Originally published in NYSBA *Bright Ideas*, Summer 2020, Vol. 29 No. 2, pp. 4-10.

Patenting Artificial Intelligence Inventions: Introduction and Selected Issues

By Michael Mauriel, Andrew Noble and Rory Radding

I. Introduction: Recent Growth of AI and AI Patenting

Artificial intelligence (AI) technology has been around for decades. Patent filings covering AI-related inventions have also been around for decades. However, it is only in the past 15 to 20 years that AI has exploded in the technology world generally. And it is only in the past 10 years that AI has exploded in the world of patent filings.

Last year, the World Intellectual Property Organization (WIPO) initiated a new series of reports, *WIPO Technology Trends*, with a 154-page report entitled *Artificial Intelligence* (the "WIPO report"). The report looks at AI patenting trends across industries and around the world. The mere fact that WIPO chose to make AI the focus of its first *Technology Trends* report suggests that AI patenting has become particularly significant. Notably, of the roughly 340,000 AI patent filings published since 1960, more than half have published since 2013.[1] (The statistics and charts below are either directly from or based on statistics in the WIPO report.)

AI is an umbrella term covering many different categories of specific techniques. The "machine learning" category has dominated AI patent filings in recent years and now appears in nearly 90% of AI-related patent filings.[2] As the chart below shows, new machine learning patent families grew steadily but modestly from 1990 to 2010. After 2010, growth accelerated dramatically.[3]

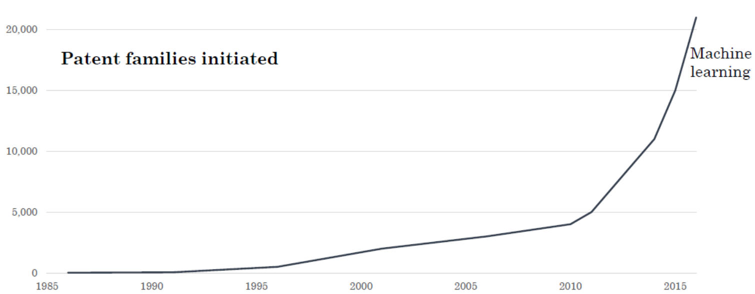

The recent AI patent filing boom is driven mainly by filings originating in China and in the United States. Although Japan led the world in AI patent filings until the late 1990s, and Korean filings have grown steadily since the early 1990s, filings in the United States and in China over the past 10 years have significantly outpaced those in other countries. Furthermore, China's AI-filing growth is now much greater than that of the United States, as shown below.[4]

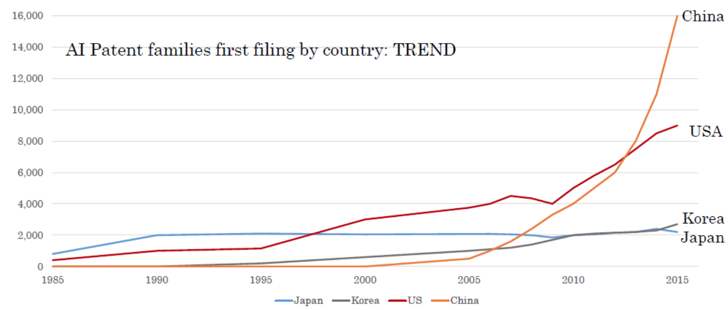

Telecommunications and transportation are the industries most often targeted in AI patent families. Each is targeted in 24 percent of AI patent families.[5] Given recent growth in autonomous vehicle technology, the prominence of the transportation industry in AI patenting is not surprising. Notably, however, the life science/medical industry area is not far behind and is targeted in 19 percent of all AI patent families.[6] AI patenting now spreads across many diverse industries. For example, some of the industry areas with the biggest recent AI patenting growth are agriculture and banking/finance.[7]

II. AI in Context

AI inventions should not be thought of as divorced from the specific context of their application. As discussed further below, although the heart of an AI system is typically realized in computer hardware and software, the particular design of such an AI system can vary significantly depending on the application to which it is applied. An AI invention is often intertwined with the context of the things it seeks to control and/or analyze data from such as a car, a scientific instrument, or a medical device. In some contexts, such as driving and medical diagnosis, the use of machine learning raises not only interesting patent questions but also important ethical and regulatory issues.

