quinn emanuel

MARCH 2026

Business Litigation Report

THE EVOLVING LANDSCAPE OF AI IN DEFENSIVE DOCUMENT DISCOVERY

For years, technology-assisted review ("TAR") was state of the art in AI-assisted discovery. That is no longer true. Generative AI tools have moved from novelty to regular use in the defensive review and production of documents in discovery.

For in-house legal departments managing complex litigation, understanding these tools, their limitations, and the disclosure and professional responsibility obligations they trigger is no longer optional.

Technology-Assisted Review

The first mainstream use of AI to facilitate defensive document discovery was TAR, sometimes referred to as "computer-assisted review" or "predictive coding." The technology running behind the scenes in a TAR review is a form of machine learning where example documents are provided (both responsive and non-responsive) and algorithms are used to identify similar documents and predict the likelihood of responsiveness of those documents. The first iteration of TAR involves simple passive learning or simple active learning, in which a human reviewer codes a sample set of documents (a "seed" set) for responsiveness that the software then uses to code the remaining documents. The advantage of using a simple passive learning approach is once the seed set has been reviewed, the model can make predictions about the remaining, unreviewed corpus. The second iteration of TAR involves continuous active learning (TAR 2.0). This iteration starts with a small seed set of documents to train the model. After this initial training, the software promotes additional documents for human review that it identifies as most likely to be responsive to the top of the workflow. As human reviewers evaluate and code those documents as responsive or non-responsive, the model is continuously refined in real-time. The advantage of a continuous active learning approach is the ability to continue training and refining the model, but it does require more human review.

TAR, in either iteration, is now a well-known, court-sanctioned method for conducting defensive document discovery. See *Rio Tinto PLC v. Vale S.A.*, 306 F.R.D. 125, 127 (S.D.N.Y. 2015) ("[I]t is now black letter law that where the producing party wants to utilize TAR for document review, courts will permit it."). However, because TAR requires examples to compare with the review corpus, it may struggle to categorize documents that are unique or complex.

Generative AI

Generative AI closely mimics how human reviewers approach a defensive document review project, both in terms of the input into the software and the output. As to the input, instead of training the software by coding individual documents and allowing the software to learn through that coding (as is done with both iterations of TAR), the software is trained in much the same way as human reviewers are trained. Written prompts are used to instruct the software

FIRM HIGHLIGHTS

Tariff Refund Task Force

On February 20, 2026, the United States Supreme Court struck down the tariffs imposed by President Trump under the International Emergency Economic Powers Act ("IEEPA"), ruling them unlawful in *Learning Resources, Inc. v. Trump*. Quinn Emanuel has assembled a Tariff Refund Task Force to help companies assess their claims, understand their options, and pursue refunds aggressively—whether through litigation in the Court of International Trade, the Customs and Border Protection administrative process, or both.

Please click here for more information.

