Protecting Inventions Relating to Artificial Intelligence: Best Practices

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

Alex Bridge - Hewlett Packard Enterprise
Thomas Burton – Siemens Corporation
Michael Carey – Chau IP Law/F. Chau & Associates
Christina Huang - Faegre Drinker Biddle & Reath LLP
Dave Kincaid – Cantor Colburn LLP
Jennifer Lacroix - SmithAmundsen LLC
Giordana Mahn - Marshall, Gerstein & Borun LLP
Ryan Phelan – Marshall, Gerstein & Borun LLP
John Pienkos – SmithAmundsen LLC

---

[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-Lead

Introduction

Artificial Intelligence (AI) has quickly emerged at the forefront of a fourth industrial revolution that is fundamentally changing the way we live, work, and interact with the environment around us. The importance of AI's role in our future cannot be overstated. As a consequence, many parties are investing heavily in developing new AI related inventions, and seeking to protect the valuable intellectual property so generated. Given the relatively recent surge in interest in AI, some practitioners and applicants may not fully appreciate the unique considerations associated with protecting AI innovation. This white paper seeks to discuss the patenting process as it relates to AI, and aid the practitioner and applicant in forming sound decisions.

I. Preparation of AI Patent Applications

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

(Mike Carey, Dave Kincaid, John Pienkos)

1. Set the Context

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

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

2. Identify the Problem

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

Interview Question: What technical challenges did you face that led to the AI invention? How did you address them?

Interview Question: Did you encounter any unexpected problems during development? How did you address them?

Interview Question: Why were conventional/existing approaches inadequate?

3. Identify Technical Advantages

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

Interview Question: What is improved?

Interview Question: Is the functioning of a computer, another technology, or technical field improved?

Interview Question: How are the improvements realized?

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

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

4. Understand System Inputs

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

Interview Question: What kind of data is used as input to the system?

Interview Question: What kind of data is used for training the system?

5. Understand the Data and Data Collection

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

Interview Question: How is the data collected?

6. Understand Pre-processing Steps

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

More particularly, when conducting an inventor disclosure interview, an effort should be made to obtain an understanding of what types of problems can arise in the data being used for development or implementation of the artificial intelligence or machine learning, and how those problems are avoided or alleviated by pre-processing. For example, is raw data pre-processed to correct or add to the data so as to eliminate known problems or deficiencies in the raw data, such as noise that may be impacting the data, or to reorganize or reformulate the data? Does any such pre-processing entail systems or components that involve more than merely a conventional computer processing device that might provide a further basis for contending that the invention is subject matter eligible and constitutes significantly more than any judicial exception?

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

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

7. Understand Post-processing Steps

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

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

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

8. Understand Network Architecture

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

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

Interview Question: are multiple high level architectural elements used in the invention?

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

Interview Question: what is the structure of the individual layers within the high level architectural components? What activation functions are used?

Interview Question: how are the layers within a high level architectural component connected?

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

Interview Question: what is the computing environment of the AI invention? Are there user devices, databases, or other external elements?

Interview Question: does the AI invention include sensors or control elements for interacting with a physical environment?

9. Understand the Training Process

The training process is one of the key elements that differentiates AI inventions from other software inventions. Training claims have both advantages (e.g., it is sometimes easier to show patent eligibility) and disadvantages (e.g., the training 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 in the specification or claims.

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

Interview Question: Does the invention use supervised learning, unsupervised learning, and reinforcement learning?

Interview Question: What kind of training data is used to train the network?

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

Interview Question: What loss function(s) are used in the training process?

Interview Question: What task does each loss function represent?

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

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

Interview Question: What hyperparameters can be used in the training process?

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

Interview Question: Are different parts of the network fixed, trained separately, or trained jointly?

Interview Question: Are parts of the network trained in multiple phases?

10. Understand the Process at Inference Time

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

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

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

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

Interview Question: What underlying architecture enables the inference?

Interview Question: What improvements result from the inference method?

Interview Question: Are two or more conventional networks combined to address a new problem?

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

Interview Question: Is the trained model updated as a result of the inference? If so, how?

Interview Question: Are the input data used by the trained model expected to change over time?

B. Ethical issues to consider when preparing AI invention disclosures

(John Kind, Ryan Phelan, Giordana Mahn, Alex Bridge)

1. Topic Overview

Artificial Intelligence (AI) has become increasingly important to companies and institutions of all sizes and industries in various different areas. It comes as no surprise, therefore, that AI is being in applied in ways that can have a disproportionate impact on people based certain demographics such as, for example, sex, gender, race, and socioeconomic status.

A disproportionate impact can arise, for example, when an AI model, such as a machine learning model, is not trained on data that takes into account such differences. That is, a failure to train a machine learning model can, in some cases, introduce statistical bias in an AI model where some groups of people are treated differently, but where they otherwise should not be.

2. Technical Overview

An understanding of how an AI model trained is illustrative in how statistical bias can arise in an AI model. Generally, in the field of computer science, the familiar phrase "garbage in, garbage out" applies especially to AI models. That is, if the data used to build an AI model is faulty, one can expect that the AI model would produce a faulty output.

This is because AI is fundamentally a data-driven technology. Training a model involves developing unique datasets as input to train AI computer models. Once trained, an AI computer model may take new data as input to predict, classify, or otherwise output results for use in a variety of applications.

Therefore, an inherent danger lies in the use of incomplete, biased, or otherwise faulty datasets in training an AI model. If such faulty datasets are used, an AI model can fundamentally be trained to output faulty predictions, classifications, or otherwise outputs or decisions that can have real-world impact on people on their livelihoods.

3. Tips for In-House Counsel on promoting good ethical practices

In-house patent counsels who handle AI inventions are uniquely positioned to evaluate the ethical implications of new products and projects early in the development process. Patent counsel have knowledge of the technical operation of AI inventions as well as insight into the business use cases of products incorporating those AI inventions. Patent counsel can efficiently issue spot for possible ethics concerns at the disclosure stage. Particularly, issues such as consumer data privacy, social bias, and exploitation of the technology by malicious customers can be evaluated based on various factors described in the disclosure. For example, ethical issues may be exposed in the type of data being used in the invention, the types of determinations being made by the invention, and other contextual factors.

At a minimum, patent counsel can gather additional information about the invention when a potential ethics issue is flagged. By asking questions of the inventors, the topic is brought to the forefront of the inventors' minds, and they can adjust their implementation to safeguard against those ethics concerns becoming a reality. Many companies, though, are opting for more robust consideration of ethical AI, including forming task forces around the topic. Patent counsel have a valuable opportunity to enmesh the patenting process with ethical AI initiatives. For example, patent counsel can connect inventors with AI ethics personnel when they determine that an invention impacts the ethical AI space. In a currently running initiative, the patent counsel flags potential ethics impacting invention disclosures, and the AI ethics partner evaluates the flagged invention disclosures. If the invention sufficiently impacts ethical considerations, the AI ethics partner reserves a portion of the patent drafting inventor interview to conduct a short conversation on the ethical implications of the invention. The conversation is not aimed at forcing change on the inventors, but rather at empowering the inventors by giving them information about the ethical considerations that the company is focused on. In some industries, the company may benefit from patent counsel or AI ethics partners taking an even more active role in the ethical conversation by guiding the product team through ethical risk measurement exercises.

One particularly sensitive consideration that can often lead to counter-productive results is the assertiveness of non-technical partners when discussing ethical AI with inventors. AI has always been a space where non-technical partners have relied on buzzwords and term ambiguity to embolden their particular interests. Whether that be from sales teams, marketing departments, executive staff, journalists, enthusiasts, or other sources, AI engineers are particularly sensitive to the distinction between those who "talk the talk" and those who "walk the walk." To be effective, ethical AI discussions with technical partners must be conducted as discussions, not as lectures, lest "ethical AI" become yet another buzzword tossed around by non-technical partners and loathed by the technical partners. Patent counsel, as semi-technical partners, are well positioned to set guardrails for the non-technical partners in order to avoid offending or otherwise suppressing the technical voices in the conversation.