Michael Mauriel, Andrew Noble and Rory Radding are partners at Mauriel Kapouytian Woods LLP (www. mkwllp.com).

III. Elements of an AI Invention

To gain insight into the kinds of things that go into AI patent applications and some of the unresolved legal issues related to AI patenting, it helps to get a feel for how AI actually works. As stated above, in the technology world and in the patent world, machine learning is far and away the most prominent AI technology category. Therefore, although there are other categories of AI technology, we focus here on machine learning.

A. Machine Learning Generally

The WIPO report defines "machine learning" as:

> [A]n AI process that uses algorithms and statistical models to allow computers to make decisions without having to explicitly program it to perform the task. Machine learning algorithms build a model on sample data used as training data in order to identify and extract patterns from data, and therefore acquire their own knowledge.[8]

To put it more succinctly, machine learning technology allows a computer to "learn" from examples. The quotation marks remind us that when we speak of a computer "learning," it is really shorthand for "appearing to learn" or "learning-like" activity.

Many things we observe and learn about in the real world can be represented as a mathematical function that maps inputs to outputs. For example, a baseball pop-fly can be represented by inputs such as velocity, height, mass, flight angle, air velocity, etc. To the extent those inputs impact an output of interest, e.g., where the ball will land, there is a probably a discoverable mathematical function that can be used to predict the output based on the set of inputs. In the pop-fly example, classical mechanics provides rules describing such a function and can, with sufficient accuracy, predict the ball's landing spot given correct inputs. Thanks to Newton and various equations derived from his laws of motion, we already know rules describing such a function. But before discovering his laws of motion, Newton needed to make many observations of objects in motion and figure out what those laws were.

Machine learning's goal (at least in the context of what is known as "supervised" learning) is to allow a computer to carry out at least part of the scientific method. The first step (the hard one) is to use real-world *known* examples of something—i.e., instances of that something for which relevant conditions and results are already known—to discover a mathematical function that maps those real world conditions (function inputs) to real-world results (function outputs). The second step is to apply that newly discovered function to *untested* examples of that something by measuring or otherwise obtaining conditions associated with those untested examples and then using those conditions as inputs to the function to predict a real-world result based on the function's output.

Below are just a few illustrative examples of useful real-world applications and "inputs" and "outputs" that might be associated with a machine learning implementation for those applications:

Application	**Inputs**	**Math function** (mapping inputs to outputs)	**Outputs**
Object recognition	Pixel values from a digital image (and/or features values computed using those pixel values)	?	Object identity (e.g., word or phrase selected from a set of thousands of words or phrases corresponding to various objects)
Tissue pathology slide analysis	Pixel values from a digital image (and/or features values computed using those pixel values)	?	Tissue classification as positive (e.g. for cancer) or negative
Speech recognition	Digital audio data (and/or feature values computed using that data)	?	Recognized words
Autonomous driving	Lidar data, image data, car velocity data, weather data, time of day data, etc.	?	Driving control instructions (e.g., for steering, accelerating, braking, etc.) and/or intermediate outputs for determining such instructions (e.g., indication that object 30 feet away is another car)
Cardiac diagnosis	EKG data	?	Arrhythmia identification / classification

In sum, machine learning strives to fill in the "function" column above by analyzing training data derived from real-world examples in which the values for the inputs and outputs are already known. Once a sufficiently accurate function is discovered, then the trained machine learning system applies that function to determine outputs for new, untested examples, assuming the inputs of the new examples can be measured and provided to the trained machine learning system.

B. How Neural Networks Find a Function From Examples

We now dip a toe in the water of some neural network details—without pretense of completeness or technical precision—to provide a feel for the technology on a small, simplified scale to give some sense of what is involved on a larger scale. Machine learning techniques, such as neural networks, essentially use a "template" function that has a structure but also has many unknown parameters. The machine learning system then uses training data and a training algorithm to try to "learn" the optimal parameter values so that the function's output can be calculated accurately for new examples given a new set of inputs.

