PROGRAM MATERIALS
Program #3684
April 21, 2026

Patent Eligibility for AI and Software: What's Changed and What Works

Celesq® AttorneysEd Center
www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919

BLUESHIFT IP

www.blueshiftip.com

Patent Eligibility for AI and Software:

What's Changed and What Works

CLE Presentation | April 21, 2026

Robert Plotkin | Managing Partner, Blueshift IP

rplotkin@blueshiftip.com

Overview

- The AI Patent Paradox: USPTO and Federal Circuit moving in opposite directions
- Recentive Analytics, Inc. v. Fox Corp. (Fed. Cir. Apr. 18, 2025)
  - First precedential Federal Circuit ruling on ML patent eligibility
- USPTO Memorandum — Deputy Commissioner Kim (Aug. 4, 2025)
  - Five examiner reminders for AI/ML claims
- Ex Parte Desjardins (PTAB Appeals Review Panel, Sept. 26, 2025, precedential)
  - ML training improvements = patent-eligible technology
- December 5, 2025 MPEP Advance Notice
  - Desjardins incorporated into MPEP §§ 2106.04(d), 2106.05(a), 2106.05(f)
- Practical Drafting Strategies: specification, claims, and language
- Intersecting Pressures: §§ 112, 103, and inventorship documentation

The AI Patent Paradox

- The USPTO and the Federal Circuit are moving in opposite directions
- USPTO (Director Squires): pro-innovation record — reduce inappropriate § 101 rejections
- Federal Circuit (Recentive): generic ML applied to new domains = abstract idea, ineligible
- The practical consequence:
  - An examiner may allow a claim that a district court later invalidates
  - The real test of durability comes years later — in litigation
- The historical parallel: 'patent everything on the Internet' → Alice (2014)
- Now: 'patent everything using ML' — Recentive is the first corrective ruling
- The decisions you make in the first draft determine whether the patent survives in ten years

Recentive Analytics, Inc. v. Fox Corp.

- Federal Circuit Case No. 2023-2437 | Decided April 18, 2025
- Panel: Judge Dyk (author), Judge Prost, Chief District Judge Goldberg
- Question of first impression:
  - Whether claims that do no more than apply established methods of ML to a new data environment are patent eligible
- Four patents directed to using ML for:
  - 'Machine Learning Training' patents — live event scheduling ('367, '960)
  - 'Network Map' patents — TV broadcast channel assignment ('811, '957)
- All four patents: specifications called for 'any suitable machine learning technique'
- District court dismissed; Federal Circuit affirmed

Recentive: Core Holdings

- Step One — Abstract Idea:
  - Claims directed to the abstract idea of using generic ML in a particular environment
  - Iterative training and dynamic adjustments are 'incident to the very nature of ML' — not an improvement
  - Field-of-use limitation (event scheduling, network mapping) does not create eligibility
  - Speed and efficiency improvements from generic ML do not create eligibility
- The decisive failing: neither the claims nor the specs describe how any improvement was accomplished
- Step Two — Inventive Concept: none found; proposed inventive concept was simply the abstract idea restated
- Bottom-line rule:
  - Patents that do no more than claim generic ML applied to new data environments, without disclosing improvements to the ML models, are patent ineligible

Recentive: Key Practice Implications

- Disclosing what ML does is not enough — you must disclose how the ML itself is improved
- Purely functional language describing desired outputs cannot survive § 101 scrutiny
- Field-of-use limitations are dead weight without technical improvement to the ML
- 'AI for healthcare,' 'ML for finance,' 'ML for scheduling' — none add eligibility alone
- Standard ML behaviors are now presumptively abstract:
  - Iterative training, real-time updating, model optimization = inherent to ML
  - Do not build an eligibility argument on them
- The door remains open: the Court specifically stated ML 'may lead to patent-eligible improvements'
- Examiner allowance ≠ court survival — build for litigation durability from day one

The USPTO Responds: Three Key Documents

- 1. USPTO Memorandum — August 4, 2025
  - Deputy Commissioner Kim → Technology Centers 2100, 2600, 3600
  - Five examiner reminders for subject matter eligibility analysis
- 2. Ex Parte Desjardins — PTAB Appeals Review Panel
  - Issued Sept. 26, 2025 | Designated Precedential Nov. 4, 2025
  - Director Squires personally authored; binds all examiners and PTAB panels
- 3. December 5, 2025 MPEP Advance Notice
  - Formally incorporates Desjardins into MPEP; effective immediately
  - Amends §§ 2106.04(d), 2106.04(d)(1), 2106.05(a), 2106.05(f)
- These documents govern prosecution — they carry zero weight at the Federal Circuit

Kim Memo (Aug. 4, 2025): Overview

- Not new policy — a reminder of existing MPEP guidance
- Signals to examiners that inappropriate § 101 rejections in AI/software have been a problem
- Scope: Technology Centers 2100, 2600, and 3600 (the software-heavy centers)
- Addresses Step 2A of the USPTO's subject matter eligibility analysis
- Critical caveat: Zero weight at the Federal Circuit
  - A claim allowed under the Kim Memo is not insulated from district court challenge
  - Use as a prosecution tool — not as a guarantee of litigation survival
- Five key examiner reminders follow

Kim Memo: Five Key Examiner Reminders

- 1. 'Recites' vs. 'Involves' (Step 2A, Prong One)
  - Claims that only involve an abstract idea — without reciting it — are eligible; no full Alice analysis needed
- 2. Mental Process Limitations Have Limits
  - AI operations that cannot practically be performed by a human mind are NOT mental processes
- 3. Holistic Claim Analysis
  - Evaluate claims 'as a whole' — do not isolate individual limitations and dismiss each separately
- 4. Recognize Genuine Improvements
  - The spec need not say 'improvement' explicitly — but must make it apparent to a POSITA
- 5. The 51% Standard
  - Only reject when ineligibility is 'more likely than not' (>50%); uncertainty alone is insufficient

Kim Memo: Prosecution Tools for Practitioners

- On 'recites' vs. 'involves': draft claims to avoid explicitly reciting the mathematical calculation
- Use functional/operational language that describes what the system does without naming the algorithm
- On mental processes: document the computational scale — parameters, data volumes, latency requirements
- Both qualitatively (impossible for a human) and quantitatively (billions of parameters)
- On holistic analysis: when responding to a § 101 rejection, demand the examiner address the claim as an ordered combination
- On improvements: use the specification to describe the mechanism of improvement, not just the result
- On the 51% standard: in borderline cases, argue the examiner has not met the preponderance standard
- Cite the Kim Memo directly in office action responses

