How to get the most out of disclosure interviews with AI inventors

CHAU | IP

2022

1. Set the Context

An AI invention may be directed to one or more inventive aspects including training method, inference method, architecture, and application. It is important to understand from the start to which of these aspects the AI invention is directed. This will set the stage for the interview. For example, an interview for an AI invention that relates to how a neural network is pruned or how a reward function is tuned will focus primarily on the training process and any underlying architecture to support the training. An interview for a novel use-case of a conventionally trained neural network will focus on the inference process and any architecture that supports the inference. Some AI inventions may touch multiple areas of AI and should be discussed accordingly. Sample interview questions will be provided below as examples for guiding the interview.

> ***Interview Question: Does the AI invention relate to data collection, data pre-processing, training, inference, or the underlying AI architecture?***

2. Identify the Problem

Once the focus of the AI invention is determined, it is helpful to frame the AI invention in terms of the problem it addresses (technical challenge) and how the problem is addressed (technical solution). This approach is useful for novelty/non-obviousness considerations under 35 U.S.C. §§ 102 and 103 as well as subject matter eligibility considerations under § 101. Like many other software-related technologies, some AI inventions can face subject matter eligibility challenges during prosecution and/or enforcement. Discussing eligibility considerations with the inventors early in the application drafting process enables the patent practitioner to draft an application that avoids or overcomes subject matter eligibility rejections later.

> ***Interview Question: What technical challenges did you face that led to the AI invention? How did you address them?***
>
> ***Interview Question: Did you encounter any unexpected problems during development? How did you address them?***
>
> ***Interview Question: Why were conventional/existing approaches inadequate?***

3. Identify Technical Advantages

Once the technical challenges are identified, it is important to understand how the AI invention provides technical improvements to those challenges. This is useful when attacking a subject matter eligibility rejection using Step 2A, prong II of the USPTO's eligibility test, namely whether the alleged judicial exception is integrated into a practical application of that exception.

> ***Interview Question: What is improved?***
>
> ***Interview Question: Is the functioning of a computer, another technology, or technical field improved?***
>
> ***Interview Question: How are the improvements realized?***

The improvement need not be an improvement over well-understood, routine, and conventional activities. See MPEP 2106.04(d). Rather, it need only be shown that the functioning of the computer, another technology, or technical field is improved.

For example, consider a neural network trained to operate (i.e., perform inference) on a specific device (e.g., an edge processing device). To the extent that the neural network is designed in such a way that it can operate on the device where a conventional neural network could not, the device is improved. Further, even if a conventional neural network could operate on such a device, the inventive neural network may still improve the functioning of the device if the neural network enables the device to operate more efficiently, with fewer resources, faster, etc. To that end, it is helpful to describe and claim architectural aspects of the invention to show non-conventionality (e.g., a recurrent neural network (RNN) used for natural language processing is a specific instance of a neural network that is trained for a specific purpose and causes the device on which it runs to perform a particular task not performed by conventional computing systems).

4. Understand System Inputs

In the context of inventions involving artificial intelligence or machine learning, it is important to recognize that data can be utilized in several different ways in such inventions and, accordingly, to obtain a clear understanding of how data is used in regard to any given invention. In particular, it is important to determine whether data is being used to develop or train a machine learning model such as a neural network, or being input for use by an already-trained machine learning model to make determinations or decisions.

> ***Interview Question: What kind of data is used as input to the system?***
>
> ***Interview Question: What kind of data is used for training the system?***

5. Understand the Data and Data Collection

An effort should be made during the inventor disclosure interview to gain an understanding of the nature of the data that is being used to develop, or being operated upon by, artificial intelligence or machine learning. How is the data organized or structured? Is the data organized in the form of objects having attributes or in some other manner? What are the actual or possible names/values associated with objects/attributes? It may also be helpful to learn more about the volume of data that is being employed, the sources from which the data is being provided, or how the data is being obtained.

> ***Interview Question: How is the data collected?***

6. Understand Pre-processing Steps

Although artificial intelligence and machine learning involve processing data, the successful development or implementation of artificial intelligence or machine learning depends upon both the quantity and quality of data that is employed. Attaining useful or desirable results by way of artificial intelligence/machine learning can critically depend upon whether the data being utilized is accurate, complete, or properly formatted or normalized. Pre-processing may not, in some cases, seem to be the focus of a given invention. Still, even though the addition of "insignificant extra-solution activity" does not amount to an inventive concept, particularly when the activity is well-understood or conventional (see Parker v. Flook, 437 U.S. 584, 588-89, 198 USPQ 193, 196 (1978)), this general rule should not inhibit the interviewer from gaining an understanding of any pre-processing that may be employed in regard to any given invention relating to artificial intelligence or machine learning.

More particularly, when conducting an inventor disclosure interview, an effort should be made to obtain an understanding of what types of problems can arise in the data being used for development or implementation of the artificial intelligence or machine learning, and how those problems are avoided or alleviated by pre-processing. For example, is raw data pre-processed to correct or add to the data so as to eliminate known problems or deficiencies in the raw data, such as noise that may be impacting the data, or to reorganize or reformulate the data?

Does any such pre-processing entail systems or components that involve more than merely a conventional computer processing device that might provide a further basis for contending that the invention is subject matter eligible and constitutes significantly more than any judicial exception?

The types of pre-processing that can be performed in any given invention can vary significantly depending upon the embodiment, circumstance, or purpose of the artificial intelligence or machine learning. For example, if the artificial intelligence or machine learning relates to image processing, pre-processing can be implemented that involves operations such as changing various geometric features, rotational orientations, or brightness or color characteristics, performing erosion, or dilation, normalizing features, or performing filtering, image segmentation, or super-resolution.