The underlying principle is more familiar than it might seem. Consider a very simple "template" function, the basic linear equation from junior high math:

$$y = ax+b$$

In the language of machine learning, y is an output value corresponding to some real-world thing; x is an input value corresponding to some real-world thing; and a and b are unknown parameter values. If you know that the relationship between x and y is linear, then you can discover the value of parameters a and b with just two known examples, each of which is represented by an input value x and an output value y.

For example, assume:

a is the speed someone travels directly away from home from a given starting point;

b is the distance from home of the given starting point;

x is the time the person spends traveling directly away from home; and

y is the person's final distance from home.

If we assume a and b are fixed values that do not change from one example to another, then we can determine those fixed values with two different "training" examples in which x and y are known. Once we know the fixed values for the parameters a (speed) and b (initial distance), we can determine the value of y (final distance from home) for any value of x (time spent traveling) us-ing the above equation.

Of course, if all real-world relationships of interest were simple linear ones, then we could get by with basic algebra and avoid the need for sophisticated machine learning techniques such as neural networks. But many important real-world relationships are very non-linear.

Neural networks provide intricate template functions to model complicated, non-linear relationships between measurable conditions (inputs) related to a real-world thing of interest and particular outcomes of that real-world thing of interest. Those template functions might have several thousand (or more) unknown parameters. A very simple "feed forward" type neural network is shown below by way of illustration.

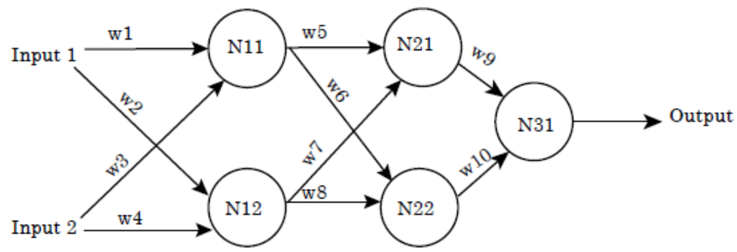

A neural network typically is arranged in "layers" of nodes known as artificial neurons. For example, in the above illustration, nodes N11 and N12 are in a first layer, and nodes N21 and N22 are in a second layer. Each node of the above neural network implements what is known as an "activation" function, which typically is a fairly simple non-linear function. For example, a commonly used activation function in recent years is known as the "ReLU" function (ReLU refers to "rectified linear units"). The rule for that function is simple: If the input is greater than 0, then the output equals the input. If the input is less than or equal to 0, then the output equals 0. The activation function determines a node's output value by applying the function's rule to a weighted sum of inputs from nodes in a prior layer. The "weights" shown above (w1, w2 . . . w10) are the parameters to be learned during training.

The above network's computations would proceed as follows: The input to node N11 is equal to (Input1)(w1) + (Input2)(w3). The input to node N12 is (Input1)(w2) + (Input2)(w4). In similar fashion, the input to node N21 is a weighted sum of the output of nodes in the prior layer, i.e., (N11output)(w5) + (N12output)(w7), as is the input to node N22, i.e., (N11output)(w6) + (N12output)(w8), and the input to node N31, i.e., (N21output)(w9) + (N22output)(w10). The output of each node is determined by applying the activation function to the input.

The above network would be "trained" for a particular real-world prediction problem by running training data with known output values through the network and making incremental adjustments to weight values to reduce the network's prediction error until it cannot be reduced any further. A simplified overview of that process is as follows:

1. Start with arbitrary values for w1, w2, . . . w10;
2. For a first training example, run the example's input data through the network to determine a predicted output;
3. Compare the predicted output to the known output of the training example to obtain an error measurement;
4. "Back propagate" the error through the network using something called the "back propagation algorithm" (which involves a series of partial derivative determinations working back through the network's processing pathways) to assess the amount of the error attributable to each weight (w1, w2 . . . w10);
5. Incrementally adjust the weight values to try to slightly decrease the error; and
6. Repeat steps 1-5 using additional training data examples until error is minimized.

