

REQUEST FOR COMMENTS
BY THE UNITED STATES PATENT AND TRADEMARK OFFICE
Comment period ends 10/11/2019

Agency: United States Patent and Trademark Office, Department of Commerce.

Document Citation: 84 FR 44889

Agency Docket No.: Docket No. PTO-C-2019-0029

Document Number: 2019-18443

Web Publication: <https://www.federalregister.gov/documents/2019/08/27/2019-18443/request-for-comments-on-patenting-artificial-intelligence-inventions>

SUBMITTER

First Name	Kim
Last Name	Rubin
Mailing Address	825 Cambridge Av. Menlo Park CA 94025
Email Address	xkrubin@gmail.com
Phone Number	650 328-3424
Fax Number	650 322-4283
Organization Name	Kim Rubin Inventor and Patent Agent
Submitter's Representative	self
Background of Submitter	See end of this document

QUESTIONS AND COMMENTS

1. Inventions that utilize AI, as well as inventions that are developed by AI, have commonly been referred to as "AI inventions." What are elements of an AI invention? For example: The problem to be addressed (*e.g.*, application of AI); the structure of the database on which the AI will be trained and will act; the training of the algorithm on the data; the algorithm itself; the results of the AI invention through an automated process; the policies/weights to be applied to the data that affects the outcome of the results; and/or other elements.

"It is of utmost importance that patents issue with definite claims that clearly and precisely inform persons skilled in the art of the boundaries of protected subject matter." MPEP 2173

Starting with the simple and clear requirement for disclosure above, we see that there is no recognition of any specific technology.

There is no such thing as a steel patent or a rubber patent. There is no such thing as a software patent or an AI patent. There are only inventions. As the US Constitution clearly states, they need only be "new and useful." Despite the fact that the US Supreme Court in several famous cases (e.g., "*Alice/Mayo*") has decided that many useful things do not legally qualify as "useful," including mathematical formulas, new methods of doing business, and games, inventing continues to occur. The individual inventors (such as myself) and organizations that fund development of new technology would like to be rewarded — and deserved to be rewarded for their efforts — with granted patents.

The world today runs on data. Sometimes it is called personal data and sometimes "big data." Data is the underlying, tangible asset that drives today's innovations and the most valuable companies in the world. Processing that data is today's version of building out of steel or putting motors into things so they become "automatic" instead of "manual." 65% of patent applications today require some portion to be implemented in software.¹

Data is composed of bits, at the lowest level. Every bit is as tangible as a compound in an alloy, or a tiny transistor on a chip, or a manufactured gene in a cell. Bits in a memory chip may be a charge on a gate of a transistor or a conducting state (v. a non-conducting state) of a transistor. The physical reality of this charge can be easily measured using common tools in the art such as voltmeters, oscilloscopes, or scanning electron microscopes. Observing these bits indirectly is even easier using consumer tools such as a computer with a screen.

A bit on a disk drive is a small magnetic domain. It may be observed with another magnet, movement under a coil of wire, and with other tools. Magnetic domains on a disk drive, although smaller, are manufactured similarly to a refrigerator magnet. Indeed, the earth itself is a magnetic domain, as easily measured by a conventional compass.

If bits, and therefore data, are "not tangible," where exactly is the dividing line between the charge of electricity in a bolt of lightning (which is tangible enough to do significant damage) and the charge of electricity on the gate of a transistor? And, where exactly is the dividing line between the earth (generally considered tangible) and a magnetic domain on a disk drive?

