# Patent

## Best Practices in Drafting AI Patent Applications

CHAU | IP

2022

## 1. 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.

## 2. 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]

[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.

## 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.