November 8, 2019

Via Electronic Delivery
Eligibility2019@uspto.gov

United States Patent and Trademark Office
Mail Stop CFO
P.O. Box 1450
Alexandria, VA 22313-1450

Attn: Andrei Iancu, Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office

RE: Comments on Patenting Artificial Intelligence Inventions

Dear Director Iancu:

Schwegman Lundberg & Woessner, P.A. ("SLW") thanks the United States Patent and Trademark Office ("the Office") for the opportunity to provide comments on patenting artificial intelligence (AI) inventions set forth at 84 FR 51522 as a series of 12 questions. We request that the Office consider the below answers to the following questions:

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.

In general, artificial intelligence (AI) is the use of machines to mimic cognitive functions associated with humans. Twenty years ago, the term AI would have encompassed decision-tree structures; ten years ago, neural networks; today, the main driver in the fast progress of AI is generally associated with machine-learning algorithms. Machine Learning (ML) is an application

of AI that provides computer systems the ability to perform tasks without explicitly being programmed, by making inferences based on patterns found in the analysis of data.

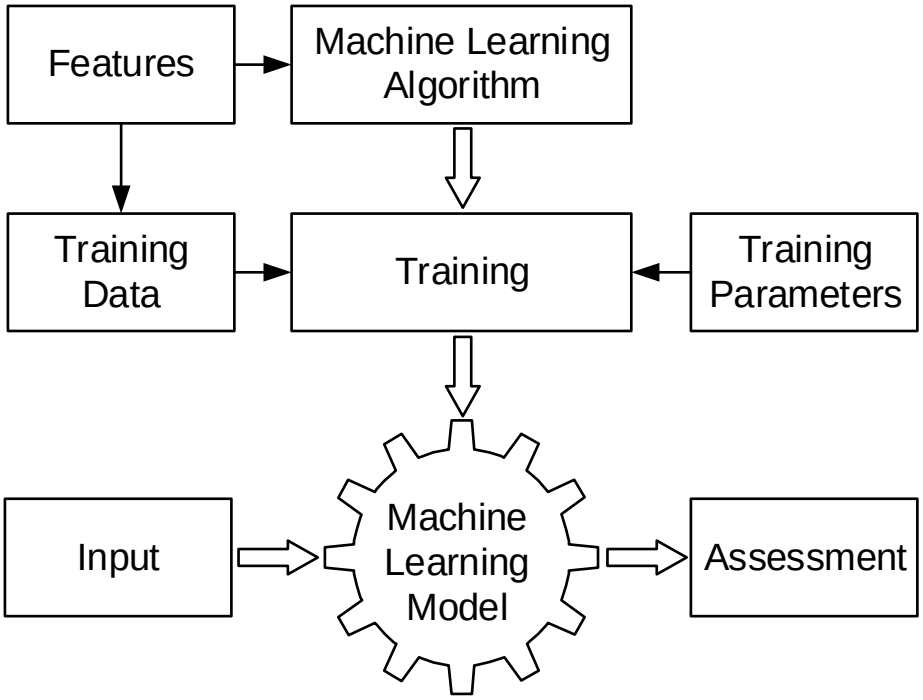

FIG. 1: The Machine Learning Model

ML systems commonly utilize a ML algorithm to analyze training data to create a ML model capable of providing assessment of a given input. Features define the structure of the data used during the training. A feature is an individual measurable property of a phenomenon being observed. The concept of a feature is related to that of an explanatory variable used in statistical techniques such as linear regression. Choosing informative, discriminating, and independent features is important for effective operation of the ML model, such as in pattern recognition, classification, and regression. Data representation refers to the method of organizing the data for storage on a computer system, including the structure for the identified features and their values. Clear and precise data representation is important for the training to identify patterns or correlations within the data.

During training, the ML algorithm, given an ML model specification with adjustable model elements or parameters (e.g., as defined by a data scientist), creates the ML model by selecting or tuning ML-model elements or parameter values to optimize a certain metric that

determines how well the model captures patterns inherent in the data. For example, in supervised machine learning, where the training data includes pairs of input data items and corresponding "labels" representing a ground-truth output of the data item (e.g., in classification problems, an assigned class), the metric may be cost function that quantifies the difference between the outputs (assessment) computed by the model and the ground-truth outputs. Training often involves analyzing large amounts of training data (e.g., from several gigabytes to a terabyte or more) based on identified features and exploring many possible functions and model parameter values to determine the best patterns or correlations within the data, which may require large amounts of computing resources and a substantial amount of time (e.g., from hours to days or longer).

