AI Patenting Handbook

CHAU | IP

2023

AI Patenting Handbook

Authors[1]
Frank Chau – Chau IP / F. Chau & Associates, LLC*
Sumon Dasgupta – Jordan IP Law, LLC*

Thomas Burton – Siemens Corporation
Michael Carey – Chau IP / F. Chau & Associates, LLC
Christina Huang - Faegre Drinker Biddle & Reath LLP
Dave Kincaid – McCarter & English LLP
Jennifer Lacroix - Amundsen Davis LLC
Ryan Phelan – Marshall, Gerstein & Borun LLP
John Pienkos – Amundsen Davis LLC
Mike Jennings – AA Thornton
Shankar Krithivasan – Qualcomm Technologies, Inc.
Alexander Korenberg – Kilburn & Strode LLP
Nicholas Malden – D Young & Co LLP
Edoardo Mirabella – Beckman Coulter, Inc.
Justin Mullen – Capital One Financial Corporation
Judy Naamat – Locke Lord LLP
David Pointer – Thomas Horstemeyer LLP
Nick Transier – Dinsmore & Shohl LLP
Wen Xie – Global IP Counselors LLP

[1] This article results from a collaboration between members of the IPO's Software Related Inventions Committee and the AI & Other Emerging Technologies Committee. Opinions stated in this article may not be consistent with the firms, corporations, or clients affiliated with the authors.
* Project Co-Chair

Table of Contents

I. AI Definitions and Technology Overview

Artificial intelligence (AI) is a rapidly expanding field of technology that has the potential to transform society in multiple ways. Accordingly, many innovative firms and individuals are making improvements in AI technology. Advances in AI technology have been accompanied by explosive growth in the number of AI-related patent applications. The following terms and definitions will be useful in understanding the nature of AI and AI inventions.

A. AI Invention Categories

(Michael Carey; Sumon Dasgupta)

AI inventions and AI patents can be categorized as inventions related to core AI technology, applications of AI technology, or the product of AI technology. In some cases, patents and patent claims can include elements of multiple categories. Each category relates to products, designs, processes, computer programs, and other inventions. To understand these categories, we present the following definitions:

<u>AI Inventions on Core AI Technology ("Core AI Technology"):</u>

Inventions that enable AI technology, relate to the operating principles of AI, or have general applicability (i.e., are not limited to specific problem domains). Core AI Technology (e.g., AI enablement technology) relates to the building blocks for application-specific tools, including software (e.g., AI training, architectures, and methodologies) and hardware (e.g., computer processors, accelerator chips, and neuromorphic chips).

<u>Applications of AI Technology</u>

Inventions employing core AI Technologies to perform specific tasks in a particular context. Applications of AI technology often integrate AI Enablement Technologies with domain specific systems that provide input data (*e.g.*, sensors) or use the outputs of the AI Technology for a specific end goal (*e.g.*, to predict outcomes or control machines).

<u>Inventions Made Using AI</u>

Inventions that are conceived of or designed by a human with the assistance of an AI technology (AI as a tool for innovation), a human in collaboration with AI (AI as co-inventor), or an AI system without a human inventor (AI as sole inventor).

B. AI Definitions

(Michael Carey; Alexander Korenberg; Edoardo Mirabella; Justin Mullen)

The following terms will be used throughout the document to describe various aspects of machine learning and artificial intelligence:

Architecture

AI architecture refers to the structure and design of an ML model. For instance, the architecture of an ANN refers to the organization of layers and the connections between the layers in the neural network.

Artificial Intelligence (AI)

Artificial Intelligence is a broad field of computer science aimed at creating machines that match or exceed human intelligence, including perceiving, synthesizing, and inferring information.

Artificial Neural Network (ANN)

An artificial neural network is an ML model that mimics the way the human brain operates. An ANN includes a network of parametric mathematical functions (also called nodes). The nodes are organized in interconnected layers so that the output of the nodes in one layer provides the input for nodes of another layer.

Data Label

Data labeling is the process of adding labels to data so that the data can be better understood by an ML model. A data label is a piece of information that is attached to a piece of data to provide context or meaning. For example, an image of a wolf might be labeled with "wolf," a piece of text might be labeled with "positive sentiment", and transaction data might be labeled with "fraud." Data labels are used to train ML models so that they learn to identify patterns in the data and make predictions.

Deep Learning

Deep learning is a type of ML that employs artificial neural networks with many layers (hence "deep") to model and understand complex patterns in datasets.

Evaluation

Evaluation is the process of assessing the quality, performance, and predictive power of an ML model, using specific metrics such as accuracy, precision, recall, and the like.

Feature

A feature is a measurable property or characteristic of an object or event. In ML, features include data used to train models to make predictions. Features can be numerical, categorical, or text-based. The quality of the features used in an ML model can have a significant impact on the accuracy of the model. In some cases, an ML model generates or modifies features using an encoder portion of the model. A decoder generates an output based on the features.

Feature Engineering

Feature engineering is the process of making raw data more useful for ML algorithms. This can be done by cleaning the data, removing errors, and filling in missing values. It can also involve transforming the data into a format that is easier for the algorithm to understand. Feature engineering is an important step in the ML process, and it can help to improve the accuracy and performance of ML models.

Fine-tuning

Fine-tuning involves taking a pre-trained ML model (often trained on a large dataset) and adjusting its parameters through further training on a smaller, specific task-related dataset, improving the model's performance on this specific task.

Hyper-parameters

Hyper-parameters are settings for the architecture and the learning process of an ML model. Hyper-parameters are usually determined before training. Exemplary hyper-parameters include the number of layers in a neural network and the learning rate (i.e., the “speed” at which the model learns).

Inference

Inference is the process of making predictions or decisions based on a trained ML model and new input data, i.e., data not used during training.

Loss Function

A loss function is an objective function for evaluating how well a specific achieves a task. The loss function is used to update model parameters during training. For example, if the predictions deviate too much from the actual results, the loss function would output a large value. The goal of optimization during the training process is to minimize this loss function.

Machine Learning (ML)

Machine Learning is a branch of AI that studies and constructs mathematical models that automatically learn and improve from experience without being explicitly programmed. While the terms AI and ML are often used interchangeably, "machine learning" more specifically refers to techniques for training machines to perform AI tasks.

Machine Learning Model

A machine learning model is a mathematical model for performing a task, such as making predictions or decisions, interacting with humans through natural language, and interpreting the visual world. Rather than being explicitly programmed to perform its task, an ML model automatically learns and improves from data during the training process. An ML model includes parameters whose values are determined during the training process.

Testing

Testing is the process of assessing the performance of a trained ML model on predetermined input data, typically to estimate how well the model will generalize to unseen data.

Training

Training refers to the process of providing an ML algorithm with training data to learn from. Training can refer to the overall process of developing an ML model, or the specific part of the development process where the parameters of the model are updated. The training process aims at finding a set of values of the model parameters (e.g., weights) that best describe training data. The algorithm obtains the set weights by feeding training data to the ML model and minimizing a loss function.

The following diagram provides an overview of a supervised training process:

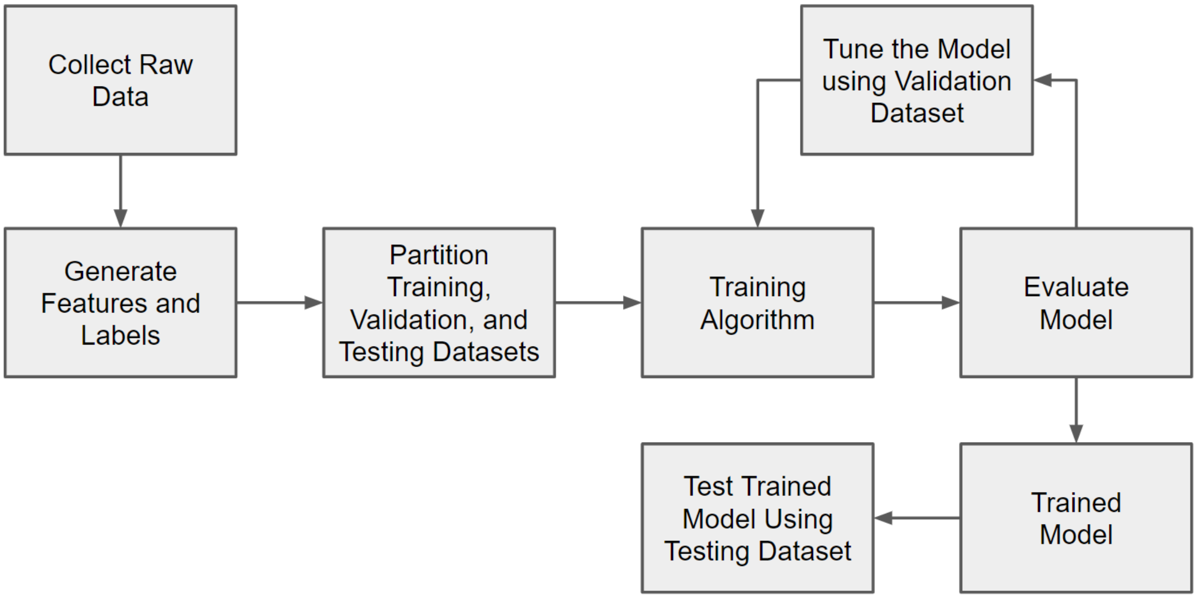

C. Types of AI

(Michael Carey; Alexander Korenberg; Edoardo Mirabella; Justin Mullen)

AI can be categorized by the architectures and applications it is used for. The following terms are useful in understanding the different types of AI:

Attention

In the context of ML, particularly in sequence-to-sequence tasks, attention mechanisms help models focus on relevant parts of the input, thereby improving the model's ability to handle long sequences effectively. Attention can be categorized as cross-attention, where the attention mechanism takes two different types of input and self-attention, where the relevance of different parts of a single input is determined.

Backpropagation

Backpropagation is an algorithm used in training neural networks. It involves efficiently calculating the gradient of a loss function with respect to the weights of a network so that the weights can be modified during training.

Computer Vision

Computer Vision is a subfield of AI that enables computers to interpret and understand the visual world, processing images and videos to detect and identify objects or features. ML models can also be trained to generate images.

Convolutional Neural Network (CNN)

A Convolutional Neural Network is a type of feed-forward neural network particularly effective for image processing and computer vision tasks. It uses convolutional layers with filtering operations to capture spatial features such as edges, corners, and other image characteristics.

Diffusion Model

A diffusion model generates an output by reversing a noise diffusion process, starting from a simple noise distribution and gradually adding details until a sample from the target distribution is obtained. Diffusion models play a central role in text-to-image generation models, where a diffusion network is trained to generate realistic images when conditioned with an embedding of a text prompt.

Embedding

An embedding is a numerical representation of an item of data (for example a word or sentence) that captures essential characteristics of the data so that items of data that are similar are close to each other in a space defined by the embedding.

Feed-forward Neural Network

A feed-forward neural network is a type of neural network where connections between nodes do not form a cycle. That is, information moves only in one direction: from the input layer, through hidden layers, to the output layer.

Generative AI

Generative AI is the subcategory of AI that focuses on generating synthetic content, such as language or images. In some cases, the content is generated based on a prompt or reference data provided by a user.

Large Language Model (LLM)

A large language model is an ML model trained on a vast amount of text data, designed to generate human-like text based on the input it is given. In some cases, an LLM takes a sequence of tokens representing input text, and iteratively predicts additional tokens representing output text.

Natural Language Processing (NLP)

Natural Language Processing is a subfield of AI focused on the interpretation and generation of natural language. In some cases, this allows machines to understand and interact with human users.

Recurrent Neural Network (RNN)

A Recurrent Neural Network is a type of neural network where connections between nodes form directed cycles. This architecture uses the internal state of nodes (i.e., node memory) to process sequences of inputs, which makes them useful for time-series and natural language processing tasks. Long Short-Term Memory (LSTM) networks are a commonly used type of RNN in natural language processing tasks.

Reinforcement Learning

Reinforcement learning is a type of ML where an ML agent learns to make decisions by taking actions in an environment to maximize a reward signal. It involves trial and error. The agent improves its policy over time based on the feedback (reward) it receives after each action.

Self-supervised Learning

Self-supervised learning is when an ML model generates its own labels from the input data for training, for example, by defining a preliminary task that the model needs to solve.

Semi-supervised Learning

Semi-supervised learning leverages a small amount of labeled data in combination with a large amount of unlabeled data during training, optimizing the benefits of both supervised and unsupervised learning.

Supervised Learning

Supervised learning is a type of ML where a model is trained using labeled data, i.e., a dataset that has both input and corresponding output values. The aim is to learn a function that maps input to output. This can be useful for making predictions on unseen data.

Transformer

Transformers are a type of neural network architecture primarily used for processing sequences, for example, sequences of words in the field of natural language processing. They rely heavily on attention mechanisms and are known for their ability to process internal dependencies in data without the need for recurrent or convolutional layers.

Unsupervised Learning

Unsupervised learning refers to a type of ML where the model is trained using unlabeled data, with the goal of discovering underlying patterns or structures, such as clustering or dimensionality reduction.

II. Drafting AI Patent Applications

Patent preparation includes a number of phases from identifying inventions to drafting applications. Accordingly, some practitioners (e.g., in-house counsel) may focus on inventor engagement, invention harvesting efforts, ranking and selection of invention disclosures to pursue, and initial review of disclosure materials. Others may hold formal disclosure interviews with inventors and draft patent documents reflecting their understanding of the invention. Regardless of the phase of preparation, using a clear framework for understanding inventions enables practitioners to distill an invention down to its core advancements over the prior art, and to present those advancements clearly in a patent application. To facilitate that process for AI inventions, the following describes a technology-specific approach to drafting AI patents.

A. Conducting Invention Disclosure Interviews

(Michael Carey; Shankar Krithivasan; Edoardo Mirabella; Justin Mullen)

An invention disclosure interview is a common way for patent practitioners to reach an understanding of an invention in preparation for drafting an application. Some of the suggested disclosure questions may be useful to include in an invention disclosure form provided prior to the invention interview. Regardless of how a practitioner interacts with inventors, leveraging this framework to educate inventors on how to think about their inventions will result in more complete and valuable patent applications.

1. Setting the Context

An AI invention may be directed to one or more categories such as data selection, data preparation, computing environment/constraints, training method, inference method, architecture, and application. Patent practitioners will benefit from determining which among these aspects comprises the inventive subject matter. In doing so, they can set the stage for an invention disclosure interview and efficiently lead discussions with inventors.

For example, an interview for an AI invention that relates to how a neural network is pruned, or how a reward function is tuned, may focus primarily on the training process and any underlying architecture that supports the training. An interview for a novel use of a conventionally trained neural network may focus on the inference process and any architecture that supports the inference. Some AI inventions may be relevant to more than one of these categories. Sample interview questions are provided for determining the context of an invention.

- Disclosure Question: Provide an overview of the invention. What did you do and how did you do it?
- Disclosure Question: Do the advancements of the AI invention relate to data collection, data pre-processing, training, inference, or the underlying AI architecture?
- Disclosure Question: Does the AI invention rely on any open-source components? If so, which ones and what functionality do they provide?

2. Identifying the Problem

Once the focus of the AI invention is determined, it is helpful to frame the invention in terms of the problem it addresses (i.e., the technical challenge). This approach will aid in addressing both prior art considerations under Sections 102 and 103 as well as subject matter eligibility considerations under Section 101. Like other software technologies, some AI inventions can face subject matter eligibility challenges during prosecution and enforcement. Discussing eligibility considerations with the inventors early in the application drafting process enables patent practitioners to draft an application that preempts potential subject-matter eligibility rejections later.