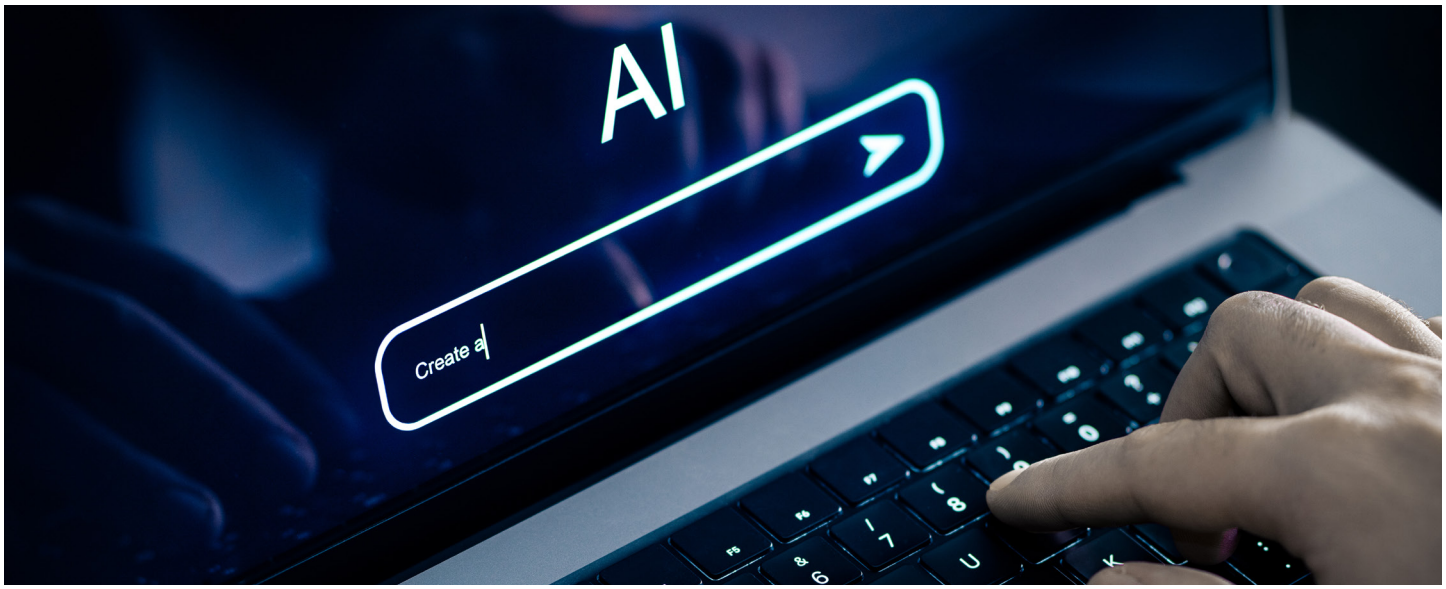

how to make responsiveness determinations. These written prompts can take a variety of forms, but closely resemble the instructions that would be given to a human reviewer in a document review protocol. The ability to mimic the human review workflow makes the entire training process more efficient. In our experience, the prompts can largely be pulled from a traditional document review protocol that is routinely prepared for instructing the human reviewers overseeing the review project or performing second level or quality-control review before documents are produced. In terms of the output, like TAR, generative AI provides a recommendation as to whether a document should be coded responsive or non-responsive. However, unlike TAR, generative AI provides the logic and reasoning, in narrative form, behind the software's responsiveness determinations.

Success in using generative AI relies heavily on how well the software is trained, which is dependent on how well the prompts are drafted. Initial prompts must be clear and precise. The process of generating the final set of prompts that are applied to the document universe can be iterative. Prompts can be continuously revised until the desired results are achieved, and results can be tested and validated using standard validation metrics. However, it is important to note that generative AI review is not an iterative process once the prompts have been run over the corpus. Thus, it is critical to test, revise, and validate the prompts before they are run against the review set. The narrative explanations that are provided make the entire process quicker and more efficient—if a non-responsive document is miscoded as responsive, or vice versa, the narrative reasoning for the responsiveness determination allows a human reviewer to identify precisely why the software miscoded the document, which makes it easy to revise the prompts. The narrative reasoning provided also helps with the validation process and facilitates subsequent human review.

In addition to the superior output in terms of receiving narrative explanations, generative AI review is also superior to TAR because it does not have the limitations of an exemplar-based review application and it allows the subject matter expert to provide the instructions directly to the model, knowing that the instructions will be followed consistently and quickly and that the results can be validated. However,

we recommend using continuous active learning workflows for human reviewers to make a responsiveness determination where the generative AI review process determined that responsiveness was borderline or a close call.

Beyond simply identifying a document as responsive or non-responsive, generative AI can be used to identify privileged documents that may need to be redacted or withheld and can facilitate the generation of descriptions of the privileged material for input into a privilege log. It can also be used to identify personally identifiable information that may need to be redacted, such as bank account or social security numbers. Further, it can categorize documents by issue and identify key documents that may require early attention. The early identification of noteworthy documents is crucial in developing case strategy, in assessing the strength and weaknesses of a case or theory, and in preparing for the taking and defending of depositions. The use of generative AI has the capability to transform document review from a purely defensive exercise into a strategic one. And, significantly, not only can generative AI perform these functions quickly, but the use of generative AI can result in substantial cost savings in document discovery, which is one of the most expensive phases in the life of a case. Although running generative AI over a set of documents does typically generate costs, those costs are a fraction of the fees typically incurred to human review those documents.

Disclosure Requirements for Use of Generative AI

The level of disclosure required concerning the use of AI in defensive document discovery is constantly evolving. Although there is a body of case law concerning the level of disclosure required when utilizing TAR, there is no set standard. One principal underlying most decisions, however, is that transparency is crucial. See, *e.g., In re Insulin Pricing Litig.*, 2025 WL 1112837, at *2 (D.N.J. Apr. 11, 2025) ("[C]ourts have mandated some level of transparency and validation of TAR methodologies in consideration of 'the complexities of TAR' and Rule 26(g)'s obligations for a producing party to undertake a reasonable inquiry in discovery."); *Progressive Cas. Ins. Co. v. Delaney*, 2014 WL 3563467, at *10 (D. Nev. July 18, 2014) ("The cases which have approved technology assisted review of ESI have required an unprecedented degree of transparency and cooperation among counsel in the review and production of ESI responsive to discovery requests."); *Berger v. Graf Acquisition, LLC*, 2024 WL 4541011, at *4 (Del. Ch. Oct. 21, 2024) (allowing use of TAR "so long as [the party is] transparent with the plaintiff about their computer-assisted review process"); Moore v. Publicis Groupe, 287 F.R.D. 182, 192 (S.D.N.Y. 2012) ("[T]ransparency allows the opposing counsel (and the Court) to be more comfortable with computer-assisted review, reducing fears about the so-called 'black box' of the technology. This Court highly recommends that counsel in future cases be willing to at least discuss, if not agree to, such transparency in the computer-assisted review process." (footnote omitted)), adopted by 2012 WL 1446534 (S.D.N.Y. Apr. 26, 2012); In *re Diisocyanates Antitrust Litig.*, 2021 WL 4295729, at * 7 (W.D. Pa. Aug. 23, 2021) ("Transparency transcends cooperation. It does not mean merely that parties must discuss issues concerning the discovery of ESI; it requires that they disclose information sufficient to make those discussions, as well as any court review, meaningful."), *adopted by* 2021 WL 4295719 (W.D. Pa. Sept. 21, 2021).

