ALM | LAW.COM
LAW JOURNAL NEWSLETTERS

THE INTELLECTUAL PROPERTY STRATEGIST

Can Artificial Intelligence Patents Overcome §112 Requirements?, Part 2

By Mark Liang, partner; Paige Hardy, associate; Grace McFee, associate
O'Melveny & Myers LLP

February 2024

As explained in Part I of this series, ***Can Artificial Intelligence Patents Survive Alice?***, while the last decade has seen a dramatic increase in the number of AI patents, such patents face difficulty in overcoming the patent-eligibility challenges under §101 and *Alice*. Section 101, however, is not the only hurdles AI patents must overcome. Section 112, with its written description, enablement, and definiteness requirements, presents additional obstacles.

Particularly in the wake of the Supreme Court's recent *Amgen Inc. v. Sanofi* decision that brought §112 back to the forefront of patent law, AI patents face multiple challenges under that provision. *See generally,* 598 U.S. 594, 614 (2023). AI software patents risk running afoul of §112 requirements if they claim black box concepts as "machine learning" or "neural networks" without elaborating the specific details of their implementation in the specification, or if they substitute a POSITA's knowledge for actual disclosure. Because AI patents might allow an inventor to obtain legal protection over a nebulous self, computer-trained synthesis of correlation and prediction, there are likely to be challenges to the scope of an AI patent's claims. *See,* Sameer Gokhale, "**AI and Written Description: When Does an AI Patent Claim Cross the Line?**," Oblon (last accessed Jan. 22, 2024). And as more and more AI patents are litigated, there are bound to be more §112 disputes concerning such patents and the AI technologies that they claim. So just as *Alice* §101 challenges have increased with the growth in AI patents, §112 challenges are likely to do the same.

While some have claimed that the recent resurgence of §112 is "a cataclysmic shift in the law of enablement," that remains to be seen — especially in the context of AI patent claims. *See,* Gene Quinn, "**SCOTUS Ushers in New Era of Enablement Law in Amgen Ruling**," IPWatchdog,

Mark Liang is a partner in O'Melveny's Intellectual Property & Technology Group, based in San Francisco. His practice focuses on patent and technology-related litigation. He has experience in all stages of patent cases, including claim construction, fact and expert discovery, dispositive motions, trial, post-trial, and appeal.
Paige Hardy is a litigation associate in O'Melveny's Los Angeles office with experience representing clients in patent litigation before the Patent and Trial Appeals Board and in federal appellate courts.
Grace McFee is a litigation associate in O'Melveny's Los Angeles office with a practice that focuses on intellectual property and technology.

(last accessed Jan. 22, 2024). This article considers how courts have treated AI-related patents in the context of §112(a) written description and enablement and §112(b, f) for means-plus-function limitations. Do not fret: We also provide illustrative hypotheticals and some practice pointers on both issues.

112(a) Written Description and Enablement

Section 112(a) sets out that the patent specification must provide a "written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same." If an inventor claims a lot but describes or enables a little, the public does not receive the benefit of the patent bargain. *See, Bonito Boats, Inc. v. Thunder Craft Boats, Inc.*, 489 U.S. 141, 150 (1989). The Supreme Court emphasized this principle in *Amgen*, holding that Amgen had failed to enable all that it claimed — in that case, an entire genus of antibodies with a particular function—because its specification disclosed only certain examples, forcing a POSITA to engage in "painstaking experimentation" to make other, undisclosed antibodies under that genus. 598 U.S. at 614. Though this decision may be interpreted to apply only to genus-species claims, it suggests that enablement and written description may be more successfully challenged by accused infringers going forward.

As detailed below, applying §112(a) and *Amgen* to AI patents, courts have found that: 1) AI patents without detailed specifications and explanations for claimed AI features are unlikely to survive a §112(a) challenge, but; 2) in some cases, expert testimony about a POSITA's knowledge of existing AI technologies might save the claims of a patent with vague or undetailed specification.

112(a) and Sparse Specifications

In applying §112(a), courts commonly invalidate patent claims that recite AI-related features without sufficient detail about their operation and function.

