

Disclosing the undisclosable: Section 112(a) challenges for artificial intelligence patents

By Sasha S. Rao, Esq., and Todd M. Hopfinger, Esq., Sterne, Kessler, Goldstein & Fox PLLC

MARCH 6, 2026

I. Introduction

Patent law's foundational bargain demands transparency. In exchange for exclusive rights, inventors must disclose how to make and use their invention and demonstrate possession. Artificial intelligence fundamentally challenges this compact. The opacity of neural-network decision-making and the complexity of multi-layered architectures stands in direct tension with § 112(a)'s written description and enablement requirements.

How can inventors demonstrate possession of an invention whose internal workings remain inscrutable even to its creators? How can specifications enable reproduction of systems whose behavior emerges unpredictably from training rather than deliberate design? The PTAB is actively drawing lines on these questions. This article examines the characteristics of AI that create disclosure challenges, then turns to three illustrative PTAB decisions revealing when AI specifications pass muster and when they fall short — and what patent practitioners can learn from each.

II. Legal framework: Section 112(a) requirements

The written description and enablement requirements of 35 U.S.C. § 112 fulfill patentees' obligation to disclose their technological advance in exchange for the grant of a patent monopoly. The written description requirement mandates that the patentee actually invented what was claimed. The enablement requirement ensures that persons of ordinary skill in the art can make and use the invention without undue experimentation.

Courts assess whether experimentation is undue using the *Wands* factors, which include: the breadth of the claims, the nature of the invention, the state of the prior art, the level of one of ordinary skill, the level of predictability, the amount of direction provided by the inventor, the existence of working examples, and the quantity of experimentation needed based on the disclosure.

The numerous unique types of inputs, the unpredictability of their outputs, and the mystery of what connects the two challenge current thinking on both disclosure requirements.

III. The AI disclosure challenge

Several characteristics inherent in AI technologies may frustrate the requisite disclosure.

The 'black box' problem. Many AI systems derive from their training data patterns too cryptic to articulate in human-understandable terms. Unlike a traditional electronic invention described by its circuits and switches, you often cannot explain why a neural network produces a particular output.

The opacity of neural-network decision-making and the complexity of multi-layered architectures stands in direct tension with § 112(a)'s written description and enablement requirements.

Training data dependency. An AI model's behavior arises from the massive datasets used to train it. But two identical architectures trained on different datasets could behave differently. How do you sufficiently describe an invention when its defining characteristics emerge from data that may be proprietary, massive, or difficult to characterize?

Non-determinism in training. Running the same code on the same training data twice can yield different models with different performance characteristics. AI is probabilistic technology, not deterministic. Thus, outputs may vary based on the same prompt. Further, models can evolve without human intervention or knowledge and develop capabilities or behaviors not intended by their developers. This non-determinism complicates the traditional patent expectation that following the specification should yield consistent, predictable results.

IV. PTAB guidance

These challenges are playing out at the Patent Office. And appeals of rejections have put the spotlight on the PTAB to opine on the sufficiency of applicants' AI disclosures.

Written description in AI at the PTAB

Two illustrative appeals on written description are *Ex Parte Kirti* and *Ex Parte Allen*, where the PTAB respectively reversed and affirmed the examiners' § 112(a) rejections of AI-related claims. *Ex Parte Rituraj Kirti et al.*, Appeal 2020-527, 2021 WL 2103251 (P.T.A.B. May 21, 2021); *Ex Parte Corville O. Allen et al.*, Appeal 2020-005211, 2021 WL 5756073 (P.T.A.B. Dec. 2, 2021).

Ex parte Kirti: Written description found sufficient

In *Kirti*, the PTAB reversed the Examiner's rejection of claims reciting "a machine learning model that is trained to determine membership in the first [or second] cluster group using the first [or second] subset of users as a training set."

*As PTAB guidance demonstrates,
the key for written description
lies in distinguishing conventional
AI applications from
novel algorithmic contributions.*

The Examiner argued that the specification did not supply the machine learning algorithm or disclose how it performs the claimed function.

The Board disagreed, finding that the specification disclosed the type of machine learning model, the training methodology, the training inputs, and the desired outputs.

The PTAB reasoned that the claims "broadly recite" a machine learning model, and the specification disclosed sufficient implementation details (e.g., training input, desired output, and how the model is used) such that a skilled artisan could build the claimed models.

The Board emphasized that "one of skill in the art can select any appropriate type of model (e.g., a neural network or a linear regression)" because "the algorithms for training known machine learning models (e.g., backpropagating a neural network or fitting a linear regression) are well known." Notably, the Board viewed this machine learning algorithm as conventional.