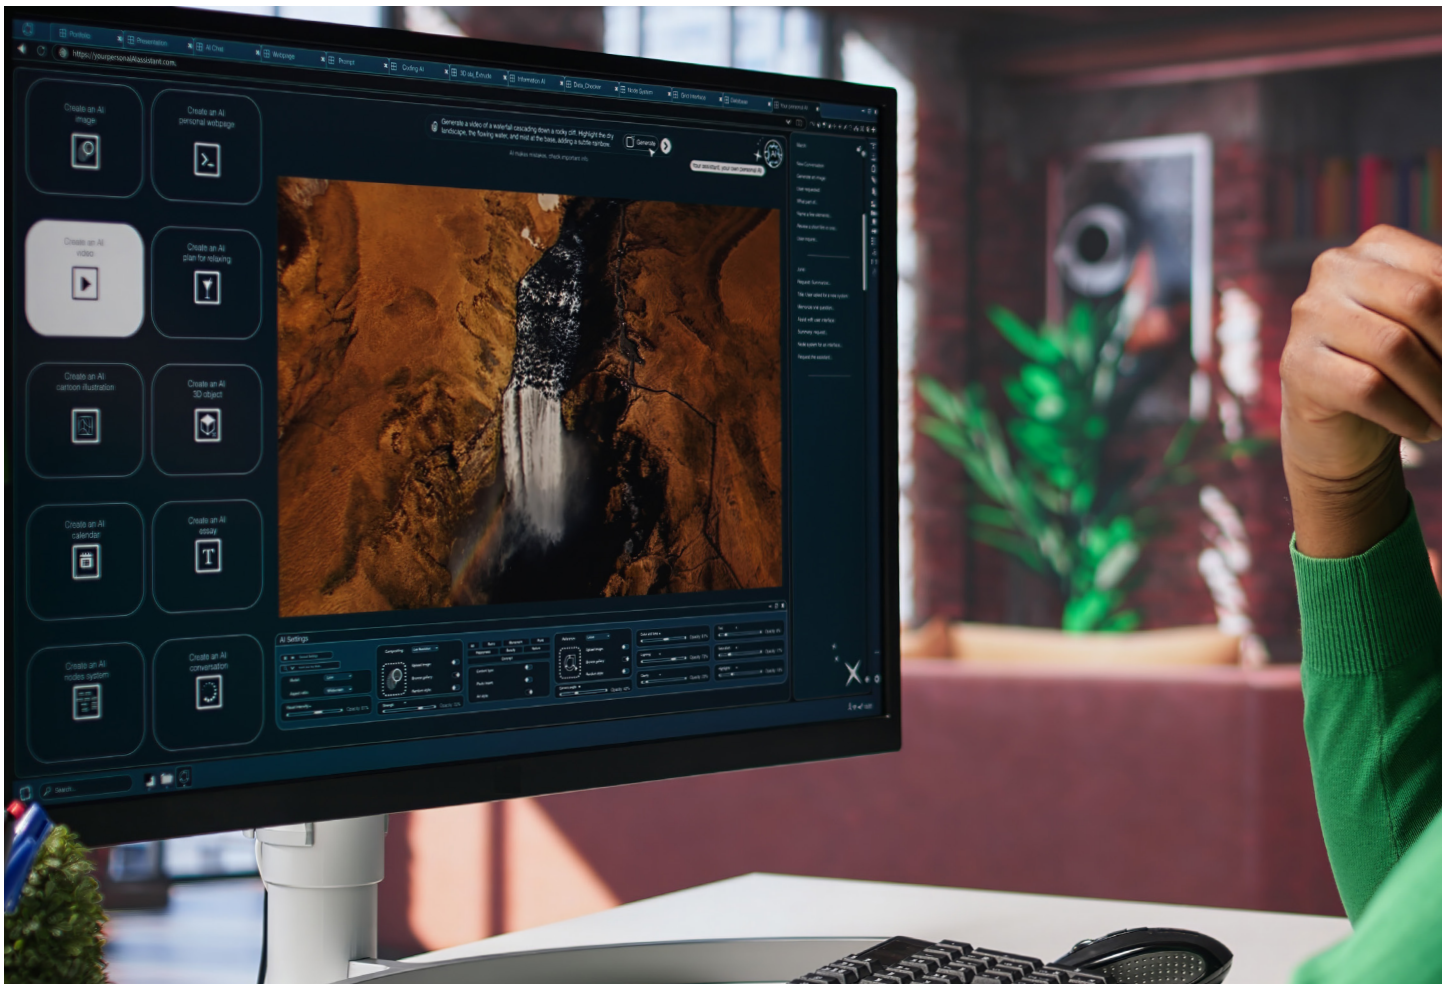

Transparency, at a minimum, includes disclosing the use of AI to the opposing party. *See, e.g., In re Valsartan, Losartan, & Irbesartan Prods. Liability Litig.*, 337 F.R.D. 610, 618 (D.N.J. 2020) (faulting party for failing to disclose "that it might use TAR . . . at th[e] time it was objectively reasonable and foreseeable"). It has also extended to disclosing details about the software and validation methodologies used. See, *e.g., Kaye v. N.Y.C. Health & Hospitals Corp.*, 2020 WL 283702, at *2 (S.D.N.Y. Jan. 21, 2020) (disclosing "detailed information regarding the collection criteria they used, the name of their continuous active learning ('CAL') software, their CAL review workflow, and how they intend to validate the review results" is "sufficient . . . to make the production transparent"); *Livingston v. City of Chicago*, 2020 WL 5253848, at *3 (N.D. Ill. Sept. 3, 2020) ("The City has disclosed the TAR software—Relativity's AL—it intends to use and how it intends to validate the review results, which in this case is sufficient information to make the production transparent.").

However, some courts have stopped short of requiring disclosure of the seed sets used to train TAR software and of requiring the opposing party's involvement in the review and validation processes. See, *e.g., In re Biomet M2a Magnum Hip Implant Prods. Liability Litig.*, 2013 WL 6405156, at *2 (N.D. Ind. Aug. 21, 2013) ("I won't order Biomet to reveal which of the documents it has disclosed were used in the seed set, but I urge Biomet to re-think its refusal."); *Livingston*, 2020 WL 5253848, at *3 ("Plaintiffs' insistence that the City must collaborate with them to establish a review protocol and validation process has no foothold in the federal rules governing discovery.").

Given the newness of generative AI as a tool for defensive document discovery, there is less clarity on the level of disclosure required. But the body of case law that has developed regarding the use of TAR suggests that a similar level of transparency will likely be required, including transparency with respect to the use of generative AI itself, the specific software used, and validation processes and metrics. There is an open question on the disclosure of the prompts used to instruct generative AI tools—are these akin to search terms, which are routinely shared and negotiated by parties? Or are these akin to the seed set of documents used to train TAR, which are sometimes voluntarily shared between parties, but

which some courts have refused to order be disclosed? Or are these more akin to document review protocols, which are widely regarded as work product and rarely shared? Courts will likely be forced to grapple with these issues as the use of generative AI becomes more prevalent and disputes inevitably arise.