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

3. Identifying Technical Advantages

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

- Disclosure Question: What technology does your invention improve? How?
- Disclosure Question: What makes your solution better than the best existing alternative?
- Disclosure Question: Is the functioning of a computer, another technology, or technical field improved?
- Disclosure Question: How are the improvements realized?
- Disclosure Question: Are there any quantifiable conditions or metrics that you have measured showing a technical improvement (e.g., reduced processing cycles, memory utilization, power use, etc.)?

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 to operate on the device where a conventional neural network could not, such as by compressing its architecture, reducing parameter count, etc., the performance and capability of the device are 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., an 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 more efficiently than would be possible using conventional computing systems).

4. Understanding the System Inputs

For AI inventions, not only can input data vary, but the data may also be utilized in different ways from conventional technology. Obtaining a clear understanding of how an invention uses data should be an important talking point for an invention disclosure interview. In particular, it is important to determine whether data is being used to develop or train an ML model or for use by an already-trained model.

- Disclosure Question: What data is used as input to the system?
- Disclosure Question: What data is used for training the system?
- Disclosure Question: What is the minimum viable type or amount of data that could be used? Could other types of data be used?
- Disclosure Question: Of all the data used, is there a particular type or part that is dispositive to providing the solution?
- Disclosure Questions: In what way is any input data manipulated to be more effectively used by the model?

5. Understanding Data Collection

An effort should be made during the invention disclosure interview to gain an understanding of how the data is obtained. How is the data 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. For example, is the training data and/or input data locally generated or is it obtained from another entity? Is the data made available to a neural network in an online/real-time manner and/or in an offline manner in batches?

- Disclosure Question: How is the data collected?
- Disclosure Question: Are there sensors or other hardware or systems involved in collecting the data?
- Disclosure Question: What is the format of the data when collected?

6. Understanding the Preprocessing Steps

Data preprocessing involves cleaning and preparing raw data for analysis. As part of the data preparation, feature engineering may be used to transform raw data into features that make the data more relevant and informative for ML algorithms.

The successful development and implementation of an AI model depends upon both the quantity and quality of data that is employed. Attaining useful or desirable results by way of ML can depend upon whether the data being utilized is accurate, complete, properly formatted, correctly normalized, and the like. Similarly, feature engineering tasks like selecting which portions of data and modifying the data to best fit the model type may support an advancement of an AI invention. Preprocessing or feature engineering 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 (1978)), this general rule should not inhibit the interviewer from gaining an understanding of any preprocessing that is part of the invention.

For example, 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 AI or ML, and how those problems are avoided or alleviated by preprocessing and feature engineering. For example, is raw data preprocessed 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 preprocessing 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? Does the invention rely on transforming the data in a particular way to allow the model to operate more effectively?

The types of preprocessing and feature engineering that can be performed in any given invention can vary significantly depending upon the embodiment, circumstance, or purpose of the AI or ML. For example, if the AI or ML relates to image processing, these operations can include 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.

- Disclosure Question: How is the raw data processed before being provided to the system?
- Disclosure Question: Did you encounter any issues in using the collected data? How does your solution address those issues?
- Disclosure Question: Is the data used for the invention proprietary, publicly available, or both?
- Disclosure Question: Does the format of the data change between being collected and used by the model? If so, how?
- Disclosure Question: Are there any statistical or aggregates of raw data that are used?
- Disclosure Question: Are input data manipulated, transformed, or the like in order to be used by your invention?

7. Understanding Post-processing Steps

In many cases, the output from an AI or ML operation requires further modifications or processing to be useful. 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. Such post-processing can be performed by the same processing device(s) that perform the artificial or ML operations, or by other processing devices or other devices or systems.

As with pre-processing, post-processing may 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 and, if so, how such post-processing is accomplished and what components, devices, or systems perform such post-processing.

- Disclosure Question: How is the output of the system processed or made usable?
- Disclosure Question: Is there any calibration or transformation of the outputs (e.g., Platt scaling)?
- Disclosure Question: Are output data manipulated, transformed, or the like in order to be used by your invention, or by downstream functions utilizing the output of your invention?

8. Understanding the 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.

- Disclosure Question: Is your model architecture "off-the-shelf" or custom? If custom, how have you modified the standard architecture to improve its performance?
- Disclosure Question: What high-level architecture is used for the invention, and how does the architecture relate to the problem being solved?
- Disclosure Question: Are multiple high-level architectural elements used in the invention?
- Disclosure Question: How do the architectural elements relate to each other?

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.

- Disclosure Question: Is there a preexisting similar architecture to the one used in your invention? How does your architecture differ, and what is the effect of those differences?
- Disclosure Question: What is the structure of the individual layers within the high-level architectural components? What activation functions are used?
- Disclosure Question: How are the layers within a high-level architectural component connected?
- Disclosure Question: How are the layers arranged with respect to other layers?
- Disclosure Question: Are other architectural models used?

In many cases, an AI invention is integrated with another device (e.g., a mobile phone, a robot, or a vehicle). In some cases, the training or deployment of the neural network may be distributed across different components, such as in a federated learning model where different layers of a neural network model are implemented in different devices or components. 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. Additionally, an AI element can be a part of a service that is connected to user devices, databases, and other computing elements.

- Disclosure Question: What is the computing environment of the AI invention? Are there user devices, databases, or other external elements?
- Disclosure Question: Can an off-the-shelf model be used?
- Disclosure Question: Why have you selected this particular model? What are the advantages of using this model?
  Disclosure Question: Which hyperparameters have been set? Why have those values been chosen?

9. Understanding 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 claims 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.

Before talking to inventors about training, it is useful to have a baseline understanding of a typical training process. ML techniques include supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning, to name a few common examples. There are also variations on each of these methods, such as autoencoder techniques and diffusion processes. Each of these methods typically includes some form of training data, although the training data is used in different ways.

- Disclosure Question: what type of learning algorithm is used for your AI invention?
- Disclosure Question: What training data is used to train the network?
- Disclosure Question: How has the data been obtained or labeled? Are any automated procedures used for labeling?

Supervised learning is a common training technique. It typically involves calculating something called a loss function that determines how well the model has performed at a given task. For example, a simple loss function could include finding the difference between an output of the model and a ground truth value. The gradient of the loss function is then calculated, and from the gradient, an optimization process called gradient descent (or ascent) is used to determine how to update the model parameters. Training hyperparameters (which may be operator adjustable), such as the number of training batches, the learning rate, and others also impact the training.

- Disclosure Question: What loss function(s) are used in the training process?
- Disclosure Question: What are the components of the loss function, and what does each component influence in terms of model performance and training?
- Disclosure Question: What task does each loss function represent?
- Disclosure Question: What optimization algorithm can be used in the training process?
- Disclosure Question: How are hyperparameters set to improve the training process?

Sometimes different parts of a model use different training methods. Also, different parts of a model can be trained simultaneously, or some portions of the model can be fixed (or “frozen”) while others are trained. Finally, portions of a model can be trained in different stages. For example, a model can be trained, tuned, pruned, and refined.

- Disclosure Question: Are different parts of the model fixed, trained separately, or trained jointly?
- Disclosure Question: Are parts of the network trained in multiple phases?
- Disclosure Question: Is the model architecture trained end-to-end or in different phases?
- Disclosure Question: Are any aspects of the model discarded after training?

10. Understanding the Inference Process

Inference refers to using a trained model (e.g., a trained neural network) to generate an output using unseen data. Consider an example in which an RNN is trained to perform speech recognition. In this 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 inputs that the model uses and outputs that the model generates. When describing the inference method in a patent application, it is also useful to describe the underlying architecture that enables the inference and any improvements that result from the inference method.

- Disclosure Question: What input(s) are received and what output(s) are generated by the trained model?
- Disclosure Question: What underlying architecture enables the inference?
- Disclosure Question: What improvements result from the inference method?
- Disclosure Question: Are existing model architectures combined to address a new problem or to improve performance?

Furthermore, 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 inference.

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

B. Drafting Claims for AI Inventions

(Frank Chau; Michael Carey; Nick Transier; Wen Xie)

The claims are a critical part of any patent application as they define the legal rights associated with the patent. AI claims should be drafted with certain AI-specific factors in mind. Although AI-based inventions cover a broad range of technical domains and concepts, most AI inventions are computer-implemented inventions. Many of them are implemented in software. As such, best practices for claiming computer and software inventions generally apply to claiming AI inventions as well.

However, recent experience with patent prosecution in the US shows that characterizing a claim as being performed by AI will generally not be sufficient. Thus, there are many AI-specific concepts that merit individual attention when considering how to best claim an AI invention.

1. Subject Matter Eligibility (§101)

To ensure that AI-related patent claims are patent eligible under 35 U.S.C. §101, it is useful to understand the test that examiners use to determine eligibility. Specifically, MPEP 2106 provides the following phases of evaluation: 1) determine if the claims fall into one of the statutorily-defined categories of patentable subject matter; 2) ask if the claims recite an abstract idea (i.e., Alice/Mayo step 2A, Prong 1); 3) consider whether the claim is integrated into a practical application (i.e., Alice/Mayo step 2A, Prong 2); and, 4) determine whether additional elements amount to "significantly more" than the abstract idea (i.e., Alice/Mayo step 2B).

Although most AI-related inventions can easily be claimed in terms of one of the statutorily defined categories of eligible inventions (e.g., a process or machine), they may be susceptible to characterization by a patent examiner as an abstract idea, such as a mental process, mathematical concept, or method of organizing human activity.

Thus, there are three basic strategies for claim AI inventions, corresponding to steps 2, 3, and 4 of the process outlined in the MPEP: 1) avoid using language that obviously invokes an abstract idea, 2) include language that integrates the claim into a practical application, and 3) include "additional elements" that are integral to the invention.

When possible, to avoid characterization as an abstract idea, avoid the following:

- Steps that sound like mental processes (e.g., predict, identify, evaluate, etc.)
- Mathematical equations and notation
- Words that invoke business or legal applications (e.g., contracts, advertising, sales, etc.)

To integrate a claim into a practical application, include steps that are directly related to improving a particular technology, such as:

- Improving the accuracy of an AI model
- Reducing the size, the number of parameters, or the number of layers of a model
- Reducing the amount of training data needed
- Enabling the use of different kinds of training data
- Improving the training speed or efficiency
- Enabling lifelong learning (e.g., utilization of previously learned parameters without complete retraining)
- Increasing the speed or efficiency of network operation/prediction
- Enabling the use or optimization of different hardware (e.g., GPU vs CPU)

To ensure that a claim includes "significantly more" than a judicial exception, include language directed to specific elements, such as:

- Specific architectural elements

- Steps that are related to specific ML architectural elements (e.g., encoding/decoding, embedding, performing a convolution, etc.)
- Steps that cannot be performed in the human mind
- Steps or components that involve physical sensors for collecting data
- Steps or components that result in a physical transformation
- Hardware components that are not part of a generic computing environment

2. Definiteness (§112)

Practitioners must also be mindful of the enablement and indefiniteness requirements under 35 U.S.C. §112. AI inventions are often claimed functionally or in such a way as to be performed by a general computing device. Examiners and courts may interpret such claims as styled in Means Plus Function format (MPF), invoking Section 112(f). In such instances, a description of structure corresponding to the claimed nonce or generic term that performs the claimed function must be present in the specification, or the claims can be rejected as failing the definiteness requirement under Section 112(b). Such description must be more than a repeat of the same generic term.

To avoid characterization as MPF, refrain from using non-structural terms that do not have a specific meaning in the art in the body of the claim. Terms considered non-structural include: Module; Component; Unit; Element; Mechanism; System; and Apparatus.

In some cases, more specific terms can be used that have a particular meaning in the field of AI, including terms such as: Machine learning model; Artificial Neural Network (ANN); Convolutional Neural Network (CNN); Generative Adversarial Network (GAN); and Reinforcement Learning model

If a non-structural term is used to achieve the desired claim scope, describe in the specification whether the element is a hardware element or if it is implemented using software (or provide a sufficient description of both implementations). If no structure is provided in the specification, then the claim may be rejected during prosecution or later held invalid. This can be the case even in situations where methods of performing the invention are well known. If the claim or the specification is deficient with respect to how the claimed element functions, the claim may be rejected during prosecution or later held invalid under Section 112(a).

3. Training Claims

Many AI inventions are enabled by trained ML models, and the training of these models is often claimable subject matter. In some cases, the training methodology is the primary invention (e.g., federated / distributed learning type inventions), whereas in other cases a conventional training methodology is a means to another inventive end.

Claims directed to training have a head-start at the USPTO because MPEP 2106.04(a)(1), example *vii* includes a training claim explicitly described as eligible subject matter. However, as described further below with reference to prosecuting patents, the USPTO has begun to reject claims similar to the MPEP's example. Accordingly, practitioners should consider strategies such as avoiding reciting explicit mathematical concepts (e.g., formulas) in the claims and describing steps of the training method within the specification as complex, computationally intense, and impractically performed in the human mind. These and other strategies will be further discussed in the sections below.

Training comes in many varieties, including supervised learning, semi-supervised learning, and unsupervised learning to name just a few. Unique aspects of each of these different methods may provide claimable content. Further, in most cases, a training method requires some objective metric to measure the progress of the training, such as a loss or objective function. Because the structure of loss functions can significantly affect the outcome of training, novel loss functions may also be claimable subject matter. In some cases, a portion or even an entire model architecture may be dedicated to training and then jettisoned once the model is trained, such as with autoencoder models. In such cases, the unique configurations of the training architecture may provide claimable subject matter.

Finally, training ML models is generally a computationally intensive task. Accordingly, special-purpose hardware (e.g., a GPU, a TPU, and/or a so-called ML or AI "accelerator") is often used to speed up training. Claiming the special-purpose hardware can help to overcome a rejection based on a "generic computer" implementation.

4. Inference Claims

AI claims directed to inferencing (i.e., using a model to generate a useful output, as opposed to training) have several distinct species, including claims directed to: (1) a standard ML model performing a novel task; (2) a novel ML model performing a standard task; and (3) a novel ML model performing a novel task. Note that while "model" is used as a convenient noun in the discussion below, it is often the case that an AI-related invention uses a processing architecture that may include one or more models as well as other processing elements.

For claims directed to a standard ML model performing a novel task, the strategy should include a focus on what the novel task is because the ML model is a "black box" unlikely to confer eligibility by itself. In other words, similar to *Alice's* rationale, the mere recitation of a generic ML model cannot transform a patent-ineligible abstract idea into a patent-eligible invention. However, citing specific elements of the operation of the ML model can still prevent rejection as ineligible subject matter even if the model is not novel.

For claims directed to a novel ML model performing a standard task, a good strategy is to explicitly claim what makes the model novel, such as a specific configuration or combination of layers or connection pathways between the layers for a neural network model. Further, the combination of different types of models (e.g., an ensemble architecture) can combine benefits of different model types in novel ways. The section below regarding drafting eligible architecture claims provides additional guidance.

For claims directed to a novel ML model performing a novel task, reciting elements of the novel task may allow for claiming the novel model more flexibly or more broadly. Because explicitly claiming specific model architectures increases the opportunity for design-around and may require disclosure of aspects of a model that might alternatively be kept trade secret, the novel task can be leveraged as a complementary tool for eligibility.