For example, in *Impact Engine, Inc. v. Google LLC*, the claims recited an AI software component described as a "compiler" or "compiling engine" in the context of patents directed to integrating ads with media for webpages. 624 F. Supp. 3d 1190, 1195 (S.D. Cal. 2022). But the specification's references to the "compiler component" were limited to "(i) a box in Figure 1 labeled compiler (116) with no corresponding reference in the text of the specification as to the purpose and function of the compiler, and (ii) Fig. 3, a flow chart depicting a method for template customization." *Id.* The court found this description insufficient, reasoning that while "[a] compiler in the computer arts at the time the patent was filed was a program that translates source code into machine or object code … [the] patent does not disclose any information or mechanism that would inform a person of skill in the art how a compiler as construed in this patent would group the claimed design layers, design elements and content containers into a collection of slides to generate a communication." *Id.* Because the "compiler" was an abstraction, with nothing disclosed about how it performed the functions required of it in the claims, the court held that the claims were invalid based on lack of enablement and written description. *Id*.

In re Starrett is an example of a court applying *Amgen* to AI patents, as the claims purported to cover broader subject matter than what was disclosed in the specification. 2023 WL 3881360 at *4 (Fed. Cir. Jun. 8, 2023). In that case, the claims recited "machine learning" functionality for certain telepathic functions. *Id.* But the court held the claims were "rife with broad, vague concepts," declaring that "[t]he application's disclosure of a broad and abstract organizational structure used to accomplish the maintenance of augmented telepathic data amounts to little more than a 'research assignment' requiring a skilled artisan to undertake undue experimentation to discover what types of devices are encompassed by the claim limitations and how they would function." *Id.*

112(a) and POSITA Knowledge

It is well established that a POSITA's knowledge should be considered, and can partially provide the basis, for written description and enablement. *See, Streck, Inc. v. Rsch. & Diagnostic Sys., Inc.,* 665 F.3d 1269, 1284-1292 (Fed. Cir. 2012) (considering inventor and expert testimony in determining whether a patent satisfied the written description and enablement requirements); *see also, Eli Lilly & Co. v. Teva Parenteral Medicines, Inc.*, 845 F.3d 1357, 1371-72 (Fed. Cir. 2017) ("Dr. O'Dwyer testified that, although 'vitamin B12' can refer to a class of compounds in other contexts, it refers specifically to cyanocobalamin when 'vitamin B12' is prescribed in the medical field …. We see no clear error in the district court's acceptance of the understanding that 'vitamin B12,' when used to refer to vitamin B12 supplementation in a medical context, refers to cyanocobalamin."). Courts have, in some cases, upheld AI-related patents over §112(a) challenges in view of expert testimony regarding a POSITA's knowledge.

For example, in *Centripetal Networks, Inc. v. Cisco Sys*, the court confirmed claims directed to correlating a plurality of log entries received and transmitted by a network device to generate a set of rules used to identify packets and provide network devices with such rules. 492 F. Supp. 3d 495, 558 (E.D. Va. 2020). Cisco and its expert argued that the specification lacked disclosure of technologies central to the claims, including "Cognitive Threat Analytics, machine learning, artificial intelligence, integrating threat feeds, [and] NetFlow." *Id.* Centripetal's expert responded that a POSITA with knowledge of AI technologies would understand the invention and its function from the specification. *Id.* ("Dr. Jaegar opined that a person skilled in the art would be able to look at column 8, lines 46 through 63 of the '176 Patent specification and determine that the invention 'utilize[s] logs to correlate packets transmitted by one or more network devices with packets received by one or more network devices.' …. Additionally, for the 'responsive to' element, Dr. Jaegar points to column 12, line 55 through column 13, line 13."). In post-trial proceedings, the court upheld the jury verdict finding validity, concluding that the jury could credit the specification's disclosures identified by plaintiff's expert. *Id*.

Similarly, in *Masimo Corporation v. Philips Electronic North America*, the court held that there was substantial evidence to support the jury's finding that Philips failed to carry its burden to show the patent did not enable the claims. 2015 WL 2379485 at *7 (D. Del. May 18, 2015). The claims at issue were directed to the processing of measured signals. *Id.* at *5. Defendant argued "the specification did not reasonably convey to one of ordinary skill that the inventor was in possession of a non-correlation canceler."

Id. But based on Masimo's expert testimony, the court reasoned that "[w]ithout any (much less clear and convincing) evidence of what a noncorrelation canceling technique is — or how this undefined technique would still be covered by the claims — the jury could reasonably have concluded that Philips failed to carry its burden on lack of written description." *Id.* at *5-6.