Ex Parte Desjardins: The Invention

- Application No. 16/319,040 | Real Party in Interest: DeepMind Technologies Limited
- Claims a computer-implemented method for training a machine learning model on sequential tasks
- The technical problem addressed: 'catastrophic forgetting'
  - When trained on a new task, a neural network loses performance on previously learned tasks
- The claimed solution:
  - After training on Task 1, compute each parameter's importance using a Bayesian posterior distribution approximation
  - When training on Task 2, apply a penalty term to the objective function that resists changing parameters critical to Task 1
  - Result: the model learns new tasks in succession while protecting prior task knowledge
- Three concrete technical improvements: continual learning, reduced storage, reduced system complexity

Ex Parte Desjardins: The ARP's Analysis

- Procedural posture: PTAB panel had entered a new § 101 ground of rejection sua sponte; Director convened Appeals Review Panel
- Step 2A Prong One — Confirmed: claims do recite an abstract idea (mathematical calculation)
- Step 2A Prong Two — Reversed: claims, evaluated as a whole, integrate the abstract idea into a practical application
- Key limitation credited by the ARP:
  - 'adjust parameters to optimize performance on the second task while protecting performance on the first task'
- This constitutes an improvement to how the ML model itself operates — not merely a mathematical calculation
- Director Squires: panels must not 'evaluate claims at such a high level of generality' that meaningful technical limitations are dismissed without explanation
- §§ 102, 103, and 112 are the 'traditional and appropriate tools' to limit patent scope — § 101 should not do that work

Desjardins: The Key Principle

- The eligibility question has shifted:
  - FROM: Does this claim automate a human task?
  - **TO: Does this claim improve how the ML model itself works?**
- A claim solves a technical problem internal to the ML system — not just in the application domain
- Recentive (Federal Circuit) vs. Desjardins (PTAB ARP):
  - Recentive: generic ML applied to a new domain = ineligible
  - Desjardins: ML training that improves the model's own operation = eligible
- These are complementary, not contradictory — the dividing line is whether the invention improves the ML technology itself
- Desjardins is now precedential and binds all USPTO examiners and PTAB panels

December 2025 MPEP Advance Notice: Key Changes

- MPEP § 2106.04(d) — New paragraph: Desjardins added as ML patent eligibility example
- MPEP § 2106.04(d)(1) — Improvements consideration clarified:
  - Spec need not use the word 'improvement' — but improvement must be apparent to a POSITA
  - Bare conclusory assertion of improvement (without technical detail) is insufficient
  - Claim must reflect the disclosed improvement — but need not explicitly recite it
- MPEP § 2106.05(a) — Two new examples of patent-eligible ML improvements (citing Desjardins):
  - Training a model to protect prior-task knowledge while learning new tasks
  - Parameter adjustments to improve ML system performance
- Anti-oversimplification mandate: examiners must not dismiss additional elements as 'generic computer components' without adequate explanation

Practical Drafting: Specification Strategies

- 1. Describe how the improvement is achieved — not merely that it exists
  - Recentive's fatal flaw: neither claims nor specs described how the improvement was accomplished
- 2. Modernize your ML boilerplate
  - Name specific architectures (transformers, CNNs, GNNs), training methodologies (federated, continual, RLHF), and deployment contexts
- 3. Build computational impossibility arguments — both qualitative and quantitative
  - Describe why the operation is beyond human cognitive capacity; include parameter counts, data volumes, latency requirements
- 4. Frame the technical problem as internal to the ML system
  - Catastrophic forgetting, gradient vanishing, mode collapse, distribution shift — these are ML problems, not domain problems
- 5. Operate at three levels of abstraction: high (broad claims), middle (key design choices), low (pseudocode/flowcharts)

Practical Drafting: Claim Strategies

- Apply the improvement-over-ML-itself test to every independent claim:
  - Does the claim improve how the ML system operates — or merely use ML to achieve a domain outcome?
- Frame AI as part of an integrated technical system — not a standalone 'using AI to do X' claim
- Show concrete interaction with hardware, sensors, data pipelines, feedback loops, control systems
- Nest your claims in three layers:
  - Broad system/method claims — infringement detection and coverage
  - Mid-level technical improvement claims — your Desjardins-style eligibility anchor
  - Narrow implementation claims — § 112 support and damages apportionment fallback
- Consider product-by-process claims for trained models when commercial value lies in the model itself
- Use continuation strategy: file strongest technical implementation first, expand to other domains later

Language to Use and Avoid

USE — Supports Eligibility

- Hardware-specific operations and component interactions
- Named ML architectures with computational specifics
- Technical constraint framing (penalty function, parameter importance scores)
- Internal-to-ML problem framing (catastrophic forgetting, gradient vanishing)
- Performance metrics (convergence rate, parameter counts, latency)
- System integration (sensors, actuators, control signals)

AVOID — Creates Eligibility Risk

- 'Analyzing data to determine [outcome]'
- 'Using machine learning to optimize [domain objective]'
- 'Applying artificial intelligence to [field]'
- 'Training a model on [data type]' — without ML-specific detail
- 'Generating [output] based on [input]' — without describing computation
- Bare assertion of improvement without technical mechanism

Beyond § 101: Intersecting Pressures

- § 101 is not the only pressure point for AI patents in 2026
- Three areas now demand attention in every AI application:
  - § 112: The functional-claiming trap — black-box invalidity risk
  - § 103: The elevated PHOSITA — a narrowed non-obviousness gap
  - Inventorship documentation — the new inequitable conduct battleground
- Design applications to satisfy § 101, § 112, and § 103 simultaneously from day one
- The specification that serves multiple masters:
  - The examiner — § 101 eligibility
  - The Board — PTAB appeal and IPR defense
  - The district court — validity and infringement
  - The damages expert — royalty apportionment

§ 112: The Functional-Claiming Trap

- The tension: AI models are becoming more powerful and more opaque simultaneously
- The law demands more algorithmic disclosure precisely when disclosure is hardest
- Courts are applying Williamson v. Citrix (Fed. Cir. 2015) to AI claims
- Functional claim language without structural support → means-plus-function treatment → narrow construction
- Practical rule: for every functional limitation in an independent claim, identify where the specification discloses the corresponding algorithm
- Step-by-step logic, pseudocode, or architecture diagrams — no exceptions
- The § 101 push for technical specificity and the § 112 enablement requirement now reinforce each other
- A rich specification satisfies both — a thin specification fails both

§ 103: The Elevated PHOSITA

- USPTO's 'AI as a tool' guidance has an underappreciated § 103 consequence
- A POSITA in 2026 can use LLMs to generate code, run automated prior art searches, and combine techniques across subfields
- What was non-obvious in 2020 may be routine in 2026
- The routine optimization gap between your invention and the prior art is narrower than ever
- Response: Build your objective indicia of non-obviousness record early
  - Unexpected results, commercial success, industry recognition, copying by competitors, long-felt need
- AI-generated prior art is a real threat:
  - Examiners and challengers can now assemble combinations no human searcher would previously have found
- Claim differentiation must be more precise, not less

