

Where AI Patent Risk Begins And How Records Can Help

By **Lestin Kenton and Todd Hopfinger** (September 16, 2026)

In-house teams, and increasingly litigators assessing a competitor's artificial intelligence portfolio, ask two questions. First, who invented this? Second, does the application explain what the model is doing well enough to claim a technological improvement rather than a result?

Both questions encounter the same difficulty: Nobody can fully explain what the model did. A contemporaneous record of the human contributions supports inventorship, written description, eligibility and definiteness. Prepared early, it does all four jobs; prepared after the first office action, it typically does none.

The stakes are not theoretical. In July, the [U.S. Court of Appeals for the Federal Circuit](#) **invalidated** deep-learning imaging claims in *Dental Monitoring SAS v. Align Technology Inc.* on eligibility grounds.[1] Five weeks later, it **vacated** an inter partes review in the same patent family because the prior art patent's provisional application may not have supplied adequate written description support.[2] Neither dispositive issue had been raised during examination.



Lestin Kenton



Todd Hopfinger

Inventorship After the November 2025 Reset

On Nov. 28, 2025, the [U.S. Patent and Trademark Office](#) replaced its 2024 inventorship guidance for AI-assisted inventions. Ordinary conception under *Burroughs Wellcome Co. v. Barr Laboratories Inc.*, [4] decided by the Federal Circuit in 1994, now governs, with AI treated as a tool. *Pannu v. Iolab Corp.*, [5] decided by the Federal Circuit in 1998, is reserved for multiple human joint inventors.

The practical effect is easy to miss: Asking whether AI was used reveals little about inventorship. The important question is which human made which inventive decision, and whether it was documented.

The revised guidance is generally pro-innovation but narrower for AI-driven discovery. Under the prior guidance, a scientist who proved out an AI-proposed candidate had a plausible path to inventorship through reduction to practice. A conception-only inquiry may

instead treat that validation as evidence that conception was incomplete.

Examiners rarely catch this on their own. Under Section 2157 of the Manual of Patent Examining Procedure, named inventors are presumed correct, and examination reaches the inventorship question only when an applicant lists an AI system as an inventor outright. Otherwise, the issue is left for litigation or postgrant review, so a defective record functions less as an examination risk than as a validity defect sitting in the file until opposing counsel finds it.[6] Title 35 of the U.S. Code, Section 256, cannot cure the absence of a human conceiver because there is no omitted inventor to add.

The Vendor Problem

When a vendor's engineers modify your architecture, loss function, training data curation or fine-tuning approach, and that work rises to conception, those engineers become co-inventors, and co-inventors are co-owners.

Under Title 35 of the U.S. Code, Section 262, each co-owner may make, use and license the patent without being accountable to the others.[7] Your model provider could then license the resulting patent to your closest competitor and keep the proceeds.

The remedy is contractual and inexpensive relative to the exposure it prevents: assignment and cooperation provisions in the services agreement, and vendor personnel identified on disclosure forms as routine practice.

The supporting record is often less burdensome than companies anticipate: a conception report prepared before the AI tool is used, describing the problem and constraints, plus a work report noting which outputs were kept or discarded and why. Keep in mind that prompt logs kept as business records remain discoverable like any other.

What Cannot Be Fixed After Filing

Even with a clean inventorship record, Title 35 of the U.S. Code, Section 112, poses a separate risk. After filing, the specification cannot be supplemented with new matter.[8] The filing date is therefore the last opportunity to provide the disclosure needed to support the claims.

Dental Monitoring presented the same problem from another posture. Align also challenged the claims under Section 112, but the Federal Circuit resolved the case under

Title 35 of the U.S. Code, Section 101, without reaching written description.

Title 35 of the U.S. Code, Section 112, asks two related questions: whether a skilled person could make and use the invention without undue experimentation, and whether the specification shows the applicant possessed it as of filing.

Possession is where AI claims most often run into difficulty. Claims drafted to a result often sweep broadly enough to cover every model that achieves that result. Both *Ariad Pharmaceuticals Inc. v. Eli Lilly & Co.*,^[9] decided by the Federal Circuit in 2010, and *Amgen Inc. v. Sanofi*, decided by the [U.S. Supreme Court](#) in 2023,^[10] require the specification to support the full scope of what is claimed.

That does not require disclosure of trained model weights, but it does require disclosure of the technical problem, the architecture, the training choices that mattered, the operation at inference and the measurable effect the claims rely on.