¹ IPWatchdog, webinar, 27 September 2019 (ipwatchdog.com)

It's time the Patent Office moved beyond the idea that somehow data, and processing data, is not tangible. Similarly the Patent Office needs to reject the idea that software is somehow "abstract." If data is abstract, then so is every cell phone, every computer, and indeed our entire wired electrical grid.² In this modern world, many of the things we interact with every day are not tangible in the 19th century sense, but nonetheless exist with definite boundaries and in specific and limited forms, are extremely useful, and can be manipulated, interacted with, used as tools, seen on a screen, rearranged, etc. The difference is only that this new type of tangibility exists as bits and is stored on a disk drive as physical magnetic domains, taking up physical space in well-defined physical locations

In *State Street Bank v. Signature Financial Group*, 149 F.3d 1368 (Fed. Cir. 1998) the CAFC ruled that a calculation that produces a number, if that number has a specific purpose, such as "a price," is "useful, concrete and tangible." The US Supreme Court, for example, in *Alice/Mayo*, did not address and did not overturn *State Street's* determination that a useful number is not abstract and is specifically patentable subject matter. There is no requirement that an element in a claim must be able to be held in one's hand.

Other patent offices in the world do not have the problem that the USPTO has believing that software algorithms are somehow not patentable subject matter. Despite the dramatic improvement of the **2019 Revised Patent Subject Matter Eligibility Guidance** (which is ignored by some examiners to this day), the USPTO still operates under the fantasy that data is not tangible and therefore data processing method steps are non-patentable subject matter. Virtually all Computer Science professionals, as well as nearly all US Patent practitioners, including academics, agree that the current examination of patent applications, and appeals, regarding eligible subject matter under 35 USC § 101 is nonsense.³

We return now the question's specific reference to AI. How exactly is an invention using AI or machine learning different than an invention that uses software? A trained neural net machine is comparable to a programmed computer. That trained neural net machine accepting inputs and generating outputs is comparable to a

² The US Supreme Court has never ruled that software is abstract. "Transitory" is not the same as abstract, and there is no reason to consider that data, which may be manufactured or converted to another form, is any more transitory than a broken bone — which most people would agree is not abstract. The treatment of data, including data collecting, data transmission, data publishing, and data structures, as abstract is primarily a figment of examiners; this error is easily corrected by proper Guidance from the Office. It may also be corrected by Congress. But why force Congress to fix an error that is easily remedied by the Office?

³ In the USPTO Roundtable, 5 December, 2016, at Stanford University, every speaker except one was adamant that the current Office examination under § 101 was absurd. (The one exception was the Electronic Frontier Foundation, but they don't think there should be any patents.)

computer executing its program: using inputs and generating outputs. Innovation, including novelty and non-obviously, should be treated the same for all applications, including both software and AI.

This is not to say that many software methods qualify for patents. The 35 USC § 102 and § 103 bars for novelty and non-obviousness are high. Despite literally millions of programs and applications (e.g., smart phone apps), only a tiny fraction of these do, or should, meet those bars. In addition, a vast number of issued patents involving software and data should have passed (but did not) the 35 USC § 112 bar, such that a person trained in the art (POSITA) would be able to “make and use” the invention. As will be discussed in more detail below, the Patent Office should require a similar level of disclosure and details for software and AI as they do for mechanical devices and chemical manufacturing methods.

The question above asks about the “problem to be addressed” in an AI invention. There are six aspects of machine learning that likely have patentable inventions: (i) selecting training data; (ii) the underlying hardware of the machine that is “learning” from the training data; (iii) software that runs on top of the underlying hardware but is part of the “learning machine;” (iv) configuration data derived from learning (e.g., “weights”); (v); hardware and software used for implementation of the end-application, or ‘deployment’; (vi) input data and output data of the end-application. An invention might be directed to more than one of the above aspects. Most applications of AI will have novelty in only one or a few of these aspects. Note that for most of the above aspects, the entire invention is directed to “data.” Therefore, if the Office expects to grant any patent for “AI inventions,” it is critical that they move past the idea that data is non-patentable subject matter.

Just as it is for “software patents,” the bar for true novelty and non-obviousness in AI-related inventions is high.

And, until practitioners change the way they write specifications, most applications would fail to pass the § 112 bar, too. How § 112 applies to AI applications is discussed in more detail below.

2. What are the different ways that a natural person can contribute to conception of an AI invention and be eligible to be a named inventor? For example: Designing the algorithm and/or weighting adaptations; structuring the data on which the algorithm runs; running the AI algorithm on the data and obtaining the results.

