

# [Patent Protection for Artificial Intelligence and Machine Learning](#)

**Go to:** [AI Basics](#) | [Trends in Patenting AI](#) | [Drafting Patent Applications Directed to AI](#) | [Prosecuting Patents Directed to AI](#)

**Maintained**

by [Kirk A. Sigmon](#), Banner Witcoff

This practice note discusses patenting artificial intelligence (AI), machine learning (ML), and related inventions. The note provides a high-level overview of AI and ML, provides tips for drafting a patent application directed to inventions relating to AI and ML, and discusses trends and strategies for handling prosecution of such inventions.

See [Patent-Eligible Subject Matter \(Section 101\) Statements of Law](#) and [Patent Law Fundamentals Resource Kit](#).

For an introduction to U.S. patent law, see Patent Fundamentals and [U.S. Patents Q&A Checklist \(Patent Fundamentals\)](#).

## **AI Basics**

As a preliminary matter, it is important to distinguish between AI in the general sense, ML, deep learning, and other commonly used terms in the field.

### **AI**

The term artificial intelligence generally refers to causing computing devices to perform human-like thinking. The phrase has been used in patent applications for decades, though historically, few computers could do anything remotely approximating human-like thinking. In fact, many patent applications seemed to use the term like a marketing mechanism, extolling the virtues of particular algorithms and/or processes.

### **ML**

ML models, a subset of artificial intelligence, are one of the latest forms of algorithms that enable computers to approximate human-like thinking. ML models are often configured (i.e., trained) through large quantities of data—often referred to as training data—to learn, through that data, to perform particular tasks. While the term machine learning is also quite old (and was used as early as the 1960s by computer scientist Arthur Samuel), it was historically somewhat infeasible, and modern computing devices permit ML model implementation on even consumer-grade hardware. Stated more plainly, the world has been trying to do ML for a long time, but modern hardware makes it significantly easier to do so.

One of the most promising implementations of ML models comprises so-called deep learning, using artificial neural networks that are intentionally designed to mimic the human brain. Such an approach is computationally costly, but can result in some amazing results: for example, the famous ChatGPT algorithm uses deep learning in a manner that allows it to answer questions realistically.

### **Natural Language Processing (NLP)**

It is not uncommon for AI and ML to be associated with natural language processing (NLP), which relates to algorithms that process (i.e., understand, output) human communications (e.g., human-written text, conversations, and the like). For example, NLP might be paired with a trained ML model such that a user can provide natural language input, that input can be processed into appropriate input data for a trained ML model, and then the input data can be provided to input nodes of the trained ML model. As another example, many NLP implementations use trained ML models for the purposes of translation, sentiment analysis, and the like. With that said, not all AI is NLP,

## Patent Protection for Artificial Intelligence and Machine Learning

and not all NLP is AI. For example, one might argue that an algorithm configured to remove stop words (e.g., "the," "is," "are") from text is an NLP algorithm, though such an algorithm does not involve AI.

### **Trends in Patenting AI**

There has been a veritable gold rush for AI-related patents in recent years, in no small part due to the rapid popularization and convenience of ML and deep learning. Along those lines, while the term "artificial intelligence" has been used in patents filed as early as 1970, modern developments—particularly, the rapid popularization of artificial neural networks that can be executed (albeit sometimes poorly) on commercial hardware—have resulted in tens of thousands of AI-related patent applications being filed per year. Moreover, many patent applications which are not outright directed to AI sometimes contain features that might be implemented using some form of AI.

This trend can be, from a patent practitioner's perspective, exciting but risky. On one hand, patent prosecutors find themselves busier than ever with AI-related work, and many find themselves specializing in the field to some degree. On the other hand, prosecuting AI-related patents is increasingly difficult, as the U.S. Patent and Trademark Office (USPTO) has seemingly applied more scrutiny to such applications.

### ***Two Major Types of AI Inventions***

You should be careful to distinguish the two key types of AI-related inventions. The two different types require surprisingly different approaches, even if both might at a high-level relate to AI.