Inventorship Documentation

- Core requirement: a human must have made a significant contribution to the conception of the claimed invention
- AI is a tool, not an inventor — the microscope analogy: what did the human see, select, and decide?
- The practical gap: clients are treating AI outputs like Google search results — use it, don't document it
- Three things to document in real time:
  - Selection — why this architecture, training approach, or training data? What judgment was exercised?
  - Refinement — how did the human modify or improve the AI output? What changes and why?
  - Prompting — the iterative reasoning that directed the AI: what was asked, what was unsatisfactory, what was the human's corrective judgment?
- Treat AI interaction logs as lab notebooks: timestamped, specific, showing the evolution of inventive thought
- Inventorship disputes will become the new inequitable conduct battleground

Summary: Key Principles

- The Recentive Rule (Federal Circuit — binding):
  - Claims that apply generic ML to new data environments, without disclosing ML improvements, are patent ineligible
- The Desjardins Opening (PTAB — precedential):
  - Claims directed to improvements in how ML systems themselves operate can be patent eligible — even when a mathematical concept is recited
- The Kim Memo / December MPEP Toolkit (USPTO — prosecution guidance):
  - Distinguish 'recites' from 'involves' | Block mental-process overreach | Assert holistic analysis | Use the 51% standard
- The Drafting Imperative:
  - Describe how the ML improvement works | Name specific architectures and methodologies | Frame the problem as internal to ML | Build computational impossibility into the record
- Write for multiple masters simultaneously — the examiner, the Board, the district court, and the damages expert

Key Authorities

Authority	Date	Key Point
Recentive Analytics, Inc. v. Fox Corp. (Fed. Cir.)	Apr. 2025	Generic ML + new domain = ineligible; must improve ML itself
USPTO Memo — Kim, Dep. Commissioner	Aug. 2025	5 examiner reminders; 51% standard; holistic analysis
Ex Parte Desjardins (PTAB ARP, precedential)	Sep. 2025	ML training that improves model's own operation = eligible
USPTO MPEP Advance Notice	Dec. 2025	Desjardins incorporated into MPEP §§ 2106.04(d), 2106.05(a)
Alice Corp. v. CLS Bank Int'l (SCOTUS)	2014	Two-step § 101 framework
Enfish, LLC v. Microsoft Corp. (Fed. Cir.)	2016	Software improvements to computer functionality are eligible
McRO, Inc. v. Bandai Namco Games Am. (Fed. Cir.)	2016	Specific computer techniques different from human processes = eligible

Additional Resources

- Contact: rplotkin@blueshiftip.com | 617-207-6141 | www.blueshiftip.com
- Key cases and documents covered today:
  - Recentive Analytics, Inc. v. Fox Corp., No. 2023-2437 (Fed. Cir. Apr. 18, 2025)
  - Ex Parte Desjardins, Appeal No. 2024-000567 (PTAB ARP Sept. 26, 2025, precedential)
  - USPTO Memorandum — Kim (Aug. 4, 2025)
  - USPTO Advance Notice of MPEP Changes (Dec. 5, 2025)

Questions?

Robert Plotkin | Blueshift IP

rplotkin@blueshiftip.com | www.blueshiftip.com

Patent Eligibility for AI and Software:

What's Changed and What Works

CLE Presentation — April 21, 2026 Robert Plotkin | Blueshift IP

Presenter Bio

Robert Plotkin is Managing Partner at Blueshift IP, a boutique patent law firm focused exclusively on software and AI patents. He has nearly 30 years of experience in patent law, holds an MIT computer science degree, and is registered to practice before the USPTO. He is the author of *The Genie in the Machine* (Stanford University Press, 2009), which predicted that AI would transform the invention process itself, and *AI Armor* (2024), which provides a strategic framework for protecting AI innovations. He is a National Law Journal IP Trailblazer and IP Super Lawyer, and maintains a LinkedIn following of over 23,000 in the patent and IP community.

Learning Objectives

By the end of this presentation, attendees will be able to:

1. Explain the Federal Circuit's holding in *Recentive Analytics, Inc. v. Fox Corp.* (April 18, 2025) and its implications for machine learning patent claims.
2. Summarize the USPTO Deputy Commissioner's August 4, 2025 memorandum and apply its five key reminders to pending applications and office action responses.
3. Identify the significance of *Ex Parte Desjardins* (PTAB Appeals Review Panel, Sept. 26, 2025, precedential) and the December 5, 2025 MPEP advance notice.

4. Apply concrete specification drafting and claim strategy techniques to improve the patent eligibility of AI-related applications.
5. Recognize emerging intersections between § 101, § 112, § 103, and inventorship documentation in the AI context.

Outline

1. Introduction: The AI Patent Paradox (5 min)
2. *Recentive v. Fox*: The Federal Circuit Draws the Line (12 min)
3. The USPTO Responds: Three Key Documents (15 min)
    1. The August 4, 2025 Kim Memo
    2. *Ex Parte Desjardins*
    3. December 5, 2025 MPEP Advance Notice
4. Practical Drafting Strategies (15 min)
    1. Specification drafting
    2. Claim strategy
    3. C. Language to use and avoid
5. Beyond § 101: Intersecting Pressures (8 min)
    1. § 112 and the functional-claiming trap
    2. § 103 and the elevated PHOSITA
    3. Inventorship documentation
6. Q&A (5 min)

SECTION I: INTRODUCTION — THE AI PATENT PARADOX

(~5 minutes)

Key Thesis

We are living through a moment of genuine paradox in AI patent law. The USPTO and the Federal Circuit are moving in opposite directions, and the gap between an allowed claim and a litigation-proof claim has never been wider.

The Divergence in Brief

The USPTO, under Director John Squires, is playing a long institutional game. The August 2025 Kim Memo, the *Desjardins* precedential decision, and the December 2025 MPEP advance notice collectively signal a deliberate effort to build a pro-innovation examination record for AI patents. The overall direction: fewer inappropriate § 101 rejections, more holistic claim analysis, and explicit recognition that improving an ML model's functioning is patent-eligible.

The Federal Circuit, in *Recentive v. Fox*, delivered a clear and cautionary ruling: generic machine learning applied to a new domain is an abstract idea. The Court's language is strict, its rule is categorical, and it has direct implications for thousands of pending and issued claims.

The practical consequence: An examiner who follows the Kim Memo may allow a claim that a district court later invalidates under *Recentive*. The real test of durability doesn't happen at the USPTO; it happens years later, in litigation.