5. Architecture Claims

Best practices for drafting AI architecture claims are similar to those for drafting training and inferencing-type claims as discussed above. However, whereas claims related to training and inferencing often include elements covering the respective processes end-to-end, an architecture claim might instead focus on a particular portion of the architecture. This is because, for example, a conventional training technique can be applied to a novel architecture to perform a conventional inference task.

In many cases, a complete claim set will include training, inference, and architecture claims that each have a different scope and rely on different patent eligibility strategies. That is, some elements that have been included in other claims may be omitted in architecture claims if other elements are included to ensure eligibility. There are multiple aspects of an AI architecture to consider, including the software (or process) architecture as well as the hardware (or compute) architecture.

In particular, AI architecture claims may be directed to specific components and processing patterns that can be recited at a high level of generality without being considered generic computing components. Examples of these architectural components include a convolutional neural network (CNN), a recurrent neural network (RNN), transformer-based models (e.g., Bidirectional Encoder Representations from Transformers (BERT) and Generative Pre-Trained Transformer (GPT), to name but a few). These components are "particular machines" that are not generic because 1) such components are not present in most typical computing systems and 2) such components are often suited for particular applications (e.g., a CNN is well-suited for image processing applications, whereas an RNN may be better suited for natural language processing (NLP)).

Referencing specific AI architectural elements in a claim can prevent a claim from being considered a judicial exception under step 2 of the USPTO's subject eligibility guidelines. For example, using training or architectural elements can provide grounds for showing that the claim is integrated into a practical application. Even where the inventive aspect relates to inference, it can also be useful to describe and claim architectural aspects of the invention to show non-conventionality. For example, an RNN used for NLP 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.

There are also ways to incorporate architectural elements into inference claims. For example, a step of a method claim could recite an architectural element that performs the step, or the step could be phrased in a way that specifies a particular kind of processing. For example, a claim that uses a CNN could claim "performing a neural network convolution" on input data.

Finally, AI inventions often include hardware aspects. For example, self-driving vehicle systems rely on the input of specific types of hardware sensors, such as cameras and ranging sensors, to provide input data to trained ML models, which in turn generate the outputs that drive (or assist with driving) the vehicle.

6. Problem/Solution Claims

Claims that do not recite specific architecture or training elements can still be crafted to avoid a subject matter eligibility rejection. Specifically, a claim can recite a problem, a technical solution, and the resulting improvement over prior art systems.

For example, if an image recognition system is used to navigate through a physical environment, the input to the system could be described in a way that identifies a problem (e.g., "receive a 2D image that depicts an obstacle in an environment"). Intermediate steps can describe the technical solution (e.g., "generate a depth map indicating a distance of the obstacle"). Finally, the result can be described as an action taken based on the ML model (e.g., "navigate through the environment while avoiding the obstacle based on the depth map").

Accordingly, in addition to including architectural and training elements, AI claims can be patent eligible if they include elements directed to improving the functioning of a computer or an ML model. Such claims should include specific terms that correspond to 1) the problem, 2) the technical solution, and 3) the result. The problem, technical solution, and result should be fully described in the specification to provide support for advocating that the claim is not directed to a judicial exception or includes "significantly more" than a judicial exception.

C. Drafting the Specification

(Frank Chau; Michael Carey; Christina Huang; Judy Naamat; Nick Transier; Wen Xie)

While the claims are important in defining AI inventions, there are a number of AI-specific techniques applicable to the written description to avoid or overcome rejections and challenges based on Sections 101 and 112.

1. Subject Matter Eligibility (§101)

Although subject matter eligibility under 35 U.S.C. §101 is primarily based on the claims, the specification can impact whether the claims are deemed ineligible. The primary reason for this is that claimed terms are understood in light of the specification. For example, the MPEP 2106.04(a)(2) III(c) states "in evaluating whether a claim that requires a computer recites a mental process, examiners should carefully consider the broadest reasonable interpretation of the claim in light of the specification."

The fact that claim terms are understood in light of the specification is particularly important because the terminology used in an AI/ML context can have meanings that differ from conventional usage. For example, the term "predict" can result in the determination that an element of a claim is directed to a mental process. If terms such as "predict", "identify", "observe", "evaluate" are used in a claim, it is important to include a technical description of these terms that describes them in an AI/ML context with a technical definition that can differentiate from a mental process.

Even terms such as "encode" could potentially be performed by a human with the aid of pen and paper, so providing technical details of how encoding is performed can help avoid a claim interpretation that causes a claim to be ineligible. For example, details can be included that clarify that "encoding" includes representing data in a high-dimensional embedding space, or that the encoding refers to a process performed with a particular ML model.

A claim that includes abstract elements can still be patent eligible if the judicial exception is integrated into a practical application. For example, the specification may include descriptions of a problem in the prior art, the technical solution, the benefits derived therefrom. The improvements may be to computer functionality, or to any other existing technology. The specification is also considered when determining whether a claim includes an additional element that amounts to significantly more than an abstract idea. Here, it is important to clearly describe any additional elements that are included in a claim and how they are related to the inventive concept.

2. Written Description, Enablement, and Best Mode (§112(a))

35 U.S.C. §112(a) includes three distinct requirements: 1) a written description of the invention;

2) a description of the manner and process of making and using the invention (i.e., the enablement requirement); and 3) the best mode contemplated by the inventor of carrying out his invention.

For many AI-inventions, the written description requirement can be viewed as relating more to architectural elements, whereas the enablement requirement relates more to training (i.e., making) and inference (i.e., use). Thus, novel details relating to training, inference, and architecture should all be described in virtually every AI patent application.

The enablement requirement states that the specification must describe the claimed invention in a manner understandable to a person having ordinary skill in the art (PHOSITA). Unfortunately, the term “a person of ordinary skill in the art” is not perfectly clear in a fast-moving field such as in AI. Terms that seem obvious to inventors might be completely unknown to examiners and judges. Furthermore, architectures, such as transformer models and diffusion models, can go from being virtually unknown to widely used in a matter of months. In section 4 below, guidance is provided for levels of descriptions of AI inventions in the specification to help avoid written description rejections.

Finally, while some would argue that the America Invents Act (AIA) eliminated the consequences for failing to disclose the inventor’s best mode, the requirement still exists. For AI applications, this means that a practitioner should describe any preferred forms of training, architecture, and inference.

3. Definiteness (§112(b))

35 U.S.C. §112(b) requires that claims must particularly point out and distinctly claim the invention. As with the patent eligibility rules, this requirement relates directly to claims, but the specification will help determine whether the claims are considered definite.

The following can be helpful for establishing definiteness under Section 112(b) with respect to a function: provision of an algorithm; a description of how a function is performed; a description of what elements are used to perform the function; a description of how exceptions are handled; and provision of examples of how exceptions are handled.

Regarding definiteness under Section 112(b), a determination of indefiniteness may be avoided by providing method steps that are performed automatically in a way that can be readily understood by a PHOSITA, as opposed to inclusion of user actions, which are more likely to render a claim indefinite. In some cases, claim terms for well-known or off-the-shelf ML and/or AI modules may avoid invocation of Section 112(f), whereas specialized ML modules would require sufficient description to provide structure, such as via algorithms, a description of the input, output, and functions performed, a description of how exceptions are handled. Examples are helpful in showing that an inventor had possession of the invention at the patent application’s priority date.

The Supreme Court, in *Amgen Inc. v. Sanofi*, has emphasized that the broader the scope of the claims, the more the applicant must enable with teachings in the specification. The teachings need to provide enough information that a PHOSITA could make or use the invention without more than a reasonable amount of experimentation. Therefore, when claiming a genus, a description needs to be provided for making or using each member of the genus. It is tempting (and sometimes wise) when applying ML to a field that is new to ML to claim broadly. Such breadth needs commensurate description, with consideration taken of how much experimentation would be needed by a PHOSITA to design each of the claimed systems or to perform each of the claimed methods.

4. Three Levels of Description

The key to describing AI and ML architecture is to become familiar with several levels of abstraction. As a rule of thumb, an AI invention should be described on three levels, including a first level based on the target function, a second level based on high-level technical description, and a third level based on a detailed technical description. By understanding these levels of abstraction, it is possible to describe how the structure performs a claimed function.

For example, when describing architecture, the highest level of description is the functional component description (e.g., an image classification network). On its own, this description may not be enough to satisfy the requirements of Section 112. Furthermore, functional terms might be interpreted as a nonce under Section 112(f). Thus, it is essential to include another level of description based on the high-level architectural paradigms described above (e.g., CNN, RNN, feed-forward network, etc.). The specification should also describe the relationship between the problem to be solved (e.g., image recognition) and the high-level architecture used to solve it (e.g., a CNN).

However, to fully satisfy the requirements of Section 112, it is important to include details that go even deeper than a recitation of the basic architecture. Thus, a third level of description can be included that provides technical details about the operation of the network at the level of layers, nodes, and activation functions. This does not require describing the actual parameter. For example, the specification could include a description of how a CNN works at the node level (e.g., describe the role of different filters of the CNN). If possible, the inventive concept should be woven into the description at each level of description.

As with the different levels of architectural description, there are also three levels of description that are useful when describing the training process. Again, at the highest level, one should include a functional description (e.g., a neural network trained to classify objects in an image). At the next level, provide details broadly descriptive of a high-level training paradigm, such as supervised learning, unsupervised learning, or reinforcement learning. Then, at the level of fine detail, provide specifics related to, e.g., the loss functions of a supervised learning process or a policy model of a reinforcement learning process.

5. Problem/Solution Description

In addition to technical details related to the architecture and training, special attention should be given to providing a description of a technical problem, the technical solution, and the resulting improvement over related technologies.

When describing a problem in the existing technology, a practitioner should provide enough description to motivate the solution provided by the invention without conceding too much background as prior art. However, it is generally useful to identify a field of technology, provide a generic name for some device or task in the field, and then describe a problem faced when implementing such a task or device.

When describing the technical solution, provide a description of the structure that performs each function recited in the claims, including how such structure performs the function. If the structure is implemented in software, then provide an algorithm for how the function is performed. Furthermore, if an AI-specific term is claimed, then provide a non-limiting definition or provide a description of an exemplary use of the claim term in the specification.

Thus, in addition to a description of the problem and a description of the technical solution, it is useful to describe how the technical solution results in an improvement over existing systems and methods. Preferably, the improvement should relate directly to the problem and be specific to a particular area of technology.

III. Prosecution and Enforcement

Many of the considerations for prosecuting and enforcing AI applications are similar to those for other areas of technology (e.g., traditional software inventions). However, several unique considerations for AI applications are described below. Of particular importance is the evolving manner that AI patents are considered under subject matter eligibility analysis.

A. Office Action Rejections

(Frank Chau; Michael Carey; Nick Transier)

1. Patent Eligibility Rejections (§101)

There are three key strategies for responding to an Office Action that includes a Section 101 rejection. The three strategies are:

- Arguing that a claim does not recite a judicial exception;
- Arguing that a claim is integrated into a practical application; and
- Arguing that additional elements (or a combination of elements) add significantly more than conventional solutions.

These strategies parallel the three strategies for *avoiding* a 101 Specifically, they correspond to steps 2-4 of the subject eligibility test. Recall that the four steps are: 1) determine if the claims fall into one of the statutorily defined categories of patentable subject matter; 2) ask if the claims recite an abstract idea; 3) consider whether the claim is integrated into a practical application; and 4) determine whether additional elements amount to “significantly more” than the abstract idea. When an examiner rejects a claim under Section 101, they must go through all four of these steps (unless the claim is rejected at step 1).

The implementation of these strategies differs somewhat from strategies for other technology due to some specifics about AI technology. Furthermore, the manner in which the USPTO implements the four-part test is rapidly changing. The following provides evidence-based strategies based on recently collected data and observations about how AI-related applications are being treated at the patent office.

2. Patent Eligibility Rejection Study

A review of 200 AI-based patent applications that were rejected under 35 U.S.C. §101 for office actions issued between 1/1/23 to 9/30/23 found that:

<u>Step 1</u>

No applications were found to be ineligible at step 1.

<u>Step 2 (2A Prong 1)</u>

- 141 out of 200 applications (71%) were rejected under mental process category (37 of these are rejected with human activity category and 23 of these are rejected with mathematical category as well)
- 84 out of the 200 applications (42%) were rejected under the human activity category (4 of these applications were rejected under the mathematical concept category as well)
- 30 out of the 200 applications (15%) were rejected under the mathematical concept category.

<u>Step 3 (2A Prong 2)</u>

- 184 out of the 200 applications (92%) were found to use generic computer components to perform the abstract idea. (e.g., “computer system/product, device, memory, processor, non-transitory computer readable medium, ML model”)
    - 111 applications included a “processor”
    - 48 applications included a “ML model” or “AI model/technique”; and
    - 11 applications included a “neural network”
- 71 out of the 200 applications (36%) were rejected based on limitation reciting data gathering that amounts to insignificant extra-solution activity.
- 66 out of the 200 applications (33%) were rejected for other reasons. For example, the claim as a whole does not integrate the abstract idea into a practical application or does not amount to an improvement in a technical field, etc.

<u>Step 4 (2B)</u>

- 193 out of the 200 applications (97%) were evaluated and rejected under this step.
- 89 out of the 200 applications (45%) were rejected for being well-understood, routine, conventional activity in the field. Of these rejected applications, most of the Examiners reasoned that either data gathering is conventional or using generic computer components is conventional.

3. Mental Process Rejections

Of the three categories of abstract ideas, "mental process" is the most cited in rejections of AI claims, about 71% in 200 rejections. Therefore, "mental process" rejections merit special attention.

One strategy for overcoming a mental process-type rejection is to amend the claims by modifying language that sounds like a mental process (e.g., change the word "predict" to "compute") or adding additional language specifying architecture-specific elements that cannot practically be performed mentally (e.g., "performing a convolution on an image" or "generating an embedding vector representing text"). Alternatively, a practitioner can simply argue that the existing language does not recite a mental process.

The MPEP explains that "[t]he courts consider a mental process (thinking) that "can be performed in the human mind, or by a human using a pen and paper" to be an abstract idea. MPEP § 2106.04(a)(2)(III) (citing *CyberSource Corp. v. Retail Decisions, Inc.*, 654 F.3d 1366, 1372 (Fed. Cir. 2011)). "[E]xamples of mental processes include observations, evaluations, judgments, and opinions." *Id*. And notably, "both product and process claims may recite a "mental process," *id*.; in other words, an apparatus configured to perform a so-called mental process would be rejected as a method claim directed to the mental process.

While claims requiring a computer to perform the alleged mental process run afoul of the rule, *id*., there is an important carve-out. Namely, "[c]laims do not recite a mental process when they do not contain limitations *that can practically be performed in the human mind*, for instance when the human mind is not equipped to perform the claim limitations. MPEP § 2106.04(a)(2)(III)(A) (citing *SRI Int'l, Inc. v. Cisco Systems, Inc*., 930 F.3d 1295, 1304 (Fed. Cir. 2019)) (emphasis added).

Note that "impractical" and "inconvenient" to perform in the human mind are not the same, and improving convenience and speed by using a computer to assist will not save a claim. *See generally* MPEP § 2106.04(a)(2)(III)(C). The MPEP provides examples of claims that do not recite mental processes because they cannot be practically performed in the human mind, including:

1. A claim to a method for calculating an absolute position of a GPS receiver and an absolute time of reception of satellite signals, where the claimed GPS receiver calculated pseudo-ranges that estimated the distance from the GPS receiver to a plurality of satellites.
2. A claim directed to detecting suspicious activity by using network monitors and analyzing network packets.
3. A claim to a specific data encryption method for computer communication involving a several-step manipulation of data.
4. A claim to a method for rendering a halftone image of a digital image by comparing, pixel by pixel, the digital image against a blue noise mask, where the method required the manipulation of computer data structures (e.g., the pixels of a digital image and a two-dimensional array known as a mask) and the output of a modified computer data structure (a halftoned digital image).