Professional Responsibilities When Using Generative AI

The use of generative AI in discovery also triggers new applications of existing professional responsibilities. The American Bar Association addressed this in detail in Formal Opinion 512, "Generative Artificial Intelligence Tools," on July 29, 2024 ("ABA Formal Opinion 512"). Among the primary professional responsibility rules implicated when generative AI is used in defensive document discovery are the following:

Competency. Rule 1.1 of the Model Rules of Professional Conduct requires a lawyer to "provide competent representation to a client." This obligation extends to understanding how to use AI tools and the accompanying risks. As this applies to generative AI, ABA Formal Opinion 512 describes this as a duty to "acquire a reasonable understanding of the benefits and risks of the [generative AI] tools that they employ in their practices" or else a duty to "draw on the expertise of others who can provide guidance about the relevant [generative AI] tool's capabilities and limitations." This also requires "an appropriate degree of independent verification or review" before relying on a generative AI tool's output. ABA Formal Opinion 512.

Confidentiality. Rule 1.6 of the Model Rules of Professional Conduct mandates the confidentiality of client information. Using paid versions of software and ensuring they have appropriate privacy protections in place to avoid using prompts or uploaded documents to train the model, thereby protecting confidentiality, privilege, and work product, is critical.

Supervision. Rules 5.1 and 5.3 of the Model Rules of Professional Conduct dictate obligations for lawyers who manage or supervise other lawyers and non-lawyers. When generative AI is being used, these obligations extend to the supervision of the use of those AI tools. "Managerial lawyers must establish clear policies regarding the law firm's permissible use of [generative AI], and supervisory lawyers must make reasonable efforts to ensure that the firm's lawyers and nonlawyers comply with their professional obligations when using [generative AI] tools. Supervisory obligations also include ensuring that subordinate lawyers and nonlawyers are trained, including in the ethical and practical use of the [generative AI] tools relevant to their work as well as on risks associated with relevant [generative AI] use." ABA Formal Opinion 512 (footnotes omitted). ABA Formal Opinion 512 also notes the importance of ensuring that generative AI providers and tools are appropriately vetted for both competency and their ability to protect confidential information.