Generally, patent counsel needs to be cognizant of the types of AI inventions coming across their desk. Inventions that ingest technical data and output technical determinations (e.g., network configurations, pharmaceutical compositions) have a very different ethical risk profile than inventions that ingest personal information and output social determinations (e.g., prison sentencing determinations, job applicant filters). Patent counsel encountering the former type of invention may be concerned more with informing the inventors of the ethical considerations. Patent counsel encountering the latter type of invention may need to assist the engineering team in navigating the ethical concerns of the particular problem space and even in determining whether the problem is best solved with the particular AI solution described in the invention disclosure. Regardless of the exact solution chosen by a particular company or patent counsel, incorporating ethical considerations into your practice is a value add at a time when ethical AI is at the forefront of many AI conversations.

C. International considerations

(Sumon Dasgupta, Christina Huang)

The patent process (e.g., patent laws, rules and regulations) vary across various countries and jurisdictions. Thus, it is important that an Applicant remain aware of such distinctions to sensibly guide the patent process from application drafting to international filings. Below is a brief discussion of how the patent process across several countries and jurisdictions, and how the patent processes may differ from the patent process in the US.

1. European Patent Office (EPO)

The Guidelines for Examination in the EPO, G-II state that “Artificial intelligence and machine learning are based on computational models and algorithms for classification, clustering, regression and dimensionality reduction, such as neural networks, genetic algorithms, support vector machines, k-means, kernel regression and discriminant analysis. Such computational models and algorithms are per se of an abstract mathematical nature, irrespective of whether they can be ‘trained’ based on training data.” In contrast, US examination practices may not find AI related claims to be abstract per se (e.g., Subject Matter Eligibility Examples: Abstract Ideas, Example 39). Rather, US patent practice involves determining *whether* AI related claims include an abstract idea (e.g., a method of organizing human activity, a mathematical relationship, etc.) and treating the claims accordingly.

The EPO clarifies that patents may still be granted when AI leaves the abstract realm by applying it to solve a technical problem in a field of technology (i.e., a technical solution to a technical problem). As such, when drafting applications that will be filed to the EPO, Applicants should bear in mind the importance of clearly defining the technical problem and the technical solution throughout the specification. Applicants may also consider that features related to the technical solution may be required in the claims in the EPO.

2. Korea

Korea appears to be less stringent with respect to patent eligibility. For example, Korea identifies patentable subject matter based on novelty and inventiveness (as long as 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.

3. Japan

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 is capable of creating 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.

4. China

In February 2020, China Patent Office 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 artificial intelligence 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 how technical features are 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.

5. Inventorship across different jurisdictions

It is no secret that AI has evolved in a myriad of ways to reach new heights. In fact, AI can now be considered an "inventor" in some contexts. This (as will be discussed herein) can be problematic in certain jurisdictions that only permit natural persons as the inventor. For example and as of the time of drafting this article, AI arguably should not be listed as an inventor in the USA, UK, EPO, JPO, Switzerland, Korea, China among others. While litigation is pending in several of the above jurisdictions to dispute this assertion, it is nonetheless difficult to ascertain the validity of such litigations and whether AI-generated inventions will be permitted.

In contrast, South Africa recently approved a patent in which AI was named as the inventor. The Federal Circuit of Australia also held AI-generated inventions as being patentable, and the PCT permitted AI to be listed as an inventor in WO 2020/079499. As such, Applicants may consider whether it is worthwhile to patent AI-generated inventions at this time (pending the outcome of ongoing litigation), and which jurisdictions are most favorable to AI-generated inventions. Applicants are also advised to carefully monitor the outcome of litigations associated with AI-generated inventions to determine whether precedent will change in the above jurisdictions.

D. Inventorship

(Thomas Burton, John Kind)

Inventorship issues can be murky, even in non-AI cases. This chapter focuses on how to identify and handle inventorship issues where AI (i.e., non-human inventors) may be viewed as the sole or contributing inventor.

1. Current Landscape of Laws/Regulations for Inventorship

a. Overview of Laws/Regulations on Inventorship

Under current laws in the United States and most other jurisdictions, inventors cannot be non-natural persons. See "Comments on Patenting Artificial Intelligence Inventions", IPO Published Letter to Director Iancu, pg. 5, available at https://ipo.org/wp-content/uploads/2019/11/IPO-Comments_Patenting-AI.pdf, (Nov. 11, 2019). In the United States, "'the threshold question in determining inventorship is who conceived the invention' as claimed in a patent." *Id.* at pg.4, (citing US MPEP 2137.01 and Fiers v. Revel, 984 F.2d 1164, 1168 (Fed. Cir. 1993)). Under U.S. laws, conception is defined to be formed "in the mind of the inventor", which makes it difficult to identify AI or a non-human as the sole inventor or co-inventor. However, recently, some jurisdictions have recognized AIs as inventors.See, e.g., https://www.ipaustralia.gov.au/about-us/news-and-community/news/commissioner-appeal-court-decision-allowing-artificial-intelligence; https://www.managingip.com/article/b1sx9mh1m35rd9/dabus-south-africa-issues-first-ever-patent-with-ai-inventor.

Recognizing an AI as an inventor presents significant practical issues. For example, to submit a patent application in the US, an inventor is required to execute an oath or declaration stating that such inventor "believes himself or herself to be the original inventor or an original joint inventor of a claimed invention in the application." *Id.* at pg. 4 (citing 35 U.S.C. 115(b)(2)). The use of a "substitute statement" under 35 U.S.C § 115 (d) in lieu of executing a declaration does not fill the gap left by 35 U.S.C § 115(b)(2) because an AI system is not an "individual" who is "(i) deceased, (ii) is under legal incapacity; or (iii) cannot be found or reached after diligent effort; or (B) is under an obligation to assign the invention but has refused to make the oath or declaration required. . . ." *Id.* (citing 35 U.S.C § 115(d)(2)).

The inability of an AI or non-human to own property and execute an inventor assignment also poses issues for naming an AI as an inventor.

2. DABUS Challenges, Current Laws/Regulations For Inventorship With AI Named As Sole Inventor

DABUS is an AI paradigm where controlled chaos combines whole neural nets, each containing simple notions, into complex notions (e.g., inventions). DABUS is argued by some to be sentient due to its characteristic that any chain-based concept launches a series of memories (i.e., affect chains) that sometimes terminate in critical recollections, thereby launching a tide of artificial molecules. The team behind DABUS has filed patent applications listing DABUS as the sole inventor in countries worldwide to challenge the inventorship laws/regulations in each jurisdiction that may inhibit AI (or a non-human) from being named as an inventor. DABUS has been pursued in many different jurisdictions with some level of success. For example, South Africa recently issued a patent with DABUS listed as the inventor, and Australia recently upheld that AI may be listed as an inventor. Numerous other jurisdictions disagree and have help that AI cannot be an inventor (e.g., USA, UK, EPO, JPO, Switzerland, Korea, China among others).

3. Conclusion

Until inventorship laws/regulations are changed, continue to pursue patent apps that only name humans that contributed to what is claimed, even where AI used as part of the “design” process. When filing patent applications, Applicants should identify aspects that may be novel and determine whether conception of those aspects may be traced back to a human inventor. This may be a tricky determination as it is not always clear whether a human conceived of a feature using an AI as a tool or whether the feature may, legally speaking, have been conceived by the AI itself. To the extent features that cannot be traced back to an act of human conception may be separated from the broader invention, these features may be better kept as trade secrets and not included in the patent application. Where this is not possible, Applicants should seek to exclude such features from the claims and focus on those features that can be clearly tied back to a human inventor to avoid having to deal with the inventorship question.