See MPEP § 2106.04(a)(2)(III)(A). Notably, the difference between these eligible examples and ineligible examples, such as "a claim to 'collecting information, analyzing it, and displaying certain results of the collection and analysis,'" *id*., seems to be related at least in part to the level of specificity in the claim. After all, the eligible GPS method (example 1, above) is an example of "collecting information" (e.g., GPS radio signals) and "analyzing it" to generate an output (e.g., a position estimate).

Relatedly, the PTAB has held, on several occasions, that when claim elements are described in the specification as being "complicated", "difficult to scale", "computationally intensive", or "computationally complex", they are thereby "impractical to perform in the human mind." *See*, *e.g*., *Ex parte Akli Adjaoute*, Appeal 2018-007443, p. 11 ("the 'classifying' steps of claims 1 and 'modules' of claim 8 when read in light of the Specification, recite a method and system difficult and challenging for non-experts due to their computational complexity. As such, we conclude that one of ordinary skill in the art would not find it practical to perform the aforementioned 'classifying' steps recited in claim 1 and function of the 'modules' recited in claim 8 mentally."); *Ex parte Jean-Baptiste Tristan*, Appeal 2018-004459, p. 6 ("when read in light of the Specification, the claimed 'identifying a particular inference algorithm' is difficult and challenging for non-experts due to their computational complexity. As such, we conclude that one of ordinary skill in the art would not find it practical to perform the aforementioned 'identifying' step mentally.")

AI-related inventions tend to be complex and computationally intensive. Furthermore, some AI concepts are not easy to claim directly because the parameters are not interpretable. However, a practitioner should describe the elements of an AI claim as complex and computationally intensive in the specification to address a Section 101 rejection later. In fact, a practitioner can explicitly describe steps of AI-related inventions as impractical to perform in the human mind. By way of example, if an ML model enables an AI-type invention and takes a month to train on a supercomputer, then a practitioner can easily point out the computational complexity and impracticality of performing such a process in a human mind.

4. Organizing Human Behavior Rejections

The second most common category of abstract idea used in rejecting AI-related applications was the 'methods of organizing human behavior', which was cited in about 40% of the cases. This rejection is particularly common in the 3600 "Business Methods" art unit, where many business-related claims are examined.

The MPEP explains that the "methods of organizing human behavior" category includes the following subcategories:

- fundamental economic principles or practices (e.g., hedging, insurance, mitigating risk);
- commercial or legal interactions (e.g., agreements in the form of contracts, legal obligations, advertising, marketing or sales activities or behaviors, and business relations); and
- managing personal behavior or relationships or interactions between people, (e.g., social activities, teaching, and following rules or instructions).

See MPEP 2106.04(a)(2). However, USPTO guidance explains that "not all methods of organizing human activity are abstract ideas (e.g., 'a defined set of steps for combining particular ingredients to create a drug formulation' is not a certain 'method of organizing human activity'), In re Marco *Guldenaar Holding B.V.*, 911 F.3d 1157, 1160-61, 129 USPQ2d 1008, 1011 (Fed. Cir. 2018). Second, this grouping is limited to activity that falls within the enumerated sub-groupings of fundamental economic principles or practices, commercial or legal interactions, and managing personal behavior and relationships or interactions between people, and is not to be expanded beyond these enumerated sub-groupings except in rare circumstances" *Id.*

In other words, examiners are not free to make up new sub-categories of the "methods of organizing human behavior" category of abstract idea. Accordingly, the primary method of overcoming an organizing human activity rejection is to avoid explicitly reciting words known to relate to these recognized methods of organizing human behavior. Many practitioners avoid using these terms in both the specification and claims to avoid getting applications sent to Art Unit 3600 in the first place.

Although many AI applications solve problems related to the types of activities considered "methods of organizing human behavior," it is usually possible to describe claims in a way that does not directly invoke these activities. For example, claims can focus on the processes performed directly by the AI model (i.e., computing vectors, text and images), as opposed to the final objective these outputs will be used for (e.g., creating a legal agreement).

5. Mathematical Concept Rejections

Surprisingly, the least common category of abstract idea used in rejecting AI-related applications was the "mathematical concept" category. It was cited in about 15% of rejections.

AI-inventions are, of course, mathematical concepts in the broadest sense, however, claims covering such inventions need not–and indeed should not–explicitly recite mathematical formulas to avoid abstract idea rejections based on the mathematical concept group. For example, in an AI-related claim directed at training an ML model, consider reciting elements of a loss function descriptively, rather than the formula representing the exact loss function.

As another example, in an AI-related claim directed at inferencing, consider claiming the specifics of the input data, the transformation to the data made by the model, and the specifics of the output data, as well as any input data preparation steps if applicable, rather than the mathematical operations that affect such transformations of input to output. As yet another example, for an AI-related claim directed at a particular ML model architecture, consider claiming the inclusion of particular types of layers and connections therebetween, as well as the function of those layers, without reciting specific mathematical operations (e.g., dot products, summations, etc.) performed by those layers.

The MPEP notes, when discussing the seminal case *Diamond v Diehr*, that "[t]he Court's rationale for identifying these 'mathematical concepts' as judicial exceptions is that a "*mathematical formula* as such is not accorded the protection of our patent laws,"' MPEP § 2106.04(a)(2)(I)) (quoting *Diamond v. Diehr*, 450 U.S. 175, 191(1981)) (emphasis added). Note the emphasis on "mathematical formula" in that rationale. Consistent with this, the MPEP explains that:

> When determining whether a claim recites a mathematical concept (i.e., mathematical relationships, mathematical formulas or equations, and mathematical calculations), examiners should consider whether the claim recites a mathematical concept or merely limitations that are based on or involve a mathematical concept. *A claim does not recite a mathematical concept (i.e., the claim limitations do not fall within the mathematical concept grouping), if it is only based on or involves a mathematical concept*.

Id (emphasis added). While the MPEP further adds that "a mathematical concept need not be expressed in mathematical symbols" to be recited in a claim, the PTAB has held on multiple occasions that a claim is not directed to a mathematical concept where the claims did not recite a specific algorithm or formula. *See*, *e.g.*, *Ex parte Akli Adjaoute*, Appeal 2018-007443 (October 10, 2019) (reversing a Section 101 and noting: "the specific mathematical algorithm or formula is not explicitly recited in the claims. As such, under the recent Revised 101 Guidance, the claims do not recite a mathematical concept."); *Ex parte Jean-Baptiste Tristan*, Appeal 2018-004459 (June 25, 2019) ("while the Specification identifies 'implement[ing] an inference algorithm' … the specific mathematical algorithm or formula is not explicitly recited in the claims. As such, under the recent Memorandum, the claims do not recite a mathematical concept."); *see also* PEG Example 38 ("The claim does not recite a mathematical relationship, formula, or calculation. While some of the limitations may be based on mathematical concepts, the mathematical concepts are not recited in the claims.”).

6. Specific Structures and Practical Applications

As noted above, over 90% of the recent AI applications were rejected under step 3 (2A, Prong 2) for using generic computer components to perform the abstract idea. In many cases, this is because examiners treat this Prong superficially.

However, it is often possible to overcome 101 rejections on this basis by:

- Arguing that the claim recites specific structures that cannot be considered “generic computer components”; or
- Arguing that the claim is integrated into a practical application.

That is, in step 3 of the eligibility test it is determined whether any additional claimed elements other than the abstract amount to significantly more or integrate the claim into a practical application. Claims stating use of “a processor” and “memory” to execute AI functions will likely draw this rejection. If such a claim is coupled with the specification that describes the claimed functions can be performed in a general-purpose computer or processing device, a rejection under this Prong is almost inevitable.

It is noteworthy that in the most recent rejections, “neural networks”, “ML models”, and “AI models/techniques” were considered generic computer components. Therefore, to argue that claims recite a structure specific enough to overcome this element of a Section 101 rejection, a practitioner will likely be required to recite more specific architectural elements.

There is some evidence that some more specific architectural elements (e.g., deep neural networks, recurrent neural networks) are starting to be considered generic. However, practitioners have found success citing even more specific architectural elements such as self-attention layers, layers with specific activation functions, U-nets, etc.

Alternatively, a practitioner can focus on the relationship between technical elements of a claim and improvements to a particular technology, Limitations that have been held to amount to significantly more or integrate a claim into a practical application include limitations that:

- Improve the functioning of a computer
- Include an improvement to the functioning of a computer can include an improvement to the functioning of an ML model. Thus, if an invention improves an ML model, the improvement should be included in the claims.
- Implement the judicial exception with a particular machine
- Include specific elements that perform the relevant steps in the claim. The USPTO has trained examiners to reject claims that appear to implement otherwise ineligible methods

using generic AI or ML models. However, specific architectures should still be considered "particular machines".

- Effect a transformation or reduction of a particular article to a different state
- Transform a particular thing. For example, ML models that transform or generate images, documents, or audio files can be considered to transform a particular article.
- Apply the judicial exception in some other meaningful way beyond generally linking the use of the judicial exception to a particular technological environment

Ultimately, the key to successfully arguing that a claim is applied to a particular technology in a meaningful way is whether the claims include technical elements that link a judicial exception to the technology in question. For example, RNNs and transformer models may utilize a sequence of ordered tokens or embeddings specific to a particular technological environment. Alternatively, citing specific sensor data, or specific outputs unique to a technological environment, along with architectural or procedural steps specific to a particular ML technology can help apply the invention to a particular technological environment.

7. Combinations of Additional Elements and Unconventional Solutions

Of the recently evaluated applications, 97% were evaluated and rejected under step 4 (2B) of the patent eligibility test. However, as with step 3 (2A Prong 2), many of these were initially superficial arguments by the examiner.

The fundamental question of step 4 (2B) of the patent eligibility framework is whether *additional elements* (i.e., not those elements indicated by an Examiner as directed to the abstract idea) amount to an inventive concept that is significantly more than the judicial exception. Some elements of the step 4 (2B) analysis are similar to step 3 (2A, Prong 2). For example, the identification of the "additional elements" carries over from the step 2 analysis, so it is critical to explicitly recite such additional elements when responding to a rejection.

However, there are two key differences that can be important in the context of AI/ML claims. First, step 4 (2B) takes into account whether an element is routine, well-understood, or conventional. Second, step 4 (2B) emphasizes consideration of the claim as a whole, including *combinations* of elements (including elements that individually could fall under a judicial exception).

The question of whether an element of routine or well-understood has been used to exclude additional elements from consideration in terms of integrating a claim into a practical application. However, it is important to incorporate unconventional elements into a claim even if they might fall under a judicial exception (e.g., unconventional mathematical techniques). Finally, an examiner is required to provide evidence if they claim that an element is conventional. See, PTO Memorandum, Changes in Examination Procedure Pertaining to Subject Matter Eligibility, Recent Subject Matter Eligibility Decision (Berkheimer v. HP, Inc.) (Apr. 19, 2018) at 3-4.

Unconventional elements can be considered in combination with other elements such that when combined, integrate the claim into a practical application. In other words, when responding to step 3 (2B) rejections, it is important to explicitly point out combinations of elements that amount to an inventive concept, even if some of the elements could be considered judicial exceptions. Furthermore, it is useful to cite passages from the specification that show how the technical solution is unconventional.

For example, if an invention uses an AI architectural element in an unconventional way (i.e., to solve a task that usually involves a different architecture), the combination of the architectural element and the output can be considered an additional element that can be used to argue that the claim includes something more than an abstract idea.

8. Recent Changes Regarding Training Claims

The manner in which the patent eligibility steps are applied has been changing rapidly. For example, as discussed above with reference to drafting patent eligible claims, the USPTO guidance provides a training claim as an example of a claim that is not directed to an abstract idea at step 2 (2A, Prong 1). However, current practice in the USPTO has begun to change and more training claims are being rejected under Section 101.

Specifically, MPEP 2106.04(a)(1), example vii (based on Example 39 of the USPTO's "Subject Matter Eligibility Examples: Abstract Ideas") is directed to an eligible method of training a neural network. The example claim is reproduced below:

A method of training a neural network for facial detection comprising:

collecting a set of digital facial images, applying one or more transformations to the digital images, creating a first training set including the modified set of digital facial images;

training the neural network in a first stage using the first training set;

creating a second training set including digital non-facial images that are incorrectly detected as facial images in the first stage of training; and

training the neural network in a second stage using the second training set.

The USPTO's original guidance explained that the above claim is not directed to a judicial exception under Step 2 because, inter alia, "the claim does not recite any mathematical relationships, formulas, or calculations" and "the claim does not recite a mental process because the steps are not practically performed in the human mind."

However, examiners at the USPTO have begun to reject claims directed to training with a structure similar to the cited example. From the study of 200 recent rejections, 30 applications that include claims directed to training were rejected. 24 out of the 30 applications were rejected under Step 2, because the training processes were claimed at a high level of generality or failed to provide improvement to computer functionality or technology. Therefore, claims directed to training should be drafted with the same considerations given above to be patent eligible as in other AI claims.

9. Written Description and Definiteness Rejections (§112)

As discussed above with reference to drafting the specification, 35 U.S.C. §112(a) includes three distinct requirements: 1) the written description requirement; 2) the enablement requirement; and 3) the best mode requirement. However, "best mode" rejections are very rare post AIA so emphasis should be placed on the first two.

The MPEP provides specific guidance for evaluating the written description requirement for computer-implemented inventions: "When examining computer-implemented functional claims, examiners should determine whether the specification discloses the computer and the algorithm (e.g., the necessary steps and/or flowcharts) that perform the claimed function … An algorithm is defined, for example, as 'a finite sequence of steps for solving a logical or mathematical problem or performing a task.' Microsoft Computer Dictionary (5th ed., 2002)." MPEP 2161.01(I).

These elements of Section 112 should be interpreted from the perspective of a person having ordinary skill in the art (PHOSITA), which may not be clear in a fast-moving field such as in AI. In general, the PHOSITA generally does not have the same level of expertise in the field as the inventors themselves. Therefore, it is important to differentiate between concepts that would be understood by an average person working in the field of AI and someone with the particular expertise of the inventors.

The primary strategy for responding to Section 112(a) rejections of AI claims is to point out portions of the specification that provide a description of the architecture of a model or the training paradigm. See the three levels of description of architecture and training recommended to be included in the specification in section II.C.4 above.

Regarding Section 112(b), an application is required to "particularly point out and distinctly claim" the subject matter which the inventor regards as the invention. In many cases, rejections under Section 112(b) relate to informalities such as mistakes regarding proper antecedent basis in a claim. These mistakes are often easily corrected.

Since AI claims are often expressed with a model performing a function, a common reason for rejections under Section 112(b) is the citation of generic structures which have been interpreted as means-plus-function elements under Section 112(f), with a subsequent determination that the specification does not provide sufficient corresponding structure. Guidance on establishing definiteness in the written description can be seen in section II.C.3.

In particular, it is worth repeating this tip: for every claimed function, provide an identification of the structure performing that function, a description of an algorithm; a description of what structural elements are used to perform the function; and a description of how exceptions are handled.

