Riding The Changing Winds For AI Innovations At The USPTO

By **Mark Campagna and Ryan Phelan** (December 15, 2025)

Mark Campagna

Ryan Phelan

In recent months, the U.S. Patent and Trademark Office has increasingly focused its comments and guidance, recalibrating how examiners and the Patent Trial and Appeal Board should review software and artificial intelligence innovations in view of Title 35 of the U.S. Code, Section 101.

In its "Reminders on evaluating subject matter eligibility of claims under 35 U.S.C. 101," released in August, the office reiterates that it is not changing the law so much as focusing the lens through which examiners view subject-matter eligibility.[1]

But for AI-assisted inventions, the focusing is significant. The reminders are directed at examiners in Technology Centers 2100, 2600 and 3600, which receive the bulk of software and AI-related applications, and have quickly become a practical road map for evaluating AI claims.

The reminders begin at Step 2A, Prong 1, where examiners decide whether a claim recites a judicial exception. The USPTO again emphasizes the three familiar groupings — mathematical concepts, certain methods of organizing human activity and mental processes — but it adds important texture around how those groupings should be deployed against software inventions.

The mental-process grouping has historically been a catch-all for AI and other computer-implemented claims. The reminders push back on that practice by stressing that a claim does not recite a mental process when it includes limitations that cannot practically be performed in the human mind.

That clarification is especially critical for AI. The guidance explains that when claim elements encompass AI "in a way that cannot be practically performed in the human mind," they fall outside the unpatentable mental-process grouping. This moves the focus away from a superficial characterization of human reasoning or classification and toward the real capabilities of the claimed system — its data volumes, processing speed and algorithmic complexity.

For innovation teams, the message is clear: Claims that highlight specifically computer-executed operations and technical aspects at nonhuman implementation or scale are less likely to be swept into the mental-process bucket.

The reminders also revisit the distinction between claims that recite a judicial exception and those that merely involve one. A claim that expressly recites a mathematical formula or named algorithm must undergo a complete patent eligibility analysis. A claim that simply uses math behind the scenes, without reciting it, may never trigger a judicial-exception finding at all.

AI-related USPTO examples illustrate the difference. In one example, a neural network is trained in stages to improve facial detection performance; because the claim describes the training functionally, without reciting any specific equations, the office treats it as not reciting a mathematical concept.

In another example, a claim that explicitly invokes backpropagation and gradient descent is treated as reciting a mathematical concept and is therefore analyzed as an unpatentable judicial exception.

How an AI technique is described — at the level of high-level function versus explicit algorithmic detail — can thus determine how much Section 101 scrutiny it receives.

Once a judicial exception is identified, Step 2A, Prong 2 becomes the battleground. Here, the reminders urge examiners to consider the claim as a whole and to ask whether it integrates the exception into a practical application.

The guidance explains that improvements to computer functionality or other technologies can provide that practical application, even if the specification never uses the magic word "improvement."

The key is whether a person of ordinary skill would recognize that the claimed combination yields a technical benefit, such as reduced memory usage, lower latency, improved model performance under specific constraints or a more efficient training pipeline.

The USPTO also clarifies that such improvements may be supported by the specification, even if they are not explicitly stated in the claim language. Conversely, claims can encode improvement concepts without reciting every performance metric.

This interplay between specification and claims will be central for AI inventions. Applicants who document concrete technical gains in their specifications and then tie those gains to claim elements are far better positioned to argue that their inventions represent patent-eligible improvements rather than abstract manipulations of data.

The reminders finally address what has become a familiar refrain in eligibility rejections: the "apply it" rationale. Examiners are cautioned not to oversimplify detailed claim language into a generic instruction to apply an abstract idea on a computer.

Where a claim recites specific steps, architectures or data flows that plausibly improve computer capabilities, examiners are told to grapple with those details rather than collapsing them into a stock "apply it" dismissal.

This is reinforced by further reminders on when to issue Section 101 rejections at all. For close calls, the guidance stresses that the preponderance-of-the-evidence standard applies; an examiner should not reject under Section 101 merely because of uncertainty about eligibility.

At the same time, the office renews its call for compact prosecution, pressing examiners to analyze every claim and to rely on Sections 102, 103 and 112 where appropriate, rather than using Section 101 as a blunt gatekeeper.

Collectively, these moves reshape how AI-assisted inventions will be examined. For companies building AI pipelines, this amounts to a more predictable and technology-focused framework for moving innovations from lab bench to patent office.