#### *Inventions Improving AI/ML Itself*

The rarest, but perhaps most pure, form of AI-related invention is the improvement to AI itself. These inventions are characterized by relating to improvements to artificial neural networks and/or ML models themselves; the hardware capable of implementing such artificial neural networks and/or ML models; the process by which nodes in those artificial neural networks and/or ML models are trained; or the like. It is often easy to identify these inventions, as they typically involve complex math, lengthier disclosure calls, and are often designed to be input- and output-agnostic (that is, an invention improving an artificial neural network is often designed to improve that artificial neural network in a variety of use cases).

As will be detailed below, these inventions can involve relatively straightforward prosecution. They typically land in an AI-related art unit, are typically examined by examiners familiar with the technology, and are somewhat less likely to face battles over subject matter eligibility under [35 U.S.C. § 101](#).

#### *Inventions Using AI/ML*

The more common but more difficult form of AI-related invention is the invention that uses AI or ML in some other field of endeavor. In other words, these inventions improve some other field of endeavor by, for example, replacing one or more steps of a process with AI. As an illustration, use of an ML model to improve video post-processing would likely fall into this category, as would use of an ML model as part of managing a control system.

These inventions have faced substantially more onerous prosecution difficulties in recent years, and some examiners at the USPTO have openly admitted that they have been instructed by their superiors to treat these inventions with greater scrutiny. These applications often end up in a non-AI-related art unit, are typically examined by examiners with limited knowledge of AI, and often face substantial pushback regarding subject matter eligibility under [35 U.S.C. § 101](#). Moreover, these applications often face quite strong rejections under [35 U.S.C. § 103](#), especially given the veritable goldrush for AI-related inventions that has resulted in a bevy of potential prior art.

### **Drafting Patent Applications Directed to AI**

## Patent Protection for Artificial Intelligence and Machine Learning

Regardless of which of the types of AI-related invention you have, many aspects of the drafting process are largely similar, and mainly hinge on ensuring that AI is described in a way that cannot be misconstrued by an examiner or potential litigation target. See [Patent Litigation Fundamentals](#).

### ***Prior Art Searching and AI Inventions***

Given the current AI boom, the value of prior art searches cannot be understated. This is especially the case for inventions using AI, as the internet is replete with so-called "wouldn't it be cool" discussions where technologists speculate as to how AI could improve various industries.

A strong prior art search for AI-related references often begins with a search of both internet sources and patent databases. For the first type of invention (inventions improving AI itself), the best sources of potential prior art include research papers and other university publications, patent applications, and documentation websites. In contrast, for the second type of invention (inventions using AI), the internet as a whole often provides more possible prior art, particularly since many companies' efforts in using AI are reflected in marketing materials, and particularly since many technologists often post so-called "wouldn't it be cool"-type articles and discussions.

You should exercise caution when relying on inventors to know about potential prior art, especially for inventions using AI. Many companies and universities have strongly incentivized inventors to seek out ways to implement AI in existing processes, and this leads many inventors to inadvertently reinvent what others may have already explored or implemented. Of course, this does not necessarily mean that both sets of inventors invented the same solution. If you do discover such a problem, you should thoroughly investigate whether the approaches are truly the same, or if there are distinctions to the approaches which might be sufficient for a showing of novelty and non-obviousness. See [Obviousness Rejections: Attacking the Prima Facie Case](#) and [Obviousness Rejections: Rebutting the Prima Facie Case](#).

Prior art searches are often a useful way to remind inventors that their inventions must necessarily be more than "do this, but with AI" and must be more than what others have invented previously. In particular, during disclosure calls, it can be helpful to openly discuss prior art issues with inventors and encourage them to explore their idea deeply. In the circumstance where you find a pertinent reference and want to encourage inventors to think about their invention more deeply, useful questions for those inventors include the following:

- What is different between your AI/ML approach and this reference?
- What difficulties did you experience in implementing AI/ML that are not remedied by this reference, but are remedied by your invention?
- If you chose a particular type of ML (e.g., using labeled or unlabeled data, using a particular type of model), what about it makes it appropriate for your problem?
- Is there something special about your input data or output data that is different from this reference?