However, the technical field of "AI inventions" should not be characterized as limited to a particular set of elements. While AI inventions can be directed to innovations in the structure of the ML model, the ML algorithm, the structure of the training data, etc., no list of elements should be considered exhaustive. The term AI continues to evolve. Any attempted enumeration of the elements of AI inventions will soon be incomplete. It is important that any framework considered by the Office be as applicable now to the existing structure as it will be to future developments without continuous amendment. It is our opinion that any rigid framework would impose undue burdens to the continued progress and protection of science and the useful arts.

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.

Title 35 of the United States Code defines the requirements for patentability as "[w]hoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title." *35 U.S.C. 101*. "Patentability shall not be negated by the manner in which the invention was made." *35 U.S.C. 103*.

Accordingly, any new and useful process, machine, manufacture, or composition of matter—*that happens to involve AI*—is subject to patentability. As described above, there are

numerous elements related to ML. Inventions conceived to improve any of these elements should remain patentable, so long as they meet the statutory requirements for patentability.

This debate about the patentability of AI inventions is similar to past debates about the patentability of software-based inventions. Some argued that common machine operations, such as receiving data, processing data, and outputting results are common computer operations, and that software-based inventions should not be patentable. However, as stated by the Federal Circuit, "We have found software-related patents eligible." *Bascom Global Internet Services, Inc. v. AT&T Mobility LLC.* (Fed. Cir. 2016). Similarly, AI inventions are patent eligible so long as they meet the requirements of Title 35.

FIG. 1 illustrates some of the numerous possible elements of an AI invention. For each element, a "natural person" may make choices or decisions that affect or improve the AI invention. In addition, a "natural person" may conceive new or improved elements. For example, a natural person may select a specific type of ML model (e.g., a particular neural network), configure model hyperparameters for a particular application (e.g., set a number of network layers), select data features to be used as input into the model, define a new ML algorithm or select and configure an ML algorithm to train the model, create a new metric to be optimized during training or used to measure performance of the trained model, implement a new practical application of a (new or existing) ML model, etc.

Consider an application to find the meaning of a text message. Each word in a dictionary may be represented as a vector, where each word contains a "1" in the location corresponding to the word and "0" elsewhere. However, vectors having large dimensions (e.g., 50,000 or more) are not easy to use. Conversions can be defined to assign a vector of a small dimension to each word (e.g., 30-50). The particular way in which the vector conversion is implemented can greatly improve the operation of the ML model. In other examples, when many features are being considered, the way in which those features are combined into vectors or matrices can require technical thought to achieve efficacy. The discovery of a new, useful, and non-obvious data representation, set or order of features, or other specific selection, combination, or implementation of AI elements should be protectable.

Consider an example application to provide recommendations for users. If the number of possibilities is very large, the ML model must be able to quickly correlate user features to possible recommendations. Further, some facts about users are not deterministic, but

probabilistic. For example, when providing travel recommendations for a user, the probability that the user has been to Paris may follow a probabilistic distribution centered around 75%. Representing and accounting for probability distributions regarding facts about a user, as well as considering a very large number of possible recommendations, can make specific problems quite complex. Solutions to specific problems will be subject to patent protection if a scientist is able to generate a workable ML model that trains quickly and is able to generate the right recommendations for each user. That is, selecting the right training set may also present technical challenges. For example, if an online service with a billion users wishes to determine which news articles to recommend to the users based on past viewings of users, the amount of historical data and the large number of users will require intelligent selection of the training data necessary to reduce the data set to a workable amount of data that can be processed by the ML algorithm. Finding the right data set, with the right set of features, will be inventive if it solves these technical challenges.

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?

We do not believe that laws and regulations need to be revised because of the use of AI tools. The question itself implies that a ML model may contribute to the conception of an invention in the same way that a natural person contributes to the conception of an invention. We respectfully believe this to be a false premise. We consider ML algorithms to be powerful tools for the development of solutions in many aspects of today's technical environment. A tool is anything used as a means of accomplishing a task or purpose. ML algorithms are tools used, for example, to help identify hidden patterns or correlations within data that may provide practical applications. Depending on where the inventive concept lies within an ML tool, credit for the inventive concept is due to the natural person that created the ML or to the natural person who used the ML tool in that particular manner and/or for a particular application. For example, the inventor of the inventive concept may be the natural person who created or set up the ML model or the ML algorithm, selected the desired features, chose the proper data representation for the training data, or devised experiments to collect the training data. However, the machine itself

deserves no credit, as without the actions of the aforementioned people, the machine would be useless.

