Response to Request for Comments on Patenting Artificial Intelligence Inventions

Michael Mauriel (mmauriel@mkwllp.com) and Andrew Noble (anoble@mkwllp.com)

patent attorneys (see https://mkwllp.com/biographies/michael-mauriel/ and https://www.mkwllp.com/biographies/andrew-a-noble/)

These comments respond to the United States Patent and Trademark Office's Request for Comments on Patenting Artificial Intelligence Inventions published August 27, 2019 in the Federal Register, Vol. 84, No. 166 at 44889 ("the Request").

We have not consulted with our clients in making these comments and these comments are not made on behalf of any of our clients. Our clients may or may not agree with these comments. The views expressed below in these comments are solely our own, not those of our firm.

INTRODUCTION

Over the past several years, we have drafted patent applications covering various AI inventions for biotech and life science applications. These inventions provide cutting edge tools promising to improve the efficiency and effectiveness of medical research and, ultimately, diagnoses and treatments. However, despite the potential criticality of these AI inventions to advances in medical and life sciences, current U.S. patent law leaves an undesirable level of uncertainty regarding the subject matter eligibility of many 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, we welcome the USPTO's effort to seek public comments on this topic and work towards improving the U.S. patent system's ability to effectively respond to the importance and recent growth of AI inventions.

In this response, we address only issues raised in questions 1.-4. of the request. We have not freed up enough time yet to prepare written comments on the important issues raised in questions 5.-12. of the Request. We would welcome the opportunity to do so at a later date. We would particularly welcome the opportunity to provide further comments on the §§ 101 and 112 issues raised in questions 5.-8. Please let us know if there is a future opportunity to provide comments on those issues.

RESPONSE TO QUESTIONS 1.-2.

Separating "AI as inventor" issues from those regarding "AI inventions" by natural persons

To avoid confusion, we recommend that the USPTO, the courts, and Congress define "AI inventions" only as "inventions that utilize AI" rather than as "inventions that are developed by AI." The latter phrase raises interesting semantic and even philosophical issues, and may even already raise practical issues as suggested by questions 3. and 4. in the Request. But such issues are quite different than those raised by the challenges of protecting the many AI inventions that are made by natural persons. We recommend separating the two types of issues to avoid confusion and inefficiency. The more speculative and provocative questions about whether non-natural entities can be "inventors" does not need to be resolved in order to address the more immediately pressing questions surrounding patenting of specific AI systems currently being invented by natural persons. Therefore, in responding to the question of *what are elements of an AI invention*, we address only AI elements invented by natural persons.

AI systems today are often inventions but are not inventors

Notwithstanding the recently well-publicized "DABUS" example of a container design generated by an AI system, most AI systems in use today are not themselves designing products from scratch. Typically, an AI application enables a computer to do something that might otherwise be done by a human. Examples include recognizing objects, recognizing speech, diagnosing a patient, driving a vehicle, vacuuming, reading pathology slides, or other tasks that a human could do. Of course, designing a product might be one such task. But it is important to keep in mind that there are many critically important applications of AI in which the AI system does not generate a product design or generate a design of anything else that might otherwise be patentable. However, in all these cases, the AI systems themselves, made by humans, are potential inventions and proper candidates for patent protection.

AI inventions by human inventors

AI inventing today typically involves selecting and arranging a particular mix and configuration of AI-related techniques to facilitate a computer learning how to classify and/or predict a particular real-world thing using data related to that real world thing. What an AI system looks like in one context might be very different than in another context. For example, the type and arrangement of processing 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 than the particular type and arrangement of 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.

Selected examples of possible AI invention elements

"AI" encompasses a wide range of techniques. Deep learning / neural network systems are not the only type of AI systems being developed today. However, their indisputable importance for many current AI implementations makes them worth using as an example for discussing appropriate elements of an AI invention. A list of possible elements for a deep learning / neural network invention follows. Of course, the below list barely scratches the surface of the various elements that might be part of a deep learning / neural network invention. However, it hopefully gives a feel for the sort of things that AI inventors have to figure out. Much thought and experimentation go into figuring out each of the below for a particular machine learning task.

a. Input data determination including raw data selection, 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 element of AI inventions.

b. Layer types

Each layer of a neural network is defined in part by its type. For example, a layer might be feed forward, recurrent, or convolutional. These layer types might themselves have sub types. For example, one convolutional layer might perform normal convolution. Another convolutional layer might perform separable convolution.