Similarly to a “software patent,” discussed above in the reply to Question 1, a natural person can contribute to an AI invention. First, by improving the operation of an AI core — same as improving the function of a computer. Second, by coming up a novel and non-obvious use of an AI core in a useful and fully disclosed application. Third, by new forms of training data. Fourth, by new forms of output data. Fifth, by finding

new applications for the AI technology, including finding new problems to solve. And sixth, by novelty of the input or output used in the end-application. In many cases, the AI training core is very different than the AI application core.

As courts have said, “sometimes identifying the problem is the invention.” Examiners should not be permitted to use impermissible hindsight of the technological problem solved and disclosed by the invention to claim obviousness. If the examiner cannot find the problem to be solved well-identified in the prior art, then at least that portion of the invention is new. For new problems, reading the “claim as a whole” (which is technically required now, but often ignored by examiners), and if the problem is sufficiently identifiable in the claim (including the preamble), the invention is patentable.

The question presented above is, “how can a natural person contribute” to the invention. The answer is: same as now. “Each joint inventor must contribute to the conception of the invention. *Burroughs Wellcome Co. v. Barr Lab., Inc.*, 40 F.3d 1223, 1227-28 (Fed. Cir. 1994). That is, an inventor is a natural person who contributes to any part of any claim. I emphasize that “contribute” refers to the conception of the idea, not the implementation of the idea.

The possibilities listed in the question above of ways that a person is eligible to be named as an inventor are all correct — assuming that the claim itself reads directly on the contribution. For example, if a claim of the invention is directed only to input data, then only a person who contributed to the novelty of the input data is an inventor.

It is worth noting, importantly, that the question above includes the text, “structuring the data.” This text is an implicit recognition that “data is a [patentable] thing,” and that “structure of data” is also patentable subject matter. After all, it is called a “data structure” because it is indeed a structure: a designed (and potentially invented) form. New Guidance must make this clear. Of course the data structure must be sufficiently defined to pass § 112, and patentable subject matter does not directly impact examination under § 102 and § 103.

3. Do current patent laws and regulations regarding inventorship need to be revised to take into account inventions where an entity or entities other than a natural person contributed to the conception of an invention?

No.

“Entities” include companies, which we know are not inventors.

The question asks if “entities other than a natural person contributed to the conception.” I argue that only natural persons can “conceive” of an idea. I think it will be a very long time before a machine tells, me, “I have an idea for a new invention.”

4. Should an entity or entities other than a natural person, or company to which a natural person assigns an invention, be able to own a patent on the AI invention? For example: Should a company who trains the artificial intelligence process that creates the invention be able to be an owner?

No.⁴

The answer to the second part of the question above is: if the invention is directed to “training the artificial intelligence process,” then the contributing natural persons are the inventors. If the bounds of the claim do not include training, there is no reason for an examiner to directly consider that aspect of innovation.

In other words, there is no need to change current law or regulations, assuming they are properly applied.

5. Are there any patent eligibility considerations unique to AI inventions?

As explained extensively in answers to the other questions herein, once patent eligibility for “software inventions” is properly resolved, there will be no need for exceptions to “AI inventions.”

6. Are there any disclosure-related considerations unique to AI inventions? For example, under current practice, written description support for computer-implemented inventions generally require sufficient disclosure of an algorithm to perform a claimed function, such that a person of ordinary skill in the art can reasonably conclude that the inventor had possession of the claimed invention. Does there need to be a change in the level of detail an applicant must provide in order to comply with the written description requirement, particularly for deep-learning systems that may have a large number of hidden layers with weights that evolve during the learning/training process without human intervention or knowledge?

There are indeed disclosure-related considerations for AI invention. Such disclosure requirements should be similar to disclosure requirements for software inventions. However, disclosure requirements for software inventions are not provided in current Office Guidance. And, a vast fraction of issued US patents for software inventions, in fact, have grossly deficient disclosure.

⁴ Note that text in the first question of Question 4 is punctuated incorrectly. It should be, “Should an entity or entities, other than a natural person or assignee, be able to own a patent on the AI invention.” See title of book: *Eats, Shoots & Leaves: The Zero Tolerance Approach to Punctuation*.