The confusion with the attribution of an invention to machines comes likely from the fact that ML-based software "programs itself" in that, by the nature of ML, the trained ML model results in part from machine operations. However, as explained above, this automatic creation uses an ML algorithm created by a natural person, which optimizes, based on training data (e.g., as provided by a natural person) an ML model whose elements and parameters are at least partially defined by a natural person. For example, in the context of neural networks, a natural person may define the structure of the neural network by specifying the number and type of layers, the data flow between layers, etc., and the neural network weights may be automatically determined using the ML algorithm. We recognize that ML inventions may be "nested," such that, for example, input features to an ML model, instead of being hand-selected by a natural person, are themselves machine-learned, or that an ML model is automatically selected or configured. In these cases, however, the feature selection algorithm, or the "auto-ML" algorithm that selects or configures the ML model, depend on human definition or input at some level. Thus, ultimately, ML tools are the product of human ingenuity.

We also note that, if one were to attribute conception to an AI tool, such attribution would raise a host of new issues that would require congressional action to address. AI tools do not qualify as natural persons in any legal context, and therefore could not sign an oath or declaration.

The same way technology is used today to find information, ML algorithms are tools designed to find and exploit patterns or correlations in data. The inventive step is not performed by the machine that found the pattern or correlation, but by the ML designers and data scientists that created and used the tools to find them.

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?

Neither the United States Constitution nor U.S. law provide for "things" to have ownership rights. Nothing, other than a natural person or a company via assignment, should be

able to own a patent. As stated above, an AI tool is not a natural person or a company and thus cannot own an invention.

Further, the example in the question asks about the *possibility* of a company who trains the AI process that creates the invention be able *to be an owner*. However, a company, in of itself, cannot be the owner of an invention at the time of its conception. It is the people of the company that create the invention; thus, the people are the initial owners of the invention that results from the training of an artificial intelligence process. Therefore, as is standard practice, the company may become the owner of the invention should the people that conceived of the invention have "hereby assigned" or agree to assign their ownership rights to the company.

The apparent question that the example seems to elicit is, "Should the company that trains the artificial intelligence process that creates the invention be an owner because of the contribution to creating the invention?" Under U.S. law, the answer is no, as the company that performed the training did not contribute to the inventive concept. The company that trains the ML algorithm may develop inventive mechanisms to make training faster or identify more accurate correlations within that training data. As such, the employees of the company performing the training will likely be the inventors and able to file patents on these innovations. However, providing a tool *alone* does not give inventive rights to a technological process that uses this tool, just as with any other software or hardware tool.

Consider the process of manufacturing a hardware processor. A processor manufacturer designs a new and improved processor and hires a semiconductor manufacturer to make the new and improved processor. The semiconductor manufacturer uses a semiconductor manufacturing tool provided by a third-party tool manufacturer. The tool is the only one in the market that is able to make the new and improved processor because of the stringent requirements of the design. When the new and improved processor is manufactured, will the semiconductor manufacturer or the tool manufacturer be considered inventors for the new processor designed by the processor manufacturer? The answer is no. The tool manufacturer could have patent rights to the innovative technology that allows them to make the processor, but the tool manufacturer has no patent rights to the products that are made using their semiconductor manufacturing tool without some other innovation of the products themselves.

In the same way, the people of companies that provide AI services will be entitled to patent rights on the inventions that they conceive, but the patent rights will not expand to the use

of their inventions by others, unless agreed by the parties through some other contractual agreements.

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

No. The current patent eligibility considerations, and in particular eligibility considerations pertaining to software inventions, should apply to AI inventions in the same way that they apply to anything else. When an invention is new, non-obvious, and useful, the invention should be subject-matter eligible.

Some may argue that ML algorithms are directed to mathematical concepts and therefore, should not be patent eligible. However, inventions incorporating mathematical concepts are patent-eligible so long as such concepts are integrated into a practical application.

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?

The current guidelines related to written description require that: "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. Information that is well known in the art need not be described in detail in the specification." *Examining Computer-Implemented Functional Claim Limitations for Compliance With 35 U.S.C. 112*, 84 Fed. Reg. 4, 61. Unless claims are directed to the structure of the "hidden layers" referred to in this question, there is no need to describe their state, and hence no written description issue. The state of such hidden layers will depend on the algorithms and training data, which are very predictable if described sufficiently in the specification. Accordingly, the same criteria used for other inventions in the predictable arts with respect to disclosure should apply to AI inventions.