Each neural node implements a simple function. However, a large network of such nodes collectively implements a very complicated non-linear function that, ideally, can be "trained" as described above to effectively model the relationship between inputs and outputs of the real-world phenomenon to which it is applied.

The above example shows just a small piece of one type of neural network. Other popular types include, for example, recurrent neural networks and convolutional neural networks. And many neural networks include multiple types of neural layers, e.g., convolutional layers, feed-forward layers, etc. and are used with other types of processing such as pooling or other techniques that reduce data size along the processing flow of the network. Moreover, the particular structure in terms of layer widths, number of layers, location of auxiliary inputs (e.g., inputs injected initially downstream of the initial network layer), and various other factors have infinite possible variations.

C. Neural Network Inventions

A neural network invention typically involves selecting and arranging a *particular* mix and configuration of neural network layer types, sizes, and depths that works well for a particular real-world problem. Neural network inventiveness can also reside in identifying the optimal input data, input pre-processing techniques, and/or training techniques that work best given the particular real-world problem to which the neural network is applied.

The possible elements of a neural network invention listed below could arguably all be characterized as mathematical, thus raising a question of whether such inventions are simply "abstract ideas" and therefore not eligible subject matter under 35 U.S.C. § 101. However, the choices made regarding the elements in a neural network invention can dramatically and concretely impact the ability of a computer to efficiently solve problems existing in the physical world. And, in some fields, that efficiency can be a matter of life and death. Cancer research, for example, is a race against time for those who have or will get diagnosed with the disease. One neural network design might help analyze genetic sequences an order of magnitude faster than another. To the extent significant improvements in processing time and/or accuracy in analysis of physical phenomena result from a neural network invention, we believe that invention is more than simply an abstract idea. However, as discussed near the end of this article, the law does not yet appear to have clearly reached that conclusion.

1. Architecture

The type and arrangement of neural network or other machine learning structures and techniques that work best for allowing a computer to use image data of biopsied tissue to predict whether the tissue contains malignant cells might be very different from the particular type and arrangement of structures and techniques that allow a computer to recognize spoken words based on captured audio data. In both cases, the individual techniques/processing elements are likely well known, but their particular arrangement and configuration in the invented AI system for performing the particular task is not well known. Thus, the types of layers, arrangement of different layers, size (width) of various layers, number of layers (network depth) and, in sum, the overall architecture of a neural network tailored to a particular problem is often at the heart of an AI invention.

2. Input Data Determination, Pre-Processing, and Feature Extraction

AI inventors distinguish between "raw" data (or processed raw data such as normalized, weighted, or encoded data), on the one hand, and "features," on the other. Both types of data are potential candidates for inputting into a neural network. "Features" are typically some values derived from the raw data. For example, in processing digital image data for input into a neural network for object recognition, the pixel values might be the "raw" data. It is possible to input all the pixel values into a neural network. However, it is also possible to extract image "features" from the pixel data and input those "features" into the neural network rather than, or in addition to, the raw pixel data. For example, a "feature" might be computed based on a change in pixel values over a portion of the image. It is also possible to use a neural network to learn what "features" are most useful to derive from a given type of raw data for a particular classification or prediction task.

The choice of what raw data to collect and use, how to pre-process it, and what features to extract from that data, if any, for input to a neural network can significantly impact how well the neural network performs a particular learning task. Significant thought and experiment go into making these choices, and they are an appropriate subject of AI inventions.

3. Training Methods

Neural networks are trained by iteratively passing training data through the network, measuring the "error" (sometimes called "loss") in the outputs relative to known values of what the output should be, and updating weight values to gradually reduce that error. Currently, many neural networks are trained using some variation of the back-propagation algorithm previously mentioned. For many neural network applications, standard training techniques are used in a routine way, and those techniques are not elements of the invention.

However, standard training techniques can be modified and/or combined with other known training techniques for specific training tasks in a manner deserving of protection. For example, known training techniques can be modified to prioritize certain training performance goals appropriate for a particular application. Such goals might include learning weights quickly that produce acceptable, but not spectacularly low, error levels or, by contrast, more finely tuning weights over a longer time period to increase predictive precision. We think such training techniques and/or the selection thereof in a particular context can be appropriate elements of an AI invention.