B. PTAB Decisions

(Ryan Phelan; Nick Transier)

In addition to Office Actions, it is also important to consider how AI is treated by the Patent Trial and appeal Board (PTAB). While there have been a number of PTAB decisions related to AI, there have been no recent decisions designated as precedential or informative. Recently, Examiners are refusing to consider PTAB decisions cited in Office Action responses unless those decisions are designated as precedential or informative, and there are very few such decisions to cite.

Ex parte Hannun, 2018-003323 (April 1, 2019), is the latest decision designated as "informative". In *Hannun*, an AI claim was found patent eligible based on reciting a practical application of human speech-to-text translation using deep learning.

It is also notable that among recent decisions, when the PTAB is confronted with a combination of apparatus, method, and computer program claims, the PTAB has chosen to select the method claim as exemplary. As method claims are more susceptible to a determination of being directed to mental processes in Step 2 (2A Prong 1), a selection of the method claim as exemplary renders all claims more vulnerable to patent ineligibility despite the applicant undergoing the careful practice of presenting apparatus and device claims. A practitioner should argue against this as it allows the adjudicating body to bypass the default requirement of a claim-by-claim analysis.[2]

For example, as shown below in *Ex parte Philip E. Vasey* (Appeal 2022-001109), Rule 41.37(c)(1)(iv) allows for this practice by stipulating that, when an applicant does not provide separate arguments for different patent claims, the PTAB may select a single claim from a group and decide the appeal on the basis of the selected claim alone. Therefore, to avoid this type of blanket interpretation by the PTAB, the burden is on the applicant to argue the patentability of each claim type separately with regard to the matter of judicial exception during prosecution and appeal practice.

When it comes to Step 3 (2A Prong 2), recent PTAB decisions remain consistent with prior decisions such as *Ex parte Wataru Kimizuka* (Appeal 2018-001081), in that technological improvements must relate directly to the functioning of a computer or device to make a finding of integration into a practical application. Specifically, user benefits are not sufficient to make a finding of technological improvements. Rather, the test requires a technological improvement. This is a high standard for many applications and may render technological improvements a more difficult hurdle to meet at the PTAB compared to the Federal Circuit, where the court has not followed the same standard.

[2] See, e.g., *A House Divided: Is the PTAB Ignoring the USPTO's Section 101 Guidance?*, available at https://ipwatchdog.com/2020/04/13/house-divided-ptab-ignoring-usptos-section-101-guidance/id=120589/

The selection of decisions below also illustrates how the PTAB addresses the question of what is well-understood, routine or conventional via assessment of the patent disclosure. For example, *Ex Parte Akira Harada* (Appeal 2022-003628) illustrates the danger of using "black boxes" to describe computer parts. An insufficiency in disclosure with respect to computer components as shown in the drawings will not only pose a problem for the applicant with regard to Section 112, but also subjects the claimed invention to being deemed as well-understood, routine or conventional.

Ex Parte Hannun (2018-003323)

In *Hannun*, the patent at issue related to "systems and methods for improving the transcription of speech into text." The claims included several AI-related elements, including "a set of training samples used to train a trained neural network model" as used to interpret a string of characters for speech translation. Applying the two-part Alice test, the Examiner had rejected the claims finding them patent-ineligible as merely abstract ideas (i.e., mathematical concepts and certain methods of organizing human activity without significantly more.)

The PTAB disagreed. While they generally agreed that the patent specification included mathematical formulas, such mathematical formulas were "not recited in the claims." (original emphasis). Nor did the claims recite "organizing human activity," at least because the claims were directed to a specific implementation comprising technical elements, including AI and computer speech recognition. Finally, and importantly, the PTAB noted the importance of the specification describing how the claimed invention provides an improvement to the technical field of speech recognition, with the PTAB specifically noting that "the Specification describes that using DeepSpeech learning, i.e., a trained neural network, along with a language model 'achieves higher performance than traditional methods on hard speech recognition tasks while also being much simpler.'"

Ex parte Aurélien Coquard and Christopher Bourez (2022-004679)

Regarding step 2 (2A Prong 1), the PTAB held that processing contract documents is considered a part of a legal interaction. Specifically, the PTAB concluded that claim 1 recites a legal interaction, which falls under the category of "methods of organizing human activity". This suggests that using language such as "contract document" in a claim makes it more vulnerable to a subject matter ineligibility rejection.

Regarding step 3 (Step 2A, Prong 2), the PTAB found arguments with respect to "an improvement in computing" unpersuasive because the Specification describes using generic computing devices to process contract documents and the Appellant did not persuasively explain why using generic computing devices to process data makes claim 1 patent eligible. In other words, the claims did not represent an improvement in the computers themselves, but merely used generic computers as tools for performing a method of organizing human activity.

Ex parte Philip E. Vasey (2022-001109)

As mentioned above, the PTAB took a method claim as representative, despite the presence of apparatus claims. The claimed steps of "evaluating the first rule based on first inputted information to generate a partially customized document comprising the compulsory content elements, the first symbol element, the second rule, and the second symbol element" and "subsequently generating a fully customized document from the partially customized document, without reference to the document template, by evaluating the second rule based on second inputted information" can be performed in the human mind.

Under step 2 (2A Prong 1), the PTAB concluded the claims recite the judicial exception of mental processes. Under step 3 (2A Prong 2) the PTAB considered the Appellant's purported improvement. Appellant argued that its claims integrate the judicial exception into a practical application by enabling a customized document to be generated on a remote computer that does not have access to the template documents. However, the PTAB found that solution was not necessarily rooted in computer technology as Appellant argued.

Ex parte Robert Kerr (2023-000284)

Under step 2 (2A Prong 1), the Appellant contended that requiring a process be stored electronically and performed by a processor overcomes the rejection because the claim excludes processing by the human mind. The PTAB found that to the extent some limitations cannot be performed mentally, this does not demonstrate error or demonstrate eligibility when at least some of the limitations recite steps that could be performed mentally. If a claim under its broadest reasonable interpretation covers performance in the mind but for the recitation of generic computer components, then it is still in the mental processes category unless the claim cannot practically be performed in the mind.

The PTAB stated that the Appellant provided insufficient detail of how the additional elements are implemented. Appellant's lack of a detailed disclosure of computer hardware or functional requirements and the lack of details describing a computer-specific implementation of the recited ML model and other functions (such as might have been indicated by inclusion of a detailed flow chart depicting unconventional computer operations and/or routines for performing each of the claimed steps), meant the details were well-understood, routine, and conventional.

Ex Parte MARTIN RUPP (2023-000022)

In analyzing step 2 (2A Prong 1), the PTAB stated the claims merely recite performing certain calculations without detailing any particular hardware circuitry performing mathematical operations. Manipulation of the input data relates to the pre-Internet activity of performing mathematical computations to convert input data into equivalent output performance data. Accordingly, the PTAB found the cited steps are mathematical steps that can be performed with pen and paper by an operator to reduce the amount of data stored in memory as set forth in the claim. Furthermore, the claims, under their broadest reasonable interpretation, recite a mental process for organizing information through mathematical.

Under step 3 (2A Prong 2) the PTAB reasoned that the Specification does not provide additional details about the general-purpose computer that would transform it into a specific computing device for converting input data from one form to another. Further, Appellant's identified improvements were held to relate to the abstract idea itself, not improvements to a technology or computer functionality. Under step 4 (Step 2B), the PTAB held that the Appellant failed to establish how converting input data into equivalent output performance data is distinguished from the conventional processor-implemented calculation of data.

Ex parte Akira Harada (2022-003628)

Under step 2 (2A Prong 1), the PTAB held that the claims were directed to a mathematical formula. Specifically, "a numerical data acquisition step of receiving numerical data derived from . . ." and "a standard error calculation step of calculating a standard error of numerical data at each data acquisition point or the data based on the numerical data" relate to mathematical relationships for processing data by receiving data and calculating the associated standard error.

Under step 3 (2A Prong 2), the PTAB held that the additional element of claim 1 is a detector in a chromatography analysis device. However, the Specification did not refer to any chromatography analysis device, and the closest device to a chromatography analysis device appeared to be the analysis device in the Specification. Because the Specification merely illustrated the analysis device as a generic box and merely referred to—but did not describe in any detail—the device, then the claimed "chromatography analysis device" was well understood, routine, or conventional in the field.

C. Appeals Court Decisions

(Ryan Phelan)

To date, the Federal Circuit has not reviewed many cases involving AI. However, in one recent case, the Federal Circuit found that a machine learning claim element lacked sufficient enablement because both the claim itself and the written description failed to describe "how" the claimed invention implemented this element. *See In re Starrett*, 2023 U.S.P.Q.2d 684 (Fed. Cir., Jun. 8, 2023). While the decision is nonprecedential, the case reveals the Federal Circuit's analysis and treatment of AI-type claims with respect to Section 112.

<u>*In re Starrett*, 2023 U.S.P.Q.2d 684. (Fed. Cir., Jun. 8, 2023)</u>

In *Starrett*, the Federal Circuit considered an appeal from the Patent Trial and Appeal Board (PTAB) regarding U.S. Patent Application 15/299,124 ("the '124 application"). The '124 application claimed an invention for maintaining "data structures representing categories of biological signals in a body such as a 'Nervous System' and a 'Sensory System.'" *Starrett,* 2023 U.S.P.Q.2d 684.

Claim 1 of the '124 patent recited, in part, a "machine learning" element directed to a specific "configuration," specifically: "[b)] configur[ing] to receive, relay, transmit, or distribute one or more signal [*sic*] wherein at least one signal comprising data representative of information about one or more biological body [*sic*] wherein the processing of biological systems data using at least one *machine learning* task intelligibly recovering perceived, experienced, remembered, or imagined imagery, sounds, or feelings as one or more computational, visual, auditory, textual, numeric, symbolic, coordinate, or haptic representation…" *Id.* (citing '141 application, claim 1.) (emphasis added).

During prosecution, the examiner rejected all claims for lacking enablement. The rejection was appealed to the PTAB, which affirmed the examiner. *Id*. The PTAB found that claim 1 was a type of genus claim that "contain[ed] forty-seven 'or' clauses, thereby allowing it to cover over 140 trillion embodiments." *Id.* In addition, while the patent applicant had argued that claim 1 was "fully enabled" by the patent application's "laboriously detailed" specification, the PTAB disagreed, finding such assertions merely conclusory. *Id.* at *2-*3. Finally, the Board noted that the patent applicant's contentions essentially amounted to "argu[ing] that if an apparatus is well-known . . ., then any function that [the inventor] claims for that apparatus is also fully enabled." *Id.* at *3 (citations omitted).

On appeal, the Federal Circuit affirmed, citing the Supreme Court's precedent regarding enablement in the recent *Amgen Inc. v. Sanofi* decision: "If a patent claims an entire class of processes, machines, manufactures, or compositions of matter, the patent's specification must enable a person skilled in the art to make and use the entire class. In other words, the specification must enable the full scope of the invention as defined by its claims. The more one claims, the more one must enable. *Id.* (citing 143 S. Ct. 1243, 1254 (2023)).

The Federal Circuit then applied this enablement principle to the '141 application, finding that "[h]ere, much is claimed, and little is enabled." *Id*. In particular, the Federal Circuit found particularly troubling the '141 patent application's failure to explain *how* the claimed features would operate without undue experimentation: The application's disclosure of a broad and abstract organizational structure used to accomplish the maintenance of augmented telepathic data amounts to little more than a "research assignment" requiring a skilled artisan to undertake undue experimentation to discover what types of devices are encompassed by the claim limitations and *how* they would function. *Id.* (citing *Amgen*, 143 S. Ct. at 1256) (emphasis added).

While the Federal Circuit did not specifically address the "machine learning" element of claim 1, it did find, more generally, that claim 1 was "rife with broad, vague concepts." *Id.* at *5. For this reason, the Federal Circuit invalidated the claim based on a lack of sufficient enablement. *Id.* The Federal Circuit also addressed the applicant's contention that the claimed features were "well-known" and, as a consequence, allegedly "fully enabled." *See id.* (discussing this aspect as part of a *Wands* factor analysis, of which consideration of "well-known" components is a part. *See In re Wands*, 858 F.2d 731, 737 (Fed. Cir. 1988)). However, the Federal Circuit found that whether a feature is "well-known" (or not) is but one factor of the Wands analysis and is not dispositive on its own. *See id.* Again, the Federal emphasized the importance of describing *how* the claim elements function. *Id.* (stating that "the Examiner's discussion of the Wands factors properly faulted the specification for failing to describe how the claim elements function.").

Ultimately, claim 1 (containing the "machine learning" element) was found non-enabled because the applicant had failed to describe how this (and other) aspects of the invention worked, and the applicant could not rely on the knowledge of a person of skill in the art to cure this defect, no matter how "well-known" such prior art elements were. As the Federal Circuit noted, "[a]lthough the knowledge of one skilled in the art is indeed relevant, the novel aspect of an invention must be enabled in the patent." *Id.* (citing *Auto. Techs. Int'l, Inc. v. BMW of N. Am., Inc.*, 501 F.3d 1274, 1283 (Fed. Cir. 2007)).

<u>*Realtime Data v. Array Networks,* 2023 WL 4924814 (Fed. Cir., Aug. 2, 2023).</u>

Enablement could become a more prominent area of focus with respect to computer-implemented inventions if, for example, Section 101 is resolved via legislation. This topic was recently previewed in *Realtime Data v. Array Networks*. 2023 WL 4924814 (Fed. Cir., Aug. 2, 2023).

In *Realtime Data*, the patent at issue generally related to computer-implemented technology (and not AI), where the claims recited methods and systems for digital data compression. The majority opinion, by Judge Reyna, affirmed a district court's decision invalidating the claims of the patents-at-issue as abstract ideas pursuant to Section 101. *Id.* In particular, the Federal Circuit agreed with the district court that the claims were directed to the abstract idea "of manipulating information using compression." *Id.* The court admonished the claims and specification, stating that "[w]e have determined that "the claim itself ... must go beyond stating a functional result" and that "the claim must "identify *'how'* the functional result is achieved by limiting the claim scope to structures specified at some level of concreteness, in the case of a product claim, or to concrete action, in the case of a method claim." *Id.* (citing *Am. Axle & Mfg., Inc. v. Neapco Holdings LLC,* 967 F.3d 1285, 1302 (Fed. Cir. 2020)) (emphasis added).

Because the claims "failed to do this," the majority opinion held the claims to be invalid pursuant to Section 101. *Id.* at *8. In particular the majority found that "none of the claims at issue specifies any particular technique to carry out the compression of data—the particular rules for producing a smaller set of data out of a larger starting set." *Id.* Rather, the claims "all take the availability of compression techniques as a given and address the threshold matter of choosing to use one or more such available techniques." *Id.* The majority further faulted the abstract nature of the claims stating that "even as to making such a selection, the claims are directed to only abstract ideas, calling for unparticularized analysis of data and achievement of general goals." *Id.*

Judge Newman dissented, arguing that the proper lens for determining was not Section 101 but Section 112, in particular, enablement. *See id.* At *12 (Newman, J., dissenting) (stating that "This is properly an enablement case."). Judge Newman did not analyze the claim under Section 112. Rather she advocated that the proper review belonged under Section 112, and not Section 101: "I write separately to note once again that § 101 was never intended to bar categories of invention in this way. This judicial exception to eligibility is an unnecessary and confusing creation of the courts. This case is an example, for the enablement requirement of § 112 is better suited to determining validity of these claims than is the distortion of § 101. I respectfully dissent, and would remand for determination of validity under § 112." *Id.* (Newman, J., dissenting)

She ended her dissent by noting that "[e]ligibility law has been called a 'morass of seemingly conflicting judicial decisions'" (citations omitted) and that "[w]e should not wade further into this morass." *Id*. "This case is another example that conforms with our flawed precedent. I respectfully dissent. I would remand for a determination of validity under § 112 and, if applicable, §§ 102 and 103." *Id.*

In view of the *Starrett* decision, practitioners should endeavor to explain sufficiently in the written description the specific aspects of ML (and other computer-implemented invention features). In particular, practitioners should endeavor to describe *how* a claimed computer-implemented invention (e.g., an AI invention) operates or otherwise works.