The Office needs to take seriously the need for proper Guidance for disclosure of software (and then, by extension, AI) patents. Source code in some form (e.g., actual source code, pseudo code, commands to identified applications, identified open source libraries, and the like) should be a disclosure requirement, in most cases. Vague block diagrams that repeat claim steps are rarely adequate.

Note however, that disclosure need only disclose ONE embodiment. Practitioners should not be scared to include source code because they believe that such source code limits the scope of the invention.⁵ New Guidance should make clear that one detailed embodiment is not limiting — only claims are limiting. Limitations imposed under 35 USC § 112(f) are discussed below.

All software patents that include algorithms⁶ are inherently method claims, as all software — including object oriented programming, exception handling and parallel execution units — executes sequential steps. All software steps inherently perform a “function.” That is, after all, the point of including that step in the software algorithm. For example, consider a simple SINE(x) execution step. Indeed, SINE(x) is even called a “function call.” There is no reason that a specification should include examples of how to compute a SINE(x) function. Guidance should go back to the old “means steps are identified by the phrase, ‘means for’ in the claim.” Plus a very small set of nonce words, like ‘module.’” Module by itself is a nonce word because there is no function call for a ‘module.’ The list of acceptable nonce words to be interpreted as equivalent to “means for” in the claim should be in the new Guidance.⁷ MPEP 2181(I) discusses nonce words triggering a § 112(f) interpretation. A nonce word is defined in Wikipedia (https://en.wikipedia.org/wiki/Nonce_word) as “essentially meaningless and disposable and are useful for exactly that.” It should not be up to an examiner to decide what words are “meaningless.” For a word to be meaningless is a high bar. New Guidance needs to limit examiner’s arbitrary decisions on construing a claim limitation as triggering § 112(f) to a specific and short list of words or phrases.

The question refers to “hidden layers” than can change attributes without additional human intervention or knowledge. The entire point of every machine is that, once built and configured, is that it can perform some kind of work. The machine’s work product is typically at least partially automated without additional human direction or knowledge. Or, the machine may produce more work than a human alone could produce. For example, consider a simple thermostat. The thermostat has “hidden

⁵ Should an applicant fear that a portion of a claim (or god-help-us, every limitation in a claim, as some examiners have done) be treated as a “means” under 35 USC § 112(f), and that source code in the disclosure will thus be the limiting construction of the claim element, the applicant should note that if there no source code then the claim as a whole will fail because there is no structure in the specification to implement the means. Thus, inclusion of additional details of an embodiment should never weaken a claim.

⁶ Algorithms as part of a claim directed to a specific, non pre-emptive application.

⁷ Such a list in Guidance meets court-created doctrine that means steps don’t have to actually include the phrase, “means for.”

layers” that adjust temperature without additional human intervention. Thus, AI and deep-learning machines are still machines. The fact that there are elements of the machines that “evolve” is the nature of many machines. After all, both internal combustion engines and shoes need to “break in” — that is, evolve. A computer that is programmed then processes data. The exact details of the data, such as individual bits or a person’s name in a database, are not known to the programmer in advance. From a patent viewpoint, it does not matter if, during training, an AI component “learns” various internal weights. The invention did not morph into a different invention. The machine as a whole still operates as designed.⁸ A fully trained AI machine may be an invention unto itself, but the fact that it was “programmed” is no different than a general-purpose computer programmed to perform a specific task. A programmed general-purpose computer is patentable subject matter separate from the un-programmed computer. If a patent is directed only to the as-programmed box, it does not matter who (or what) did the programming.^{9, 10}

That is, a system with an AI component is still a machine. No special considerations are needed just because it uses an AI component. As said above, “There is no such thing as a software patent or an AI patent. There are only inventions.” What is required is proper application of current patent laws for disclosure.

The question asks, “Does there need to be a change in the level of detail an applicant must provide?” The level of detail to meet § 112 requirements has always been both a function of both the field of art and the level of knowledge in the art. If a particular type of neural net is required, for example, that type must be disclosed. If that type is known in the art, simply identifying it clearly is sufficient. If the invention is directed to a novel architecture within the neural net, then that new architecture IS the invention and the architecture must be disclosed so that another person in the art could make and use that architecture without undue experimentation.