IV. Challenging Patent Law Issues for AI Inventions

AI inventions are everywhere these days. And many areas of AI innovation are consequential to advancing particularly high-stakes endeavors. For example, these inventions provide cutting-edge tools that promise to improve the efficiency and effectiveness of medical research and, ultimately, diagnoses and treatments. However, despite the potential criticality of AI inventions to advances in medical and life sciences, current U.S. patent law leaves an undesirable level of uncertainty regarding section 112 written description and enablement requirements and section 101 subject matter eligibility requirements for AI inventions. This uncertainty risks incentivizing leading companies to keep the inner workings of important AI inventions secret rather than seek patent protection. Such decisions by those on the cutting edge of applying AI to medicine and other fields risk reducing the open exchange of information that is fundamental to the patent system's constitutional purpose of promoting the progress of science and useful arts. In light of these risks, below we offer thoughts and practice tips for how patent attorneys might think about the application of sections 112 and 101 to AI inventions.

A. Written Description and Enablement: Is AI an "Unpredictable Art?"

The law regarding both written description and enablement under section 112 distinguishes between so-called "predictable arts" and "unpredictable arts." The former typically requires less detailed disclosure in order to support relatively broad claim scope, while the latter typically requires greater disclosure detail. Traditionally, many mechanical, electrical, and computer-related inventions are treated as "predictable arts" for section 112 purposes, whereas many chemical, life science, and medical treatment-related inventions are treated as "unpredictable arts."

As discussed further below, AI inventions present potential challenges to this traditional division between "predictable" and "unpredictable." Notably, the USPTO's 2019 Request for Comments on Patenting Artificial Intelligence Inventions[9] raises the question of whether AI inventions should be treated as unpredictable arts for purposes of section 112.

1. Written Description

To meet the written description requirement under section 112, a patent specification must describe the claimed invention in sufficient detail that a person of ordinary skill in the art can reasonably conclude that the inventor had possession of the claimed invention. Generally, disclosures for relatively new, complex, and/or unpredictable arts require a heightened level of detail to satisfy the written description requirement.[10] As a new field evolves, the balance between what is known and what is added by each inventive contribution also evolves.[11]

The level of detail required to satisfy the written description requirement varies depending on the nature and scope of the claims and on the complexity and predictability of the relevant technology.[12] Computer-related inventions have traditionally been treated as "predictable" in the sense that it has been assumed that one skilled in the art would understand what the inventor has invented if the specification provides at least high-level disclosure of the underlying processing. In other words, the inventor does not need to spell out in detail all variations on how a particular solution might be implemented in order to claim the solution with reasonably broad scope.

However, in the context of deep-learning technology, the reasons why one neural network design performs better than another is not necessarily clear to one skilled in the art or even to the inventor. Whether a particular solution will work can, in some cases, be as much a matter of trial and error as it is a matter of discovering principles that underly the efficacy of that solution. Therefore, assessing whether undisclosed variations on the primary embodiments were in the possession of the inventor at the time of the application's filing might be more challenging in the context of AI inventions than in the context of other computer-related inventions. Notably, the USPTO has raised the question of whether AI inventions might, under the written description requirement, require more detail in the specification's disclosure than other inventions require.[13]

2. Enablement

In order to meet the enablement requirement, an invention must be disclosed in sufficient detail to allow one of ordinary skill in the art to both make and use the invention without undue experimentation.[14] *In re Wands* sets forth several factors for determining whether undue experimentation is required: (1) the breadth of the claims; (2) the nature of the invention; (3) the state of the prior art; (4) the level of one of ordinary skill; (5) the level of predictability in the art; (6) the amount of direction provided by the inventor; (7) the existence of working examples; and (8) the quantity of experimentation needed to make or use the invention based on the content of the disclosure.[15]

For "predictable arts" a person of ordinary skill in the art is generally assumed to have reasonably high mastery of the basic tools and techniques disclosed in the specification for making and using the invention. Such an assumption allows for a relatively high-level description of subject matter, e.g., generic computer components, network connections, etc., which are considered to be well-known by a person of ordinary skill.[16] Therefore, an inventor might satisfy the enablement requirement by describing only limited or incrementally inventive portions of a computer-implemented application in greater detail.