> ***Interview Question: How is the data processed before being input to the system?***

7. Understand Post-processing Steps

In many embodiments or circumstances, the output from artificial intelligence or machine learning operation requires further modifications or processing to be useful. Again, the types of post-processing that can be employed in any given embodiment or circumstance can vary widely depending upon the ultimate purpose or use of the results of the artificial intelligence or machine learning process. Such post-processing can potentially be performed by the same processing device(s) that perform the artificial or machine learning operations, or by other processing devices or other devices or systems.

Again, as with pre-processing, post-processing may not in some cases seem to be the focus of a given invention. Nevertheless, an effort should be made during an inventor disclosure interview to determine whether any post-processing is performed based upon the artificial intelligence/machine learning output and, if so, how such post-processing is accomplished and what components, devices, or systems perform such post-processing.

> ***Interview Question: How is the output of the system processed or made usable?***

8. Understand Network Architecture

A neural network, or parts of a neural network, can often be classified under one of several high-level architectural categories. For example, a convolutional neural network (CNN) includes at least one convolution layer (which finds local features for each data element that take into account neighboring data points), and is often useful for image processing. By contrast, a recurrent neural network (RNN) includes skip connections between different layers, and is often useful for natural language processing, and processing of time series data. In some cases, more than one high level architecture is used.

> ***Interview Question: what high level architecture is used for the invention, and how do these elements relate to the problem being solved?***
>
> ***Interview Question: are multiple high level architectural elements used in the invention?***

In some cases, there is an improvement to the functioning of a particular high level architecture. For example, such an improvement could include the size or number of layers, or the way in which the layers are connected. Furthermore, in many cases, a neural network layer includes a combination of linear (or affine) functions and non-linear activation functions. Thus, an invention could include structure or constraints related to the function of individual layers.

> ***Interview Question: what is the structure of the individual layers within the high level architectural components? What activation functions are used?***
>
> ***Interview Question: how are the layers within a high level architectural component connected?***

In many cases, an AI invention is integrated with another device (e.g., a mobile phone, a robot, or a vehicle). Thus, it is important to understand the context or system architecture of the AI components. For example, an AI system could include various sensors and control systems. Alternatively, an AI element could be a part of a web service that is connected to user devices, databases, and other computing elements.

> ***Interview Question: what is the computing environment of the AI invention? Are there user devices, databases, or other external elements?***
>
> ***Interview Question: does the AI invention include sensors or control elements for interacting with a physical environment?***

9. Understand the Training Process

The training process is one of the key elements that differentiates AI inventions from other software inventions. Training claims have both advantages (e.g., it is sometimes easier to show patent eligibility) and disadvantages (e.g., the training can be difficult to enforce due to detection and split infringement issues). However, it is important to understand the training process before deciding whether to include training claims in the specification or claims.

Before talking to inventors about training, it is important to have a baseline understanding of a typical training process. Machine learning techniques include supervised learning, unsupervised learning, and reinforcement learning. There are also variations on each of these methods, such as the autoencoder technique (which is a variant of supervised learning). Each of these methods typically includes some form of training data, although the training data is used in different ways.

> ***Interview Question: Does the invention use supervised learning, unsupervised learning, and reinforcement learning?***
>
> ***Interview Question: What kind of training data is used to train the network?***

Supervised learning is perhaps the most common training technique. It typically involves calculating something called a loss function that determines how well the network has performed at a given task. For example, a simple loss function could include finding the difference between an output of the network and a ground truth value.

> ***Interview Question: What loss function(s) are used in the training process?***
>
> ***Interview Question: What task does each loss function represent?***

The gradient of the loss function is then calculated (e.g., via a process called back-propagation), and from the gradient, an optimization process called gradient descent is used to determine how to update the network parameters. Training hyperparameters (which may be operator adjustable), such as the number of training batches, the learning rate, and others also impact the training.

> ***Interview Question: What optimization algorithm can be used in the training process?***
>
> ***Interview Question: What hyperparameters can be used in the training process?***

Sometimes different parts of a network use different training methods. Also, different parts of a network can be trained simultaneously, or some layers can be fixed while others are trained. Finally, portions of a network can be trained in different stages. For example, a network can be trained, then pruned, and then refined.

> ***Interview Question: Are different parts of the network fixed, trained separately, or trained jointly?***
>
> ***Interview Question: Are parts of the network trained in multiple phases?***

10. Understand the Process at Inference Time

Inference refers to using a trained model (e.g., a trained neural network) to generate an output using real-world (non-training) data. Consider an example in which an RNN is trained to perform speech recognition. In such an example, inference occurs when the trained RNN receives real-world data such as a sound clip of a person speaking and outputs a textual representation of the person's speech.

It is important to understand the input(s) that the model uses and output(s) that the model generates.

> ***Interview Question: What input(s) are received and what output(s) are generated by the trained model?***

When describing the inference method in a patent application, it is useful to describe the underlying architecture that enables the inference and any improvements that result from the inference method.

> ***Interview Question: What underlying architecture enables the inference?***
>
> ***Interview Question: What improvements result from the inference method?***
>
> ***Interview Question: Are two or more conventional networks combined to address a new problem?***

Also, consider whether the trained model is updated as a result of the inference. This can occur in various ways, such as retraining an existing model or combining predictions from the existing model with a new model created using outputs of the inference. In such cases, the trained model can be improved over time using the inference.

> ***Interview Question: Is the trained model updated as a result of the inference? If so, how?***
>
> ***Interview Question: Are the input data used by the trained model expected to change over time?***