II. Prosecution

A. How to draft AI applications to prevent or overcome rejections under 35 U.S.C. §101 and §112 before the USPTO

(Frank Chau, Mike Carey, John Pienkos, David Kincaid)

1. Introduction

Artificial intelligence (AI) is a broad term that includes many types of applications. Machine Learning (ML) is the dominant category in AI applications, followed by subcategories such as Image Processing, Natural Language Processing (NLP) and Speech Processing. ML involves gathering and processing data to feed a neural network (NN), training the NN, and inference methods to determine the outcome. Image processing relates to manipulating an image in order to enhance it or extract information from it, and can be used for interpreting and analyzing images for a wide range of applications, such as improving or restoring images, face recognition and authentication, and detecting and recognizing objects and patterns in images and videos.

Given that AI technology fundamentally involves software, AI-based patent applications often face some of the same types of challenges that are faced by other software-related patent applications. In particular, as with many other software-related patent applications, AI-based patent applications face heightened scrutiny as to whether the claimed inventions concern patent-eligible subject matter under 35 U.S.C. §101. Additionally, AI-based patent applications can face additional concerns involving whether the claims are sufficiently definite or whether there is sufficient support for the claims of an application in the specification of the application under 35 U.S.C. §112. Notwithstanding these concerns, due to the subject matter of AI-based patent applications, there do exist strategies for achieving allowability of AI-based patent applications that are especially well-suited to those applications even if those strategies are unavailable or less applicable with respect to other software-related patent applications.

2. Drafting Patent Eligible AI Applications Under 35 U.S.C. §101

a. Alice/Mayo Eligibility Test

As in other technology areas, USPTO examiners use the Alice/Mayo test to determine the eligibility of AI patents under 35 U.S.C. §101. To apply the test, it is necessary to: first determine if the claims fall into one of the statutorily-defined categories of patentable subject matter;[1] next, ask if the claims recite an abstract idea (step 2A, prong 1);[2] then consider whether the claim is integrated into a practical application (step 2A, prong 2);[3] and, finally, determine whether additional elements amount to "significantly more" than the abstract idea (step 2B).[4] Although these same steps apply to AI inventions, there are special considerations that can improve the likelihood that an AI patent will be eligible.

Recent Federal Circuit cases show a trend of declaring claims ineligible if the inventive concept is described using functional language, even if the claim includes concrete elements. For example, in *American Axle & Manufacturing, Inc. v. Neapco Holdings LLC*[5], a claim for a method of manufacturing a shaft assembly included a functional limitation related to tuning a shaft liner.[6] The claim was held to be ineligible because it was directed to Hooke's law. Similarly, in *Yu v. Apple*[7], the claims included concrete camera components (e.g., image sensors).[8] Still, the CAFC found the claims ineligible because these elements were considered "generic processes and machinery."

When drafting AI claims, it can be tempting to fall back on functional language because key elements of the invention are performed by a "black box" (i.e., a system with so many parameters that describing the precise structure is impractical). However, functional claiming can lead to a rejection under 101 as in *American Axle* and *Yu.* Therefore, AI claims should be carefully crafted to avoid this fate.

b. Claiming Strategies

One way to avoid 101 rejections is to recite AI-specific architecture or training elements. Claims directed to the architecture and training of neural networks are likely patent eligible because: (1) they do not fall within the judicial exceptions enumerated in the 2019 PEG, and (2) they are integrated into a practical application. Furthermore, these claims can avoid falling into the trap of being held ineligible like those of *American Axle* or *Yu* for reciting functional limitations without structure.

AI claims that do not recite specific architecture or training elements should be drafted carefully to avoid being ineligible. For example, as described below, they can be drafted in a problem/solution style that clearly shows how the claim recites an improvement over existing technology.

Also, AI-based patent applications involving particular types of subject matter can be significantly less likely to be rejected under 35 U.S.C. §101, especially if the AI-based patent applications are assigned to particular art units at the USPTO.

i. Architecture and Training Claims

In *Alice*, the Supreme Court stated that "the mere recitation of a generic computer cannot transform a patent-ineligible abstract idea into a patent-eligible invention." Courts and USPTO examiners have often cited this to reject software claims that include steps described at a high level of generality that are performed on a computer.

However, AI functional claims are typically performed using particular architectural components (i.e., processing patterns) that can be recited at a high level of generality. The architectural components are not simply a “generic computing environment”. Examples of these architectural components include a convolutional neural network (CNN), and a recurrent neural network (RNN). 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 a RNN may be better suited for NLP). Thus, referencing specific AI architectural elements in a claim can provide grounds for showing that the claim is integrated into a practical application under step 2A, prong 2 of the *Alice*/Mayo test.[9]

Even where the inventive aspect relates to inference, it is also useful to describe and claim architectural aspects of the invention to show non-conventionality. For example, a recurrent neural network (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. At a minimum, such an approach can provide a helpful “backup” position when undergoing §101 scrutiny.

In addition to claims reciting architectural elements, AI and machine learning claims can avoid being held patent ineligible if the claims recite a training process. This is because the training of neural networks involves specific processes tailored for particularized machines, and these processes are not conventional processes performed by generic computer systems.

In fact, the USPTO itself has provided guidance directly related to AI training claims. Specifically, the MPEP listed examples 37 to 42 for USPTO examiners to use in conjunction with the 2019 PEG. Example 39 relates to training a neural network. In the analysis of example 39, the guidance clearly states that the claim does not recite a judicial exception under Step 2A, prong 1 of the eligibility test.[10] As a result of this guidance, USPTO examiners in at least some art units have been instructed to consider AI training claims to be patent eligible *per se*.

ii. Problem/Solution Claims

Claims that do not recite specific architecture or training elements can still be crafted to avoid a §101 rejection. Specifically, a claim can recite a problem, a technical solution, and the resulting improvement over prior art systems, or improvement in computer/AI functionality, to satisfy the *Alice/Mayo* test.[11]

The USPTO's Patent Eligibility Example 42 shows how to draft problem/solution claims. The claim in Example 42 relates to a method of transmitting notifications when medical records are updated. One limitation recites language indicating a problem to be overcome (i.e., "information in a non-standardized format"). Another describes a technical solution to the problem ("converting... into the standardized format"). The final limitation of the claim describes the improvement ("each user has immediate access to up-to-date patient information").

AI claims can also be structured in this manner. 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"). Machine learning steps correspond to 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 machine learning model (e.g., "navigate through the environment while avoiding the obstacle based on the depth map").

Thus, AI claims can be crafted with elements or combination of elements that reflect an improvement specific to AI technology. Examples of AI-specific improvements include:

- Improving the accuracy of predicted annotations
- Reducing size or layers of a neural network
- Improving inference accuracy
- 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)
- Reducing the number of network parameters
- Increasing the speed or efficiency of network operation/prediction
- Enabling the use or optimization of different hardware (e.g. GPU vs CPU)

Accordingly, in addition to including architectural and training elements, AI claims can be patent eligible if they include claim elements directed to improving the functioning of a computer or AI technology. Such claims should include specific terms that correspond to 1) the problem, 2) the technical solution, and 3) the result.[12]

iii. Art Unit Selection

Because AI encompasses so many applications, AI inventions can fall into any of a large number of art units at the USPTO. The treatment of AI applications in various art units varies, and there are some types of applications that are more likely to avoid subject matter eligibility rejections. It may thus prove useful to draft an AI application so that is clearly directed to the art unit in which it is most likely to succeed.