Practical Impacts for Companies

The USPTO's doctrinal recalibration is not happening in a vacuum. On Sept. 26, the office's appeals review panel, led by Director John Squires, **issued** a closely watched decision in Ex parte Desjardins that put flesh on the bones of the new guidance.[2]

In that case, the claims were directed to training a machine learning model on multiple tasks in a way that preserved performance on prior tasks, addressing the well-known problem of catastrophic forgetting.

The ARP vacated a Section 101 rejection and found that the claimed invention provided a concrete technical solution: It reduced storage needs, simplified the overall system and

maintained performance across tasks.

In doing so, the panel warned against overbroad Section 101 rejections that risk stifling innovation in emerging technologies such as AI. It criticized earlier reasoning that effectively equated all machine learning with unpatentable algorithms on generic computers.

Instead, the ARP framed the claimed continual-learning approach as a technological improvement, consistent with U.S. Court of Appeals for the Federal Circuit precedents that uphold computer-implemented claims where they enhance the functioning of the computer or another technical system.

Equally important, the ARP took the opportunity to remind examiners that the Patent Act already has the right tools for cabining claim scope. The panel explicitly pointed to Sections 102, 103 and 112 as the provisions best suited to defining the proper limits of protection.

In Desjardins, the claims at issue remained subject to a Section 103 rejection on prior-art grounds, underscoring that questions of novelty and nonobviousness, rather than eligibility, should be the focus of examination once an invention is framed as a genuine technological advance.

The Ex parte Desjardins decision was later designated as precedential, meaning that future PTAB panels will have to apply its reasoning, which will undoubtedly have a trickle-down effect on examiners, e.g., as more examiners are reversed when there are disputes over Section 101 rejections in view of AI innovations.[3]

For in-house teams, the practical implications are immediate.

Research and development efforts aimed at AI should be designed with technical improvement in mind from the outset. That means building experiments around measurable performance gains — whether in accuracy, robustness, speed, memory consumption or resource efficiency — and preserving those results in forms that can later be translated into patent specifications.

Specifications should describe those improvements in concrete, technical language rather than relying on high-level business benefits or outcome-oriented rhetoric.

At the drafting stage, patent counsel will want to ensure that key aspects of an AI system's architecture, training regimen and deployment configuration are tied to the claimed performance benefits. Claims that merely gesture at using a machine learning model are likely to struggle, but claims that connect specific training sequences, model structures or update mechanisms to identify technical outcomes are squarely in the zone that the reminders and Desjardins treat as eligible.

The specification can carry much of the explanatory weight, but at least some claims should encode the contours of the improvement, so that those benefits are not dismissed as an after-the-fact attorney argument.

Portfolio strategy should also adjust. Companies that previously abandoned AI cases in the face of aggressive mental-process or "apply it" rejections may find those files look different under the new guidance. Continuation filings or appeals may be worth revisiting where the invention genuinely improves computer performance or system architecture.

At the same time, as Section 101 recedes somewhat as a gatekeeper, practitioners should anticipate more robust scrutiny under Sections 102, 103 and 112. In-house groups responsible for budgeting and staffing prosecution will need to plan for a reality in which AI applications are more likely to survive eligibility only to face conventional prior art and written description challenges.

Balancing Innovation and Risk

The fact that AI inventions are now on firmer eligibility footing does not mean that every AI advance should be patented. Modern uses of AI span private, internal models trained on highly sensitive data, models embedded in consumer-facing devices, and cloud-based systems that quietly guide business decisions.

For each of these, companies must decide not only whether an invention is patent-eligible, but whether patenting is the right protection strategy compared to trade secrecy.

A helpful way to think about this is to walk through a typical AI development and deployment pipeline. Training begins with gathering and storing raw data — potentially including sensor outputs, user behavior logs or confidential customer information. That training data must be preprocessed, often extensively, to clean, normalize or augment it for model consumption.

Developers then select or design a model architecture, sometimes combining multiple models or modalities. Over time, the model is updated with new or additional data, giving rise to successive versions that may be more accurate or robust than their predecessors.

The outputs of the trained model may themselves be post-processed, routed to downstream systems or aggregated. Finally, the model's inference behavior — how it generates outputs when presented with new data — is deployed into production, either inside devices at the edge or on centralized servers.