⁸ The full capabilities of a machine are often not known in advance of implementation. For example, detailed flight characterizes of an airplane are not known until an actual plane is tested. Similarly, the success of a languages translation machine may not be fully known until after training.

⁹ The issue with copyrights may be different than patents. If a programmed computer writes poetry (some do), who owns the copyrights? Fortunately, courts have already decided that only “natural persons” (e.g., not animals) can own copyrights. Is a “creative work” that a copyright protects different than the “conception” of an idea that a patent protects?

¹⁰ Even of some aspect of the as-programmed box was generated by other than a “natural person,” there must have been some natural person who contributed to the invention. Since anyone who contributes any portion of a claim is an inventor, it is moot if some “non-natural persons” also contributed. Only the natural persons need be named as inventors. It is hard to conceive that a patentable invention is 100% conceived by a machine. And if so, did the machine itself sign away its inventorship rights to an assignee? Can an assignee be another machine? Ultimately, it may be that inventions created entirely by machine are in the public domain, as they cannot get patents. This is probably good. Who wants patent trolls that are entirely machines?

The issue of “what level of detail” is required for disclosure comes up with every new technology. The people in the art — POSITAs — know what level of detail is required to make and use the invention. A more relevant question for the Office may be, “how do we find POSITAs to be new examiners?” However, that particular question has not been put out for public comment.

7. How can patent applications for AI inventions best comply with the enablement requirement, particularly given the degree of unpredictability of certain AI systems?

Every new technology introduces some level of unpredictability. Nobody knew how fast or how high the first airplanes could fly until they were tested. That did not make airplanes magic. New classes of drugs do not have fully predictable effects until extensive human trials are done. That does not mean the drugs are magic. Similarly for AI, the exact performance of machines is not known prior to training. The AI machine may produce some unexpected output. But then, drugs may produce unexpected side effects and planes may have unexpected failure modes.

The inventors have to disclose what they have invented. They do not need to disclose details of the performance of the invention. As examiners are fond of saying, “the invention is the structure, not the benefit.”

One might argue that a common AI platform performs one way with a first set of training data and another way with a different set of training data. Well, a computer performs one way when programmed with a first set of instructions and a different way with a second set of instructions. As discussed above, AI inventions have numerous parallels to software inventions.

An AI invention might be directed to the training data; or to the architecture of an AI core; or to an overall application requiring at least three elements: an AI core, training data, and a deployment (use after training). The level of disclosure and nature of disclosure depend hugely on to which aspect of AI the invention is directed.

Because of the above multiple aspects of AI, it is neither appropriate nor useful to refer, generally, to an “AI patent.” Just as, to repeat, it is not particularly useful to refer to a “steel patent” or to a “software patent.”

8. Does AI impact the level of a person of ordinary skill in the art? If so, how? For example: Should assessment of the level of ordinary skill in the art reflect the capability possessed by AI?

To the first question: yes, as it does for every new technology. To the second question: a person of ordinary skill in the art needs to be trained in the art. AI is a

new field of art, and there are now extensive courses, degrees, careers, devices, benchmarks, conferences and experience in that field. Becoming trained in the art follows the same basic education and experience arc of other technical fields.

The third question discusses the “capability possessed by AI.” It is easy to be wowed by the new capabilities of any new technology, whether it be airplanes, computers, or the internet, for example. It is not trivial to “assess” a level of skill in an art whether that art is old or new. This is not a new problem. As always, such assessment is best performed by those already in the field. The Office might well ask how is the level of skill in an art assessed of examiners in that art unit. The Office works diligently to expose neither this question nor the answer.

9. Are there any prior art considerations unique to AI inventions?

No.

However, in all new fields of art, the rate of innovation, which then rapidly becomes prior art, is fast. Most of that newly-formed knowledge will not be in the form of published patents. Finding prior art requires reading and working within the industry. For AI and machine learning, as for all new technologies, examiners who look simply at patent publications will not be effective.