For example, AI applications that are classified as relating to image processing (e.g., AI applications falling into Class 382 concerning Image Analysis) recently have been less likely to face subject matter eligibility rejections than some other types of AI applications.[13] Accordingly, in general it is desirable that any AI application relating to image processing be appropriately recognized and categorized as such. In this regard, the patent application should have a title that makes clear that the application concerns image processing or similar concepts such as image transformation or pattern recognition. Also, the patent application should include claims (at least independent claims) that emphasize how the invention concerns image processing or similar concepts, including in the preambles of the claims.

---

[1] Specifically, whether the claims recite a process, machine, manufacture, or composition of matter.

[2] The three categories of abstract ideas (also referred to as "judicial exceptions") outlined in the MPEP, and in the 2019 Revised Patent Eligibility Guidance (2019 PEG), are mathematical concepts, methods of organizing human activity (business methods; fundamental economic practices), and mental processes (concepts performed in the human mind).

[3] The 2019 PEG lists, as one example of "a judicial exception integrated into a practical application," the following: claim elements reflect an improvement in the functioning of a computer, or an improvement to other technology or technical field.

[4] According to the MPEP, elements that the courts have found to qualify as "significantly more" include:

improvements to the functioning of a computer; improvements to any other technology; applying the judicial exception with a particular machine; effecting a transformation or reduction of a particular article to a different state or thing; or unconventional steps that confine the claim to a particular useful application.

[5] American Axle & Manufacturing, Inc. v. Neapco Holdings LLC, et al., 2018-1763 (Fed. Cir. July 31, 2020).

[6] Claim 1 of the patent at issue (U.S. Patent No. 7,774,911) included the limitation "*tuning at least one liner to attenuate at least two types of vibration transmitted through the shaft member*."

[7] Yu v. Apple, 2020-1760 (Fed. Cir. June 11, 2021).

[8] Claim 1 of the patent (U.S. Patent No. 6,611,289) included the limitation "*a digital image processor...producing a resultant digital image from said first digital image enhanced with said second digital image*."

[9] The 2019 PEG also lists, as another example of "a judicial exception integrated into a practical application," the following: an additional claim element implements or uses a judicial exception in conjunction with a particular machine.

[10] Further for example, Section 2106.04(a)(1) of the MPEP identifies the following as an example a claim that does not recite an abstract idea: "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."

[11] These considerations are applicable to both step 2A and step 2B of the Alice/Mayo test. According to the Section 2106.04(d) of the MPEP: "Step 2A Prong Two is similar to Step 2B in that both analyses involve evaluating a set of judicial considerations to determine if the claim is eligible. See MPEP §§ 2106.05(a) through (h) for the list of considerations that are evaluated at Step 2B. Although most of these considerations overlap (i.e., they are evaluated in both Step 2A Prong Two and Step 2B), Step 2A specifically excludes consideration of whether the additional elements represent well-understood, routine, conventional activity."

[12] The problem, technical solution, and benefit/result should be fully described in the specification to provide support for advocating that the claim is directed to a "non judicial exception" or a "significantly more" inventive concept.

[13] See "Determining the Likelihood that an AI Patent Application Will Be Allowed at the USPTO" by Kate Gaudry et al., IPWatchdog, October 25, 2020.

B. Drafting AI-Based Patent Applications with Section 112 in Mind

(Frank Chau. Michael Carey)

1. Avoiding 112 Rejections

Considerations similar to those of the patent eligibility determination apply to the disclosure requirements of 35 U.S.C. §112. Namely, despite the fact that AI and machine learning inventions often depend on the specific values of many internal parameters, these inventions should not be described as a black box. Rather, as in the case of the claims, an AI patent specification can include details related to architectural elements, training elements, and problem/solution elements.

2. Architecture Description

The key to describing AI and machine learning architecture is to become familiar with several key levels of abstraction. By understanding these levels of abstraction, it is possible to describe how the structure performs a claimed function without the need to describe the value of parameters in a neural network. The highest level of description is the functional component description (e.g., an image classification neural network). But on its own, this description is not enough to satisfy the requirements of §112. 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.) For example, the specification should 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 satisfy the requirements of §112, it is important to include details that go deeper than high level architecture models. 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 values of the parameters. As an 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.

3. Training 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 is the 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, for example, the loss functions of a supervised learning process or a policy model of a reinforcement learning process.

By thinking of an AI architecture and training on these three levels, the specification will have enough depth to describe HOW the invention is accomplished, thereby satisfying the disclosure requirements of §112.

4. Problem/Solution

In addition to technical details related to the architecture and training, AI applications can include a description of a technical problem, a technical solution, and the resulting improvement over related technologies.

When describing a problem in the existing technology, one 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.

With 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 software/AI implemented, provide an algorithm for how the function is performed. Furthermore, if an AI specific term is claimed, provide a non-limiting definition or provide a description of an exemplary use of the claim term in the specification.

In addition to a brief but clear description of the problem, and a detailed description of the technical solution (as recited in the claims), it is also 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.

C. How to draft AI applications to withstand challenges brought before the CAFC

(Sumon Dasgupta, Ryan Phelan)

As of August 2021 (at the time of writing), AI is a relatively newer technology and as such, court cases analyzing AI-related patents are infrequent. The Federal Circuit has yet to analyze an AI-related patent for patent-eligibility. It is generally expected that the Supreme Court's two-part Alice test will be used to analyze general AI-related inventions at the Federal Circuit, as the Federal Circuit has used this test for more general software-related inventions, of which AI-related inventions are a subset.

For example, the USPTO and District Courts and US Court of Appeals, have followed the Supreme Court's two-part Alice test to analyze general software AI-related inventions. As such with respect to software AI-related subject matter, patent practitioners may adopt a similar approach as has been used with respect to computer-related inventions in general (e.g., strive to avoid reciting an Abstract idea in the claims and describe how the claimed invention provides an improvement to a technical field).

It is further apparent that at least some judges of the Federal Circuit have signaled caution in getting questions about "important inventions" (such as AI) wrong, and are reluctant to grant case dispositive motions. In *Smart Sys. Innovations, LLC v. Chicago Transit Auth.*, Judge Linn described the uncertainty of patent eligibility review using the two-part Alice test and the dangers inherent to new age technology, such as Artificial intelligence, of arriving at the wrong result. 873 F.3d 1364, 1378 (Fed. Cir. 2017) (Linn, J., dissenting and concurring in part) ("Despite the number of cases that have faced these questions and attempted to provide practical guidance, great uncertainty yet remains. And the danger of getting the answers to these questions wrong is greatest for some of today's most important inventions in computing, medical diagnostics, artificial intelligence, the Internet of Things, and robotics, among other things."). Similarly, in *Athena Diagnostics, Inc. v. Mayo Collaborative Servs.*, Judge Moore, writing in dissent, expressed her concern as to the impact of *Alice* (and related case *Mayo*). LLC, 927 F.3d 1333 (Fed. Cir. 2019) (per curiam). The *Athena Diagnostics* court had denied a petition for rehearing en banc regarding claims directed to a "method for diagnosing neurotransmission or developmental disorders." The court, issuing the opinion per curiam, expressed its obligation (and concern) in holding the claims ineligible. *Id.* (Lourie, J.) ("I concur in the court's decision not to rehear this case en banc. In my view, we can accomplish little in doing so, as we are bound by the Supreme Court's decision in Mayo. … Some of us have already expressed our concerns over current precedent").