Each step in this pipeline raises its own legal and business questions that bear on patent strategy. Any of these layers might contain potentially patentable inventions, but not all of them should necessarily be exposed in a published patent.

One key factor that tends to push toward patenting is public exposure. When a trained model is deployed as part of an AI-based device offered for sale, or when model outputs are provided to customers or third parties, the risk that competitors will detect, infer or independently recreate the underlying techniques increases dramatically. In such settings, patents provide a way to protect value even when the model weights or infrastructure are not disclosed.

Other considerations cut in the opposite direction. Some AI inventions are difficult or impossible to detect from outside the organization. Internal recommendation engines, scheduling optimizers or fraud-detection tools that never surface their inner workings — and whose outputs are not unique or revealing — may be well suited to trade secret protection.

The sensitivity of training data also matters. If patenting a model would effectively disclose collection methods or data relationships that must remain confidential for regulatory, ethical or competitive reasons, that weighs in favor of trade secret protection.

There is also a tension between the disclosure demanded by patent law and the instinct to maintain broad swaths of an AI pipeline as proprietary know-how. AI patents must explain the invention enough to enable others to make and use it, which will often require revealing aspects of model architecture, training flows and deployment choices.

For some companies, that level of disclosure may be an acceptable tradeoff for 20 years of exclusivity. For others, the better choice may be to patent only select, high-level aspects of their AI systems while keeping the most sensitive details in-house.

In practice, most sophisticated AI strategies blend both approaches. Companies may seek patents on core architectures, training techniques and deployment frameworks that could be reverse-engineered or that are broadly reusable across product lines, while maintaining trade secrecy around bespoke data pipelines, tuning heuristics and operational refinements.

The critical task is to evaluate, for each project, how detectable the innovation is in the wild, how sensitive the underlying data and methods are, and how the balance of disclosure versus exclusivity plays out over the life of the technology.

Keep in mind that the effective life of an AI invention is unlikely to exceed the 20 years of patent protection due to the rapid pace of AI technology innovation, which diminishes some of the advantages of trade secret protection.

Strategic Outlook

Looking ahead, the USPTO's evolving approach to AI and Section 101 is likely to influence, but not dictate, the broader IP landscape.

Even so, the record created in prosecution — shaped by the software-related invention reminders and decisions like Desjardins — will frame how eligibility disputes are litigated. Applications that clearly articulate technical improvements, avoid over-recitation of mathematical minutiae in the claims, and document concrete performance gains will present a more compelling case for eligibility when they reach the courthouse.

On the transactional side, a more permissive and predictable eligibility environment should increase the number and quality of AI-related patents available for licensing and cross-licensing. As portfolios mature under the new guidance, we can expect AI patents to play a larger role in negotiations, standards-setting and collaborations, particularly in sectors such as healthcare, transportation and infrastructure, where AI is becoming deeply embedded in physical systems.

At the same time, greater patenting activity is almost certain to trigger more challenges before the PTAB and in district courts, testing the durability of these assets under heightened scrutiny.

Internationally, the USPTO's guidance joins a growing body of AI-specific frameworks at

other major offices. The European Patent Office and China's China National Intellectual Property Administration, among others, have issued their own directions on AI-related inventions, generally moving toward a shared emphasis on technical character, concrete technical effects, and sufficient disclosure of training and implementation.

For global companies, this means that a single, well-crafted story of technical improvement can increasingly be adapted across jurisdictions, even if the doctrinal labels differ.

In the end, the USPTO's new guidance does more than remind examiners of existing law. It invites applicants, in-house counsel and policymakers to reengage with Section 101 as a technology-oriented inquiry rather than an abstract, policy-driven veto.

For AI innovators, that shift is both an opportunity and a challenge: an opportunity to secure meaningful protection for genuine technical advances, and a challenge to be more deliberate and strategic about how those advances are conceived, documented and disclosed.

Mark Campagna is an attorney. He has previously been in-house counsel at Medtronic Inc., TiVo Corp., Google LLC and DexCom Inc.

Ryan Phelan is a partner at Marshall Gerstein & Borun LLP.

[1] Reminders on evaluating subject matter eligibility of claims under 35 U.S.C. 101,

[2] 2024-000567 - Ex Parte Desjardins et al Rehearing Decision Sep 26 2025.

[3] https://www.uspto.gov/subscription-center/2025/ptab-designates-precedential-appeals-review-panel-decision.