Looking simply at patent publications is not particularly effective prior art searching even for well-established arts. Other patent offices in the world find ways to search and review publications other than patents. The USPTO is well behind other patent offices in this regard. US examiners, for example, almost never refer to college textbooks or conference proceedings, and are even less likely to refer to Masters or Ph.D. theses.

10. Are there any new forms of intellectual property protections that are needed for AI inventions, such as data protection?

Yes.

There are indeed new forms of intellectual property protections that are sorely need for AI inventions. However, the requirement for new forms of IP has been around for a while. Distributed inventions, such as those that use computers, communications, the internet, or distributed elements also need new forms of IP.¹¹

¹¹ The Office has created a hole in IP protection in between design patents and utility patents. A device might well deserve protection both for its ornamental design and its novel structure for increased utility. However, the Office requires disclosure to examiners if the “other type” of patent is pending on the same device. Examiners then use this as an excuse to deny both patents. The design patent examiner says the design is “useful,” as evidenced

It is hard to put square pegs in round holes. The “round holes” refers to old, traditional inventions — the kind you could buy in a hardware store and hold in your hand. New inventions might comprise thousands of hand-held devices sending data to an arbitrary server, which then performs computations and sends results to yet another party. Such a system might well have been invented (conceived and programmed) by a small group of people, who got exactly what they hoped for: a new machine with vast benefits and commercial value. However, the Office is unable to issue a patent on this invention because they can’t get this new “system” square peg to fit into a 19th century round hole.

An even more challenging example is commonly referred to as a “platform.” Interface to the platform may be through an “application programmer interface,” or API. Sometimes such APIs are called “frameworks.” For platform inventions, there are two levels of users of the invention. First are the “application developers,” who use the platform’s resources to build a new Invention (aka, a “program” or “app.”) Second are the users of that app.

The US Patent Office does not currently know how to examine such “system” inventions. This is due, in part, because the underlying structure of patent claims today is unsuitable for system-type inventions.

There are three fundamental types of US patents: utility, plant, and design. Each of these has completely separate rules. Within utility patents there are three main claim structures: the well-known device and method claims, and the lessor-known recipe or kit structure (elements are listed but relationships between the elements are not). There is also a “system” claim structure, identified in the preamble, that examiners, with remarkable inconsistency, don’t know how to examine. Are system claims device claims or are they method claims? Can they be a mixture of the two? Can they refer to other claims in the body of the claim? If they do refer to other claims, are they dependent or independent? Examiners typically (but not always) reject claims that have a mix of “device” limitations and “method steps.” Yet, a system is, by its nature if not definition, precisely such a mix. The system has well-defined physical elements (such as smart phones, servers, cell-towers, and the like) and critical methods of communicating between these elements (such as internet packets, encrypted wireless, and the like). The novelty of the system requires considering the “claim as a whole.” However, in this case, the “claim as a whole” is the “system as a whole” which includes BOTH device elements and method steps.

solely by the utility patent application on the same device (but with different claims). The utility examiner insists that the novel elements of the utility claims are all “design choices,” with no more evidence than the existence of an application for a design patent, and without actual examination of the utilitarian structures. Arguments to both examiners that innovation “deserves” a patent fall on deaf ears: “each type of patent is handled independently,” we are told. This is obviously not true. The Office needs to correct this deficiency in examination policy.

Such a “system” claim structure does not currently exist. That is, it is not currently possible to apply for a proper system invention at the USPTO.¹²

Such “system” inventions also have a raft of issues regarding ownership and infringement. These topics would normally be for another Public Comment. However, in question 10 above, the Office has opened the Pandora’s box of “a new form of intellectual property.”