In another example, in *Singular Computing LLC v. Google LLC*, 2020 U.S.P.Q.2d 10708 (D. Mass. 2020), the court denied a 12(b)(6) motion to dismiss, where the defendant alleged that the asserted claims were patent-ineligible per the 2-step Alice test. The patent-at-issue generally related to "a processor or other device" that included "processing elements designed to perform arithmetic operations . . . on numerical values of low precision but high dynamic range [i.e., LPHDR]." The patent owner had argued that such architecture allows "for more efficient use of a computer's transistors, which improves computer performance in certain applications such as artificial intelligence software." In denying the motion, the court focused on the alleged improvement (via LPHDR technology) to existing computer architecture, finding that such improvement made the claimed invention plausibly inventive, which was at least enough for the patent claims to survive at the pleading stage where allegations as to inventiveness were accepted as true.

In some cases, the court can find AI-related patents invalid at the outset, especially where the respective claims lack any AI-specific elements or architecture. For example, in *Applied Predictive Techs., Inc. v. Marketdial, Inc.*, the court granted a Rule 12(b)(6) motion to dismiss where the claims lacked AI-specific elements and were instead "directed to the abstract concept of optimizing parameters for business initiative testing." 2020 Us Dist Lexis 221981 (D. Utah Nov. 25, 2020). In *Applied Predictive Techs.*, even though the specification described the possible use of a "neural network" to perform the claimed "testing," the court found that "no specific, technical details regarding the exact sort of analysis" was provided, and, therefore, the claims were considered to cover "history data analysis" broadly. "To grant a patent on the concept as it is articulated in the [patent-at-issue] would be to preempt any method of analyzing historical data to determine how certain parameter settings influence that data." Thus, the court granted the motion finding that claims-at-issue patent ineligible.

While not binding on federal courts, decisions by the Patent Trial and Appeal Board (PTAB) can be illustrative, and perhaps persuasive, in future litigation regarding AI inventions. As an example, the Patent Trial and Appeal Board (PTAB) more recently applied the 2019 PEG (as revised) in an ex parte appeal involving an artificial intelligence invention. See ex parte Hannun (formerly Ex parte Linden), 2018-003323 (April 1, 2019) (designated by the PTAB as an "Informative" decision). 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 with the Examiner's assessment. While the PTAB generally agreed that the patent specification included mathematical formulas, such mathematical formulas were "not recited in the claims," and claims did not recite any feature relating to "organizing human activity." That is, the PTAB found that at least some of the claims were directed to a specific implementation comprising technical elements, including AI and computer speech recognition.

Lastly, 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.’”

D. Drafting AI Claims that are detectible or Trade Secret

(Sumon Dasgupta, Thomas Burton, Alex Bridge)

Introduction

What are “detectable features”?

This section will focus on detectability related issues surrounding AI patents. As discussed below, AI related subject matter poses unique challenges from a detectability standpoint. Generally, a valuable patent is one which is detectable. Thus, consideration relating to detectability may guide aspects of preparation and prosecution.

1. What are “detectable features”?

At the outset, it is beneficial to briefly discuss how features are detectable. Generally, a detectable feature is viewable and discernable from a product, process and/or method embodying the AI invention. In a best case scenario, such a detection may be readily apparent from commonly viewed products (e.g., “off the shelf” products). In other more difficult scenarios, infringement may be detected through reverse-engineering (although such an endeavor may prove to be time consuming and costly) even in cloud based situations. In some cases, infringing actions may be disclosed in publicly available documents (e.g., market materials, user specification and guides, websites, online video, online forums, etc.) simplifying the infringement analysis.

An often perplexing and difficult question is how to determine which features of an invention are detectable. For example, to satisfy the requirements of patentability (e.g., enablement and written description), many patent practitioners will harvest significant details from the inventors. Not all of these details will be detectable in an AI context. The parties most readily able to differentiate between detectable features and non-detectable features are generally the inventors. Thus, it is important for patent practitioners to carefully discuss such concerns with the inventors, and inquire as to which aspects are able to be most readily discussed and disclosed in product literature, websites, etc., and easily discernible from an infringing product itself. It may also be helpful for the patent practitioner to avoid (if possible while still meeting statutory requirements) disclosing features which will never be detectable and maintaining such features as a trade secret.

2. Claimed features that are viewable when product/method embodying AI invention is operated

A software or AI related invention disclosure that has a key feature that is viewable when the software or AI is operated should be claimed in a patent application since such an infringing use by a competitor or competitor's customer is easily detectable. For example, an AI related invention may require a user interface to prompt a user for key inputs for the AI to perform a task and display corresponding AI intermediate or final results, such as to identify user specific criteria or data sets for training a speech recognition or transcription AI model or an AI tool for learning how to detect faults in a particular product.

The US PTO 2019 Subject Matter Eligibility Examination Guidance provides an example of eligible claims for a software related invention that is detectable when the underlying software is operated (see "Example 37- Relocation of Icons on a Graphical User Interface" at https://www.uspto.gov/patents/laws/examination-policy/subject-matter-eligibility). In this example, the claimed step of "determining, by a processor, the amount of use of each icon over a predetermined period of time;..." may be detectable even if implemented with AI. Moreover, the claimed feature may be detectable even if AI processing is performed in a cloud but requires user inputs via a browser user interface (i.e., client) or a mobile device.

In another example, AI that is part of an autonomous device (e.g., vehicle, robot, game, etc.) or a device that requires learning (e.g., a thermostat that learns your schedule based on detected occupancy) may be also be detectable in the operation of the device. For example, a robot may be trained to perform a specific task based on user specified preferences that may be a key claim element. A learning thermostat may learn a user's preferences for heating and cooling based on the user overriding pre-programmed schedules in the thermostat. AI claims should focus on those observable features.

3. Claimed features very likely to be disclosed in competitor's publicly available documents

AI claimed features may also be detectable based on publicly available documents pertaining to the competitor's infringing product, system or service (including a cloud software-as-a-service). Even if the AI claimed feature is directed to an AI model/architecture improvement, a competitor may mimic the AI model/architecture key features in advertising materials to match or beat your business's AI marketing materials for the same features. For example, a competitor's marketing materials may emphasize that their AI "uses a 7 layer neural network model that provides the following advantages…" or their AI "uses deep learning in a particular way to provide better, faster results. Thus, AI claims should focus on those model/architecture features that may be detectable in competitor's marketing materials.

Similarly, if the AI invention is directed to an AI commissioning/training process, a competitor's user guide or technical manual is likely to identify AI key features for installing, setting up, and commissioning the AI to train the AI for a specific customer. For example, a company that develops AI based data analytics tools for energy management of a building or home may need to specify certain details in its user documents to guide a technician in commissioning the tool for a particular building or home. Most buildings and homes are different. Thus, an AI user guide may identify key steps for installing, commissioning, and training the AI tool for the specific building or home. Likewise, and AI related medical device may require setup to properly train the AI on data for specific type of patients. AI claims should focus on those commissioning and training features.

4. Cloud-Based and As a service (aaS) Documentation Paragraph

One particular concern when attempting to detect ML features using publicly available documentation is the industry shift toward cloud-based [1] and as-a-Service (aaS)[2] provision of machine learning solutions. Rather than publishing documentation targeted at a customer's IT staff and software developers that describes, in intricate detail, how the system works,[3] the documentation for cloud-based and aaS solutions merely need to describe how to interact with the system.[4] Since the provider manages and maintains their own systems, the details required to troubleshoot issues, build customizations for particular workloads, and otherwise implement novel features are held as trade secret, only published to a limited audience. Such features will not be detectable using documentation.