Practitioners should, however, be aware that a POSITA's knowledge is not a cure-all for a sparse specification. Courts place little weight on assertions about a POSITA's knowledge where it lacks evidentiary support. For example, in *In re Starrett*, the court found "Starrett's arguments on enablement conclusory and unresponsive." 2023 WL 3881360 at *5. It reasoned that "[a]lthough a skilled artisan's familiarity with the components of a claimed invention is relevant … it is not dispositive of enablement on its own." *Id.* Starrett's arguments on appeal did not address how the patent's disclosures enabled novel functions of allegedly well-known components and were not supported by any evidence other than an interested inventor's own assertions. *Id*.

Takeaways

AI patent claims are often broad and a bit nebulous by nature, and AI models are largely developed through self-training, making it difficult for patentees to fully and adequately describe them. Courts, however, are unlikely to give AI patents leniency in complying with written description and enablement requirements. Patentees must describe AI inventions in sufficient detail to enable persons skilled in the art to reproduce the invention without undue experimentation. For example, rather than broadly claiming the use of AI concepts like "machine learning" or "deep learning," the specification should describe specific algorithms, training datasets and techniques, and neural network attributes (e.g., number of layers, activation functions, weights, biases, etc.) that are used. On the other hand, a specification that lacks such details may still overcome §112(a) challenges if (1) the specification relies on AI models and techniques that were well-known or widely available at the time of the patent's priority date (and ideally, describes them as such) and/or (2) the patentee-plaintiff can supplement the specification's disclosures with credible expert testimony and supporting evidence confirming that any disclosed AI models and techniques were well-known to a POSITA. Still, relying on expert testimony is likely to fare less well if a patent claims to invent or employ new and improved AI models.

Claims With §112(b, f) Means-Plus-Function limitations

Patent drafters can opt to use means-plus-function claiming, which may appeal to AI patent drafters, who may disfavor reciting the specific structures, training datasets, and algorithms that an AI model uses, either to avoid narrowing the scope of the claims or because of the cost or difficulty in describing such details in the claims. Means-plus-function limitations allow AI patent claims to more broadly recite functions performed by an AI model, subject to the condition that such limitations cover only the "means identical to or the equivalent of the structures, material, or acts described in the patent specification." *WMS Gaming, Inc. v. Int'l Game Tech*., 184 F.3d 1339, 1349 (Fed. Cir. 1999).

But a means-plus-function claim risks being found invalid as indefinite if there is an "absence of structure disclosed in the specification to perform [the recited] functions." *Aristocrat Techs.*

Australia Pty Ltd. v. Int'l Game Tech., 521 F.3d 1328, 1333, 1338 (Fed. Cir. 2008). For software, a means-plus-function limitation's corresponding structure cannot be a general purpose computer because "general purpose computers can be programmed to perform very different tasks in very different ways." *Id.* at 1338. Instead, the corresponding structure must be a step-by-step algorithm. *See, e.g., Ergo Licensing, LLC v. CareFusion 303, Inc.*, 673 F.3d 1361, 1365 (Fed. Cir. 2012) (holding that "if special programming is required for a general-purpose computer to perform the corresponding claimed function, then the default rule requiring disclosure of an algorithm applies.").

Courts have established two doctrines for software means-plus-function patents that are especially relevant to AI. First, drafters cannot disclose a black box or an abstraction in lieu of a step-by-step algorithm. Second, in contrast to §112(a), patent drafters cannot substitute an algorithm with the knowledge of a POSITA or rely on such knowledge to fill in any gaps in the disclosed algorithm.

Impermissible Black Box Algorithms

If patent drafters claim AI models using means-plus-function limitations, the specification must go beyond describing any corresponding algorithms by merely identifying the algorithms' inputs and outputs. Instead, the specification must break down, step-by-step, how an AI model reaches such outputs. Otherwise, the AI model will be deemed an indefinite "black box."

The Federal Circuit in *Augme Technologies, Inc. v. Yahoo! Inc*. explained that specifications cannot rely on black-box disclosures of algorithms for software means-plus-function limitations. 755 F.3d 1326, 1329 (Fed. Cir. 2014). In that case, the patentee sought to enforce patents related to adding media or other functionalities to a webpage. *Id*. One claim covered a "means for assembling" a "computer readable code module." *Id*. at 1337. The patentee pointed to a figure and its accompanying text as the necessary algorithm. The court found that the identified portion of the specification only "discloses inputs to and outputs from the code assembler instructions, but does not include any algorithm for how the second code module is actually assembled." *Id*. at 1338. So, the court concluded that the purported algorithm was insufficient as "[s]imply disclosing a black box that performs the recited function is not a sufficient explanation of the algorithm required to render the means-plus-function term definite." *Id*.