Although new layer types are occasionally developed, most neural network inventions apply known layer types to solve particular problems. Even if a layer type is known, the application of that layer type to a particular real-world data problem might not be known or obvious. Moreover, a particular multi-layer neural network might combine different known layer types in a unique way for optimal performance. Also, even if a particular type of layer or combination of layers is known, the number of layers, dimensions of the layers, or other layer parameter combinations might be uniquely suited to a particular real-world data problem.

c. Layer dimensions

Generally, the size of the input layer corresponds to the number of distinct data values to be input into the network and the size of the output layer corresponds to the number of possible classes into which the thing of interest (i.e., the thing to which the input data relates) might be categorized. However, the size of the one or more "hidden" layers in between the input and output layer is not fixed by the input or the output size.

"Size" is an imperfect word to describe a layer (and can mean different things for different types of layers), but we use it anyway for its intuitive value. What matters practically is how much computation is being done at each layer and how many parameters are to be learned at each layer. For example, in a feed forward layer, the more neurons are in the layer, the more computation the layer requires and the more individual weight values there are to learn. In a convolutional layer, one or more "filters" are applied to the input. A filter is a moving window of weight values that are moved together, step-by-step, over the input values and linearly combined (via matrix multiplication) with input values in a current window location of the filter. In such layers, the filter size and/or the number of filters determines how many weights must be learned. However, other parameters, such as the stride (how big each "step" is as the filter is moved through the input between) also affect how many calculations are performed each time input data is passed through the layer.

Size is also controlled, particularly in convolutional networks, by "pooling" layers used between one or more other layers of the network. Pooling layers are used to reduce the amount of data that is passed between, for example, two convolutional layers. Placement and configuration of such pooling layers is another object of thought and experimentation for AI inventors in optimizing a particular neural network for a particular problem.

d. Activation functions (with caution)

A neuron in a neural network takes several inputs and generates an output. Generally, the output is determined in a two-phase process. The first phase is typically a linear combination (e.g., via matrix multiplication) of the inputs with a set of current weights to determine a single value. That value is then input into the second phase of the neuron's processing which applies an "activation function" to the first phase's output. An "activation function" is typically a non-linear rule for determining what the neuron's output should be for a given result of the first-phase linear combination.

Most AI inventions use well-known activation functions. Currently, the "rectified linear unit" or "ReLU" is widely used across many different neural network applications.

Other activation functions are also widely used. In most cases, the choice of activation function is <u>not</u> part of the inventiveness of a new neural network.

However, different layers in the same network can use different activation functions. Therefore, the choice of and/or combination of activation functions can be inventive as applied to a particular problem solved by the neural network, even if the activation function itself is not new. Also, new activation functions will almost certainly be adopted in the future when their general usefulness (or usefulness for particular types of problems) is discovered through future research.

At the same time, because many structurally different neural networks use the same activation function for solving very different types problems, granting patent protection to the general use of a new activation function in the context of any neural network could have undesirable preemption consequences. We do think care should be taken before granting protection for claims to use of a new activation function in neural networks generally. However, it is quite possible that use of a new activation function in the context of a particular neural network architecture solving a particular learning problem might be protected in a limited way that sufficiently minimizes preemption effects.

e. Overall network architecture

The arrangement of layers of various types and sizes and the number of each such layers in the network defines the overall architecture of the network. Significant thought and many months (or even a few years) of research and trial and error can go into developing a new neural network having a particular architecture or range of architectures for solving a particular real-word machine learning problem. Training data sets are run through many different possible network architectures to find the architectures that best solve the learning problem for a particular application. Each architecture might have different performance levels in terms of speed and/or accuracy in handling a given learning task.