---

[1] https://www.datasciencecentral.com/profiles/blogs/6-cloud-based-machine-learning-services

[2] https://www.altexsoft.com/blog/datascience/comparing-machine-learning-as-a-service-amazon-microsoft-azure-google-cloud-ai-ibm-watson/

[3] https://developer.nvidia.com/deep-learning-software

[4] https://docs.microsoft.com/en-us/azure/machine-learning/how-to-deploy-and-where?tabs=azcli

5. AI claimed product/method features can be reverse engineered or can't be maintained as a trade secret

While there are many different implementations and frameworks for machine learning solutions, some general principles can be distilled out that can be used to better detect features using reverse engineering or other technical methods. Machine learning provides value mostly at scale, as recognized by the popularity of the term "Big Data." Big data usually does not come from a small handful of devices controlled by the machine learning provider, but from a large number of broadly deployed devices, such as mobile devices, edge-deployed sensors, internet of things (IoT) devices, medical equipment, and other such devices. In some applications, the machine learning solutions use publicly available data and/or publicly defined data structures to train on, exposing both the inputs and the outputs of the solution to inspection.

Features can be detected at many stages of the process, and 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. Usually, the edge devices and the communication between the edge devices and the centralized cloud compute resources are much easier to capture and investigate via technical methods than the centralized compute resources themselves.

Of particular advantage to the investigator is the need for ML services to forward the input data, in some form, to the centralized compute resources for further action. Given network constraints on most enterprise networks, raw data dumps are less desirable, as they tend to saturate the network if the data being collected is large enough. Therefore, many different feature engineering, data processing, and even training and model application steps may take place at the edge device. The ability to easily (compared to centralized compute resources) acquire, analyze, and monitor such edge devices allows such investigation to uncover what is being done to the inputs before sending information to the centralized compute resources.

The types of information that can be gathered at the edge device is highly dependent on the particular ML architecture, but information such as raw inputs, data pre-processing methods, feature engineering algorithms, and even trained ML models may be discoverable at the edge device with the proper reverse engineering methods. However, as with all reverse engineering, detecting these features will likely require high effort and high technical competence, which makes it unsuitable for casual infringement analysis. For such casual analysis, scouring the documentation is still likely the best bet.

When claiming ML inventions, pay particular attention to the dataflow between centralized compute resources and edge devices. Inventions that are purely targeted at features contained within the centralized compute resources may be nearly impossible to detect, but the after-effects of such features may be detectable based on information communicated to edge devices or actions taken by the edge devices once they receive such information. Also pay attention to steps in the process that are accomplished by the edge devices themselves, as these devices are usually not able to do substantial computation without "phoning home" to the centralized compute resources, and claims can be targeted to the types of information included in that "call back home". Further, thinking in terms of the types of data being used and the types of inferences being made can help capture broader infringement of the invention rather than only the productized embodiment. There are many paths to the same type of inference, and it is important, when scoping out a specification and claims, to understand why certain datapoints are being used, how other datapoints could be substituted, and what the inference actually represents. Certain types of ML models are easier to perform such an analysis on than others. For example, deep neural networks are notoriously shrouded, and figuring out which datapoints contributed to which intermediary signals in which way is a problem being addressed by university research. Other ML models may be more easily inspected.

6. Claim features that are standard compliant

Standards are needed for the adoption and compatibility of new technologies. Standard bodies, such as American National Standards Institute (ANSI), European Telecommunications Standards Institute (ETSI), Institute of Electrical and Electronics Engineers (IEEE), Internet Engineering Task Force (IETF)[2], 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.

When it is possible to request a license from a standard compliant implementer under FRAND, it may be worth contributing patented technologies to the standard.

Usually, some parts of a standard are considered essential and some parts are considered optional (for example a product may omit an optional feature or functionality, but it cannot omit an essential one). Detecting infringement of an essential feature of a standard compliant product is straightforward. When a product is standard compliant, it usually includes all its essential features. If a patent claim can be mapped to an essential feature of the standard, then it is all it takes to demonstrate infringement for a corresponding standard compliant product. When it comes to optional features, an additional step may likely be required to demonstrate infringement, in that the patent holder may need to demonstrate that the optional feature is included in or supported by the standard compliant product.

There exist a few ways in which patented technology can make its way into the standards. First, an entity may actively push its proprietary technology into a standard by having their own standard representatives who contribute in defining the standard. Second, standards users may patent proprietary technology while developing their standardized product e.g. the patent covers an aspect that is not described by the standard and that is left free to the implementer, but that may become standardized at some point in the future. Third, the standardization body may unknowingly adopt a technology that was previously patented.

[2] https://www.wipo.int/patent-law/en/developments/standards.html

AI is currently being deployed within many new technologies, and standardization bodies are not an exception. Among others, IEEE is working on standards for AI affecting human well-being[3]. ANSI proposes opportunities for standardization for AI in healthcare[4]. IETF has an initiative[5] aiming at applying AI technologies to networks. The third generation partnership project (3GPP) reports[6] that operators and vendors are now turning their attention to Artificial Intelligence and Machine Learning (AI/ML) to address different challenges in 5G, the fifth generation of wireless telecommunications. ETSI is working on a first standard for securing artificial intelligence[7,8].

Those who have patented AI technologies that become part of standards may very well be able to reap the benefits of their effort.

[3] https://transmitter.ieee.org/new-ieee-standards-artificial-intelligence-affecting-human-well/

[4] https://www.ansi.org/news/standards-news/all-news/2021/03/3-29-21-ansi-releases-report-standardization-empowering-ai-enabled-systems-in-healthcare

[5] https://www.ietf.org/mailman/listinfo/IDNET

[6] https://www.3gpp.org/news-events/2201-ai_ml_r3

[7] https://www.etsi.org/newsroom/press-releases/1871-2021-01-etsi-report-paves-the-way-for-first-world-standards-in-securing-artificial-intelligence

[8] https://www.etsi.org/committee/1640-sai

III. Enforcement

(**Ryan Phelan, John Pienkos, Jennifer Lacroix)**

In the last quarter of 2020, the United States Patent and Trademark Office (USPTO) reported that patent filings for Artificial Intelligence (AI) related inventions more than doubled from 2002 to 2018. *See* Office of the Chief Economist, Inventing AI: Tracking The Diffusion Of Artificial Intelligence With Patents, IP DATA HIGHLIGHTS No. 5 (Oct. 2020). However, because AI is still a relatively newer technology, cases involving AI-related patents have been few in number. Nonetheless, the developing body of cases makes it clear that patents involving AI and other software-related technologies are often invalidated in early stages of litigation as not being directed to patentable subject matter.

A. Early Court Treatment of AI (following Alice)

In the early years following Alice (since mid-2014), courts analyzing AI-related patents (and, more generally, software-related patents) typically conducted Alice analyses in the early/pleading stages of a case. Indeed, in the few months following Alice, Judge Mayer, writing in a concurring opinion, encouraged this approach, stating that "[f]rom a practical perspective, there are clear advantages to addressing section 101's requirements at the outset of litigation" and that "unnecessary litigation … could have been avoided [here]" if the claims-at-issue had been invalidated under the two-part Alice test. *I/P Engine, Inc. v. AOL Inc., 576 Fed. Appx. 982* (Fed. Cir. 2014) (per curiam; unpublished) (Mayer, J., concurring). In that case, the claims described a system for filtering "information for relevance to a user's query using combined content and collaborative data," and were found to be invalid under 35 U.S.C. §101. *Id*.

District courts also followed suit, and many cases invalidated patents under 35 U.S.C. §101 at pre-trial stages for reciting high-level mathematical formulas and/or mental processes as performed by generic computing "modules" or by general-purpose computers.