Finally, in view of the *Realtime Data* decision, practitioners can get a preview of what may yet come in the event of a legislative change to Section 101 (or a ruling from the Supreme Court bringing about the same), where the new invalidity battleground is not Section 101, but instead Section 112.

D. District Court Decisions

(Ryan Phelan; David Pointer)

The following set of AI-related cases are from district courts of different geographic locations.

Health Discovery Corp. v. Intel Corp., 577 F. Supp. 3d 570 (W. D. Tex. 2021)

In this case, Intel filed a motion to dismiss Health Discovery Corp (HDC)'s complaint on the grounds that the claims are invalid under 35 U.S.C. § 101. HDC's complaint alleges that Intel infringed HDC's patents related to using learning machines, such as Support Vector Machines (SVM), to identify relevant patterns in datasets and to identify a selection of features within the datasets that best enable data classification. The asserted patents were directed to feature ranking, selection, and reduction using SVM to facilitate a Recursive Feature Elimination (RFE) process on a large dataset.

Under step 1 of the eligibility test, the Court reasoned to follow the guidance provided from prior cases that analyzed patents with similar subject matter. After selecting two similar cases, the District Court asserted that the specification merely describes improving a mathematical analysis used by conventional systems. The specification explained how conventional systems reduce a feature size in data sets by ranking and eliminating features according to correlation coefficients. Similarly, the asserted patents involve ranking and eliminating features using SVM-RFE, which the court characterized as a purportedly novel, but mathematical technique. As such, the District Court reasoned that the claims are directed to an abstract mathematical concept of SVM-RFE.

Under step 2, the District Court stated that HDC's complaint failed to sufficiently allege an inventive concept. The Court stated that improving data quality was an unpersuasive argument, and the Court also considered that the claims were not limited to a particular field of invention. Accordingly, the Court granted the motion to dismiss.

Pavemetrics Sys. v. Tetra Tech, 2021 U.S. Dist. LEXIS 117651 (C.D. Cal. 2021)

In this case, Tetra Tech moved for a preliminary injunction to enjoin Pavemetrics from importing, using, and selling their "Laser Rail Inspection System" (LRAIL) products on the grounds of patent infringement. Tetra Tech's '293 patent generally relates to a three-dimensional railway track inspection and assessment system that collects and processes data during and/or after a high-speed assessment of a railway track. In 2018, Pavemetrics began developing AI-based deep learning algorithms using convolutional neural networks to identify defects in railway tracks in their LRAIL products.

Ultimately, the District Court concluded that Tetra Tech could not demonstrate a likelihood of success for a preliminary injunction because of substantial questions regarding infringement and invalidity. With regard infringement, the Court noted that substantial questions remain because Tetra Tech relies the "gradient neighborhood" limitation being reflected on a prior design of Pavemetrics' LRAIL product from two years prior to the issuance of the asserted patent. At that time, Pavemetrics significantly changed how LRAIL processed data when it switched to detecting missing features using deep neural networks. The Court agreed that Tetra Tech had not provided sufficient evidence to indicate Pavemetrics' "deep neural network" design within the current product meets the "moving gradient neighborhood like a sliding window over the 3D elevation data suing the processor," as recited in claim 1.

With regard to invalidity, the Court determined there are substantial questions because Pavemetrics alleged that one of its prior designs anticipates claim 1 of Tetra Tech's '293 Patent. For these reasons, the Court denied the motion for preliminary injunction.

IBM v. Zillow Grp., Inc., 2022 U.S. Dist. Lexis 41831 (W. D. Wash. 2022)

In this case, Zillow filed a motion to dismiss IBM's claims of patent infringement on the grounds of ineligible subject matter under 35 U.S.C. § 101. The '676 patent generally relates to a method of annotating response sets via an adaptive algorithm and the supplied annotations are used for a visualization system that presents resource response results. In the specification, the '676 patent states that the "*s*ystem discovers contexts and context attributes among users which can be used predictively," by using "a highly specialized and optimized combination of supervised & unsupervised logic along with" automated entry of learned results. Col. 19, Lines 39-44.

Although the specification of the '676 Patent appears to describe improvements relating to computer and/or search engine functionality, the Court focused on the subject matter recited in the claims for step 1. The Court found that the claimed subject matter was directed to an abstract idea because the claim language was result-oriented and recited a process that could be performed with a pen and a paper. The claims failed to recite the inner functionality of the invention.

Under step 2, the Court concluded that the claims failed to satisfy the second Prong of the test for a couple of reasons. First, the specification failed to provide a description of the alleged inventive concepts offered by IBM. Second, the Court determined that the claims in the '676 patent do not provide a specific, discrete implementation of the abstract ideas for applying an ordering and annotation function, mapping the user context vector with the resource response set, or generating an annotated response set. Accordingly, the Court concluded the '676 Patent is invalid under 101 and the related patent infringement count is dismissed for failing to state a claim upon which relief can be granted.

E. Detectability

(Michael Carey; Sumon Dasgupta; Thomas Burton)

In addition to issues that arise in court, AI related patents pose unique challenges related to detecting infringement. Generally, a patent is more valuable if infringement is easily detectable. However, many AI patents relate to software, and in some cases, to a configuration of software parameters that is not easily discernible or interpretable. This can create challenges in detecting infringement.

Detectable features are observable from the use, appearance, or construction of a product or process embodying an AI invention. In some cases, infringement may be detected through using a publicly available product, reviewing publicly available documents (e.g., academic publications, market materials, specifications, etc.), through reverse-engineering, or through application to technological standards.

When drafting an AI-related patent application, significant attention should be paid to identifying and claiming any features that are observable at inference (i.e., during use). For example, an AI related invention may include prompting a user for key inputs, and providing an output based on those inputs. The relation between the inputs and outputs may be observable even if the inner workings of the system are not observable. For example, AI features that are part of an autonomous device (e.g., vehicle, robot, game, etc.) or a device that involves learning (e.g., an intelligent thermostat) may be detectable in the operation of the device.

AI features may also be detectable based on publicly available documents. For example, because AI is a rapidly developing field that involves many academic contributions, many inventors publish their work in academic journals, conferences, and archives. In another example, companies may disclose features of an AI model in other ways such as user guides and advertising materials.

AI features may also be detectable via reverse-engineering a product. This is made more challenging by the prevalence of Software-as-a-Service (SaaS) AI products. That is, it may be challenging to reverse-engineer a product that is not publicly available. However, in some cases, AI products use publicly available data or publicly defined data structures that expose the inputs or the outputs to inspection.

AI features can be detected at several different stages including data collection, training, testing, and operation. In some cases, different features may be exposed by attempts at distributing the process via federated learning, distributed data processing, distributed training, edge ML, and other emerging technologies that alter the workload balance between edge devices and centralized (e.g., datacenter, cloud) compute resources. In some cases, communication between the edge devices and the centralized cloud compute resources is much easier to capture and investigate via technical methods than the centralized compute resources themselves. Therefore, when claiming ML inventions, pay particular attention to the dataflow between centralized compute resources and edge devices.

Finally, some features of AI patents may be detected by virtue of relating to technological or communication standards. In many cases, standards are needed for the adoption and compatibility of new technologies. Standard bodies, such as American National Standards Institute (ANSI), Institute of Electrical and Electronics Engineers (IEEE), or the Third Generation Partnership Project (3GPP) have different patent policies. For example, these patent policies may request patent holders to disclose information regarding patents or patent applications that are relevant to the standard. Some may require that any license issuing from these patents be granted under fair, reasonable and non-discriminatory terms (FRAND license). Some standard bodies may even request that a license must be royalty free. Thus, it may be worth contributing patented technologies to a standards body standard.

IV. Other Considerations

A. Ethics for AI Inventions

(Thomas Burton; Sumon Dasgupta)

With the emergence of AI-enabled products and tools such as ChatGPT and Bard, the US government and other countries are introducing guidelines and legislation to address other ethical concerns related to AI. For example, in October 2022, the US government published the "Blueprint For An AI Bill of Rights: Making Automated Systems Work for the American People "(hereafter, the "Blueprint") to provide companies guidance to address potentially inherent ethical risks of AI enabled systems. The Blueprint is not law. However, it can be a signal of potential government action (e.g., laws, rules, regulations).

As a corollary, patent practitioners now have other ethical issues to consider when preparing a patent application directed to an AI related invention. This section provides patent drafting practice tips for avoiding such ethical issues. For example, patents should not employ discriminatory language to avoid ethical issues, such as using "binary" pronouns like he or she that may not be viewed as being inclusive of all. Such language can have negative public impact.

AI principles under the Blueprint and other models, such as the recent G7 code of Conduct include:

1. Safe and Effective

An "automated system" that uses AI to "determine outcomes, make or aid decisions, inform policy implementation, collect data or observations, or otherwise interact with individuals and/or communities" (including "people connected by affinity, identity, or shared traits") should be safe and effective based on their intended use. In particular, such automated systems should be "developed with consultation from diverse communities" as well as "domain experts" and undergo pre-deployment testing, risk identification and mitigation, and ongoing monitoring" to support that they are safe and effective.

To illustrate why this principle is important, the Blueprint highlights real examples of AI enabled automated systems that violated this principle. For example, a proprietary model system was developed to predict the likelihood of sepsis in hospitalized patients and was implemented at hundreds of hospitals around the country. An independent study showed that the model predictions underperformed relative to the designer's claims while also causing "alert fatigue" in certain patients by falsely alerting likelihood of sepsis. Thus, this AI enabled prediction model was ineffective and caused harm to certain patients that were falsely identified.

Practice Tip: In preparing a patent application, patent practitioners should always avoid making any definitive statements in the patent application that the underlying invention "is safe", "error free" or any other statement that could be used to support a product liability claim against your client (the patent owner) or lead to public scrutiny of the effectiveness of your client's invention. This is especially true of an AI enabled invention where the boundaries of the AI invention are not yet fully tested, such as the case of the above AI enabled sepsis prediction model system. To avoid such product liability issues, the patent practitioner should counsel his client not to make any definitive claims of the AI system's effectiveness.

Instead, provide facts related to the data sets used to train the AI system and use safety terms such as "inhibits safety issues" (such as with outbreaks of sepsis in a hospital in which the AI system is deployed) and/or "enhances the likelihood of an accurate prediction." Such patent preparation practices are aligned with this Blueprint's principle for Safe and Effective Systems.

2. Algorithmic Discrimination Protections.

Users of "automated systems" should be protected against algorithmic discrimination. Such systems should be designed to include "proactive equity assessments" and "use of representative data" to protect against bias and ensure equitable treatment.

To illustrate the necessity of this principle, the Blueprint provided the following real example. The U.S. Transportation Security Administration (TSA) deployed certain body scanners at airport checkpoints that required the operator to select a "male" or "female" scanning setting based on the passenger's sex. But this manual setting selection introduces bias into the body scanner system based on the operator's perception of the passenger's gender identity. These scanners are more likely to flag transgender travelers as requiring extra screening done by a person. TSA has recently announced plans to implement a gender-neutral algorithm while simultaneously enhancing the security effectiveness capabilities of the existing technology.

Practice Tip: In preparing a patent application directed to an AI enabled system (such as a body scanner that purports to implement a "gender-neutral algorithm", the patent practitioner should counsel the inventor(s) on the need to ensure the underlying algorithm is not discriminatory. The practitioner should ask if the system assesses data associated with people (as opposed to data associated with a substance, an object, an animal, a car, a building, etc.) and, if so, ask how the system takes "proactive" measures to "enhance the effectiveness" of the system to operate equitably so that the patent application can disclose those measures (which may aid in identifying potential elements to claim) and inhibit third party claims of algorithmic discrimination. Moreover, the practitioner should question the inventor(s) on the data sets used to train the applicable AI enabled system so that such data sets could be disclosed to further aid in preventing third party claims of bias.

3. Data Privacy.

People should be able to control how their data is used, and they should not be subjected to abusive data practices. For example, the Blueprint describes ensuring that data collection conforms to reasonable expectations and that only data strictly necessary for the specific context is collected. Designers, developers, and deployers of automated systems should notify and seek permission with regard to data usage, and respect decisions regarding collection, use, access, transfer, and deletion of private data in appropriate ways and to the greatest extent possible. Systems should provide clarity about user choices, and not obfuscate user choice or burden users with defaults that are privacy invasive. For example, consent requests should be brief and understandable. Enhanced protections and restrictions for data and inferences related to sensitive domains (e.g., health, work, education, criminal justice, and finance), and for data pertaining to youth should only be used for necessary functions.

Practice Tip: Be cautious about describing the types of data that are used in some types of AI applications. It is almost second nature to the patent practitioner to broaden the scope of an application, but caution should be exercised with respect to data, and particularly in sensitive domains. Unwanted implications of unauthorized and/or unexpected data usage should be avoided in patent applications. Furthermore, bear in mind that while the Blueprint is non-binding, other information-related restrictions are in place. For example, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) is a federal law that describes standards to protect sensitive patient health information from being disclosed without the patient's consent or knowledge.

4. Notice and Explanation.

Users should know why and how an AI system made its determination. People should have an understanding that an automated system is being used and understand how and why the automated system contributes to outcomes that impact them. For example, it may be necessary to provide generally accessible plain language explanations (including clear descriptions of the overall system functioning and the role automation plays), notice that such systems are in use, who is responsible for the system, and explanations of the outcomes. In other words, people should timely receive notice as to how and why an outcome that impacts them was determined by an automated system.

An example provided in the Blueprint describes that "a lawyer representing an older client with disabilities had been cut off from Medicaid-funded home health-care assistance couldn't determine why, especially since the decision went against historical access practices. In a court hearing, the lawyer learned from a witness that the state in which the older client lived had recently adopted a new algorithm to determine eligibility. The lack of a timely explanation made it harder to understand and contest the decision."

Practice Tip: One of the many applications of AI can include the determination of a decision that impacts a user. In such cases, it may be prudent to describe in the patent application, that the user is notified in some fashion of the decision of the AI, that the AI is being used, and so forth. It may also be prudent to inform your client of the above guidance to mitigate potential issues in the future.

5. Human Alternatives, Consideration, and Fallback.

Finally, the Blueprint sets forth the principle that an "automated system" should provide people a choice to opt out of AI decision-making and have "fall back" access to a human if the system has an error, fails, or they want to challenge a decision made by the system. The Blueprint highlights that "automated systems with an intended use within sensitive domains, including, but not limited to, criminal justice, employment, education, and health, should additionally be tailored to the purpose" in addition to incorporating "human consideration for adverse or high-risk decisions."

To illustrate the problems this principle seeks to address and protect against in AI systems, the Blueprint identified several examples, including the following. A fraud detection system for unemployment insurance distribution incorrectly flagged entries as fraudulent, leading to people with slight discrepancies or complexities in their files having their wages withheld and tax returns seized without any chance to explain themselves or receive a review by a person. A patient was wrongly denied access to pain medication when the hospital's software confused her medication history with that of her dog's. Even after she tracked down an explanation for the problem, doctors were afraid to override the system, and she was forced to go without pain relief due to the system's error. A large corporation automated performance evaluation and other HR functions, leading to workers being fired by an automated system without the possibility of human review, appeal or other form of recourse.

