

Response to USPTO request for comments on patenting artificial intelligence inventions  
Randall Davis\*  
EECS, MIT  
MIT Computer Science and Artificial Intelligence Lab

I believe that the issues as phrased make the questions considerably more difficult to answer than is necessary. The perspective in the document also takes a rather narrow view of AI and machine learning, producing as a consequence questions at the wrong level of detail.

This note outlines a rationale for rethinking the issues and rephrasing the questions, showing that when the irrelevancies are removed, simpler questions emerge that still get at all the essential issues. The simpler questions are also often familiar questions.

### Preliminaries

One general technique useful in simplifying the questions is to see what happens when we substitute other technologies for *AI*. I suggest that if the core issues the question raises are unchanged, then the essence of the question has nothing to do with AI. Recognizing this can make the issue far easier to grapple with.

It's also useful to recognize that AI is a set of computational processes, typically embodied in software. AI has a particular set of computational techniques that it uses, but for our purposes here it's the general notion of a *computational process* that matters. Some of these questions are better phrased as about computational processes, that is, the AI focus is at times irrelevant.

Recognizing AI as a set of computational processes helps in seeing through some of the questions below. Some of the purported difficulties arising around an AI invention – e.g., disclosure and enablement – arise not because of the AI character, but because it involves a complex computational process. The same disclosure and enablement issues can arise around *any* sufficiently complex computation process, even if it has nothing to do with AI. AI in and of itself has little or nothing to do with the difficulties in disclosure and enablement.

It's useful to avoid the term *algorithm* as much as possible, or at least be explicit every time it's used. An algorithm is a specification indicating how to compute something and there are frequently lots of them in use in AI. Consider one AI technique, neural nets. Training the net is commonly done using the *backpropagation algorithm*. Using it produces a trained network that accomplishes some computation, perhaps image recognition, generating movie recommendations, etc. That trained network is *also* an algorithm, *wholly different* from the backpropagation algorithm. Either or both of the algorithms – the one used to train the network, and the trained network that results – might be novel, but they are quite distinct and must be evaluated separately on their own terms. Hence statements about “the algorithm” are almost surely going to lead to confusion.

The questions also seem based on the far too limited perspective that AI *is* the use of neural networks; there is far more to AI than that. There is also far more to machine learning than just neural nets, so even if this is about patents in the context of machine learning, the current phrasing is too restrictive.

## Question 1

*I. Inventions that utilize AI, as well as inventions that are developed by AI, have commonly been referred to as “AI inventions.”*

The phrasing here, while likely in wide use, is part of the problem. It is confusing and obscures distinctions that are important and useful to make. Consider: would we say that “inventions that use lasers are laser inventions?” I suspect not. We’d say they are inventions that use lasers. There can also be inventions that are *about* lasers, e.g., new techniques for producing coherent light.

A better phrasing for the subject at hand would make three distinctions:

- a) There inventions that are *about new techniques of AI*. I suggest this is what the phrase *AI invention* ought to be used for. These are inventions focused on AI technology and may have application in a wide variety of fields. Generative adversarial networks, for example, are/were plausibly considered an *AI invention* in this sense.<sup>1</sup>
- b) There are *inventions that use AI technology*, without advancing it, just as there are potentially novel devices (e.g., the first laser rulers) that use laser technology without advancing it. A current example is in 10,431,207. It incorporates some state of the art AI technology (e.g., a recurrent neural net), but the claimed novelty lies in the “method for operation of a spoken language understanding (SLU) system.” AI technology is a component of the invention, just as laser technology is a component of the laser ruler, but the claimed novelty is in the capability developed, not in the AI techniques used.
- c) Inventions developed by the use of AI.

One simple example of this might be neural nets trained to do a useful task, e.g., detect tumors in pathology slides. Here the usual tests apply, i.e., novelty, etc. I suspect this is the sort of AI invention that the author of the RFC had in mind.

The difficulties in *creating* the net include selecting the architecture, training set, training regimen, etc. But the end result of this process – the trained net – is just a very large arithmetic expression. The matters because it means that claimed invention (the net) does not seem to present any novel issues with respect to invention and patenting.

A second example of an invention developed by the use of AI comes from a program developed in 1995 by one of my graduate students. When given a rough sketch of a design for an electro-mechanical device (e.g., a circuit breaker) it used a set of AI techniques to methodically and cleverly explore possible variations of that design. Its output was a set of novel circuit breaker designs. The potential invention(s) here were new circuit breaker designs. The AI program functioned roughly as a human assistant might have, but it was not the invention.

Combinations of these three are of course possible: if both the program in c) *and* its output exhibited sufficient novelty they might both be inventions, but they would be distinct inventions with different subject matter foci.