The Federal Circuit reached a similar conclusion in *Blackboard, Inc. v. Desire2Learn, Inc*. The claim term at issue was a "means for assigning a level of access to and control of each data file based on a user['s] … role in a course." 574 F.3d 1371, 1382-83 (Fed. Cir. 2009). The court found the claim indefinite, because there was no corresponding algorithm. The patentee attempted to remedy the issue by pointing to a component called the "access control manager" in the specification. But the court held "that [i.e., access control manager] is not a description of structure; what the patent calls the 'access control manager' is simply an abstraction that describes the function of controlling access to course materials, which is performed by some undefined component of the system. The ACM is essentially a black box that performs a recited function." *Id*.

Generic references to "software" will also be deemed insufficient to disclose an algorithm. In *Finisar Corp. v. DirecTV Group, Inc.*, the patentee pointed to a passage in the specification describing "software 132 (executed by CPU 130)" that "generates a hierarchical set of indices referencing … database 112 and embeds those indices in the information database." 523 F.3d 1323, 1340–41 (Fed. Cir. 2008). The Federal Circuit held that the alleged algorithm was an abstraction, no more than a rephrasing of the claimed function. *Id.* Specifications must explain step-by-step how software works, instead of relying on general words like "analyzing" or "processing." *Napco, Inc. v. Landmark Tech. A, LLC*, 2023 WL 5000756, at *15 (M.D.N.C. Aug. 4, 2023) (explaining that "the specification merely provides functional language and does not contain any step-by step process for how the data … is 'analyzed' or 'processed' … analyzing is a function. Analyzing is not a how … It's just functions all the way down with no ending to the answer of the how am I going to process and analyze these answers." (internal citations omitted)).

While citing actual source code for an algorithm structure would seemingly satisfy any level-of-detail requirements, that approach is risky — the source code needs to work. In *Media Rights Tech., Inc. v. Capital One Fin. Corp*, the Federal Circuit held that a patent related to the control of data output and management of data output paths was indefinite, even where it cited C++ source code, because the code was non-functional. 800 F.3d 1366, 1374-1375 (Fed. Cir. 2015). Specifically, there was "unrebutted expert testimony that this code only returns various error messages. The cited algorithm does not, accordingly, explain how to perform the diverting function, making the disclosure inadequate." *Id*.

Impermissible Reliance on a POSITA's Knowledge

Just as a patentee cannot merely refer to an abstraction or a black box for an algorithm, a patentee also cannot refer to a POSITA's knowledge of an algorithm. In the *Blackboard, Inc*. case, the patentee argued "that the process of putting together control lists through software is well known to a person of ordinary skill in the art." 574 F.3d at 1385. The Federal Circuit concluded that argument "conflates the definiteness requirement of section 112[(b) and (f)] and the enablement requirement of section 112[(a)]. . . A patentee cannot avoid providing specificity as to structure simply because someone of ordinary skill in the art would be able to devise a means to perform the claimed function." *Id*.

This outcome mirrored the Federal Circuit's reasoning in *Medical Instrumentation & Diagnostics Corp. v. Elekta AB*, where a patent was directed to software that compares two images and calculates the distances between the objects and their surface areas. The court explained that "[t]he correct inquiry is to look at the disclosure of the patent and determine if one of skill in the art would have understood that disclosure to encompass software for [the claimed function] and been able to implement such a program, not simply whether one of skill in the art would have been able to write such a software program. . . . It is not proper to look to the knowledge of one skilled in the art apart from and unconnected to the disclosure of the patent." 344 F.3d 1205, 1212 (Fed. Cir. 2003).

Applying these holdings, the Federal Circuit invalidated an AI patent, relating to detection of driver drowsiness, because of the specification's overreliance on a POSITA's knowledge. *Ibormeith IP, LLC v. Mercedes-Benz USA, LLC*, 732 F.3d 1376 (Fed. Cir. 2013). Though the specification included a table with factors the algorithm for drowsiness detection was to consider, the specification did not disclose how best to weigh those factors, relying instead on a POSITA's knowledge to fill that gap in the specification's disclosure. *Id*.

Takeaways

The Federal Circuit's jurisprudence on means-plus-function limitations suggests that AI patent claims using such limitations are highly susceptible to §112(b, f) indefiniteness challenges. Although it may be difficult to explain in detail how an AI model works, caselaw demands such disclosure in the specification. In particular, the *Ibormeith* a nd *Media Rights* cases arguably require the specification to fully disclose all aspects of an AI model —the specific algorithms used, the training data, the attributes of any neural network, and potentially even evidence that the AI model works as intended—since none of that can be assumed or substituted with the knowledge of a POSITA. And even if the AI model's details are sufficiently disclosed in the specification, a patentee risks having the means-plus-function limitations construed in a way that limits their scope to the specification's disclosures.