Client Communications. Rule 1.4 of the Model Rules of Professional Conduct mandates that a lawyer "reasonably consult with the client about the means by which the client's objectives are to be accomplished." ABA Formal Opinion 512 advises that "lawyers should consider whether the specific circumstances warrant client consultation about the use of a [generative AI] tool" and that, even if not required, explaining the use of a generative AI tool "may serve the interest of effective client communication." When the use of generative AI is being evaluated or used, we coordinate closely with our clients in assessing the benefits, risks, costs, level of disclosure, and expected results.

As of early 2026, a significant and growing number of states have also issued some form of AI guidance for practitioners, creating a patchwork of jurisdiction-specific obligations that counsel must navigate—particularly in multi-jurisdictional litigation where multiple states' ethical rules may apply. In the next section, we provide a set of best practices for complying with these rules if generative AI is utilized.

Best Practices

Best practices when it comes to using generative AI to facilitate defensive document discovery are likely to evolve as the technology develops and becomes more widespread, and as disputes arise, prompting court decisions that will provide further guidance. Even in these early stages, however, there are a number of best practices to ensure a party using generative AI for defensive document discovery is well-situated for success in using the software and obtaining the desired results, for compliance with professional obligations, and for compliance with any potential disclosure requirements that may arise during the life of a case.

Involvement and oversight by a senior attorney with subject matter expertise. Appoint a senior attorney with a thorough knowledge of the case to oversee the entire review process. Specifically, this attorney should own the prompt development process—drafting initial prompts, analyzing the AI's narrative reasoning on sample sets, revising prompts based on error patterns, and signing off on final validation metrics before the prompts are applied at scale. The attorney who should be put in this role is akin to the attorney typically entrusted with developing a document review protocol and overseeing human reviewers in a manual human review of documents.

Awareness of limitations of any software used. Promptly after deciding to utilize generative AI for defensive document review, identify any documents the software may not be able to process (e.g., image files, Excel files, audio files, or documents with a large amount of data) and any other limitations of the software being used (e.g., limitations in reviewing documents in multiple languages). If there are limitations to the software being utilized, develop a separate workflow and strategy for the review of documents that are incapable of being read by the generative AI tools, or are otherwise ill-suited to the generative AI workflow.

Well-documented processes and procedures. Ensure that the processes and procedures used in employing generative AI are well-documented. Given the uncertain disclosure requirements when it comes to utilizing generative AI in defensive document discovery, and considering the focus on transparency when it comes to TAR, this is crucial to ensuring compliance with any possible disclosure requirements and for dealing with possible challenges to those processes by an opposing party or tribunal. In our experience, this is best accomplished by drafting a memorandum that memorializes, at a minimum, (a) the specific tool used (including the specific version and all settings applied), (b) the complete prompt history showing each iteration along with the final prompts applied across the review universe, (c) the validation methodology and results (including precision, recall, and F1 scores, where applicable) along with any sampling or quality control steps taken; (d) the level of human oversight and the attorneys providing that human oversight; and (e) any documents or file types excluded from the AI workflow with the basis for exclusion and the alternate workflow utilized for reviewing those documents.

Early discussions and disclosures with opposing counsel. Disclose any intention to use generative AI for defensive document discovery to opposing counsel early in the life of a case. Although the disclosure requirements remain uncertain, early disclosure is the most prudent path forward. In our experience, these disclosures naturally arise during Rule 26(f) conferences (and state equivalents) or otherwise in the context of negotiating an ESI (electronically-stored information) protocol. Coming to an agreement between the parties on the use of generative AI, the level of disclosure required, and the validation metrics to be utilized, among other things, can help ward off discovery disputes, which can be costly, time-consuming, and lead to surprises and unanticipated results.

* * *

When used responsibly and with appropriate oversight, generative AI has the ability to substantially compress the time required to complete defensive document discovery. But realizing those benefits requires in-house legal departments to engage proactively—understanding the tools their outside counsel are deploying, asking the right questions about validation and documentation, and ensuring that ESI protocol negotiations address AI use early in the life of a case before disputes arise. Firms and in-house legal departments that invest in understanding and deploying these tools today will be best positioned to manage the cost, speed, and quality demands of complex discovery for years to come.