The solution is not overwhelmingly complex and it does not require new patent law. The solution is to adapt the current “system” claim type to be both more comprehensive and more flexible. It should specifically recognize the concept of a “distributed” invention and be tolerant of distributed implementation, control, and ownership of the elements of the system, including both the physical elements (e.g., a smart phone) and the communication methods (e.g., cellular data packets). The application must make clear the bounds of the invention. In the case of system patents, the inventor should be quite clear where the novelty lies. Examiners, now, are likely to look only at the elements individually, where every individual element is well-known prior art, and simply conclude without any analysis of the system as a whole that it is obvious under 35 USC § 103 because the individual elements are known. Virtually all patent claims consist of elements, which in isolation, are known. Examiners should not be permitted their current tomfoolery for software method claims, and system claims, which they would never apply to a novel machine for sorting gravel, for example.

Currently in the US, patent claims do not, within the claim, identify the point of novelty. This is unfortunate and is not the case in the EPO. For example, the preferred “Jepson” claim structure identifies the prior art in the preamble and then expresses the point of novelty following. For my new proposed “system” claim, the Jepson format should be required. Otherwise, it would be very hard for an examiner — or any reader of the patent — to tell exactly what the invention lies.

(US practitioners dislike the Jepson format because they fear that ANY concession of prior art will harm some hypothetical breadth. This is unfortunate, as the inventor is supposed to describe the metes and bounds of “her invention.” If the patent applicant did exactly that — “patent the invention” — there should be no such fear. Currently, both examiners and practitioners waste vast amounts of time and money arguing the obvious — that 10 out of 12 elements in the claim (for example) are “well-known.” Examiners sometimes simply ignore (at least in the first Office action) the actual points of novelty. Why should an examiner have to search and argue about 10 elements of obvious prior art? If the practitioner simply pointed out in advance (such as in a Jepson preamble) the prior art, then everyone could concentrate on examining the novelty of the invention, rather than repeatedly

¹² At least not without violent contortions, such as listing method steps “wherein” they use structural elements, or listing structural elements “wherein” they communicate using a method. When there are only two shapes of holes, large hammers are needed.

wasting time pointing out the obvious. Ideally, the USPTO would require that ALL patent claims be in Jepson (or equivalent) format. But that is not going to happen.)

I now discuss how such a new “system” claim is appropriate for AI inventions. By their nature, AI inventions are often, but not always, such systems. For example, it is likely that the dataset used for training is not owned by the inventor. It is also likely that the user or beneficiary of the invention is removed from both the hardware of the invention and the methods of the invention. In modern AI, the hardware used for training is typically far different than the hardware used for deployment. An AI invention might be directed to only the training side; or it might be directed to only the deployment (use after training) side; but the complete AI system likely includes components of both training and deployment.

Therefore, the only rational and appropriate way to examine and issue AI claims is to directly consider all of these AI system elements, how they are related, and exactly what parts are the nexus of the invention.

As one example, consider an invention that identifies flower type from an image of a flower. The training set may start with a widely-available corpus of flower images. However, this corpus of images has to be carefully filtered before use in training. Currently, any such filtering operation would be “non-patentable subject matter” according to the USPTO, as examiners simply say that “data processing is abstract” and that there is nothing novel about one collection of flower images over another collection. Additionally, the examiner will say data is not patentable, so then neither the method of filtering nor the created corpus are patentable subject matter. This is wrong. Data is not abstract, and a new method of sorting data is patentable subject matter (although it may not be novel). An examiner will move down the line, arguing that every step is either “well-understood, routine, and conventional” or is not patentable subject matter. If, on the other hand, the applicant had to describe exactly which elements and method steps were individually known prior art (perhaps all of them), she could then focus on the point of novelty.

For these proposed new “system” claims, which are appropriate for AI inventions, the applicant must clearly specify, ideally in the claims but at least in the disclosure, exactly what is novel. The applicant should say (indeed, should be required under § 112) how the output or use of the invention produces a result not previously available.

For AI inventions, it may not be appropriate to describe exactly the “technical problem to be solved.” For example, the problem for the above example is, “fast, automatic identification of flowers from a just-taken smart phone photo.” An examiner might think this is not a “technical problem.” An examiner today is likely to argue that such an invention “could be performed manually, and is therefore not patentable subject matter.” The fact that no human could memorize 100,000 flower types, or provide this capability to anyone in the world on a hike within a few seconds, falls on deaf examiner ears. However, if the invention can deliver this capability, and this capability did not previously exist, the invention should be

patentable subject matter, even if all of the individual elements were previously known or could, in some extreme sense, be done manually.