However, for AI inventions, the USPTO has asked whether the unpredictability of certain AI systems raises challenges for applicants to provide sufficient detail to avoid "undue experimentation."[17] While no special considerations apply to AI for enablement at this point, we note that AI inventions share some characteristics with inventions in the traditionally "unpredictable arts" category, i.e., chemical, pharmaceutical, biotech, and other life-sciences related inventions. AI inventions, like chemical and life-science inventions, often require significant research and trial and error over a long period of time to discover solutions that work for a particular problem. Just as it might be difficult to understand exactly why a particular chemical formulation leads to effective results for addressing a particular problem, it might be difficult to understand why certain neural network structures and not others work well for a particular prediction problem. Because the underlying reasons for success might not be well understood, it might be difficult for one of ordinary skill to generalize the disclosed embodiments to implement a variety of other solutions that are within the scope of a claim but differ significantly from the exact embodiments disclosed in the specification.

Treating AI as an "unpredictable art" for the purposes of enablement would have significant implications for preparing and prosecuting AI patent applications. As the Federal Circuit's predecessor court stated in *In re Fisher*, "In cases involving unpredictable factors, the scope of enablement obviously varies inversely with the degree of unpredictability of the factors involved."[18] In the chemical arts, the guidance and ease/difficulty in carrying out an assay to achieve the claimed objectives may be considered in determining the amount of experimentation needed in an enablement analysis.[19] Likewise, with AI technology, enablement compliance might require significant detailed guidance in the specification regarding how to make or use the invention in order to avoid a finding that "undue experimentation" is required.

It is too early to know whether AI inventions will ultimately be treated as within the "predictable" or "unpredictable" arts for purposes of the written description and enablement requirements. But, at the very least, patent practitioners who draft patent applications for AI inventions should keep in mind the current uncertainty regarding written description and enablement requirements for AI inventions and be prepared to defend the sufficiency of their disclosures during prosecution.

3. Subject Matter Eligibility

USPTO guidance for subject matter eligibility under section 101 does not yet include AI-specific examples, and a robust body of section 101 jurisprudence on AI-specific questions does not yet exist. However, the USPTO's 2019 Request for Comments suggests that the USPTO is actively considering how to best treat AI inventions under section 101.[20]

In the meantime, the USPTO has continued to update its section 101 eligibility guidance and examples, which provide something, at least, for practitioners to go on when trying to apply section 101 law to AI. In 2019, the USPTO further revised its section101 guidance with a January revision (the "2019 PEG") and request for comments and provided a further update in October[21] (the "October Update").

The 2019 updates added a new "integrated into a practical application" prong to Step 2A of the existing USPTO framework for evaluating subject matter eligibility. This new prong makes clear that even if a claim recites a judicial exception (abstract idea, law of nature, or natural phenomenon) the claim is not considered to be "directed" to that judicial exception if "the claim as a whole integrates a judicial exception into a practical application."[22]

In general, the judicial exception is integrated into a practical application if it does something concrete with the exception's output. The October Update, at least in the medical science context, clearly distinguishes between data *input gathering* activity and data *output utilization* activity. The latter appears to satisfy this new "integrated into" prong, while the former, by itself, does not. For example, according to the October Update, a claim reciting vaccinating cats using different vaccination schedules and then analyzing results to determine a lowest risk schedule would *not* integrate the judicial exception into a practical application.[23] The vaccinating step in that example is merely "in order to gather data."[24] However, if the relevant claim goes on to recite using the identified lower risk schedule to then vaccinate other cats, that *would* integrate the judicial exception into a practical application, and the claim would be subject-matter eligible.[25]