Practice Tip: Preparing a patent application directed to an AI system trained for a particular use case (e.g., unemployment insurance fraud detection) typically would not require the applicant to disclose or include an "opt out" feature for an individual to contact a human to question an output of the AI system. But where the output of an AI system is based on accessing data that is personal to an individual, a patent practitioner should be alert to the harm that the AI system may cause the individual when the output is based on data that is not error-free but includes "slight discrepancies or complexities" that could lead the AI system to output a false positive output.

Accordingly, where an invention is directed to an AI system that processes personal data (e.g., AI system for insurance fraud detection, determining patient medication prescription, or automated employee performance evaluation), a patent practitioner should advise his client that such an AI system may be viewed by the public and government agencies as providing “high-risk” decisions for persons. To mitigate any negative sentiment from the public based on the patent application being directed to such a “high-risk” AI system, the patent practitioner may recommend that the applicant/inventor add potential further features or embodiments to the AI System for an individual to address any “high-risk” decision relating to the individual. Although likely not patentable subject matter, the patent practitioner may recommend that the applicant also include an embodiment where an individual may alternatively contact the owner or operator of the AI System to dispute any output or decision of AI system. Such an embodiment may be positively viewed by the public as in line with the objective of the Blueprint’s “Human Alternative” or “Fallback” principle.

6. G7 Code of Conduct

In addition to the Blueprint, AI ethics are being considered by a number of organizations around the world. For example, the Group of 7 (G7) countries recently adopted a voluntary AI Code of Conduct. The G7 Code is largely directed at companies that operate in G7 countries.

This 11-point ethical code seeks to ensure that AI is both safe and dependable on a global scale. It is intended to furnish voluntary guidance for organizations engaged in the development of cutting-edge AI systems, encompassing advanced foundation models and generative AI systems. The primary objective of this code is to harness the advantages of AI while effectively addressing the associated risks and challenges. It places a notable emphasis on urging companies to implement measures that identify, assess, and mitigate risks throughout the entire lifecycle of AI. Furthermore, it compels companies to confront and rectify incidents and patterns of misuse that may arise after the deployment of AI products in the market.

Additionally, the code encourages companies to disseminate public reports detailing the capabilities and limitations of AI systems, as well as their usage and potential misuse. Furthermore, it underscores the importance of investing in robust security controls to ensure the responsible development and deployment of AI technologies.

7. Summary

Although these “ethical risk mitigation” principles set forth in the Blueprint are not mandatory for companies or individual inventors to follow when implementing an AI system, patent practitioners would be wise to address these principles with their company or individual clients while preparing a patent application directed to AI system that may be viewed as having ethical concerns. By following the above practice tips, a patent practitioner can prepare a patent application for an AI system that aid the client to address any potentially inherent ethical risks in the AI system that the five principles seek to protect the public against.

B. Considerations for the EPO

(Alexander Korenberg; Nicholas Malden; Edoardo Mirabella)

The EPO considers AI and ML algorithms to be mathematical methods per se. According to the European Patent Convention (EPC), mathematical methods claimed per se are included in the non-exhaustive exclusion list of Art. 52 EPC and, hence, are excluded from patentability.

1. EPO Standards of Evaluating AI Inventions

The EPO assesses inventions involving items of the exclusion list and mathematical methods such as ML and AI algorithms, by using the so-called two-hurdle approach. The first hurdle aims to establish whether the claimed subject matter is excluded from patentability or not, i.e., whether it is directed to a mathematical method *per se*. This is relatively easy to overcome, because, in practice, the presence of any technical means (e.g., a computer, a network, or the internet) is sufficient to overcome this first hurdle.

The second hurdle constitutes the substantive examination and, among other goals, aims to establish whether the patent application provides an enabling disclosure of the invention or not and whether the claim set is clear, unitary, novel, and inventive. According to EPO practice, novelty is assessed by taking into consideration both technical and non-technical features. Non-technical features are, for instance, mathematical methods and AI and ML algorithms per se.

The assessment of inventive step is carried out via the so-called COMVIK approach, which prescribes how to apply the problem-solution approach in the presence of technical and non-technical features. According to the COMVIK approach, inventive step is assessed by assessing the differences from the closest prior art and only those that contribute to technical character are considered for inventive step. In practice this means that non-technical features are often entirely ignored, although may be considered if, in the context of the invention, they contribute to the technical character of the invention by producing a technical effect.

A mathematical method may contribute to the technical character of an invention, i.e., may serve a technical purpose, by its application to a field of technology and/or by being adapted to a specific technical implementation. Where an ML algorithm serves a technical purpose, the steps of generating the training set and training the algorithm may also contribute to the technical character of the invention, if they support achieving that technical purpose.

Application to a field of technology:

The application to a field of technology shall be specific, e.g., a generic purpose such as "controlling a technical system" is not sufficient. Moreover, the claim shall be functionally limited to the technical application, i.e., shall not encompass further non-technical applications.

Examples of technical contributions of a mathematical method are (cf. Guidelines for Examination in the European Patent Office, G-II 3.3 and G-II 3.3.1): controlling a specific technical system or process, e.g. an X-ray apparatus or a steel cooling process; digital audio, image or video enhancement or analysis, e.g., de-noising, feature detection in a digital image, or estimating the quality of a transmitted digital audio signal; classifying digital images, videos, audio or speech signals based on low-level features, e.g. edges or pixel attributes of images; and using a neural network in a heart monitoring apparatus for the purpose of identifying irregular heartbeats.

Examples of non-technical purposes include classifying text documents by using their textual content; and classifying abstract data records (or even telecommunication network data records) without any indication of a technical use of the resulting classification.

AI and ML algorithms may contribute to the technical character of the invention when they are adapted for a specific implementation, i.e., their design is motivated by technical considerations of the internal functioning of the computer system or network, e.g., producing the technical effect of an efficient hardware implementation of the algorithm. For example, the implementation of ML techniques in a computing platform comprising a GPU and a CPU, so that training steps are executed by the GPU and preparatory steps by the CPU, contribute to the technical character of the invention.

2. EPO Case Law

Case law on ML inventions at the EPO is limited, but a few principles can be gleaned. ML per se is part of the trend of technology and hence obvious (T 2246/18, T 0161/18). An invention therefore needs to specify more than just using ML for solving a problem. And the problem must be technical; solving a non-technical problem with ML per se does not contribute to an inventive step (T0872/19) unless the ML has been designed with the functioning of the computer in mind (T1358/09), an example of which would be enabling or improving parallel processing (T2330/13, T2910/19). Any improvement to be used as the basis of an inventive step must be present for all embodiments covered by the claim scope (T0702/20), must be credible from the content of the application (T0702/20) and must be more than merely encoding / automating human expert knowledge (T1635/19). While ML case law is limited, the EPO treats ML in the same way as any other mathematical method and so the case law concerned with mathematical methods more widely is also relevant.

3. Claim drafting best practice for AI-related European Patent Applications

The above outline of EPO practice in assessing inventions relating to AI and ML demonstrates that patent eligibility is typically trivial to achieve in a claim filed at the EPO. The challenge lies in the assessment of inventive step, in which claim features will be assessed as "technical" or "non-technical". Clearly it is vital that any features which are key to the invention are categorized as "technical", or they will essentially not be considered in the assessment of inventive step, and they will not be available for contributing to the "technical effect", the vital and central component of a successful demonstration of non-obviousness before the EPO.

For AI-related inventions, a notable aspect for patenting is that machine-learning technologies invariably consist of two phases: a training phase (in which a system "learns") and an inference phase (in which that "learning" is put into effect). Separate claims can therefore be crafted to cover each phase. Moreover, the technical features in each phase should form the core structure of the claim in each case, so that the technical purpose is evident in both, and the technical effect of the claimed features can be demonstrated. In both cases, even if it not desirable to include such elements in the claims themselves, sufficient specific technical detail of, e.g., training data, training parameters, input data for inference, and inference parameters should be provided in the description, to sufficiently enable the claimed invention(s). Disclosed technical features that are not claimed but could potentially be used as fallback positions to help patentability should be disclosed in claim ready language in a suitably generalized context, for example in the summary section, given the restrictive EPO rules on added matter.

In the claims themselves, avoidance of non-technical terms (where possible) can help to avoid some features being disregarded out of hand as not being technical, and this includes terminology which may be too easily associated with mathematical methods or business methods in the eyes of an EPO examiner. Furthermore, the claim features, including those relating to the AI or ML, should be readily viewed as part of a causal chain leading to the technical effect being asserted. This aim can be supported by ensuring that sufficient interconnections or interactions between claim features are present, such that the causal chain is easily highlighted.

C. Considerations for Japan

(Sumon Dasgupta; Christina Huang)

While Japan Patent Office (JPO) did not revise the Patent Examination Guideline specifically for AI, JPO has published several AI invention examples that raised awareness on the enablement and disclosure requirements in 2019. According to Patent Act Article 36(4)(i), "[t]he statement of the detailed explanation of the invention shall be clear and sufficient as to enable any person ordinarily skilled in the art to which the invention pertains to work the invention." Patent Act Article 36(6)(i) provides that a claimed invention shall be disclosed in the description.

The JPO provided Comments to Patenting Artificial Intelligence Inventions in September 2019, "[i]n order for the AI-applied invention to satisfy its enablement requirement, the description that the invention can achieve a certain degree of accuracy in estimation processing should be in Specification, that is, the capacity to create a learned model with a certain degree of accuracy in estimation processing is required for the description in Specification. . . Therefore, if there is any relationship between input and output data in the training data used to create the learned model, we consider that the AI algorithm can create a learned model that performs accurate estimation processing based on the above-mentioned input and output data relationship."

As an example, in Case Example 49, the description discloses that (i) a feature value representing a face shape of a person is a face-outline angle, which is defined between a tangent line to a jaw and a tangent line to a cheek, and (ii) there is a statistically significant correlation between a cosine of a face-outline angle and BMI (defined as a body weight divided by the square of a body height) of a person. However, the description only discloses that any feature value other than a face-outline angle representing a face shape may be obtained from a face image and used. It does not disclose a correlation or the like between (i) a feature value other than a face-outline angle representing a face shape and (ii) a body height, weight, and the like of a person and BMI based on these. As such, the application fails to meet the support requirement or the enablement requirement in Example 49.

D. Considerations for China

(Sumon Dasgupta; Christina Huang)

In February 2020, China National Intellectual Property Administration (CNIPA) enacted a new revision of the Patent Examination Guideline ("Guideline"). More specifically, a new section, Section 9.6, is added, which includes relevant rules on the examination of invention patent applications that include algorithmic features or business rules and method features, such as AI related patents.

First, a claim as a whole should be considered. For patent applications involving AI and big data, the claims are often related to rules and methods of intellectual activities such as algorithms, business rules, and methods. This revision clarifies that during the examination, technical features, algorithms, business rules, and process steps should not be evaluated separately, but all the features recited in the claims should be considered as a whole. The rationale is that when evaluating claim limitations separately, the substantive contribution of the invention cannot be evaluated objectively.

Second, some patentability examples have been provided. Specifically, if the claims include limitations different from an existing technology from the perspective of algorithm adjustments, the claims possess inventive steps. In Example 2, since the claim recites a solution to solve a technical problem of a convolutional neural network (CNN) capable of processing only images of fixed sizes and reaches the technical effect of the train CNN that can process images of various sizes, the claimed invention is patentable subject matter.

Third, the written description requirements have been updated. According to the Guideline, the specification should clearly provide technical features supporting and interacting with each other in their respective functionality to solve a technical problem. For example, for a patent application including an algorithm, the specification should combine the algorithm with the practical application, such as, including at least one input parameter and at least one output related to the specific technology area. As another example, for a patent application including business rules and method features, the specification should include detailed explanations on the entire process of how to solve the technical problem. Additionally, if the improvements to user experiences are objective (not subjective), the specification may provide how the user experience improvements are implemented using technical features, and how features are interacting with each other to generate such effects.

E. Considerations for Korea

(Sumon Dasgupta; Christina Huang)

Korea appears to be less stringent with respect to patent eligibility. For example, Korea identifies patentable subject matter based on novelty and inventiveness (if technical ideas are embodied in a computer). It is worthwhile to note that a technical idea embodied within a "general purpose computer" may be sufficient to satisfy patent eligibility if software and hardware operate together. Some inventions may fail to meet patent eligibility (e.g., economic laws, mathematical formula, mental activity, etc.) that do not satisfy the above (e.g., lack of software and hardware together, lack of technical idea, etc.).

Moreover, the enablement requirement in Korea may require that the application includes a description of a relationship (e.g., a correlation) between input data and output data from a trained model to implement AI-related inventions. For example, a correlation may be met when learning data is described, correlations between learning data and solution to a technical problem, a description of a learning model/method based on input data, how a trained model for solving a technical problem is generated based on input data and methods.

F. Considerations for the UK

(Mike Jennings)

For most inventions that involve specific technical applications of AI/ML (and other computing technology such as computer simulation), the UK courts and UKIPO practice are now largely similar to the EPO, applying a relatively narrow interpretation of the UK's exclusions from patentability for computer programs, mathematical methods and mental acts. However, the UK test for patentability remains somewhat different from the EPO test, so outcomes can differ for some types of AI.

Specifically, section 1(2) of the UK Patents Act corresponds to Article 52(2) of the European Patent Convention and lists certain subject matter that is not to be considered an invention:

"*(a) a discovery, scientific theory or mathematical method;*

(b) a literary, dramatic, musical or artistic work or any other aesthetic creation whatsoever;

(c) a scheme, rule or method for performing a mental act, playing a game or doing business, or a program for a computer;

(d) the presentation of information;"

However, these exclusions only apply "*to the extent that a patent or application for a patent relates to that thing as such*", and this wording has been interpreted relatively narrowly since 2012.

As at the EPO, claims that are functionally limited to a specific technical application of AI/ML are treated positively, but claims to abstract mathematical methods and applications of AI/ML technology for non-technical business methods will be refused. Therefore, it continues to be important to include within an AI-related patent application a description of all of the inputs, training, outputs, and technical applications – i.e. all interactions with physical systems outside the computer – as these may be important for patentability. It is also necessary to include a detailed description of the system/software architecture for sufficiency reasons, as in all jurisdictions.

In its September 2022 guidance on AI patenting, the UKIPO endorsed the EPO Board of Appeal's insufficiency decision T161/18, which criticized a lack of detailed disclosure related to the training of an ANN in an application that relied on that training as its point of novelty.

On 21 November 2023, in the UK's first substantive patentability judgement for an application of AI (in *Emotional Perception AI v Comptroller-General of Patents [2023] EWHC 2948 (Ch)*), the UK High Court overturned an earlier UKIPO refusal of a patent application for a system and method for providing media/text file recommendations using an Artificial Neural Network (ANN) that was trained to perceive semantic similarity between pairs of files. The Judge found that the claimed system made a technical contribution to the art, at least because a selected file having certain attributes such as similarity characteristics is output to an end user by a system that had established the identification system and implemented it. The recommended media/text file was identified as being semantically similar by the application of technical criteria, which the system had worked out for itself. This was considered to be a technical effect, which contributed to patentability.

It should be noted that a claim to a system outputting a file recommendation based on semantic similarity was considered to be more than a mathematical method or computer program, regardless of the system and its training involving a computer program implementing an algorithm. It was also considered that the ANN that resulted from the training process is capable of being an external technical effect, which prevents the computer program exclusion from applying.