The [Patent Trial and Appeal Board](#)'s decisions illustrate both requirements. *Ex parte Lev*, decided in 2024, affirmed an enablement rejection because the specification never progressed beyond using a generative model to detect anomalies,^[11] while *Ex parte Allen*, decided in 2021, affirmed a written description rejection for the opposite reason, claiming a score without describing the algorithm producing it.^[12]

Ex parte Kirti, reversing on written description in 2021, provides the better model: The specification described the model type, training inputs, methodology and desired outputs.^[13]

Functional claim language raises a related but distinct problem. A limitation reciting a machine learning module configured to determine some result is functional, and module is not a structural term, so *Williamson v. Citrix Online LLC*, decided by the Federal Circuit in 2015, makes the presumption against means-plus-function treatment easy to overcome.^[14]

Once Section 112(f) applies, the corresponding structure for a computer-implemented function must be a disclosed algorithm; a specification stating only that a trained model performs the function discloses none. The claim is indefinite, not merely narrow.

A Predictable Shift

Since *Ex parte Desjardins* was designated precedential on Nov. 4, 2025, many commentators have argued that the USPTO would ease Section 101 scrutiny of AI claims and rely more heavily on Sections 102, 103 and 112.[15]

Data shows that this is not the case. In CPC class G06N, the core machine learning art, Section 101 rejections appeared in roughly 66% of office actions before November 2025 and roughly 57% after, while Section 112(a) rejections moved only from 6.8% to 6.6% and have held between 6% and 8% for 15 straight months.

The Section 101 decline also began earlier than the story requires, peaking near 72% in mid-2025 and reaching 59% by October, well before the precedential designation. Broadening the dataset to include the natural language and computer vision classes changes little: Section 101 falls from 51.6% to 45.7%, while Section 112(a) barely moves.

We did not directly measure Sections 102 and 103, but one possible explanation is examiner workload rather than doctrinal change. Prior art rejections are supported by established search tools, while a written description rejection requires an examiner to identify precisely what the specification omits.

A flat Section 112 rate therefore does not establish compliance; it may show only that the issue was not tested. Examination statistics measure allowance risk, not enforcement risk. Scrutiny may have shifted not from Section 101 to Section 112 within the USPTO, but from examiners to defendants with the resources and incentive to challenge the specification.

Two Frameworks, One Portfolio

Eligibility asks whether a claim improves technology or merely applies it, and the USPTO and the Federal Circuit are answering that question somewhat differently right now.

Desjardins arose from the Appeals Review Panel's review of a DeepMind application, holding that the PTAB had assessed the claim at too high a level of generality under Step 2A, Prong Two, of the USPTO's patent subject matter eligibility guidance, and pointing to *Enfish LLC v. Microsoft Corp.*, decided by the Federal Circuit in 2016, as the right framework for claims improving computer functioning.[16]

The decision is now in the MPEP, and it is one of the more favorable signals that the USPTO has sent in some time, though it rewards specificity rather than a general assertion that AI was used.

Recentive Analytics Inc. v. Fox Corp., decided by the Federal Circuit last year, points the other direction.[17] Applying an established machine learning technique to a new data environment is not eligible without more, since iterative training and real-time updating are inherent to machine learning; the claims recited a result without explaining how it was achieved.

The Federal Circuit applied similar reasoning in Rensselaer Polytechnic Institute v. Amazon.com Inc. in February, **holding** that a conventional application of case-based reasoning, even in a new setting, is abstract.[18]

Dental Monitoring illustrates the point. The claims required training a deep-learning device on more than a thousand dental arch images, and the court held that training on a particular data subset is inherent to machine learning and supplies no technological solution by itself. The court relied in part on the specification's description of the device as selectable from known commercial neural networks. Language intended to broaden the claims ultimately undermined them.

Taken together, these decisions mean allowance no longer marks the point at which risk is resolved. A patent allowed on a Desjardins-consistent record can still be vulnerable to an early dispositive motion under Recentive before discovery begins, leaving a portfolio that looks stronger on paper than its enforcement position warrants.

One further development bears watching, since the USPTO has promoted it in an April 30 memorandum.[19]

Under Title 37 of the Code of Federal Regulations, a sworn Rule 132 declaration tied to the claims and supported by benchmark data can help, but it cannot supply disclosure that is otherwise absent: It supports an adequate specification and does not remedy an inadequate one. It is also a permanent statement.

Three Touchpoints

Inventorship depends on capturing the human conception story before filing. Section 112 depends on disclosure that matches the scope of the claims. Section 101 depends on a technological improvement rather than a generic application of AI to information. Viewed in isolation, these can appear to be distinct prosecution issues.