I suggest that a more careful phrasing that respects these distinctions will be very useful throughout.

---

<sup>1</sup> For ease of discussion I skip over questions of subject matter, novelty, etc., here. The issue is first to get clear what we’re talking about, then the additional criteria can be applied.

### Question 1 continues:

*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.*

The phrasing of this question is problematic in several ways:

- a) It's written with the sharply limiting perspective that AI *is* the use of neural networks, without recognizing how much more there is to the field. Even if the focus is machine learning, there is considerably more to it than neural networks. The phrasing is thus far too limited a view to be useful. The document needs to understand and acknowledge the breadth of AI ideas and techniques.
- b) It's focused (on purpose?) on the tools that go into creating the invention and hence seems to be at the wrong level. Why ask what "elements" go into building a neural net? It's a little like asking "what are the elements of a biotechnology invention?" The answer will likely be of the form DNA, various enzymes (ligases, polymerases, restriction enzymes, etc.), etc., but what do we get from this? The list of tools is not of the essence. (Or do I misunderstand the question?)
- c) As above, it fails to distinguish the different meanings of "AI invention"

Setting aside neural nets, we can ask the question more carefully, using the distinction outlined above:

- a) *What are the elements of an invention of that produces new techniques of AI?* This seems fairly straightforward: it's the new techniques, e.g., generative adversarial networks. Note that "elements" is used here to mean the *components of the invention*, not the tools that were used to create it.
- b) *What are the elements of an invention that uses AI?* Again straightforward – it's whatever is novel in the invention, possibly without the AI.
- c) *What are the elements of an invention that was developed using AI techniques?* Again straightforward – it's whatever is novel in the thing developed, independent of how it was arrived at.

All three of these questions make sense, are well phrased from a technical point of view, and as a result relatively straightforward to answer in any given instance. As with any technology, difficult cases may arise, but by being more precise about an "AI invention" we make the questions technically sensible and for the most part unproblematic. Like all such they may still call for careful judgment, but the questions are for the most part familiar.

### Question 2

*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.*

First, as above, the question is too narrow, as a consequence of the overly restricted view that sees AI as neural nets.

But taking it on its own terms, and restating for clarity, the question seems to be “what are the key elements that lead to a trained neural net (that a person might contribute to)? The list above of “for examples” is correct. (The resulting net may or may not be novel, of course.)

Note that there are many sorts of neural nets; the list of key elements may vary depending on the particular variety at issue. If the question is properly considered more broadly – as one about machine learning – there is a yet wider set of technologies, whose key elements may vary. Again, neural nets is not all of AI and it is not even all of machine learning.

Second, consider what happens if we simply omit “AI” in the question:

*What are the different ways that a natural person can contribute to conception of an invention and be eligible to be a named inventor?*

That seems a familiar question with lots of precedent to rely on. Being a named inventor arises in part from contributing to the novelty, and if we know what novelty we’re talking about, we can evaluate the extent to which a person contributed to it. That’s another reason why it’s important to be clear about what an AI invention might actually be.

### Question 3

*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?*

If, as in question 4, “other than a natural person” is meant to refer to a company, see below.

If “other than a natural person” is meant to include programs, the answer is No. Software has for decades assisted in creating things, some of these things were novel and inventive. As a pioneering example in the world of art (and hence not patent subject matter), a program named AARON by Harold Cohen generated artistic works for forty years. Its capabilities were constantly being refined over that time and the results are in major art museums internationally.

This segment of an article about Cohen in 2016 is informative:

Cohen wrote extensively about AARON, looking at the questions a computer-based artistic system raised both within the computing and art worlds. Was AARON creative? Cohen certainly thought that it was not as creative as he had been in creating the program. *Who was the artist—Cohen or AARON?* Cohen compared it to the relationship between Renaissance painters and their studio assistants. *Was the fact that AARON created art works evidence of computer intelligence?* On that, Cohen seemed non-committal, saying that AARON never worked to improve itself, a sure sign of intelligence, but at the same time defending the fact that AARON did just what human artists did, taking knowledge of forms and applying them to the creation process.

<https://computerhistory.org/blog/harold-cohen-and-aaron-a-40-year-collaboration>

Note in particular:

*Cohen compared it to the relationship between Renaissance painters and their studio assistants.*

Yes, just so. All of us – artists and engineers – have assistants aiding us in doing our work. Some of those assistants may be human (studio assistants), some may not (e.g., AARON, or the techniques of advanced calculus that show us how to solve a thorny problem on the way to invention). But Cohen was the artist, and the humans doing the inventing are the inventors.