The Historical Parallel

This isn't the first time we've seen this dynamic. In the late 1990s and early 2000s, the patent bar patented everything "on the Internet." Examiners allowed broadly. Courts eventually said no, and *Alice Corp. v. CLS Bank* (2014) was the result. Now the pattern is repeating: patent everything "using machine learning." *Recentive* is a major corrective ruling from the Federal Circuit.

For practitioners who lived through the first cycle, the lesson is clear: **the examiner's allowance is not the finish line.** The decisions you make in the first

draft of a patent application now determine whether that patent will survive in ten years. This presentation is about how to make the right decisions.

SECTION II: *RECENTIVE ANALYTICS, INC. v. FOX CORP.*

(~12 minutes) **Federal Circuit Case No. 2023-2437, Decided April 18, 2025**
Panel: Judge Dyk (author), Judge Prost, Chief District Judge Goldberg

Background

Recentive Analytics, Inc. owned four patents directed to the use of machine learning for (1) generating and optimizing schedules for live events and (2) creating "network maps" (the assignment of TV programs to broadcast channels across geographic markets). Recentive sued Fox Corp. and related entities for infringement. The district court dismissed, finding all four patents ineligible under § 101.

The Federal Circuit framed its task as a **question of first impression**: whether claims that do no more than apply established methods of machine learning to a new data environment are patent eligible.

Spoiler: They are not.

The Four Patents

Machine Learning Training Patents ('367 and '960)

Both titled "Systems and Methods for Determining Event Schedules."
Representative Claim 1 of the '367 patent recites a method that:

1. Receives event parameters and target features (e.g., venue availability, ticket prices, event attendance goals);

2. Iteratively trains a machine learning model on historical data to identify relationships between those inputs and desired outcomes;

3. Generates an optimized event schedule via the trained model;

4. Detects real-time changes to input parameters; and

5. Automatically updates the schedule in response.

The specification acknowledged that the ML model could be "any suitable machine learning technique" — gradient boosted random forest, regression, neural network, decision tree, support vector machine, Bayesian network, or any other type.

Network Map Patents ('811 and '957)

Both titled "Systems and Methods for Automatically and Dynamically Generating a Network Map." Representative Claim 1 of the '811 patent recites a method that:

1. Receives broadcast schedules for multiple live events;

2. Uses a machine learning technique to generate a network map optimizing overall TV ratings across events and channels;

3. Automatically updates the map in real time based on changes to the schedule; and

4. Uses the map to determine what each station displays.

Again, the specification specified "any suitable machine learning technique."

The Federal Circuit's Core Holdings

Step One (Abstract Idea)

The Court focused on "the character of the claims as a whole" and asked whether they were directed to an "improvement in computer capabilities" or instead to "a process that qualifies as an abstract idea for which computers are invoked merely as a tool."

The answer was unambiguous. The claims were directed to **the abstract idea of using a generic machine learning technique in a particular environment.** Three key sub-holdings:

1. Iterative training is not an improvement — it's inherent to ML. Recentive argued that iterative training and dynamic real-time adjustments to the model

made the claims technologically special. The Court rejected this decisively, pointing to Recentive's own admissions:

- "The way machine learning works is the inputs are defined, the model is trained, and then the algorithm is actually updated and improved over time based on the input." *(Recentive's own words at oral argument)*
- "Using a machine learning technique necessarily includes an iterative training step." *(Recentive's brief)*

In other words: **iterative training is incident to the very nature of machine learning.** You don't get credit for doing what ML always does.

2. Applying ML to a new field of use does not confer eligibility. The Court was blunt: "An abstract idea does not become nonabstract by limiting the invention to a particular field of use or technological environment." Applying existing ML to event scheduling or TV network mapping — domains where human experts had previously made the same decisions manually — is not a technological improvement. The Court compared this to its earlier cases rejecting claims that applied existing algorithms to "novel databases."

3. Speed and efficiency improvements from generic ML do not confer eligibility. Even if the ML-driven system produced event schedules or network maps faster than humans, that is not sufficient. The Court has consistently held that performing existing human tasks more quickly with a computer does not, without more, satisfy § 101.

The decisive failing: "Neither the claims nor the specifications describe how such an improvement was accomplished." The claims described *what* the ML did (optimize a schedule), not *how* the ML technology was improved to do it.

Step Two (Inventive Concept)

At Alice step two, the Court found nothing in the claims, individually or as an ordered combination, that would transform the abstract idea into a patent-eligible application. Recentive's proposed inventive concept (using ML to dynamically generate optimized maps based on real-time data) was simply the abstract idea restated.

The Court's Bottom-Line Rule

> "Today, we hold only that patents that do no more than claim the application of generic machine learning to new data environments, without disclosing improvements to the machine learning models to be applied, are patent ineligible under § 101."

Note what the Court did not hold: it explicitly stated that "machine learning is a burgeoning and increasingly important field and **may lead to patent-eligible improvements in technology**." The door remains open. The question is what it takes to walk through it.

Key Practice Implications from *Recentive*

1. **Disclosing merely what ML does is insufficient. You must disclose how the ML itself is improved.** Purely functional claim language describing desired outputs, without implementation details, cannot survive § 101 scrutiny.

2. **Field-of-use limitations are dead weight.** "AI for healthcare," "ML for financial services," "machine learning for event scheduling" — none of these domain identifiers add eligibility. They must be coupled with genuine technical improvement to the ML architecture, training, or deployment.

3. **Standard ML behaviors are now presumptively abstract.** Iterative training, real-time updating, model optimization — the Court treats these as inherent to ML. Do not build your eligibility argument on them.

4. **Examiner allowance ≠ court survival.** A claim allowed at the USPTO under post-*Recentive* guidance may still fail under *Recentive* in district court. Durability requires building claims that can survive both.

SECTION III: THE USPTO RESPONDS — THREE KEY DOCUMENTS

(~15 minutes)

The USPTO has issued three important documents since *Recentive* that practitioners must understand together as a coordinated response.

A. USPTO Memorandum, August 4, 2025

From: Charles Kim, Deputy Commissioner for Patents

To: Technology Centers 2100, 2600, and 3600

Subject: Reminders on Evaluating Subject Matter Eligibility of Claims under 35 U.S.C. § 101

What It Is and Is Not

The Kim Memo is explicitly **not** new policy. It is a reminder of existing examination guidance contained in the MPEP. Its significance is behavioral: it signals to examiners that inappropriate § 101 rejections in software and AI applications have been a problem and identifies specific corrective measures.