### ***Specification Drafting for AI Inventions***

When drafting an AI-related application of either AI invention type, it is important to not reinvent the proverbial wheel. It is all but guaranteed that your inventors did not invent the concept of AI or ML models; as such, laborious re-explanation of these concepts often serves very little purpose and is often ignored by USPTO examiners. Practically speaking, this means that it can be helpful to provide a brief description of AI or ML in a patent application, but it might be somewhat wasteful to spend numerous pages laboriously re-explaining the concept to the USPTO.

The converse applies as well: you should not under-explain AI/ML in your patent application. A common mistake made by many attorneys drafting AI-related patent applications (especially the second type of invention, those involving inventions that use AI) is that they under-explain how AI actually works in their system, effectively treating AI as a black box that does little more than pop out desired answers upon demand. This can be fatal to the patent application, as it effectively invites a host of rejections. For example, it makes it easier for the invention to be rejected with prior art under [35 U.S.C. §§ 102 and 103](#), it makes the black box look similar to a human mind in a

## Patent Protection for Artificial Intelligence and Machine Learning

way that invites a [35 U.S.C. § 101](#) rejection, and could potentially garner rejections under [35 U.S.C. § 112](#). To avoid such issues, the specification should be drafted to clearly explain (to the extent possible) how the AI is implemented (e.g., how the ML model is trained, what type of model(s) are used, what sort of data is used in training, whether such data is labeled or not), what type of input is provided to the AI (e.g., an example of the data structure provided to the trained ML model), and what sort of output is expected from the AI (e.g., a Boolean value, some sort of selection within the data, some sort of percentage figure).

Another common mistake made during drafting is to describe AI doing human-like thinking. This mistake, in conjunction with similar language in the claims, often invites subject matter eligibility rejections under [35 U.S.C. § 101](#), as the tenor of the specification leads the examiner to conclude that the AI is just performing steps that a human being could do mentally. Along those lines, it can be very helpful to review the draft specification and replace phrases that suggest thinking on the part of the AI (e.g., "the artificial intelligence may decide whether to proceed," "the artificial intelligence may weigh various considerations") with phrases that are more technically precise (e.g., "the artificial intelligence may output a Boolean value and, based on that Boolean value, the system may proceed," "the artificial intelligence may be provided, as input, considerations such as . . ."). Remember, although the point of AI might be to cause computers to think like a human being, practically speaking they are implemented using data and often perform quite rudimentary steps.

### ***Claim Drafting for AI Inventions***

In part due to the newness and perceived complexity of AI, many claims are unfortunately imprecise when describing it. This can be a fatal flaw in a variety of ways: it can invite extremely easy rejections under [35 U.S.C. § 103](#) and, in some cases, can make proving infringement quite difficult.

For the most part, you should not draft claims that describe AI as little more than a black box algorithm. For example, the following illustrative claim language, without more, arguably renders the ML model as little more than a black box algorithm:

*determining, via a machine learning model, whether to trigger an alarm;*

Even if the concept of triggering such an alarm based on ML output is entirely new, this structure makes a [35 U.S.C. § 103](#) rejection very easy for the examiner. After all, the examiner need only show that some reference teaches "*determining . . . whether to trigger an alarm*" based on some algorithm and combine such a reference with a second reference indicating that ML models might be used instead of conventional algorithms. Stated differently, in circumstances where it is very easy to treat ML models as a conventional algorithm, examiners will do exactly that.

You should also avoid drafting claims that potentially describe AI performing steps that AI cannot do without additional actions (that is, "AI magic"). For example, the following illustrative claim language, without more, assumes that a trained ML model is capable of performing a step that it likely cannot do:

*triggering, by a machine learning model, an alarm based on input data;*

The key distinction here is that, while the output of an ML model certainly might be usable to trigger an alarm, the ML model itself likely cannot do so; at most, it probably outputs a Boolean or similar value indicating whether an alarm should be triggered, and that Boolean is likely used by some computer to ultimately cause triggering of the alarm.

Another potential mistake is to try to describe ML concepts in a single step. Cramming all aspects of an ML model into a single step (e.g., such that the active step focuses on output data, and the input and training steps are reduced to wherein clauses) might be desirable in some circumstances, but it can downplay key aspects of the ML model that define over conventional algorithms. Put differently, explicitly reciting multiple active steps involving AI can help avoid an examiner hand-waiving involvement of that AI in an overall process.