Edison claimed to have tested thousands of possible materials for the filament of his light bulb. More precisely, he and his associates tested them. Do the assistants qualify as inventors for having done the long and repetitive work of trying one material after another? Evidently not. So it is with some of the programs that do our bidding, in some cases trying out thousands – or maybe millions – of possibilities from a set we (the inventors) have defined.

Looking far (far!) ahead, programs someday may begin to learn on their own, choosing what to investigate in the manner of students who no longer need their teachers to guide them. Then we may have a deeper quandary, but, crucially, it's not wholly unfamiliar: when a human invents, how much credit do his/her teachers deserve? Ought they to be named inventors? Either answer is plausible and the correct answer surely depends on the specifics of the case. The important point is that the question is to a large degree familiar and, in that distant day, ought to be approached as such.

If the question is whether a program ought to be able to be judged in the *legal sense* to be an inventor or owner of a patent, that gets into questions about agency and as such is more a philosophical question than a legal one. (It also seems quite problematic. If the program has enough agency to be an owner of a patent, can the program in turn be owned? If the program is the owner/inventor, who gets any royalties? This seems a not very useful road to go down, except perhaps as a thought exercise to keep philosophers occupied.)

#### Question 4

*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?*

I suggest not, but this is an issue of legal definition, with relatively little reliance on the technical perspective.

#### Question 5

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

No obviously so, but best to keep an eye out here, and to be clear-eyed about any claims of such. As attempted here, ask whether the allegedly unique considerations are in fact unique and whether they are unique to AI inventions.

### Question 6

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?

As phrased, the issue is complexity, not AI. Surely there are current patented inventions sufficiently complex that the *disclosure* and *ordinary skill in the art* criteria have become challenging. Even an algorithm with no connection to AI can still be sufficiently complex to make disclosure problematic. AI is not the issue, complexity is.

Once again, the question is largely familiar: how can you explain a complex device so that someone else can recreate it? For software a more specific version of the question is, how do you explain a complex computational process in sufficient detail that someone can recreate it.

This is an instance of the issue noted above: AI is a set of computational processes and it's often properties of the computational process that matters, not the AI aspect.

As one concrete possibility, for very large neural nets there the possibility of requiring publication of the network architecture and the weights.

### Question 7

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

First, the question is unclear. What is the perceived difficulty?

Second, there's nothing new in the notion of uncertainty (stochastic behavior) in a method. Some optimization methods for example use stochastic selection of starting points. There is no problem here as long as the nature of the uncertainty is disclosed via well understood methods, e.g., indicating that a selection was from a distribution with clearly specified parameters.

### Question 8

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?

The question is unclear, but I suspect it's actually asking whether *ordinary skill* means ordinary skill in the AI technology and/or ordinary skill in the area of the invention.

This seems like a familiar problem, and would seem to have nothing to do with AI. If a patent described a method of doing medical diagnosis by using a device built with superconductors,

ordinary skill in the art would presumably require knowledge of both the art superconductors and the art of medical diagnosis.

#### Question 9

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

Not obviously so. Not clear why examining prior art for an invention that uses AI techniques would be anything other than the usual processes.

#### Question 10

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

Have no comments on this one.

#### Question 11

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

As a general rule (much indulged in this note), the main issue is to be clear about what the invention is, and constantly ask whether inventions that use AI are any different with respect to these issues. Almost all the time the answer is that they are not.

If some exception to this occurs, that would be extremely interesting to examine and work out, but I suggest keeping a high bar on the question of whether the issues that crop up are anything other than the usual questions applied to yet another technology.

That said, the dynamic character of some AI systems might be problematic. Some systems constantly incorporate (i.e., learn from) additional examples/experience while in operation. Do the claims need to be constantly updated as well?

#### Question 12

*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?*

Don't have any comments on this one.

---

\* Randall Davis received his PhD in artificial intelligence from Stanford in 1976 and came to MIT in 1978. He is a Professor of Computer Science and has been working on a variety of areas of artificial intelligence.

He has also been active in the area of intellectual property and software. In 1990 he served as expert to the Court in *Computer Associates v. Altai*, a case that produced the abstraction, filtration, comparison test for software copyright. A 1994 paper in the *Columbia Law Review* analyzed the difficulties in applying intellectual property law to software and proposed a number of remedies. From 1998-2000 he served as the chairman of the National Academy of Sciences study on intellectual property rights and the information infrastructure, entitled *The Digital Dilemma: Intellectual Property in the Information Age* (National Academy Press, 2000). He has served as an expert in a variety of cases involving software copyrights and patents, including the investigation by the Department of Justice of the Inslaw matter, where he investigated allegations of copyright theft and cover-up by the Federal Bureau

---

of Investigation, the National Security Agency, the Drug Enforcement Agency, the United States Customs Service,  
and the Defense Intelligence Agency