Ex parte Allen: Written description found insufficient

In *Allen*, the PTAB affirmed the § 112(a) rejection for claims directed to a method of processing patient medical data that used natural language processing to identify medication-

related content, generate a medication listing for the patient, and determine with a scoring rubric whether certain medications should be removed from that listing.

The specification failed to disclose the algorithm used, how scores were generated, or how they should be weighted and compared. The specification stated only that "[t]he weights applied to the various scores may be specified by a subject matter expert, may be learned through a machine learning process, or the like."

The Board found this inadequate, emphasizing that simply describing a desired result does not satisfy the written description requirement. Rather, the specification must explain how the claimed function is achieved. Here, the claims recited a specific desired result without the specification providing the algorithm. Unlike well-known machine-learning classification techniques, the specific scoring mechanisms claimed were not conventional. They were the inventive steps.

The difference in outcomes comes down to this: When claims invoke well-known, conventional techniques, describing inputs, outputs, and general methodology may suffice. But when claims recite specific computational operations that are novel to the invention, the specification must disclose how those operations work.

Enablement in AI at the PTAB

The PTAB's enablement analysis in *Ex parte Lev* provides key insights. *Ex Parte Guy Lev, Matan Ninio, and Oren Sar Shalom*, Appeal 2023-001664, 2024 WL 3220638 (P.T.A.B. June 28, 2024). The claimed invention was a method for unsupervised anomaly detection using generative adversarial networks (GAN). The PTAB affirmed the Examiner's enablement rejection, finding the specification lacked critical implementation details.

The Board acknowledged that GAN technology was known and mature but emphasized the invention's complexity, noting the specification itself called the task "extremely complex." The appellant's argument that only the novel data-manipulation aspect required disclosure was rejected.

The decision provides three practical takeaways.

First, the Examiner detailed an 18-item list of missing disclosure elements — e.g., loss functions, pseudocode, parameter/weight/coefficient, architecture layout, unlabeled data, and convergence and stability analysis. This essentially functions as a checklist of what the PTAB and examiners may expect to see in a specification for a machine learning invention. There were also no working examples of the claimed invention in the specification.

Second, *Lev* rejects the argument that AI components could be treated as a "black box" and that enablement should focus only on the novel data-manipulation aspects rather than the underlying network architecture. The Board's rejection of that argument sends a clear signal that applicants cannot rely on the general knowledge of skilled artisans to fill in the structural

and operational details of complex AI systems, even well-known ones like GANs.

Third, *Lev* grapples with the tension between a mature, well-published art (over 10,000 Google Scholar results for GANs) and a finding that enablement is still lacking. The Board acknowledged the sophisticated state of the art and high skill level both weighed in the appellant's favor yet concluded these couldn't compensate for the specification's failure to provide implementation-level guidance. This may be a common scenario: the underlying techniques are widely published but applying them to a specific claimed invention still demands concrete disclosure.

About the authors



Sasha S. Rao (L) is counsel in **Sterne, Kessler, Goldstein & Fox's** trial and appellate practice group. He has experience in a range of matters from initial fact investigation through trial in technologies including artificial intelligence (AI), pharmaceuticals, materials science, chemical arts, and automotive technology. He can be reached at SRao@sternekessler.com. **Todd M. Hopfinger** (R) is a director in the firm's electronics practice group. He specializes in strategic intellectual property (IP) counseling and patent prosecution for technology companies at all development stages, from innovative startups to major global enterprises. His areas of expertise

include fields such as artificial intelligence (AI), bioinformatics, data science, blockchain, 5G, consumer technology, transportation and logistics (including autonomous vehicles), digital advertising, and cybersecurity. He can be reached at thopfinger@sternekessler.com. The authors are based in Washington, D.C.

V. Conclusion

AI's inherent opacity need not defeat patent protection, but it demands creative disclosure strategies. As PTAB guidance demonstrates, the key for written description lies in distinguishing conventional AI applications from novel algorithmic contributions. For the latter, patentees must illuminate not merely what their inventions achieve, but precisely how they achieve it. For enablement, AI's black-box nature is not enough to escape the requirement. Examiners and the PTAB will look for specific elements of the AI system that will allow the making and use of the invention without undue experimentation.

Sasha S. Rao and Todd M. Hopfinger are regular contributing columnists on AI and intellectual property for Reuters Legal News and Westlaw Today.

This article was first published on Reuters Legal News and Westlaw Today on March 6, 2026.