NOTED WITH INTEREST

Trends in IPR Institution Rates Under Director Squires

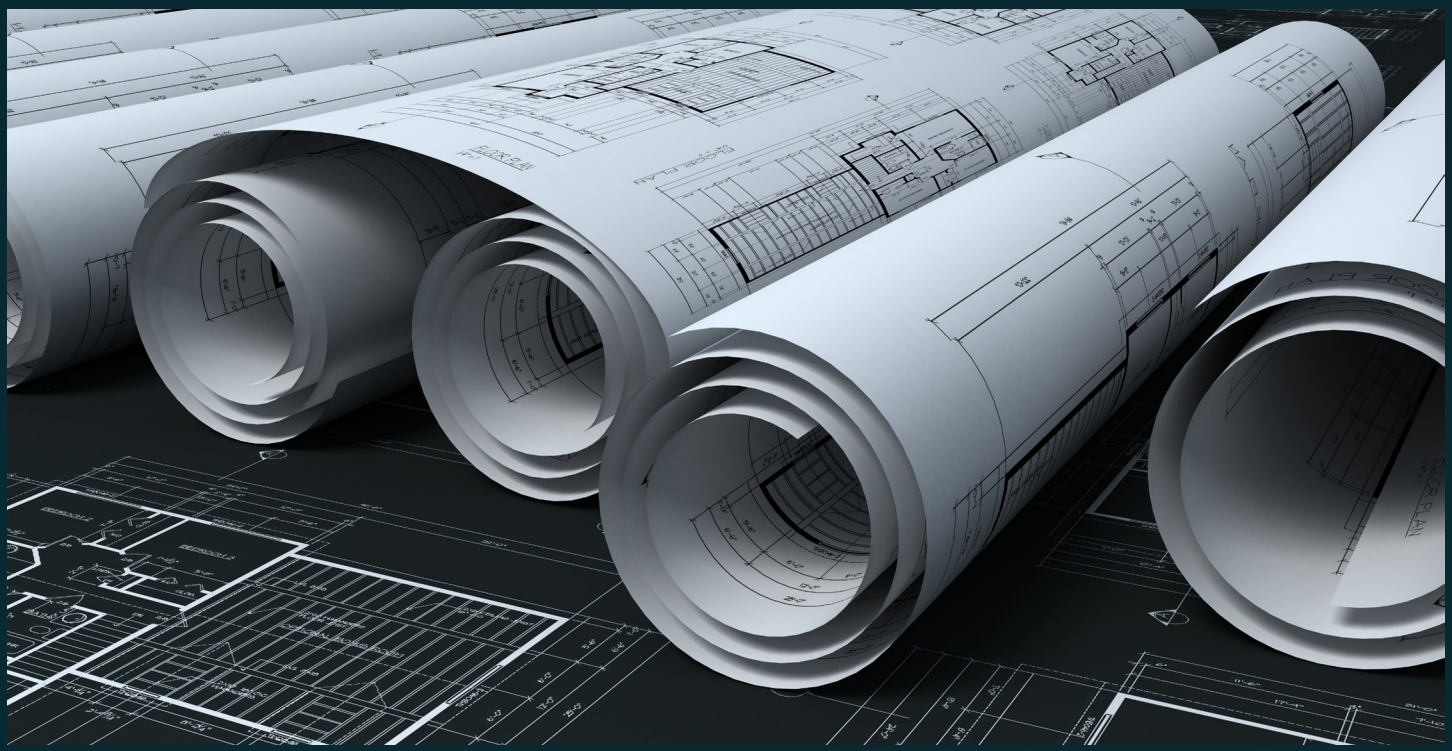

Inter partes review petitions—commonly referred to as "IPRs"—have, since the America Invents Act ("AIA") took effect in 2013, been a reliable tool for patent challengers to invalidate patents asserted against them in patent infringement proceedings. In prior years, petitions were instituted at a high rate, with over two-thirds of all petitions instituted in 2022, 2023, and 2024. In 2025, the U.S. Patent and Trademark Office ("USPTO") had a change in leadership. In January 2025, Coke Stewart took over as Acting Director.

Under Acting Director Stewart, the Patent Trial and Appeal Board ("PTAB") instituted new policies regarding discretionary denial, a procedure which allowed the Director to deny IPR petitions in her discretion before the merits of IPR petitions are considered. With this new procedure in place, overall institution rates dropped in 2025. However, as this article explains, IPRs considered on their merits continued to be instituted at a high rate. IPRs, therefore, remain an integral part of overall patent litigation strategy and should always be considered in the context of the current discretionary denial landscape and an accused infringer's litigation goals.

Acting Director Stewart denied 60% of all petitions on a discretionary basis, issuing summary decisions with minimal analysis. In addition to reviving the *Fintiv* judicial economy factors, which had largely been avoidable by stipulation under prior Directors, Acting Director Stewart began analyzing other factors, including "settled expectations" which had not previously been a consideration for IPR petitioners. And when, on September 23, 2025, John Squires was sworn in as Director of the USPTO, he continued Acting Director Stewart's approach to discretionary denial. This has led to a bevy of articles proclaiming

the death of IPRs, but a close review of Director Squires's decisions tells a different story. Indeed, if the discretionary denial process can be successfully navigated, the chances of IPR institution on the merits are high.