It carries **zero weight at the Federal Circuit.** A claim allowed under the Kim Memo's guidance is not insulated from district court invalidation under *Recentive*. Understand the memo as a tool for prosecution, not as a guarantee of litigation survival.

The Five Key Reminders

1. "Recites" vs. "Involves" (Step 2A, Prong One)

The eligibility analysis under Step 2A, Prong One only applies when a claim *recites* a judicial exception — meaning it *sets forth or describes* an abstract idea. Claims that merely *involve* an abstract idea (i.e., that rely on or are based on one) without reciting it in the claim language itself are patent-eligible and do not require further Alice analysis.

Practical impact: Careful claim drafting that avoids explicitly reciting mathematical calculations, algorithms, or other abstract ideas — while still capturing the functional behavior — can potentially sidestep Step 2A Prong One entirely. (See the drafting strategies in Section IV.)

2. Mental Process Limitations — Boundaries Matter

Examiners are reminded not to expand the "mental process" abstract idea category to cover limitations that cannot practically be performed by a human mind. The memo is explicit: "Claim limitations that encompass AI in a way that cannot be practically performed in the human mind do not fall within this grouping."

Practical impact: AI operations that require massive computational resources, process billions of parameters, or operate at speeds and scales impossible for human cognition are not mental processes. Specifications should affirmatively describe these computational demands — qualitatively (the process is impossible for a human to perform) and quantitatively (parameter counts, processing volumes, real-time latency requirements).

3. Holistic Claim Analysis

Examiners must evaluate claims "as a whole," considering how all limitations interact and combine, rather than isolating individual limitations and separately characterizing each as abstract or conventional.

Practical impact: When responding to § 101 rejections, argue that the claim as an ordered combination integrates any abstract idea into a practical application. Do not let the examiner dissect your claim into component parts and dismiss each in isolation. This is expressly prohibited by the memo.

4. Genuine Improvements Consideration

Examiners must recognize genuine improvements to computer functionality or other technical fields when present. The specification does not need to explicitly use the word "improvement," but it must describe the invention in enough detail that the improvement would be apparent to a person of ordinary skill in the art (POSITA). Conversely, a bare assertion of improvement without supporting detail is insufficient.

Practical impact: Your specification is your primary vehicle for the improvements argument. Describe *how* the AI system is technically better, not just that it produces better outputs, but that the computation itself is more efficient, the model architecture is more capable, or the training process is more effective than prior approaches. (See the drafting strategies in Section IV.)

5. The 51% Standard — When to Make a Rejection

Examiners should only issue a § 101 rejection when it is **more likely than not (greater than 50% probability)** that the claim is ineligible. Uncertainty alone is not grounds for rejection. Unpatentability must be established by a preponderance of the evidence.

Practical impact: In close cases, examiners now have affirmative direction to allow rather than reject. This is the most immediately useful tool in prosecution. When a rejection issues in a borderline case, cite the preponderance standard in your response and argue that the examiner has not met it.

B. *Ex Parte Desjardins*

PTAB Appeals Review Panel Decision, Appeal No. 2024-000567

Issued: September 26, 2025 | **Designated Precedential:** November 4, 2025

Panel: Director John Squires (author), Acting Commissioner Valencia Martin Wallace, Vice Chief APJ Michael W. Kim

Why This Decision Is Significant

Desjardins is a PTAB Appeals Review Panel decision (a special panel convened under the authority of the USPTO Director) and it is now **precedential**. It addresses ML patent eligibility directly, and the Director personally wrote the opinion. Its precedential designation means all USPTO examiners and PTAB panels are bound to follow it. It is not binding on the Federal Circuit, but it now governs prosecution and PTAB appeals.

The Invention

The application (App. No. 16/319,040, filed January 18, 2019) claims a computer-implemented method of training a machine learning model on multiple sequential tasks in a way that addresses a well-known challenge in machine learning called **"catastrophic forgetting."** The problem is that when a neural network is trained on a new task, it tends to lose performance on previously learned tasks, effectively "forgetting" them.

The claimed method addresses this by:

1. Training the model on a first task and determining each parameter's *importance* to that task (using a Bayesian posterior distribution approximation);

2. Obtaining training data for a second, different task;

3. Training on the second task in a way that *adjusts parameter values to optimize performance on the new task while protecting performance on the first task* — accomplished by an objective function that penalizes changes to parameters that were highly important to the first task.

The real party in interest was DeepMind Technologies Limited.

Procedural History

A PTAB panel had (1) affirmed a § 103 rejection of all claims, and (2) *sua sponte* entered a new ground of rejection under § 101, finding the claims directed to a mathematical concept. On request for rehearing, the panel denied relief. The Director then convened the Appeals Review Panel.

The ARP's Analysis

Step 2A, Prong One — Confirmed: The claims do recite an abstract idea (a mathematical calculation — specifically, computing an approximation of a posterior distribution). Prong One was undisputed.

Step 2A, Prong Two — Reversed: The ARP found that the claims, evaluated as a whole, integrated the abstract idea into a practical application by improving how the machine learning model itself operates.

The key limitations the ARP credited:

- "adjust the first values of the plurality of parameters to optimize performance of the machine learning model on the second machine learning task **while protecting performance of the machine learning model on the first machine learning task**"

The specification supported this with three concrete technical improvements:

1. **Effective continual learning**: the model can learn new tasks in succession while preserving knowledge of prior tasks (solving "catastrophic forgetting")

2. **Reduced storage**: a single model serves multiple tasks rather than requiring separate model instances with separate parameter sets

3. **Reduced system complexity**: one set of parameters maintained rather than multiple

The ARP concluded: these are improvements to how the machine learning model *itself operates*, not merely improvements in the application domain or business outcome.

The Director's Policy Commentary

Director Squires added pointed commentary that practitioners should read carefully:

> "Categorically excluding AI innovations from patent protection in the United States jeopardizes America's leadership in this critical emerging technology."

The original PTAB panel had "essentially equated any machine learning with an unpatentable 'algorithm' and the remaining additional elements as 'generic computer components,' without adequate explanation." The ARP called this analysis "overbroad."

> "Examiners and panels should not evaluate claims at such a high level of generality."

The Director also signaled an important structural preference: §§ 102, 103, and 112 are "the traditional and appropriate tools to limit patent protection to its proper scope. These statutory provisions should be the focus of examination." In other words: § 101 is not the right tool for doing the work of novelty and nonobviousness analysis.

The Key Principle from *Desjardins*

The eligibility question has shifted from **"does this automate a human task?"** to **"does this improve how the model itself works?"**

A claim that solves a technical problem *internal to the ML system* — not just a problem in the application domain — can satisfy § 101 even when the core mathematical calculation is abstract, as long as the specification discloses the technical improvement with sufficient detail and the claim reflects that improvement.