In practice, they often reflect different consequences of the same underlying failure to document what the humans contributed and how the technology works. Three tests, one record, which is why the inconsistency we watch for most closely is a file that treats AI as merely a tool when inventorship is at issue, and as the entire inventive contribution when eligibility is at issue.

The discipline that avoids this is unglamorous: Describe how the model actually works, and keep every argument anchored to what a human contributed. A file built that way tends to hold up years later, when someone reads it looking for a reason it should not.

Practice Pointers

- Build the conception record before the model is used, documenting the problem, the goal and the constraints, and recording which outputs were kept or discarded and why.
- Address vendor relationships before development begins, negotiating assignment and cooperation terms rather than confronting co-ownership afterward.
- Draft to the mechanism at the level of detail the claims reach for.
- Treat allowance as a milestone, not a clean bill of health, and audit claims for functional and means-plus-function language before a defendant does.
- Keep in mind that Rule 132 eligibility declarations become a permanent part of the record.

[Lestin Kenton Jr.](#) and [Todd M. Hopfinger](#) are directors at [Sterne Kessler Goldstein & Fox PLLC](#).

The opinions expressed are those of the author(s) and do not necessarily reflect the views of their employer, its clients, or Portfolio Media Inc., or any of its or their respective affiliates. This article is for general information purposes and is not intended to be and should not be taken as legal advice.

[1] [Dental Monitoring SAS v. Align Technology Inc.](#), No. 2024-2270 (Fed. Cir. July 7, 2026) (nonprecedential), available at <https://www.cafc.uscourts.gov/opinions-orders/24->

[2270.OPINION.7-7-2026_2719362.pdf](#).

[2] Dental Monitoring SAS v. Align Technology Inc., No. 2025-1752 (Fed. Cir. Aug. 10, 2026), available at https://www.cafc.uscourts.gov/opinions-orders/25-1752.OPINION.8-10-2026_2736022.pdf.

[3] Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54636 (Nov. 28, 2025), rescinding 89 Fed. Reg. 10043 (Feb. 13, 2024).

[4] [Burroughs Wellcome Co. v. Barr Laboratories Inc.](#), 40 F.3d 1223 (Fed. Cir. 1994).

[5] [Pannu v. Iolab Corp.](#), 155 F.3d 1344 (Fed. Cir. 1998).

[6] MPEP Section 2157 (9th ed., rev. as cited), available at <https://www.uspto.gov/web/offices/pac/mpep/index.html>.

[7] 35 U.S.C. Section 262.

[8] 35 U.S.C. Section 132(a).

[9] [Ariad Pharmaceuticals Inc. v. Eli Lilly & Co.](#), 598 F.3d 1336 (Fed. Cir. 2010) (en banc).

[10] [Amgen Inc. v. Sanofi](#), 598 U.S. 594 (2023).

[11] Ex parte Lev, Appeal No. 2023-001664 (P.T.A.B. June 28, 2024).

[12] Ex parte Allen, Appeal No. 2020-005211 (P.T.A.B. Dec. 2, 2021).

[13] Ex parte Kirti, Appeal No. 2020-000527 (P.T.A.B. May 21, 2021).

[14] [Williamson v. Citrix Online LLC](#), 792 F.3d 1339 (Fed. Cir. 2015) (en banc).

[15] Ex parte Desjardins, Appeal No. 2024-000567 (P.T.A.B. Appeals Review Panel Sept. 26, 2025) (designated precedential Nov. 4, 2025). See also USPTO, Subject Matter Eligibility, <https://www.uspto.gov/patents/laws/examination-policy/subject-matter-eligibility>.

[16] [Enfish LLC v. Microsoft Corp.](#), 822 F.3d 1327 (Fed. Cir. 2016).

[17] [Recentive Analytics Inc. v. Fox Corp.](#), 134 F.4th 1205 (Fed. Cir. Apr. 18, 2025), cert. denied, No. 25-505 (U.S. Dec. 8, 2025), available at https://www.cafc.uscourts.gov/opinions-orders/23-2437.OPINION.4-18-2025_2500790.pdf.

[18] [Rensselaer Polytechnic Institute v. Amazon.com Inc.](#), No. 2024-1725 (Fed. Cir. Feb. 24, 2026) (nonprecedential), available at https://www.cafc.uscourts.gov/opinions-orders/24-1725.OPINION.2-24-2026_2651991.pdf.

[19] USPTO, Memorandum, Best Practices for Submission of Rule 132 Subject Matter Eligibility Declarations (Apr. 30, 2026) (superseding the Dec. 4, 2025 memorandum), listed at <https://www.uspto.gov/patents/laws/examination-policy/subject-matter-eligibility>.