The UKIPO quickly confirmed that it is changing its examination practice to reflect the court's judgement, stating that: "*Following the Emotional Perception judgment, patent examiners* *should not object to inventions involving ANNs under the "program for a computer" exclusion*". Accordingly, the UKIPO Hearing Officer had already accepted that, even if an ANN and a method of training an ANN is based on a mathematical algorithm, "*its specific application here as part of a file recommendation engine is, in my opinion, enough to dispense with the mathematical method as such objection*".

The UKIPO is currently reviewing its examination guidelines and updates are expected. However, even before the positive court decision on 21 November 2023, the UKIPO's September 2022 guidance on patentability of AI was already quite encouraging for applicants. Also, the UKIPO and EPO recently worked closely on a recommendation for member states of the European Patent Organisation with the aim to achieve greater harmonization of patent practice for computer-implemented inventions (within the limits of what is possible while taking account of national case law). The result is that the majority of patent applications for technical applications of AI/ML will be accepted as patentable in the UK (as well as in the EPO) if the claimed invention is new, inventive, and sufficiently described.

G. AI Inventorship

(Sumon Dasgupta; David Kincaid; John Pienkos; Jennifer Lacroix)

According to United States patent law, the threshold question in determining inventorship is who conceived the invention (MPEP 2109). Conception can mean "the complete performance of the mental part of the inventive act" and it is "the formation in the mind of the inventor of a definite and permanent idea of the complete and operative invention as it is thereafter to be applied in practice." (MPEP 2138.04).

AI systems pose unique challenges when determining inventorship. A fundamental inquiry is whether the AI systems are used as a tool to help natural persons conceive of an invention, or whether the AI systems conceive of the invention. If the AI system is simply used as a tool, then the inventor would be the natural person. That is, actions by a natural person(s) that qualify as a contribution to the conception of an invention are unaffected by use of an AI system. Such actions have been broadly interpreted, and in most instances (nearly all to date), a natural person can be fairly called an inventor. As an example, actions including designing the architecture of the AI system, choosing input data to provide to the AI system, or developing an algorithm to enable the AI system to process data, may be adequate to qualify a natural person as a contributor to conception.

Still, there are circumstances (and possibly an increasing number of circumstances) in which a human has only minimal interactions with an AI system and questions can arise as to whether any human at all should properly be considered an inventor and, if not, whether in some sense the AI should be an inventor.

There are different hypothetical scenarios in which AI might be considered to constitute an inventor with respect to an invention when human involvement in the inventive process is or seems minimal. Consider generative AI models. Generative AI models can generate new and original content (e.g., computer code, designs, architecture, art, drugs, etc.). In one particular example, it is possible that an AI-enabled drug development process results in a new and novel drug being identified. Since AI is extensively used in the drug development process, there is a possibility that human intervention decreases to a point where the AI could be the only meaningful contributor to the origination of the invention. As a different example, consider that some generative AI models can generate code based on a request received from a natural person.

It is potentially possible that the request can frame a problem (e.g., a code to solve an existing problem), but does not provide any possible clues or suggestions as to how to solve the problem. In such an instance, could the natural person be considered an inventor if the resulting code is novel and patentable? Thus, at what point does the AI model cease to be merely a tool, and possibly rise to the level of an inventor?

Notwithstanding such considerations, as will be explained below, under most current patent laws, AI cannot be listed as an inventor for a patent application. It remains to be seen whether, as AI continues to evolve and expand, these laws remain the same or change over time in regard to the fundamental question of who or what can be considered an inventor.

1. **AI Inventorship in the US**

Within the past five years, the USPTO has considered and confronted many issues regarding the use of AI in innovation. For example, the USPTO has published two notices in the Federal Register seeking comments regarding the use of AI and inventorship. The first was in August of 2019, and requested comments on patenting AI inventions ("Request for Comments on Patenting Artificial Intelligence Inventions," 84 FR 44889 (August 27, 2019)). The result was a report titled "Public Views on Artificial Intelligence and Intellectual Property Policy," published by the USPTO in October of 2020. The report indicated that the comments received included very mixed views regarding whether AI is merely a tool that cannot invent without human intervention, or whether AI could contribute to the creation of inventions, jointly with humans or even on its own.

The most recent notice was published in February of 2023, and the time period for comments closed in May of 2023 ("Request for Comments Regarding Artificial Intelligence and Inventorship, 88 FR 9492 (February 14, 2023)). In the 2023 notice, the USPTO recognized that it "plays an important role in incentivizing and protecting innovation, including innovation enabled by artificial intelligence (AI), to ensure continued U.S. leadership in AI and other emerging technologies (ET)" and sought "stakeholder input on the current state of AI technologies and inventorship issues that may arise in view of the advancement of such technologies, especially as AI plays a greater role in the innovation process" *Id.* In May of 2023, IPO leadership responded to the USPTO's February 2023 request for comments.[3] In its response, the IPO stated that AI is used by inventors in the invention creation process as a tool; however, the IPO does not believe that AI is currently involved in the conception of inventions. The IPO further maintained, in view of Federal Circuit's decision in *Thaler v. Vidal* (see below) and the Supreme Court's subsequent denial of certiorari, that the Patent Act only requires the listing of natural persons as inventors, thus precluding listing of AI as an inventor/joint inventor. The USPTO has not yet released its report based on the February 2023 request for comments.

The USPTO also held its inaugural AI/ET Partnership meeting in June of 2022, which discussed issues such as whether AI could actually "conceive" of inventions. Despite the USPTO's curiosity and willingness to consider the issues regarding AI, the law to date in the United States remains clear that AI cannot be named as an inventor. In July 2019, Thaler filed U.S. Patent Application Nos. 16/524,350 and 16/524,532, naming an AI system (Device for Autonomous Bootstrapping of Unified Sentience (DABUS)) as the sole inventor. In April of 2020, the USPTO issued decisions denying the applications, and concluding that the Patent Act defines "inventor" as being limited to natural persons. Thaler appealed to the U.S. District Court for the Eastern District of Virginia. However, that court granted summary judgment in favor of the USPTO, agreeing that the Patent Act requires an "inventor" to be a natural person. Thaler then appealed to the Federal Circuit, which affirmed (*Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022)). The Federal Circuit explained that the Patent Act expressly provides that inventors are "individuals," and that the term "individuals" means a human being. *Id.* at 1211 (relying on *Mohamad v. Palestinian Auth*., 566 U.S. 449, 454 (2012)). Accordingly, the Federal Circuit concluded that, "Here, Congress has determined that only a natural person can be an inventor, so AI cannot be" (43 F.4th at 1214). Thaler filed an appeal to the United States Supreme Court in March of 2023, which the Supreme Court denied.

2. **AI Inventorship in Other Jurisdictions**

From a global perspective, there is also a significant amount of consideration and debate regarding AI-generated inventions. Thaler has provided the primary case study, since the patent applications were filed in many countries throughout the world. For example, Thaler filed patent applications in Australia, South Africa, the European Patent Office (EPO), the United Kingdom, Germany, and Israel. South Africa currently stands alone as the only global jurisdiction to recognize AI as an inventor.

South Africa's Companies and Intellectual Property Commission granted Thaler's application naming AI (DABUS) as an inventor. However, South Africa's patent system does not conduct a substantive examination of patent applications and merely checks to see if formal requirements are satisfied. For this reason, some have questioned the validity of this outcome. (See, e.g., "South Africa was wrong to patent an AI's 'invention'" by Mhangwane et al.). However,

[3] Available at https://ipo.org/index.php/ipo-comments-to-uspto-regarding-artificial-intelligence-and-inventorship/

others have championed the outcome as a progressive and pro-science stance. *See*, e.g., "AI inventorship: The right decision?" by Thaler et al.

Other jurisdictions have held that AI cannot be an inventor. For example, the Australian Patent Office initially refused Thaler's application for naming AI (DABUS) as an inventor, and Thaler successfully appealed. The Federal Court found that an AI system or device can be recognized under the Australian Patent Act 1990 as an inventor, but the court held that "a non-human inventor can neither be an applicant for a patent nor a grantee of a patent." *Thaler v Commissioner of Patents* [2021] FCA 879. The Australian Patent Office then appealed the decision, and the Full Court of the Federal Circuit ruled against Thaler, reasoning that "the origin of entitlement to the grant of a patent lies in human endeavor." A subsequent request for special leave to appeal to the High Court of Australia was denied.

The EPO has also rejected AI as an inventor. In the EPO, Thaler's applications were rejected because the European Patent Convention (EPC) requires an inventor to be a natural person. The EPO stated, "AI systems or machines have at present no rights because they have no legal personality comparable to natural or legal persons." *See* Grounds for Decision 18275147.3 dated 27.01.2020. Thaler appealed, and the EPO Legal Board of Appeal agreed with the EPO's decision.

The respective authorities in Germany, the UK, and Israel have also refused to grant Thaler's applications naming AI (DABUS) as an inventor. In Germany, the Federal Patent Court found that AI generated inventions are not excluded from patentability, but a human inventor must be named. An applicant can identify the AI being involved in the description of the invention. In the UK, Thaler's applications were initially denied, and Thaler appealed. The UK Supreme Court held oral arguments regarding whether an AI machine can constitute an inventor under UK law. The decision of the UK Supreme Court is pending as of the time of this writing. The Israel Patent Office also denied Thaler's applications.

3. Tips For AI-Assisted Inventions

Even if AI cannot be named as an inventor in most global jurisdictions, there can arise additional questions about who should be listed as an inventor regarding a patent application when the claimed subject matter of the patent application involves AI. The extent to which such other questions arise can depend upon the manner in which any given invention may involve or relate to AI. Indeed, as articulated by others, an invention may involve or relate to AI in any of several ways, including whether the invention relates to a technical improvement of an AI system or method of implementing AI (e.g., an improved method of training an AI system, or selection of a particular training data set), whether the invention performs a process that entirely or in part relies upon AI to make a determination or take an action that is part of or influences the process, or whether the invention relates to a data output or result provided by an AI system (e.g., a chemical formula) that has utility in other contexts.

As currently articulated, the law generally (notwithstanding South Africa) continues to view AI through the prism of AI constituting a tool that, even if highly sophisticated, nonetheless remains a tool. When AI is viewed from this perspective, human beings can be considered inventors in any of several manners, and it is appropriate to list human beings who fulfill any of these types of roles. First, when a patent application concerns an invention relating to an improvement to an AI system or training methodology conceived of by a human being, it seems generally to be appropriate that the human being should be listed as an inventor, just as if the human being developed an improved motor or gear system for an electric drill. Second, when a patent application concerns an invention relating to an improved process envisioned by a human being in which AI is used to perform or implement the process so as to make a determination or take an action—but in which the AI being used is a conventional AI system (that is, one not involving any particular innovative feature contributed by that human being or any other human collaborator)—then it would appear that the human being who should be listed as an inventor for that patent application would be solely that human being who envisioned the improved process (but not any human being involved with the development or implementation of the AI system used to perform the process).

Further, when a patent application concerns a useful data output or result generated by the operation of an AI system, where the usefulness of the data output or result is something that is recognized by a human evaluator, the human evaluator should be the inventor in that context. Of course, if there are multiple human beings who collaborated and jointly contributed to the claimed subject matter of a given patent application in any of the above manners, then it may be appropriate for all of those individuals to be listed as joint inventors. Indeed, for example, if a patent application concerns the useful data output or result generated by the operation of an AI system in which one human collaborator recognized the usefulness of the data output/result and another human collaborator provided an improvement to the AI system allowing for that data output or result to be generated, then in that circumstance it may be appropriate to list both of those individuals as joint inventors.

Issues surrounding who should be considered inventors in relation to patent applications that involve or relate to AI will undoubtedly continue to be the subject of debate in the years ahead as AI technology continues to advance. Additional nuanced approaches may be developed in terms of determining whether any given human being should be included as an inventor in regard to any given patent application involving or relating to AI. As to whether laws in the U.S. and around the world will more widely over time come to view AI itself, a non-human entity, as potentially being an inventor in regard to patent applications, this remains to be seen. Changes in laws to that effect would necessarily reflect profound changes in our understanding of what constitutes an acceptable basis for attributing inventorship, in terms of concepts such as whether conception is key to inventorship and what inventorship truly entails in terms of consciousness or sentience. Further, changes in laws on this subject may also be the subject of vigorous debate given that such changes could have enormous ramifications upon what persons or entities control or own the fruits of innovation efforts as AI plays a greater and greater role in these efforts.

H. Proposed Legislation

The *Patent Eligibility Restoration Act of 2023* introduced by Senators Thom Tillis and Chris Coons will no doubt impact the assessment of subject matter eligibility under the current *Alice* and *Mayo* framework, particularly with respect to the implementation of the two-part test to identify claims that are directed to a judicial exception (Step 2A) and to then evaluate if additional elements of the claim provide an inventive concept (Step 2B) (also called "significantly more" than the recited judicial exception).

Importantly, the proposed bill in its current form addresses 101 by addressing Step 2A via the elimination of judicial exceptions, and Step 2B by eliminating Alice's "well-understood, routine and conventional" test. However, two things are notable with respect the proposed bill's approach:

1) Despite the elimination of all court-created judicial exceptions under Section 2 Part A, mathematical formulas and mental processes are still designated as being explicitly excluded from categories of inventions eligible for patent protection under Section 2 Part D.
2) The proposed bill negates the well-understood, routine and conventional test to determine what is "significantly more" than an abstract idea. This negates an oft-used test that is primarily implemented by the Federal Circuit to the detriment of patentees. Yet the proposal remains silent on the USPTO's practical application analysis implemented in the 2019 Guidance that is followed by the PTAB.

The combination of these two factors will likely significantly impact the Federal Circuit's approach towards subject matter eligibility. However, by failing to address the practical application analysis of the USPTO's Step 2A Part 2, the proposed bill will still leave the PTAB's current approach of subject matter eligibility mainly intact. That is, AI and ML method claims are still equally subject to a determination of being directed to a mathematical formula or a mental process. More importantly, when confronted with a combination of apparatus, method and computer program claims, the PTAB's predominant practice is to adopt the method claim as exemplary, thereby subjecting non-method and apparatus claims as being subject to being found as being directed to a mental process along with the method claim.

Namely, the PTAB's practice of utilizing a method claim as exemplary rests in Rule 41.37(c)(1)(iv) which specifies that, when an applicant does not provide separate arguments for different patent claims, the Board may select a single claim from a group and decide the appeal on the basis of the selected claim alone. This practice was implemented as recently as the PTAB's decision in *Ex parte Philip E. Vasey* (Appeal 2022-001109) dated July 5, 2023.

As such, method claims presented with non-method claims can render the non-method claims to be vulnerable unless the applicant explicitly presents different 101 arguments with respect to each different type of patent claim. The current rule places the burden on the applicant to preemptively address representative claim treatment during prosecution or in the brief.

Therefore, the need to present problem/solution language in claim drafting may remain a dominant practice in AI and ML claim drafting despite the "technological improvement" analysis being a progeny of the well-understood, routine and conventional test set up by the Federal Circuit in *Berkheimer v. HP*. While the bill's elimination of this test should theoretically eliminate the technological improvement test, the USPTO's adoption of the technological improvement analysis under the 2019 PEG moved the assessment of improvements to a question of practical application and out of the realm of determining what is well-understood, routine and conventional. As the bill in its current form is silent on the question of practical application, the USPTO framework for Prong 2A Part 2 may still be largely implemented by examiners and the PTAB like, thereby rendering the need for applicants to maintain this practice of presenting problem/solution claims.