## Patent Protection for Artificial Intelligence and Machine Learning

A preferred approach is to describe AI in a way that contextualizes it and that involves steps that are unique to AI. For example, in the context of an ML model, it can help to include as many of the following steps as possible:

1. **Training an ML model.** As discussed above, ML models are a form of AI that can be distinguished from conventional algorithms in that they are trained in a variety of ways (e.g., supervised or unsupervised learning, using labeled or unlabeled data). Describing how these ML models are trained helps underscore the fact that the ML models are not conventional algorithms.
2. **Providing formatted input to the trained ML model.** Once the ML model is trained, it can receive data (e.g., via input nodes). It can be very beneficial to describe this step explicitly and to provide sufficient detail regarding the particular format of the input data, as doing so can be valuable ammunition against a [35 U.S.C. §§ 102 or 103](#) reference relying on a particular algorithm.
3. **Receiving particular output from the trained ML model.** After input is provided, a trained ML model can provide some form of output, such as an identification of some subset of the input data, a Boolean value, or the like. Remember, a trained ML model is not magic. The output from such a model is typically not a lofty concept (e.g., a detailed natural language analysis of why some input data is fraudulent), but is instead usually more discrete and objective (e.g., a likelihood, reflected by a percentage value, that data is similar to previous fraudulent data).
4. **Utilizing the output.** Output from the ML model should rarely be the last step of the claims, as this would mean that the claims merely culminate in receipt of data without contextualizing why such data is useful. Instead, it can help to do something with the output data above and beyond receiving/displaying it. Where possible, it can be particularly helpful to show how the output data causes some real-world change: triggering of an alarm, modification of speed, movement of some object, etc.
5. **Retraining the ML model.** Recent interviews with numerous examiners across numerous art units suggests that this is the new "gold standard" for strong ML-based claims. Specifically, examiners like to see that an ML model is further trained based on later activity in a claim, such as user feedback regarding whether the ML model's output was correct. This effectively creates a feedback loop using the ML model that is far beyond what is doable with conventional hand-programmed algorithms, making it quite persuasive against a [35 U.S.C. §§ 102 or 103](#) rejection. Moreover, this provides further ammunition against a [35 U.S.C. § 101](#) rejection, as it allows you to argue that one benefit of the claims is that the ML model—that is, a computer-implemented algorithm—is improved over time.

Given the above, a better way to rewrite the above claim step might be as follows:

*[1] generating a trained machine learning model by training, using training data comprising a history of alarms and associated operating status data, a machine learning model to output, based on input operating status data, an indication of whether an alarm should be triggered, wherein training the machine learning model comprises modifying one or more weights of one or more nodes of an artificial neural network;*

*[2] providing, to the trained machine learning model, input data comprising current operating status data;*

*[3] receiving, from the trained machine learning model, output data, based on the input data, comprising a value that indicates that an alarm should be triggered;*

*[4] triggering, based on the output data, the alarm;*

*[5a] receiving, via a user interface, user input associated with the alarm; and*

*[5b] further training, based on the user input, the trained machine learning model.*

Note that the above is rough and illustrative, and might be modified in a variety of ways based on the particularities of the application. For example, the above claim is very broad regarding the retraining step—it might be preferable to be far more specific in certain circumstances. As another example, it might be desirable in some circumstances to broaden the "generating" step, and in any event some examiners might object to the structure of the "generating" step (particularly the description of what the ML model is being trained to do) as possibly unclear or conclusory.

## Patent Protection for Artificial Intelligence and Machine Learning

In addition to the above, prior to drafting a claim directed to any form of AI, it can be helpful to review [Example 39 of the USPTO's Subject Matter Eligibility Examples: Abstract Ideas](#) and to, where possible, emulate the structure and/or approach of the claim. This example is particularly valuable as a defense against a [35 U.S.C. § 101](#) rejection, and similarities to the example claim in question can help persuade an examiner to withdraw such a rejection.