Hypothetical Patent Claims

To illustrate the application of the caselaw discussed, consider this hypothetical:

An inventor drafts a patent specification and related claims for a generative AI tool that answers questions using large language models (LLMs), akin to BLOOM, Claude, GPT, LLaMA, or Gemini. The specification provides a detailed, step-by-step example of training the underlying AI model's neural networks using self-supervised learning, but omits any mention of other training methods except to say that other methods known in the art could be used and listing unsupervised learning as another example. The specification references how, after training, the model could be applied to answer questions as part of a chatbot application, but provides no details on that application.

The inventor drafts an independent claim (Claim 1) that recites a method that states in the relevant part:

A method for determining a response to a user query using a machine learning module, the method comprising:

…

training said machine learning module to respond using data obtained from a plurality of online sources …

And a further dependent claim (Claim 2) states:

Wherein said machine learning module comprises a neural network pretrained using self-supervised learning and semi-supervised learning, and further trained using reinforcement learning.

The inventor drafts another independent claim (Claim 3), which includes two means-plus-function limitations that read:

… a first code module for training a neural network using data obtained from a plurality of online sources; and

a second code module for generating responses to user queries using said neural network after said neural network has been trained using said data . . .

If the patent were later asserted in litigation and challenged on §112(a) and (b, f) grounds, at least Claims 2 and 3 would likely be invalidated based on the relevant case law and facts discussed.

112(a) Analysis

Claim 1 presents a close case under §112(a) because the specification provides an example of training its "machine learning module" using supervised learning. The analysis and outcome would heavily depend on how detailed the description of that example is in the specification — identifying the specific supervised learning algorithms, for example, or the specific attributes of the neural network. This might include details like activation functions, weights and biases, and training data and techniques. The analysis could also depend on facts and assumptions made by the parties and court regarding a POSITA's knowledge, the state of the art of supervised machine learning as applied to LLMs, and the patent specification's own statements regarding the state of the art. Additionally, credible expert testimony, backed up by extrinsic evidence, about the state of the art for supervised learning of LLMs using neural networks may also help the claims overcome a §112(a) challenge.

Claim 2, however, is deficient under §112(a). Reciting "reinforcement learning" (a type of machine learning) without any description and listing "semi-supervised learning" (another type of machine learning) without more detail present essentially the same problems as the patent claims rejected in *Starett* and *Amgen*. Put simply, this claim covers more embodiments than the patent actually enables. A POSITA would have to perform undue experimentation to determine how to train the machine learning module using semi-supervised or reinforcement learning given the specification's limited disclosures.

112(b, f) Analysis

Claim 3 would likely be found to recite at least one indefinite means-plus-function limitation. The first means-plus-function limitation, "a first code module for training a neural network using data obtained from a plurality of online sources," may not be problematic based on the specification's step-by-step example of training an LLM using supervised learning. As with §112(a), the analysis here would depend on the level of detail provided for that algorithm. But unlike with §112(a), analysis, assumptions, facts, and expert testimony about a POSITA's knowledge cannot be considered and used to help a claim satisfy the requirements of §112(b, f).

On the other hand, the second means-plus-function limitation of "a second code module for generating responses to user queries using said neural network after said neural network has been trained using said data" is more likely to be found indefinite. If the specification does nothing more than paraphrase the function (*i.e.*, "generating responses to user queries…" after training), much like in *Finisar*, then a court would likely find a lack of a supporting step-by-step algorithm, thereby rendering Claim 3 invalid under §112(b, f).

Conclusion

It can be difficult to describe how AI models and technologies work, much less enable others to produce an AI model within the four corners of a patent — not to mention, without also confining the scope of the patent's claims to the specific model and supporting algorithms, training data, and techniques disclosed in the specification. Yet, particularly in light of the recent *Amgen* ruling, AI patents are likely to be embroiled in an increasing number of §112 disputes. For patentees, there is no easy way to overcome these challenges. For AI patents, like other software patents, the primary cure for §112(a) and §112(b, f) issues is a detailed specification. As the Supreme Court concluded in *Amgen*, "[t]oday's case may involve a new technology, but the legal principle is the same." 598 U.S. at 616.