The requirement for disclosure is not only to prove that the inventor had possession of the claimed invention, but also to allow others to practice the invention without undue experimentation. For an invention that can be practiced with several ML algorithms, the invention may reside in describing the selected features and the training data used. In such case, providing details on the configuration of the ML training may not be necessary if a person skilled in the art would be able to practice the invention without further details. For an invention that resides in a new ML algorithm or in the design of a new cost function that is to be maximized or minimized, describing the training process in detail may not be necessary, and describing the new ML algorithm or cost function and its practical application may be enough. For an invention residing in the configuration of a deep-learning system, a description of the layers of a neural network may be a requirement. On the other hand, if the inventor uses an existing, well-known neural network, a description of such layers may not be necessary.

Regarding the evolution of weights during training, generally speaking, the inventor will need to describe the setup of the training (e.g., training parameters) and the training method used. Describing how weights evolve would not be necessary if such evolution is a function of the algorithm. Similarly, intermediate values of variables during the execution of a software program would not be required, especially if such intermediate values would depend on the inputs provided.

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

AI inventions, tools, and systems should not be compared to the unpredictable arts. AI inventions are inherently no more unpredictable than the underlying ML algorithm on which they rely. ML algorithms implement series of known steps, and while the intermediate results of each step may be unknown, the end result (e.g., the output) is knowable. The enablement requirement for AI inventions should be the same standard as applied for all inventions: the invention must be described in enough detail to allow the person skilled in the art to practice the invention without undue experimentation. AI inventions are no exception, nor are they a special case that requires a heightened standard for written description or enablement.

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?

Such question infers that AI inventions possess some capability that transforms them into something other than a logical construct developed by a natural person. We submit that AI inventions do not currently possess any capability different or greater than that of a natural person using other existing tools, other than the ability to quickly process and "memorize" large amounts of data. Such ability does not change the level of ordinary skill in the art for a natural person. Implementation of the AI invention does not change the analysis of the level of ordinary skill in the art.

The person of ordinary skill in the art is different for each area of technology; no single person is of ordinary skill in all technologies. Just as there are different requirements for persons skilled in the art in the areas of semiconductor manufacturing, microprocessor design, medical diagnosis, rocket science, chemical arts, software development, hardware development, database design, network communications, engine design, etc., there will be different requirements for the person of ordinary skill in the AI arts. These requirements include the ability to practice and understand AI inventions, including ML inventions.

For inventions that improve ML algorithms, the person of ordinary skill in the art may require an understanding of the operation of ML algorithms and a high level of understanding of mathematics (e.g., probability, statistics, calculus, etc.). For inventions that use ML algorithms for a novel particular applications, the person of ordinary skill in the art may require an understanding of the capabilities of different ML algorithms, how to configure them, how to prepare the training data, how to execute the ML algorithms, etc., but not necessarily having to understand the particular details of the operation of the ML algorithms. The level of skill required of a person of ordinary skill in the art will depend on the focus of the AI invention.

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

We believe that the current prior-art considerations are sufficient for AI inventions. There is a danger, especially among those that are not versed in ML technology, to categorize any algorithm performed by a computer as a ML algorithm. We suggest Examiner training to reinforce the concepts of AI in general and ML in particular. For example, ML algorithms require

training data and the identification of features describing the data. Thus, a software program following a procedure that does not include learning from the data will not anticipate a ML algorithm.

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

No new forms of intellectual property protections are required for AI inventions. Improvements in AI itself are patentable as improvements in the software and computing hardware that implements the improved AI. Improvements in other technology or technological fields resulting through use of AI are patentable as improvements of such technology or technological fields. The protection of new, useful, and non-obvious AI tools should be patentable so long as they are integrated into a practical application.

Data used to train a ML algorithm is protectable as a trade secret, since sharing the underlying training data is not required to make use of the resulting ML model. Alternatively, to the extent that the training set is the product of human judgment (e.g., to train a ML algorithm to distinguish between "good" and "bad" poetry), the training set may be protectable under copyright.

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

AI inventions will be ubiquitous to all technology areas and should be properly examined as inventions made in any other way.

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?

The EPO Guidelines for Examination, in Section G-II-3.3.1, discuss the examination of patent claims related to AI, and require that claims recite a technical purpose to be patentable. For claims that recite the application of AI, this requirement is easily met and is appropriate. For claims related to improvements to AI (e.g., novel structures of neural networks or machine-learning layers), the EPO would reject the claims as being "per se of an abstract mathematical nature." We recommend that the USPTO not follow this approach and consider that

improvements to AI are patent-eligible as improvements in computer technology, so long as the claims are sufficiently tied to hardware.

Respectfully submitted,

/Steven W Lundberg/

Steven W Lundberg
Managing Principal
Schwegman Lundberg & Woessner, P.A.