### ***Special Consideration – Multiple ML Models***

For the most part, a single ML model is trained to provide specific output based on specific input. As such, absent extraordinary circumstances, it is generally a mistake to describe a system whereby the same ML model is expected to perform entirely different tasks (e.g., determine whether an alarm should be triggered and then identify dogs in photos).

Since ML models tend to be task-specific, it is not uncommon for inventions to use multiple ML models, for example, different models for different parts of an overall process. In such circumstances, you should be careful to distinguish between these models in the claims. Use of ordinals such as "first training data" and "second training data," "first output" and "second output" can help immensely. Inadvertently suggesting that two different ML models provide the same input, the same output, and/or are trained using the exact same data could, in extreme circumstances, result in claims that are, in practice, easy to design around.

Nonetheless, use of multiple ML models can be valuable ammunition against a wide bevy of USPTO rejections. Even if an examiner can find a reference describing use of AI in a general sense, the use of multiple ML models is much less likely to be described in such a reference. Moreover, the use of multiple ML models can, in the context of [35 U.S.C. § 101](#) rejections, be valuable storytelling to explain why the claims are far more than a rudimentary implementation of a mental process.

### ***Special Consideration – AI as an Inventor***

One common topic among the technology world at large is the idea that technology might itself invent and create. At the present moment, this concept is largely theoretical, though it has already been litigated.

Title [35 U.S.C. §§ 100–101](#) define an "inventor" as someone who "invented or discovered the subject matter of the invention." The Federal Circuit has, when considering this requirement, already concluded that an AI cannot be an inventor, and instead all inventors must be natural persons. [Thaler v. Vidal, 43 F.4th 1207 \(Fed. Cir. 2022\)](#). As of the writing of this guidance, it appears that plaintiff Thaler is considering whether to petition for a writ of certiorari to the Supreme Court regarding this issue. This ruling is approximately in line with [the U.S. Copyright Office's recent guidance on the registration of AI-generated copyright](#), which generally maintains that copyright protects only the material that is the product of human creativity.

Notwithstanding a potential overturning of the *Thaler* decision by the Supreme Court, you should exercise extreme caution when human inventors assert that an AI was an inventor. Practically speaking, such assertions are often based on the idea that an AI was somehow used by the inventor(s) during the process of invention, meaning that the AI was little more than a helpful tool used during ideation. In the same way that an integrated circuit designer is not required to credit their microchip design software for helping them design a new processor, an inventor is not required to credit AI as a co-inventor simply because it was used at some stage of the ideation process.

In short, while inventors might be excited about the idea of AI being a co-inventor from a novelty perspective, this assertion should be discouraged because it is likely inaccurate and could potentially cause the patent application to be rejected.

### **Prosecuting Patents Directed to AI**

While AI inventions were once somewhat easy to prosecute before the USPTO, they have become increasingly difficult to prosecute in recent years. Interviews with examiners indicate that they may have explicit instructions to

## Patent Protection for Artificial Intelligence and Machine Learning

treat AI-related inventions with increased scrutiny, which is likely a trend borne of the aforementioned gold rush for AI-related patents. In turn, for many AI-related inventions (especially those of the second type, inventions that use AI), you should expect a battle.

***Handling Subject Matter Eligibility Rejections***

One of the most common and perhaps most frustrating rejections faced by AI inventions is a [35 U.S.C. § 101](#) rejection. These rejections can be extremely difficult to address, in no small part because some examiners use them as a way to prevent allowance of a patent even for the narrowest of claims.

One of the biggest issues facing AI-related inventions is definitional: AI is generally designed to mimic human thinking, and courts have explicitly held that "mental processes," including "concepts performed in the human mind (including an observation, evaluation, judgment, opinion)" are abstract ideas that satisfy the first prong of the USPTO's eligibility step 2A for a [35 U.S.C. § 101](#) rejection. See MPEP § 2106.04(a). In other words, many examiners reviewing AI-related inventions will consider the invention in view of steps that could be performed in the human mind. This approach often dooms AI-related inventions to at least a perfunctory [35 U.S.C. § 101](#) rejection, particularly where the recitation of an AI algorithm in a claim could be analogized to the involvement of a human being.