In *Vehicle Intelligence & Safety LLC v. Mercedes-Benz USA, LLC*, the court applied the two-part Alice test to find patent claims invalid on a Rule 12(c) motion for judgment on the pleadings. 78 F. Supp. 3d 884 (N.D. Ill. 2015). The patent-at-issue generally described "expert systems" used "to screen equipment operators for impairments, such as intoxication, physical impairment, medical impairment, or emotional impairment" and to "control the equipment (e.g., automobiles, trucks, [etc.]) if impairment of the equipment operator is determined." While the patent-at-issue did not expressly define an "expert system," the court acknowledged that an expert system exemplified an "application of [AI]." Nonetheless, because the claims failed to "recite any new or improved computer technology or provide new physical components" beyond the mere replication of an "expert system" and its methods could be "performed entirely in the human mind." the claims were found to be directed to mere "abstract ideas." The claims also did not possess an "inventive concept" because the remaining claim elements comprised mere generic computer components without "significantly more." *Id*.

In *Blue Spike, LLC v. Google Inc.*, 2015 WL 5260506 (N.D. Cal. Sept. 08, 2015), the court also found the claims to be invalid via the two-part Alice test on a Rule 12(c) motion for judgment on the pleadings because the claims "merely discuss[ed] using routine computer components," despite the patent owner arguing that such invalidation could render "breakthroughs in artificial intelligence" categorically unpatentable. *Id*.

In *Neochloris, Inc. v. Emerson Process Mgmt. LLLP, 140 F. Supp. 3d 763* (N.D. Ill. 2015) the court found the claims to be invalid via the two-part Alice test on summary judgment, although at least one claim recited "an artificial neural network module," the court found that "it is not even clear [from the specification or claim itself] what [that term] refers to besides a [generalized] central processing unit—a basic computer's brain." *Id*.

In *Purepredictive, Inc. v. H2O.AI, Inc.*, the court applied the two-part Alice test to find patent claims invalid on a Rule 12(b)(6) motion to dismiss. 2017 WL 3721480 (N.D. Cal. Aug. 29, 2017). The patent-at-issue related to "an automated factory for predictive analytics." The claims recited a method that included generating "a plurality of learning functions" and involved selecting the most effective learned functions to create a rule set for further data input. The court found the claims to be directed to "a mental process" (an abstract idea), merely using "mathematical algorithms to perform predictive analytics." Because the claims failed to "described specific system architecture" and instead made reference to generic "modules," the court found the claims lacked an "inventive concept."

In *Power Analytics Corp. v. Operation Tech., Inc.*, 2017 BL 475384 (C.D. Cal. July 13, 2017), the court invalidated the claims via the two-part Alice test on a motion for partial summary judgment. Certain claims recited a "machine learning engine," but the court found the claims merely "describe[d] desired functions or outcomes, but [did] not, individually or in combination, constitute 'inventive concepts.'" *Id*.

B. Treatment of AI within the Last Five Years

The Federal Circuit has yet to analyze a patent by way of the two-part Alice test in which the patent at issue can be viewed as truly AI-focused. At least some of the judges have signaled caution when addressing questions about "important inventions" (such as AI). However, numerous majority opinions of the Federal Circuit continue to affirm early dispositive motions based on subject matter eligibility.

In *Smart Sys. Innovations, LLC v. Chicago Transit Auth.*, the majority affirmed a judgment on the pleadings finding the patent claims to be invalid because they were directed to the abstract idea of "the formation of financial transactions in a particular field (i.e., mass transit) and data collection related to such transactions." Judge Linn recognized that, "Despite the number of cases that have faced these [patent eligibility] questions and attempted to provide practical guidance, great uncertainty yet remains. And the danger of getting the answers to these questions wrong is greatest for some of today's most important inventions in computing, medical diagnostics, artificial intelligence, the Internet of Things, and robotics, among other things." 873 F.3d 1364, 1378 (Fed. Cir. 2017) (Linn, J., dissenting and concurring in part)).

Similarly, in *Athena Diagnostics, Inc. v. Mayo Collaborative Servs., LLC*, Federal Circuit denied a petition for rehearing *en banc* regarding claims directed to a "method for diagnosing neurotransmission or developmental disorders." 927 F.3d 1333 (Fed. Cir. 2019) (*per curiam*). In doing so, the court expressed its obligation (and concern) regarding following Supreme Court precedent in holding the claims ineligible. *Id*. (Lourie, J.) ("I concur in the court's decision not to rehear this case en banc. In my view, we can accomplish little in doing so, as we are bound by the Supreme Court's decision in Mayo. … Some of us have already expressed our concerns over current precedent"). In dissent, Judge Moore added: "I do not fault my colleagues, who under protest have concluded that they have no choice but to hold the claims in Athena ineligible because of Mayo." *Id*. Judge Moore also cited several amicus and senate hearings, including a discussion concerning AI-related inventions in the biotech space. *See id*. (citing The State of Patent Eligibility in America, Part II, 116th Cong. 9 (2019) (written testimony of Henry Hadad, President, IPO) ("[C]onfusion about what is patent-eligible discourages inventors from pursuing work in certain technology areas, including discovering new genetic biomarkers and developing diagnostic and artificial intelligence technologies. [This] uncertainty disincentivizes the enormous investment in research and development that is necessary to fuel the innovation cycle.")).

Just this year, in *Yanbin Yu v. Apple Inc.*, 2021 U.S.P.Q.2D 632 (Fed. Cir. June 11, 2021), the Federal Circuit affirmed a finding of invalidity made on Rule 12(b)(6) motion to dismiss. The majority found that the claims were directed to the abstract idea of "taking two pictures and using one picture to enhance the other." Because the claims only recited a two-lens, two-image-sensor configuration, while the Specification stressed the novelty of a four-lens, four-image-sensor system, the majority found that the claims' level of generality "merely invokes well-understood, routine, conventional components" and was not "sufficient to transform the claimed abstract idea into a patent-eligible invention." *Id*. Judge Newman issued a strong dissent, criticizing the majority for conflating the concepts of novelty and patent eligibility. *Id*. ("In the current state of Section 101 jurisprudence, inconsistency and unpredictability of adjudication have destabilized technologic development in important fields of commerce. Although today's Section 101 uncertainties have arisen primarily in the biological and computer-implemented technologies, all fields are affected. The case before us enlarges this instability in all fields, for the court holds that the question of whether the components of a new device are well-known and conventional affects Section 101 eligibility, without reaching the patentability criteria of novelty and nonobviousness.")

Recent district court decisions show a mixed bag of results when addressing the patent eligibility of computer related technologies, including AI.

For example, in *Singular Computing LLC v. Google LLC*, 2020 U.S.P.Q.2d 10708 (D. Mass. 2020), the court denied a 12(b)(6) motion to dismiss where the defendant alleged that the asserted claims were patent-ineligible via the two-part Alice test. The patent-at-issue generally related to "a processor or other device" that included "processing elements designed to perform arithmetic operations . . . on numerical values of low precision but high dynamic range [i.e., LPHDR]." The patent owner argued that such architecture allows "for more efficient use of a computer's transistors, which improves computer performance in certain applications such as artificial intelligence software." In denying the motion, the court focused on the alleged improvement (via LPHDR technology) to existing computer architecture, finding that such improvement made the claimed invention plausibly inventive, which was at least enough for the patent claims to survive at the pleading stage where allegations as to inventiveness were accepted as true.