It is important not to confuse the "integrated into" analysis under the new prong of Step 2A with the "significantly more" analysis of Step 2B of the USPTO guidelines. In contrast to additional elements under the "significantly more" analysis of Step 2B, the additional elements relied on for the "integrated into" analysis *can* be "routine and conventional." The October Update's Example 46, regarding a livestock management invention, makes this particularly clear:[26] Claim 1 recites gathering livestock data via monitors (e.g., video cameras), analyzing the data to determine whether an animal's data appears to be aberrant, and displaying the results on a display.[27] The USPTO considers this claim *not* eligible.[28] By contrast, Claim 3 of the same example adds the step of controlling a sorting gate to separate animals with aberrant behavior from those with normal behavior.[29] This additional step renders the claim eligible.[30] Note that although controlling a sorting gate, by itself, is presumably "routine and conventional," that is okay because, in context, it integrates the alleged exception into the "practical application" of separating the livestock based on behavior, which goes beyond simply identifying the behavior.

In the context of AI inventions, this answers some questions but not others. For example, a claim to an AI invention for an autonomous vehicle would presumably be eligible if the claim recited using the AI data output to control the vehicle in some way. However, we still lack official guidance that tells us whether or when an intricately designed neural network processing system, tailored to a specific real-world problem, can be patent-eligible if it produces a useful data output but that output does not trigger some further concrete action. Because machine-learning applications often make predictions or accurately identify things without necessarily taking further actions based on those predictions or identifications, further AI-specific guidance is needed. Such further guidance would help AI innovators and patent practitioners make more effective decisions regarding patenting AI.

Endnotes

1. WIPO report at 13.
2. *Id*. at 31.
3. Chart based on WIPO report at 42, Fig. 3.4.
4. Chart based on WIPO report at 90, Fig. 5.6.
5. WIPO report at 49
6. *Id*.
7. *Id*. at 51
8. *Id*. at 146.
9. 84 FR 44889 (Aug. 27, 2019).
10. *Capon v. Eshhar*, 418 F.3d 1349, 1359 (Fed. Cir. 2005).
11. *Id*. at 1358.
12. *Ariad Pharm., Inc. v. Eli Lilly & Co*., 598 F.3d 1336, 1351 (Fed. Cir. 2010) (*en banc*); *Capon v. Eshhar*, 418 F.3d 1349, 1357-58 (Fed. Cir. 2005).
13. *See* Request for Comments on Patenting Artificial Intelligence Inventions, 84 FR 44889 (Aug. 27, 2019, question #6: "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?"
14. *In re Vaeck*, 947 F.2d 488, 495 (Fed. Cir. 1991).
15. *In re Wands*, 858 F.2d 731, 736-37 (Fed. Cir. 1988).
16. *In re Fisher*, 427 F.2d 833, 839 (C.C.P.A. 1970) ("In cases involving predictable factors, such as mechanical or electrical elements, a single embodiment provides broad enablement in the sense that, once imagined, other embodiments can be made without difficulty and their performance characteristics predicted by resort to known scientific laws. In cases involving unpredictable factors, such as most chemical reactions and physiological activity, the scope of enablement obviously varies inversely with the degree of unpredictability of the factors involved.").
17. *See* Request for Comments on Patenting Artificial Intelligence Inventions, 84 FR 44889 (ug. 27, 2019), question #7: "How can patent applications for AI inventions best comply with the enablement requirement, particularly given the degree of unpredictability of certain AI systems?"
18. *In re Fisher*, 427 F.2d. at 839.
19. *In re Wands*, 858 F.2d at 737.
20. Request for Comments on Patenting Artificial Intelligence Inventions, 84 FR 44889 (Aug.27, 2019), question #5: "Are there any patent eligibility considerations unique to AI inventions?"
21. 2019 Revised Patent Subject Matter Eligibility Guidance, 84 FR 50 (Jan. 7, 2019) ("2019 PEG") and October 2019 Patent Eligibility Guidance Update, USPTO (Oct. 17, 2019), https://www.uspto.gov/patent/laws-and-regulations/examination-policy/subject-matter-eligibility ("October Update").
22. 2019 PEG at 54.
23. October Update at 14.
24. *Id*.
25. *Id*. at 15.
26. Appendix 1 to October Update at 30.
27. *Id*. at 31.
28. *Id*. at 33.
29. *Id*. at 32-33.
30. *Id*. at 40.