IPR Merits Institution Rates Remain High

There is no question that the new policies around discretionary denial have resulted in a substantial decrease in IPR institution rates. That does not mean, however, that IPRs should be completely discarded. They remain a useful tool for a comprehensive patent infringement defense if discretionary denial can be navigated.

Under Director Squires, discretionary denial and merits institution decisions are delivered en masse in a single order that addresses multiple proceedings. The decisions organize the proceedings into four groups: (1) discretionary denial; (2) further review for merits consideration; (3) institution grant on the merits; and (4) institution denial on the merits. In rare instances, the Director has also reconsidered prior referral decisions, converting these proceedings to a discretionary denial based on intervening developments. Through the Director's February 24, 2026 discretionary denial order, Director Squires has discretionarily denied 64% of all petitions he has considered, on par with Acting Director Stewart. And, for those petitions designated for merits review, the Director has instituted at a nearly 75% rate. In 2026 alone, the merits institution rate is just under 70%.

This data demonstrates that, if the discretionary denial threshold can be overcome, merits institution is highly likely. Just as before, a strong petition is a critical component to a successful IPR; now, however, petitioners must also closely scrutinize the schedule of the co-pending district court case, the scope of any *Sotera*-type prior art stipulation, the age of the patents, and any prior discussions between the parties, among other considerations. Carefully navigating these criteria may not be possible for every petitioner in every case and, as a result, some cases will be less likely to survive discretionary denial than others. But if these criteria can be navigated, there is still a strong possibility of IPR institution.

Recent Decisions Provide a Roadmap for Navigating Discretionary Denial

Much of the pessimism surrounding IPRs in 2025 was due to the unexpected nature of the new discretionary denial procedures. Prior to Acting Director Stewart, discretionary factors were considered by the panel considering the merits of the case and typically focused on the *Fintiv* time-to-trial factors which assessed whether it would serve judicial economy for the PTAB to consider a petition where the validity of the challenged patent would be evaluated by a jury in a similar timeframe. And this *Fintiv* analysis could be avoided entirely by entering into a *Sotera* stipulation, stipulating that if instituted, the petitioner would not pursue written prior art in the co-pending district court case—i.e., stipulating to apply the estoppel that would attach upon issuance of a final written decision. The new Acting Director changed this *status quo*.

When Acting Director Stewart took over, petitions had been pending for months under this old regime. Over the next eight months, decision after decision reinforced the importance of factors such as settled expectations which assesses the age of the patent being challenged. The idea is that the longer a patent

has been in force, the more settled the patentee's expectations that it would not be challenged at IPR. But although Acting Director Stewart's decisions were short, they gave important guideposts for petitioners about when settled expectations would apply. These guideposts have allowed, over time, petitions to be more tailored to account for discretionary factors such as settled expectations.

Under Director Squires, discretionary criteria have crystallized even further, providing a roadmap for navigating the discretionary denial phase of an IPR proceeding. Recent decisions designated informative or precedential have illustrated the Director's focus on inconsistent positions taken between the IPR and the co-pending district court case. In *Revvo Technologies, Inc. v. Cerebrum Sensor Technologies, Inc.*, IPR2025-00632, Paper 20 (Nov. 3, 2025) (precedential Nov. 3, 2025), the Director explained that an IPR petitioner must explain why it takes a different claim construction in the petition and in district court. Prior to *Revvo*, it was common for a petitioner to adopt the patent owner's district court claim construction. After *Revvo*, this is insufficient—the construction applied by the petitioner in the district court and the IPR should be the same. On the other side of the coin, in *Padagis US LLC v. Neurelis, Inc.*, IPR2025-00464, Paper 12 (July 16, 2025) (informative Jan. 9, 2026), the Director explained that material error during prosecution—i.e., a mistake that impacted the grant of the patent—can be sufficient to avoid discretionary denial. *Revvo*, *Padagis*, and other decisions provide guidance petitioners can use to craft a strong petition that is more likely to advance past discretionary denial to merits review.

Conclusion

Although recent changes have significantly impacted IPR strategy, the process remains viable for well-positioned petitioners. A party accused of patent infringement should consider filing IPRs where the facts, including whether there was material error during prosecution, warrant it. At the same time, IPR counsel should carefully coordinate with their district court counterparts to ensure there are no inconsistent positions that would make discretionary denial more likely..

PRACTICE AREA UPDATES

Post-Grant Patent Update:

Ex Parte Desjardins: Paradigm Shift in AI Eligibility Under § 101

On September 26, 2025, the Appeals Review Panel of the United States Patent and Trademark Office ("USPTO") issued a precedential decision in Ex parte Desjardins, Appeal No. 2024-000567. This decision has already begun reshaping patent eligibility analysis for artificial intelligence and machine learning inventions. Authored by newly confirmed USPTO Director John A. Squires, the decision vacated a previous appellate rejection under 35 U.S.C. § 101, holding that claims directed to training machine learning models can be patent eligible under certain circumstances. Designated as precedential on November 4, 2025, and incorporated into updates to the Manual of Patent Examining Procedure ("MPEP") on December 5, 2025, *Desjardins* has already become a seminal decision in § 101 jurisprudence and will carry significant influence for years to come.