C. December 5, 2025 MPEP Advance Notice

From: Charles Kim, Deputy Commissioner for Patents **Subject:** Advance Notice of Change to the MPEP in Light of *Ex Parte Desjardins*

What Changed

The December memo formally incorporates *Desjardins* into the MPEP. These changes are effective immediately and supersede the November 2024 MPEP edition. The substantive MPEP amendments are in §§ 2106.04(d), 2106.04(d)(1), 2106.05(a), and 2106.05(f).

Key MPEP Changes

MPEP § 2106.04(d) — New Paragraph Added: *Desjardins* is now cited as authority for the proposition that an ML method patent can satisfy § 101 by reflecting genuine improvements to how the model itself operates — even when the claims recite a mathematical concept at Prong One.

MPEP § 2106.04(d)(1) — Revised (Improvements Consideration): Two important clarifications:

1. The specification need not *explicitly* state that an improvement exists, but must describe the invention such that the improvement would be apparent to a POSITA.

2. If the specification asserts an improvement only "in a conclusory manner" — a bare assertion without technical detail — that is insufficient. The examiner should not credit it.

3. The claim must *reflect* the disclosed improvement — meaning the claim must include the components or steps that provide the improvement described in the specification. However, **the claim does not need to explicitly recite the improvement** (e.g., it need not say "thereby reducing storage capacity").

Desjardins is now listed alongside *Enfish*, *McRO*, *Visual Memory*, *Finjan*, and *SRI International* as an example of claims directed to patent-eligible improvements to computer functionality.

MPEP § 2106.05(a) — New Examples Added: Two new examples of claims that may show an improvement in computer functionality (citing *Desjardins*):

> "xiii. An improved way of training a machine learning model that protected the model's knowledge about previous tasks while allowing it to effectively learn new tasks."
>
> "xiv. Improvements to computer component or system performance based upon adjustments to parameters of a machine learning model associated with tasks or workstreams."

Critical Anti-Oversimplification Reminder (MPEP § 2106.05(a)): Examiners should be "careful to avoid oversimplifying the claims" and must "not dismiss additional elements as mere 'generic computer components' without considering whether such elements confer a technological improvement." The *Desjardins* ARP language is directly quoted: "Examiners and panels should not evaluate claims at such a high level of generality that potentially meaningful technical limitations are dismissed without adequate explanation."

MPEP § 2106.05(f) — "Apply It" Consideration: *Desjardins* is added to the list of cases illustrating claims that amount to more than mere "apply it" instructions — specifically, "the claims reflected a specific improvement that addressed the technical problem of 'catastrophic forgetting' in continual learning systems."

SECTION IV: PRACTICAL DRAFTING STRATEGIES

(~15 minutes)

The *Recentive*/*Desjardins*/Kim Memo triangle defines the current landscape. Here is what it means for daily practice.

A. Specification Drafting

The specification is your primary vehicle for patent eligibility. Draft it for multiple audiences simultaneously: the examiner, the PTAB, the district court, and eventually a damages expert. A thin specification that optimizes for speed at filing will cost the client far more in litigation than it saved in prosecution.

1. Describe *How* the Improvement Is Achieved, Not Just *That* It Exists

Recentive's decisive failing was that "neither the claims nor the specifications describe how such an improvement was accomplished." *Desjardins*'s success was that the specification detailed precisely how the training method solved the catastrophic forgetting problem — step by step, at an architectural level.

Do not write: "The present invention provides improved machine learning performance."

Do write: "The training process of the present invention addresses the problem of catastrophic forgetting in continual learning systems. By computing a posterior distribution approximation over parameter values after training on the first task and incorporating a penalty term based on parameter importance into the objective function for subsequent task training, the model retains performance on earlier tasks with [X%] degradation across [N] sequential tasks, compared to [Y%] degradation in conventional sequential fine-tuning."

2. Modernize Your ML Boilerplate

Generic boilerplate describing "a machine learning model" or "a neural network" no longer suffices. Update your standard specification language to include:

- **Specific modern architectures**: transformers, convolutional neural networks (CNNs), graph neural networks (GNNs), diffusion models, mixture-of-experts models — identified by name, not just by functional description
- **Parameter scales**: where relevant, describe model scale (e.g., billions of parameters). A model with billions of parameters cannot be mentally performed by a human — this directly addresses the mental process concern from the Kim Memo
- **Advanced training methodologies**: federated learning, continual/incremental learning, few-shot learning, meta-learning,

reinforcement learning from human feedback (RLHF), contrastive learning — these are meaningful technical choices, not generic ML

- **Deployment contexts**: edge computing, distributed inference, real-time processing, model compression/quantization — these show that the system is necessarily rooted in specific computational infrastructure
- **Data pipeline specifics**: preprocessing steps, data augmentation methods, how the model ingests and transforms inputs before the learned computation occurs

3. Build Quantitative and Qualitative Impossibility Arguments

The Kim Memo states that limitations "that encompass AI in a way that cannot be practically performed in the human mind do not fall within" the mental process abstract idea grouping. Build this argument into your specification.

Qualitatively: Describe why the computation is beyond human capability by its nature — not because it is fast, but because it requires simultaneous optimization over millions of interdependent parameters, probabilistic reasoning across high-dimensional spaces, or pattern recognition in data volumes no human can process.

Quantitatively: Include concrete metrics — the number of parameters, the volume of training data, the latency requirements, the scale of simultaneous inference. These create a factual record for the mental-process argument.

4. Describe the Technical Problem Internal to the ML System

Desjardins turned on the fact that the claims solved a problem *internal to how ML systems operate* (catastrophic forgetting) rather than just a problem in the application domain (scheduling, network mapping). Look for the technical ML problem your invention addresses:

- Does it improve training efficiency (fewer epochs, less data needed)?
- Does it address model degradation over time?
- Does it reduce inference latency without accuracy loss?
- Does it improve model robustness to distribution shift?
- Does it reduce storage or computational requirements for a given performance level?

- Does it solve a problem arising from the interaction of multiple ML components?

If it does any of these, describe the problem clearly, explain how prior approaches failed to solve it, and detail precisely how your invention solves it. This is the *Desjardins* template.

5. Multi-Level Abstraction — Operate at High, Middle, and Low

Draft specifications at three levels of abstraction simultaneously:

- **High level** (for broad claims): the functional result the system achieves
- **Middle level** (for dependent claims and prosecution flexibility): the architectural approach and key technical design choices
- **Low level** (for § 112 enablement and anti-black-box purposes): pseudocode, flowcharts, step-by-step algorithmic logic, or architecture diagrams showing how the model processes inputs to generate outputs

Every functional term in a claim should have a corresponding algorithmic description in the specification. In the post-*Alice* world, and especially given the § 112 pressures discussed below, there is no safe "I'll trust the examiner to understand what I mean." Document the mechanism.