Examiners also commonly evaluate inventions involving AI by removing AI-related portions of the claims and evaluating whether the remainder of the claim is sufficiently technical. For example, examiners sometimes assume that AI steps could be replaced by a human being and evaluate whether, in view of such a substitution, the remainder of the claim recites anything more than conventional computing hardware. Such an analytical approach can be devastating to many AI-related inventions, as most implementations of AI (e.g., ML models) are designed to be run on standard commercial hardware.

In turn, if an application claims AI doing nebulous, human-like thinking without sufficiently being rooted in a technological environment, that application will likely be rejected under [35 U.S.C. § 101](#) both because (1) the AI can be analogized to a mental process performed by a human being and (2) the remainder of the claim recites little more than conventional computing hardware. Traversing such [35 U.S.C. § 101](#) rejections often hinges on how well claims (1) describe AI in a way that excludes human beings and conventional algorithms and (2) define AI in the context of an overall technological environment.

For inventions involving improvements to AI itself (the first type of AI invention discussed above), overcoming [35 U.S.C. § 101](#) is often quite straightforward: the claims must clearly indicate how specific steps improve the functioning of AI itself, rather than some overall decision-making process. In the context of ML, claim amendments intended to overcome a [35 U.S.C. § 101](#) rejection might preferably be focused on specifically describing aspects of the ML model (e.g., nodes, weights), how those aspects are interrelated (e.g., specific mathematical functions, specific training approaches), and how the overall process improves the functioning of the ML model (e.g., by making the model more accurate, faster, more efficient, or the like).

For inventions using AI (the second type of AI invention), the process becomes significantly harder. The AI must generally be:

- Used in a manner that cannot be readily analogized to human thinking or simplistic algorithms –and–
- Placed into an overall technological context

***Avoid Analogy to Human Thinking and Simplistic Algorithms***

To explain how claims directed to inventions using AI are used in a manner that cannot be readily analogized to human thinking or simplistic algorithms, it is particularly helpful to focus on the unique technical aspects of AI. Examples in the ML context include the weighting of nodes in an artificial neural network, the retraining of an existing ML model, and the deploying of an already-trained ML model to different computing devices. After all, while it might be easy to argue that a human mind can be trained to identify dogs in photographs, it can be quite difficult

## Patent Protection for Artificial Intelligence and Machine Learning

to argue that a human mind can weight nodes using training data, use those nodes to identify dogs, and then re-weight those nodes based on subsequent indications of whether the dogs were correctly identified.

### *Place Claims in Technological Context*

To place claims directed to inventions using AI in an overall technological context, it often helps to focus on steps both preceding and following use of AI. For example, input data provided to an AI should be placed into context: it can help to explain what generated the input data, what the input data contains, how the input data was preprocessed for consumption by the AI, and so forth. As another example, output data generated via an AI should also be placed into context: the claims should clearly detail how the output data is being subsequently used, even if such use is little more than output via a user interface.

Where possible, it also helps to reference [Example 39 of the USPTO's Subject Matter Eligibility Examples: Abstract Ideas](#). Admittedly, this approach has limits. The USPTO places significant weight on the fact that the example claim involves an "iterative training algorithm" involving two stages of neural network training to improve neural network accuracy by minimizing false positives, meaning that examiners might not be persuaded that this example applies to claims that do not involve similar, multistep algorithm improvements. All the same, Example 39 is a valuable weapon against stubborn examiners who refuse to concede the subject matter eligibility of AI as a whole.

### **Strategies to Overcome [35 U.S.C. § 101](#) Rejections**

In reality, overcoming a [35 U.S.C. § 101](#) rejection of claims directed to an AI-related invention often involves substantial amounts of storytelling, rather than technical argument. Most examiners are preliminarily quite skeptical of AI-related inventions, and tend to be much more comfortable with an invention once they understand its overall context. Along those lines, helpful strategies include the following:

- **Conduct examiner interviews.** Ideally, these interviews should not be formulaic (e.g., "let's walk through the subject matter eligibility analysis step-by-step") but thematic (e.g., "this is why this AI invention is new, cool, and computer-oriented"). These interviews are also ideally conducted with every office action, and ideally after receiving an office action and preparing remarks but before any amendment(s)/remarks are filed. This approach builds a friendly rapport with the examiner that allows you to make a persuasive case and understand the examiner's concerns more fully without turning prosecution into an aggressive battle of briefs.