As a second example, consider an invention that allows a person to turn off lights with a simple voice command. A novelty might be that the locations and names of the lights do not have to be manually specified. Today, an examiner is likely to say, “turning off lights is not novel and neither is voice recognition.” Yet, the benefit of the invention as a whole did not previously exist.

AI inventions should be patentable subject matter if they produce a result not previously disclosed. Today, no such argument for the applicant exists.

This approach will not produce a flood of patented uses of AI. “Me too” applications are still subject to obviousness rejection under § 103. For example, an AI invention that identifies flowers might be patentable, but then the same approach used to identify leaves, or wild animal scat, is probably not.

For AI inventions, overcoming rejections may involve “secondary evidence,” such as commercial success. However, both examiners and the courts have deprecated any such secondary evidence to the point that it is rarely worth the trouble to bring it up.

Consider “commercial success.” MPEP 716.03 describes how commercial success must be clearly tied to the nexus of the invention. Often this link is hard to definitively prove. However, for AI inventions, it may not be so hard. For example, one might need only to show that you have 100,000 users of an app. The Office should update the use of secondary evidence for patentability to be more receptive to real-world evidence of novelty. “Proving” that any success, such as number of users, number of “likes,” or sales revenue, is solely a result of and only a result of the nexus of the invention is nearly impossible. The Office, through updates to all of MPEP 716 (or Guidance in advance of MPEP revision) should make examiners more receptive to real-world success of an AI — or any “system” — invention, as legitimate arguments to overcome a § 103 rejection.

It is worth noting that determination of obviousness for AI inventions will likely be more challenging than for inventions in other art units.

11. Are there any other issues pertinent to patenting AI inventions that we should examine?

Invite and request industry trade groups to adopt formal recommendations for patent applications and patent examination. Both the USPTO and the courts regularly complain that they have lack of industry guidance on decision-making. Questions relating to POSITA in the art of AI, including training, reference material, and benchmarks for improved performance — are best answered by formal

recommendations published by recognized trade groups or ad-hoc groups of credible professionals.

Most trade associations and technical professionals (engineers) do not believe they have any impact on patent policy. They feel they have to “live with a broken system.” Credible outreach, such as a presentations or panel participation at conferences will go a LONG way to getting meaningful input from leading people in the art.

Have the Silicon Valley office hold more events —they should be technical input events, not social events. Have senior representatives from the USPTO from Washington D.C. at the events.

Fix examination of “software inventions,” first.

Train examiners to actually follow the 2019 Guidance.

Have technical specialists in the art unit (not examiners or SPEs) that SPEs can call on to help out on cases that require strong technical input.

Hire outside experts to give training classes just for the AI art examining corps.

Allow “commercial success” to be a much stronger secondary consideration on overcoming both § 101 and § 103 rejections. “Long felt need” is not meaningful for new technologies and new markets. Similarly, “failure by others” is also not meaningful for new technologies and new markets. Currently, the bar for using secondary considerations is too high to be useful for nearly all technology patents. “Industry recognition” should also be a secondary consideration of new technologies.

12. Are there any relevant policies or practices from other major patent agencies that may help inform USPTO's policies and practices regarding patenting of AI inventions?

Yes. The EPO has their act together on software and AI inventions. It appears that the KIP does too.

--

BACKGROUND OF SUBMITTER

Mr. Rubin graduated from U.C. Berkeley with a degree in electrical engineering and computer science. He has been an entrepreneur for 30 years, cofounding five companies.

As a successful independent inventor, Mr. Rubin has over 300 inventions, including 18 issued patents in the vehicle-to-vehicle safety communication space.

Mr. Rubin has been a full-time patent agent for eight years, with approximately 100 issued patents for clients.

Mr. Rubin has written three books, does public speaking, has taught college Computer Science courses, and has written two international IEEE/ANSI Standards.