On the other hand, in *Applied Predictive Techs., Inc. v. Marketdial, Inc.*, the court granted a Rule 12(b)(6) motion to dismiss finding patent claims invalid where the claims lacked AI-specific elements and were instead "directed to the abstract concept of optimizing parameters for business initiative testing." 2020 US Dist Lexis 221981 (D. Utah Nov. 25, 2020). Even though the specification described the possible use of a "neural network" to perform the claimed "testing," the court found that "no specific, technical details regarding the exact sort of analysis" was provided, and, therefore, the claims were considered to cover "history data analysis" broadly. "To grant a patent on the concept as it is articulated in the [patent-at-issue] would be to preempt any method of analyzing historical data to determine how certain parameter settings influence that data." *Id.*

In *Kaavo Inc. v. Amazon.com Inc.*, 323 F. Supp. 3d 630 (D. Del. 2018), the claims were invalidated for lack of subject matter eligibility using the two-part Alice test on a motion for summary judgment. Certain claims-at-issue recited methods of "forecasting" but lacked any AI-specific elements, even though the specification explained that such "forecasting" may be performed by "neural networks" or by "Load Forecasting"/"Pricing" modules, but the court found that "claim language [did] not require the use of these modules." *Id*.

Interestingly, the District of Delaware has been addressing multiple motions to dismiss based on patent ineligibility, in separate cases, in a single hearing, and issuing a single Order. For example, in *F45 Training Pty Ltd. v. Body Fit Training USA Inc.*, C.A. No. 20-1194-LPS, 2021

BL 250008 (D. Del. July 02, 2021), the District of Delaware gave mixed rulings on three such motions. In the first case, the court denied the motion, finding under step one of the Alice analysis that the claims were "directed to the abstract idea of storing, sending, and retrieving information over a network" and could be performed by a human, but under step two that there is was "a dispute as to whether the combination of elements, including use of the generic computer components, is well understood, routine, and conventional as of the date of the invention." *Id*. In the second case, the court also denied the motion, finding in step one of the Alice analysis that the defendant's ever changing proposed abstract idea oversimplified the claims directed to a physical device, namely "a portable, self-contained device for monitoring movement of body parts during physical activity." *Id*. In the third case, the court granted the motion, finding in step one that the claims were "directed to the abstract idea of collecting and storing milestone data for different locales and then presenting that data using standardized terminology," which "invokes computers as a tool, claiming only desirable results rather than a specific improvement to a computer-specific problem." *Id*. In step two, the court found that the claims failed because when the "only 'inventive concept' is the application of an abstract idea using conventional and well understood techniques, the claim has not been transformed into a patent-eligible application of an abstract idea." *Id*. (quoting *BSG Tech LLC v. Buyseasons, Inc.*, 899 F.3d 1281 , 1290-91 (Fed. Cir. 2018)).

C. Treatment of AI in the USPTO

The USPTO has also addressed AI related claims. Along with its 2019 Revised Patent Subject Matter Eligibility Guidance (the "2019 PEG"), which has been incorporated into the Ninth Edition, Revision 10.2019 (revised June 2020) of the Manual of Patent Examination Procedure (MPEP), the USPTO provided several example patent claims and respective analyses under the two-part Alice test. *See* Subject Matter Eligibility Examples: Abstract Ideas. One of these examples ("Example 39") demonstrated a patent-eligible artificial intelligence invention. In particular, Example 39 is labeled "Method for Training a Neural Network for Facial Detection" and includes claim elements for training a neural network across two stages of training set data so as to minimize false positives for facial detection. The USPTO's analysis informs that the claim of Example 39 is patent-eligible (and not "directed to" an abstract idea) because the claim does not recite any mathematical concept, mental process, or fundamental economic concept.

Additionally, the Patent Trial and Appeal Board (PTAB) has applied the 2019 PEG (as revised) in an *ex parte* appeal involving an artificial intelligence invention. *See Ex Parte* Hannun (formerly *Ex Parte* Linden), 2018-003323 (April 1, 2019) (designated by the PTAB as an "Informative" decision). 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 the PTAB 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.'"

While the USPTO's examples and decisions may provide a useful guide, caution should be exercised in relying on such guidance for litigation matters. More specifically, the Federal Circuit has already stated that such guidance, while useful, is not precedential. For example, in *Cleveland Clinic Found. v. True Health Diagnostics LLC*, the Federal Circuit addressed an appeal from the Eastern District of Texas, in which two patents had been found invalid. 760 Fed. Appx. 1013, 1014-15 (Fed. Cir. 2019). The Cleveland Clinic argued "that the district court failed to give the appropriate deference to [the USPTO Guidance]," but the Federal Circuit disagreed. *Id.* at 1020. The court stated: "While we greatly respect the PTO's expertise on all matters relating to patentability, including patent eligibility, we are not bound by its guidance." *Id.* As a result, the court concluded "that the district court did not err in its [lack of deference to] the PTO's subject matter eligibility guidance." *Id.* at 2021. The Federal Circuit has since reaffirmed its position in additional cases. *See, e.g., In re* Rudy, 956 F.3d 1379, 1382 (Fed. Cir. 2020) ("To the extent the Office Guidance contradicts or does not fully accord with [Federal Circuit] case law, it is [Federal Circuit] case law, and the Supreme Court precedent it is based upon, that must control.").

CONCLUSION

In preparing patent applications covering AI inventions, there are general inventor interview best practices as well as AI-specific factors that have been shown to be useful in distinguishing inventions. Questions should be directed to all aspects of AI including training method, inference method, architecture, and application. Some of the most fruitful disclosure to gather include: identifying technical problems and advantages, system inputs, data collection, pre- and post-processing steps, the network architecture, the training process, and the particular ways models are used to deliver results and are updated at run-time.

A major challenge of AI claim drafting is establishing patent eligibility. AI claim drafting best practices include claiming AI-specific architecture or training elements, and to show how a particular invention improves the efficiency, accuracy, speed, use of resources, or otherwise enhances, in a non-preemptive way, the computing systems or AI system on which it operates. Selection of the art unit in which a patent is prosecuted can also help improve outcomes.

In satisfying enablement, it is necessary to go beyond high level architecture ("image classification") and describe in detail lower levels of abstraction such as the operation of the network at the level of layers, nodes, and activation functions, or for training, the high-level training paradigm, such as supervised learning, unsupervised learning, or reinforcement learning. Because AI systems are often remotely implemented in a private setting, detectability is a concern when drafting claims. Documenting and claiming how an inventive system interacts with and acquires data from sources and endpoint devices, how it interacts with end-users, and how data flows are arranged can contribute to enablement as well as downstream detectability.

Because many AI systems are heavily dependent upon training data sets, they are susceptible numerous forms of statistical and other bias, which can lead to ethical risk issues and real-world bad ethical outcomes, such as user discrimination, privacy violations, or exploitations by malicious customers. Counsel detecting this during invention discovery may need to connect their client with an AI Ethics Partner, when there is evidence that stakeholders outside the inventive process (project managers, marketing personnel, etc.) may not grasp the full legal and ethical consequences of their design choices.

Different countries and jurisdictions are at different stages of development, awareness and jurisprudence regarding AI patent protection. Many of these have provided examples of acceptable AI-related patents, or guidelines for dealing with these issues. While potentially helpful, these patent-office decisions are generally not given deference in court.

Regarding inventions that are created by employing artificial intelligence, applicants are advised to trace conception of potentially novel aspects back to a human inventor, and to pursue patent applications that only name humans that contributed to what is claimed, and to avoid having to deal with the inventorship question until current inventorship laws and regulations are further developed.

Because AI is still a relatively newer technology, there have not been many cases involving AI-related patents in district courts. Nearly all the patent cases involving AI have been invalidated in early stages of litigation by the district courts as not being directed to patentable subject matter.

The Federal Circuit has yet to analyze a patent by way of the two-part Alice test in which the patent at issue can be viewed as truly AI-focused. Nevertheless, some of the judges have signaled caution when addressing questions about "important inventions" (such as AI).

At the USPTO, examiners have granted numerous patents involving AI since the publication of the 2019 Revised Patent Eligibility Guidance (2019 PEG). However, the Federal Circuit has stated that such guidance, while useful, is not precedential.