For resources that help a patent prosecutor respond to an office action in a pending patent application in the USPTO, see [Patent Office Action Response Resource Kit](#).

- **Do not waste time on weak arguments.** Chances are, examiners will not deviate from the general concept that AI could be performed in the human mind, and thus will not be willing to budge under the first prong of the USPTO's eligibility step 2A for a [35 U.S.C. § 101](#) rejection. Rather than nitpicking this argument, it is much better to remind the examiner that "[a]t some level, all inventions embody, use, reflect, rest upon, or apply laws of nature, natural phenomena, or abstract ideas," and that the Supreme Court has cautioned "to tread carefully in construing this exclusionary principle lest it swallow all of patent law." [Alice Corp. Pty. Ltd. v. CLS Bank Int'l, 573 U.S. 208, 216, 110 \(2014\)](#). See also [MPEP § 2106.04](#); [Enfish, LLC v. Microsoft Corp., 822 F.3d 1327 \(Fed. Cir. 2016\)](#).
- **File [37 C.F.R. § 1.132](#) declarations.** In certain circumstances, it can be helpful to have an expert or inventor prepare an affidavit or declaration under [37 C.F.R. § 1.132](#). These declarations, admittedly somewhat rare during conventional patent prosecution, can be helpful to address rejections and provide a venue for a non-attorney to present persuasive arguments against a rejection. MPEP § 716. For example, if an examiner expresses skepticism regarding the idea that an AI could not be replaced by a human being, and in circumstances where such an explanation is not readily available in the specification, it can be helpful to have an expert, inventor, or other third party prepare and file a declaration explaining why the AI could not be replaced by a human being.

### **Strategies for Addressing Prior Art**

## Patent Protection for Artificial Intelligence and Machine Learning

A large swath of references in the computer science and electrical engineering fields already disclose AI, though the quality of such disclosures can vary widely. For instance, it is not uncommon for many patent applications to include generic phrases asserting the use of ML models or AI, though the applications rarely explain how such a use would in fact be implemented. Examiners commonly use these references because such broad references to AI can provide a justification for combining those references with other AI-related references under [35 U.S.C. § 103](#).

Simply put, inventions that boil down to little more than "do this, but with AI" are not likely to survive art rejections. Even if an examiner cannot find some reference in a particular field that explicitly discusses the use of AI (or ML or deep learning), the examiner could likely argue under [35 U.S.C. § 103](#) that a pertinent algorithm could be replaced with AI of some kind.

In the context of ML, there are often a number of different strong strategies for overcoming art-based rejections:

- **Focus on the training of the ML model.** Differently trained ML models can produce widely different results. Examiners nonetheless often argue that ML models are equivalent despite being trained in very different ways. Do not allow examiners to make this assumption. After all, your approach might be similar, but might be sufficiently different so as to result in significantly more accurate output, it might be faster, or the like. As a simple example, if your claims involve training an ML model based on a history of real-life network outages and a history of applications executing on a server, an ML model trained based on documentation of those applications would not be the same.
- **Focus on the input and output data.** Perhaps obviously, different formats of input to an ML model can produce different forms of output. Examiners often overlook this distinction, instead focusing on the net result of ML. For example, examiners commonly argue that different ML models directed to virus detection are the same, even when the ML models consider widely different variables and provide extremely difficult output. Make sure that you explicitly refute any assertion that two different types of input or output are analogous.
- **Focus on context.** For the second type of AI invention (inventions using AI/ML), it is particularly important to focus on the overall context via which AI/ML was used. In the context of ML models, it can be particularly helpful to emphasize how the ML models' output is used: for example, how it is used to effectuate some change. Where applicable, it can be especially helpful to concentrate on retraining steps, which help focus the examiner not only on the output itself, but how the output is ultimately used.

It is not unusual for the above arguments to require significant claim amendments, even when the examiner's art is weak. As with many other inventions, it can be helpful to regularly check in with inventors to confirm that amendments remain faithful to the original invention.

*Current as of: 04/19/2023*