted above, private contracts can solve those problems to the extent that a solution is necessary. *Cf. id.* at 1993–95 (predicting that the software designers likely will not do this).

---

<sup>4</sup> A grant-back clause is a provision in a patent license agreement that provides that the licensee will grant back to the licensor some rights, and perhaps full ownership, in downstream inventions based on the licensed patent.

### **C. Questions 5 and 8–9**

As discussed above, under current law it is likely that many AI-generated inventions are not patentable at all, by operation of either eligibility under § 101 or obviousness under § 103. It may be asked whether this is problematic because unpatentability of AI-generated inventions may mean diminished incentives to design generative AI systems in the first place. *See, e.g.,* Ryan Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. Rev. 1079, 1104–05 (2016).

However, there may not be a need to incentivize the developer of a generative AI technology by giving that developer a patent on all resulting downstream inventions produced by that AI. As noted above, the developer will already receive the federally prescribed incentive of patents on the AI technology itself. And once the generative AI is in the public domain and available to all, the possibility of patents on products of that AI would likely lead to unproductive patent races and a disincentive to develop new inventions beyond what the AI can produce; those detrimental consequences are avoided if the AI's products are not patentable.

In the hypothetical public-domain generative AI described in the discussion of question 5 under section I, it may seem unfortunate that a seeming invention produced by that generative AI does not confer patent rights; however, the alternative may be worse. Since multiple people are free to operate the system, inventorship would inhere randomly in whoever happened to observe the AI's output, purely as a matter of coincidence and without any inventive effort on the part of the “inventor.” (This is perhaps why many argue that the generative AI system itself should receive the patent, but that change to the law creates even more problems, as described above.) But, in the case where the generated AI technology is in the public domain, keeping its output in the public domain avoids these odd possibilities, and also ensures that human inventors have proper incentives to pursue inventions that require inventive skill beyond what the AI can do.

### **D. Questions 8 and 9**

The USPTO asks whether the standards for obviousness ought to be changed to account for AI-assisted invention. As noted above, the current law of obviousness appears capacious enough to account for the capabilities of AI. That being said, there are possible clarifications.

As one commentator argues, the standard for obviousness ought to account for an additional factor: the technology available to the skilled artisan. *See* Abbott 2019, *supra*, at 6. Thus, the obviousness analysis should likely at least contemplate the level of skill of machines. AI has faster computing power than humans and access to expansive data sets, so something that is difficult for a human to create could be obvious to a computer—and thus obvious to a human who has access to that computer. Over time, the use of AI will continue to grow, leading to AI being responsible for more, or perhaps most or all, innovation. The obviousness standard must correspondingly “evolve if it is to continue to reflect real-world conditions.” *Id.* at 9. Without such a change and without recognition of the capabilities of AI in assisting inventors, the standards for patentability may end up too lenient, leading to over-patenting with anticompetitive results. *See id.* at 5.



### **III. Conclusion**

The commenter thanks the USPTO for providing the opportunity to submit these comments. If there are any remaining questions relating to the matters presented herein, the undersigned would be happy to provide further information as necessary.

Respectfully submitted,

Charles Duan  
*Director, Technology & Innovation Policy*  
*USPTO Reg. No. 65,114*  
R STREET INSTITUTE  
1212 New York Ave. NW, Suite 900  
Washington, DC 20005