B. Claim Strategy

1. The Improvement-Over-ML-Itself Test

Apply this test to every independent claim before filing: *Does this claim describe something that improves how the ML system itself operates, or does it merely describe using ML to achieve a domain-specific outcome?*

Claims that pass: training methodologies that solve specific ML problems (catastrophic forgetting, mode collapse, gradient vanishing), model architectures that achieve specific computational properties (reduced inference time, improved calibration, robustness), deployment systems that achieve performance properties impossible with standard ML toolkits.

Claims that fail: "using ML to optimize [scheduling / pricing / routing / diagnosis / any other domain task]" without disclosing the ML improvement.

2. Frame AI as Part of an Integrated Technical System

Avoid standalone "using AI to do X" claims. Instead, show how the AI component integrates with and interacts with other technical system elements:

- Hardware sensors or actuators
- Domain-specific data pipelines with particular preprocessing requirements
- Feedback loops between the model and the physical environment
- Other specialized software modules with defined interfaces
- Output processing that translates model outputs into technically meaningful system actions

The more the claim describes how the AI component is *embedded* in a technical system — rather than floating abstractly above one — the stronger the eligibility argument.

3. Nest Your Claims for Multiple Purposes

Structure your claim set with three layers:

- **Broad system/method claims**: for detecting infringement and maximizing coverage
- **Mid-level technical improvement claims**: specifically reciting the ML architectural feature or training improvement that distinguishes the invention; these are your *Desjardins*-style claims
- **Narrow implementation claims**: the specific algorithmic steps, data structures, or component interactions; these serve § 112 purposes and provide royalty apportionment fallback positions

Do not write a single broad independent claim and a series of trivially narrow dependents. Build a lattice.

4. Consider Product-by-Process Claims for Trained Models

If the commercial value of your client's AI system lies in the trained model itself (e.g., they license or sell trained models), consider whether a product-by-process

claim covering the model is viable. Data structures are patentable subject matter in the U.S. If a competitor sells a trained model and you have a claim covering that model (defined by the innovative training process that produced it), each sale is an act of infringement — potentially producing greater damages than a method claim on the training process alone.

5. Avoid These Claim Drafting Patterns

Pattern to Avoid	**Why**	**Alternative**
Generic ML terminology without specifics ("using machine learning to optimize X")	*Recentive* — this is the abstract idea, not the improvement	Identify the specific ML architecture or training innovation and recite it
Field-of-use limitations without technical integration ("for use in healthcare")	*Recentive* — field of use does not create eligibility	Show how the technical system interacts with domain-specific data structures, hardware, or operational constraints
Functional results without implementation details ("wherein the model achieves improved accuracy")	*Recentive* — claims must describe *how*, not just *that*	Recite the mechanism by which accuracy is improved
Mental/mathematical terminology in claim language ("analyzing," "calculating," "determining" a value)	Triggers mental process or math concept categorization at Prong One	Use concrete operational language: "generating," "adjusting," "applying a weight vector to"
Bare assertion of improvement in spec ("the present invention provides improved performance")	MPEP § 2106.04(d)(1) — conclusory assertions are insufficient	Describe the specific technical mechanism and include comparative data or analysis
Entire method in a single independent claim including all system components	Creates infringement detection problems and apportionment vulnerability	Separate training process claims, model claims, and deployment/execution claims

C. Language to Use and Avoid

Language That Supports Eligibility

- Hardware-specific operations: “the neural network processor applies the attention mechanism to the input tensor”
- Computational specificity: “computing a Fisher information matrix approximation of the posterior distribution over parameter values”
- Technical constraint framing: “wherein the parameter adjustment is constrained by a penalty function weighted by the parameter importance scores”
- System integration: “receiving sensor outputs from [specific hardware], generating [specific data structure] using [specific ML component], transmitting control signals to [specific actuator]”
- Performance specificity: “training convergence within [N] epochs on [X] trillion parameters”
- Internal-to-ML problem framing: “to overcome catastrophic forgetting in sequential task learning,” “to reduce gradient vanishing in deep networks,” “to prevent mode collapse in generative adversarial training”

Language That Creates Eligibility Risk

- “analyzing data to determine [outcome]”
- “using machine learning to optimize [domain objective]”
- “applying artificial intelligence to [field]”
- “training a machine learning model on [type of data]” — *without specifics*
- “generating [output] based on [input]” — *without describing the computation*
- “the system automatically [performs human cognitive task]”

SECTION V: BEYOND § 101 — INTERSECTING PRESSURES

(~8 minutes)

§ 101 is not the only pressure point for AI patents in 2026. Practitioners must now design applications to simultaneously satisfy § 101, § 112, and § 103 — while building an inventorship record that will survive scrutiny.

A. § 112 and the Functional-Claiming Trap

The Black-Box Problem

AI models are simultaneously becoming more powerful and more opaque. The law demands more algorithmic disclosure precisely when the ML architecture is most difficult to describe. The irony is genuine and the tension is real.

The post-*Alice* push for technical specificity in claims — to satisfy § 101 — can create § 112 vulnerability if the specification does not support those specific limitations. Every functional term in a claim should have a corresponding algorithm — step-by-step logic, pseudocode, or architecture diagram — in the specification. Courts are applying *Williamson v. Citrix* (Fed. Cir. 2015) to AI claims and finding means-plus-function treatment where practitioners did not intend it.

Practical Rule

Before filing: for every functional limitation in your independent claims, identify where in the specification the corresponding algorithm is disclosed. If you cannot identify it, either add the algorithm to the spec or revise the claim. No exceptions.

B. § 103 and the Elevated PHOSITA

The USPTO's AI-as-a-tool guidance has an underappreciated § 103 consequence: the baseline of what a person of ordinary skill can do in 2026 is dramatically higher than it was in 2020. A POSITA in 2026 can use large language models to generate code, run automated prior art searches, and combine techniques from disparate subfields — all with minimal effort.

This means the "routine" optimization gap between your invention and the prior art is narrower than it used to be. You cannot assume that combining Feature A from one reference with Feature B from another is non-obvious simply because the combination requires technical sophistication. In 2026, that sophistication may be routine.

Build your objective indicia of non-obviousness record early. Collect evidence of: unexpected results, commercial success, industry praise, copying by competitors, long-felt but unmet need. These are your best defenses against a combination rejection based on an elevated PHOSITA.

AI-generated prior art is also a real and emerging threat: examiners and challengers can now assemble combinations of references that no human searcher would previously have assembled. The prior art search has become more comprehensive, which means your claim differentiation must be more precise.