The particular combination of elements that make up the overall architecture of a given neural network is often very unique relative to that of other neural networks and it often requires months, and sometimes years, of thought and experimentation to develop the architecture of a particular neural network well-suited for a particular learning task. Therefore, we think that there is value in granting patent protection on neural networks having particular architectures, especially when those architectures are used to solve particular real-world problems. Furthermore, for the protection to have value commensurate with the inventive effort, characteristics of a neural network's architecture will likely need to be claimed within a range of values rather than being limited to an exact number. For example, an inventor might discover that a neural network having five convolutional layers, each with three different filters followed by three feed forward layers of sizes ten, five, and three neurons respectively offers the best performance for a particular learning task. That inventor might also have discovered variations on that architecture within a range of layer types, numbers of layers, and layer size that produce acceptable performance.

f. 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. This uses iterative back propagation of error measurements to compute partial derivatives of each weight with respect to the current error level and then gradient descent plus a "learning rule" value is used to make individual weight adjustments. 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 quickly learning weights that produce acceptable 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, could conceivably be claimed as appropriate elements of an AI invention.

Weight values themselves are NOT appropriate elements of an AI invention

Conversely, we think weight values are generally NOT an appropriate element of an AI invention. Optimal weight values, discovered during training, do not need to be known to the natural person who has designed the AI invention. The neural network is designed, and training is done, so that those weight values can be "discovered." But they are not invented. Moreover, the weight values themselves are not useful apart from the particular network (as defined by its architecture of layer types and sizes) used to discover them. Therefore, there is no need to independently protect weight values if the neural network itself has other protectable elements.

RESPONSE TO QUESTIONS 3.-4.

Notwithstanding the DABUS example, our view is this: Unless and until the scientific community declares that AI has allowed computers to achieve "consciousness" such that a computer is capable of "conceiving" something, any question about what to do when an AI system "invents" something is purely hypothetical.

However, we do appreciate that questions 3. and 4. in the request reflects potentially important and challenging intellectual property questions made real by DABUS-type AI accomplishments. We would perhaps phrase those questions somewhat differently: If an AI system generates a product design, is that design something that should be protectable under some future version of our intellectual property laws even if it is not protectable by our current laws? To that end, we do think a modified version of question 10. might be well formed for exploring this issue, i.e., are there new forms of intellectual property needed for protecting the products of AI?

Our intellectual property laws have evolved over centuries to serve us well in protecting certain human inventions, discoveries, and creative works. We do think there is a need for those laws to continue evolving to protect new products of human innovation, including AI. Notably, in view of the challenges today's computer-related inventions pose to

101 jurisprudence, we are long past due for an update to the 67-year old statutory language defining eligible subject matter.

However, we would caution against trying to stretch or otherwise re-define our current laws to treat anything "invention-like" generated solely by AI under the same regime of protection we give to the products of human invention. To the extent AI starts generating new and useful designs for a process, machine, article of manufacture, or composition of matter that previously would only have been resulted from human innovation, we suspect it will be most useful to treat protection of such products of AI differently than we treat the products of human innovation.

<u>Invention and discovery revisited</u>

We caution above that any "invention-like" designs resulting from AI should be treated differently than human inventions. But in closing, we also offer the following thoughts for further consideration on how, if at all, the two might be reconciled:

The Constitution refers to "inventors." But, when naming the thing which Congress is empowered to protected, it uses the word "discoveries" rather than the word "inventions." Perhaps this word choice echoes earlier times when the two words were closer in meaning. The Shorter Oxford English Dictionary Fifth Edition lists a "now rare or obsolete" definition of "invent" that emphasizes finding or <u>discovering</u>. In contrast, modern definitions emphasize creativity and original thought.

The modern definition is consistent with our current legal definition of "invention" which emphasizes "conception." However, perhaps the older definition, which better aligns with the Constitution's language, could serve us well in the future as we try to adapt our intellectual property laws to the world of AI.

If the object of a utility patent's protection was a "discovery" (provided that the thing discovered was not naturally occurring) rather than a modernly defined "invention" (which must be "conceived" or created in the mind of the inventor), then the problem of AI-produced product designs would perhaps cease to be a problem. The human creator(s) of DABUS, for example, could be seen as having "discovered" the recently publicized container design even if those human creator(s) didn't "invent" it in the modern sense. Thus, there would be no need to address the problem of naming DABUS itself as an inventor. And the human creators of DABUS would be the owners of the container design by virtue of having discovered it, albeit with a little help from their AI friend.