C. Inventorship Documentation

The Current State of the Law

USPTO guidance issued in November 2025 revised the treatment of AI-assisted inventions. The core requirement remains: a human must have made a significant contribution to the conception of the claimed invention. AI is a tool, not an inventor. The microscope analogy captures it well: no one asks whether the microscope co-invented the discovery. The question is what the human saw, selected, and decided.

The Practical Gap

Clients are not keeping conception logs. They are treating AI outputs the way they treat Google search results — use it, don't document it. This is a significant risk.

Three things clients should document, in real time:

1. **Selection**: why did the human choose this model architecture, training approach, or set of training data? What judgment was exercised in selecting from among AI-generated alternatives?
2. **Refinement**: how did the human modify, improve, or direct the AI output? What changes did the human make and why?
3. **Prompting**: the iterative reasoning that directed the AI system. What did the human ask for, what did the AI produce, what was unsatisfactory, and what was the human's judgment in redirecting?

Advise clients to treat their AI interaction logs the way they treat lab notebooks: timestamped, specific, showing the evolution of inventive thought.

The Recursive Dimension

There is an uncomfortable irony worth naming directly: many practitioners use AI tools to assist in drafting the very patent applications that make inventorship arguments. The documentation standards we advise clients to follow apply equally to us. If you are using AI to generate claim language, specification sections, or prosecution arguments, you should be able to explain precisely what you directed the AI to produce, what you reviewed, and what judgment you exercised in accepting, modifying, or rejecting the AI's output. Inventorship disputes will become the new inequitable conduct battleground. Get the documentation right now, before it becomes a litigation weapon.

SECTION VI: Q&A PREPARATION

(Anticipate and be ready for these)

Q: Does Recentive effectively kill AI patents? A: No — the Federal Circuit was explicit that "machine learning may lead to patent-eligible improvements in technology." *Recentive* killed one category of AI patent claims: generic ML applied to a new domain without any improvement to the ML technology. *Desjardins* shows what survives: claims directed to improvements in how ML systems themselves operate. The path is narrower but it is open.

Q: What is the relationship between Recentive and Desjardins? They seem to point in different directions. A: They are complementary, not contradictory. *Recentive* (Federal Circuit): generic ML applied to new domains = abstract idea, ineligible. *Desjardins* (PTAB Appeals Review Panel, precedential): ML training methods that improve the model's own operation = eligible technical improvement. The dividing line is whether the invention improves the ML technology itself or merely applies existing ML to a new field.

Q: How do I handle a pending § 101 rejection citing Recentive? A: First, distinguish *Recentive* on the facts: does your claim describe how the ML technology is improved, not just applied? If so, argue *Desjardins* and the MPEP examples now codified in the December 2025 advance notice. Second, invoke the Kim Memo's 51% standard — if this is a close case, the examiner should allow. Third, argue holistic claim analysis under the Kim Memo: the examiner may be improperly isolating individual limitations rather than reading the claim as an ordered combination. Consider an amendment that moves more of the technical implementation detail into the claim language.

Q: Can I amend my existing claims to survive Recentive? A: Yes, if the supporting detail exists in the specification. You can add claim limitations that recite the ML technical improvement. The constraint is written description: you can only claim what the specification discloses. This is another argument for writing rich specifications at the time of filing — you preserve the flexibility to amend later.

Q: How much technical detail does the specification really need? A: More than it did three years ago, and the requirement is going up, not down. The *Desjardins* specification worked because it described, at a mechanistic level, *how* the training method protected prior task knowledge: the Fisher information matrix approximation, the importance scores, the penalty term in the objective function.

That is the bar. Generic references to "advanced training techniques" will not suffice under *Desjardins* and will be flagged as conclusory under the revised MPEP.

Q: What about trade secret strategies for ML inventions that can't be patented? A: Post-*Recentive*, some firms are reconsidering the patent/trade secret boundary for routine ML applications — particularly applications that straightforwardly apply existing ML to new domains. Trade secret protection has real advantages for training data, model weights, and hyperparameters that may not be disclosable in a patent application. But patents remain the superior tool for genuine technical innovations in ML architecture, training methodology, or deployment, and trade secret protection is vulnerable to reverse engineering and independent development. The right answer is fact-specific and should be part of every IP strategy conversation with AI clients.

SUMMARY OF KEY PRINCIPLES

The *Recentive* Rule

Claims that do no more than apply established methods of machine learning to new data environments, without disclosing improvements to the machine learning models themselves, are patent ineligible under § 101.

The *Desjardins* Opening

Claims directed to improvements in how ML systems themselves operate — where the specification discloses the technical mechanism of improvement and the claim reflects that improvement — can be patent eligible even when a mathematical concept is recited.

The Kim Memo Tools

- Distinguish "recites" from "involves" at Prong One
- Block mental-process overreach for computations beyond human capability

- Demand holistic, not atomistic, claim analysis
- Assert the preponderance standard in borderline rejections
- Build the improvements argument through the specification

The Drafting Imperative

Write for multiple masters simultaneously — the examiner, the PTAB, the district court, and the damages expert. A specification that is thin at filing is a ticking clock. The decisions made in the first draft determine whether the patent is worth anything in ten years.

TABLE OF KEY AUTHORITIES

Document	Date	Forum	Status	Key Holding/Guidance
Recentive Analytics, Inc. v. Fox Corp., No. 2023-2437	April 18, 2025	Fed. Cir.	Binding precedent	Generic ML applied to new domain = ineligible; must improve ML itself
USPTO Memo (Kim), § 101 Reminders	August 4, 2025	USPTO	Internal guidance	5 examiner reminders; 51% standard; holistic analysis
Ex Parte Desjardins, Appeal No. 2024-000567	Sept. 26, 2025	PTAB Appeals Review Panel	Precedential (Nov. 4, 2025)	ML training that improves model operation = eligible; Enfish framework applies
USPTO Advance Notice of MPEP Changes	December 5, 2025	USPTO	Effective immediately	Incorporates *Desjardins* into MPEP §§

				2106.04(d), 2106.05(a), 2106.05(f)
Alice Corp. v. CLS Bank Int'l, 573 U.S. 208 (2014)	2014	SCOTUS	Binding precedent	Two-step framework for § 101 eligibility
Enfish, LLC v. Microsoft Corp., 822 F.3d 1327 (Fed. Cir. 2016)	2016	Fed. Cir.	Binding precedent	Software improvements to computer functionality are eligible at Alice Step One
McRO, Inc. v. Bandai Namco Games Am. Inc., 837 F.3d 1299 (Fed. Cir. 2016)	2016	Fed. Cir.	Binding precedent	Specific computer techniques different from human processes = eligible