The patent application at issue in *Desjardins* relates to methods for training machine learning models on multiple tasks sequentially. The claimed invention addresses the technical problem of "catastrophic forgetting"—the tendency of neural networks to lose knowledge of previously learned tasks when trained on new ones. The specification identifies concrete advantages to the claimed method, including reduced storage requirements and system complexity, and the ability to learn tasks in succession while protecting prior knowledge.

During prosecution, the examiner rejected the claims under § 103 for obviousness but did not raise a § 101 issue. On appeal, in March 2025, the Patent Trial and Appeal Board ("PTAB") affirmed the § 103 rejection and, acting *sua sponte*, entered a new ground of rejection under § 101, finding the claims directed to an unpatentable "mathematical algorithm" with only "generic computer components" as additional elements. Rehearing was denied in July 2025, with the panel citing the Federal Circuit's decision in *Recentive Analytics, Inc. v. Fox*

Corp., 134 F.4th 1205 (Fed. Cir. 2025). In August 2025, an Appeals Review Panel ("ARP") was convened *sua sponte* to review both decisions.

Applying the two-step Alice framework, the ARP agreed that the claims recited at least one abstract idea—a mathematical calculation involving computation of a posterior distribution approximation. But, critically, and citing *Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327 (Fed. Cir. 2016), the panel found that the claims reflected concrete improvements to how the machine learning model itself operates. Specifically, the ARP identified the limitation requiring the model to "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" as constituting an improvement to the technology, not merely a mathematical calculation applied in a generic computing environment.

Newly-minted Director Squires used the opinion to deliver pointed and precedential guidance. The decision cautioned that "[c]ategorically excluding AI innovations from patent protection in the United States jeopardizes America's leadership in this critical emerging technology." The ARP criticized the lower panel for evaluating claims "at such a high level of generality" that it "essentially equated any machine learning with an unpatentable 'algorithm' and the remaining additional elements as 'generic computer components,' without adequate explanation."

Further, Director Squires signaled a reorientation of examination priorities, stating: "This case demonstrates that §§ 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." This language strongly suggests that the USPTO sees § 101 as a threshold inquiry, with the substantive work of limiting patent scope performed by novelty, nonobviousness, and written description requirements.

Desjardins provides clear guidance for practitioners with pending AI patent applications in their portfolio. Applications should consider identifying a concrete technical problem and explain in detail how the invention provides a technical solution within the specification. Practitioners should consider framing AI innovations as improvements to the functioning of the model or system itself, rather than applications of generic techniques to new domains. And given the MPEP updates incorporating *Desjardins*, practitioners should cite the decision directly when responding to § 101 rejections. Indeed, *Desjardins* has already impacted subsequent decisions: in *Ex parte Carmody*, No. 2025-002843 (PTAB, Dec. 30, 2025), the PTAB reversed a § 101 rejection where claims recited a modular machine learning architecture with defined training datasets that improved how the system operated—reasoning closely tracking the *Desjardins* framework.

Desjardins represents a significant shift in the USPTO's approach to AI inventions. It establishes that purely software-based improvements to AI models can be patent-eligible and clarified that examiners should rely on §§ 102, 103, and 112 rather than § 101 as the primary tools for limiting patent scope. Although questions remain about how courts will harmonize this guidance with the Federal Circuit's *Recentive* line of authority, *Desjardins* provides the strongest foundation in years for securing patent protection for AI innovations. Its rapid incorporation into the MPEP, precedential designation, and immediate influence on subsequent PTAB decisions underscore that the USPTO has shifted its guidance on and approach to AI eligibility under § 101.

Investment Fund Litigation Update:

Enforcement of Forum Selection Clauses Under California Law Following EpicentRx

In July 2025, the California Supreme Court issued its long-anticipated decision in *EpicentRx, Inc. v. Superior Court*, 18 Cal. 5th 58, 572 P.3d 1 (Cal. 2025), bringing needed clarity in holding that forum selection clauses involving California parties should be enforced even if they contain a waiver of the right to a jury trial. The unanimous decision explained that "a forum selection clause is not unenforceable simply because it requires the parties to litigate in a jurisdiction that does not afford civil litigants the same right to trial by jury as litigants in California courts enjoy." In the months since, courts in California and elsewhere have applied *EpicentRx* to enforce previously contested forum selection clauses, changing the landscape for forum selection litigation involving California parties.

The California Supreme Court reaffirmed in *EpicentRx* that forum selection clauses are presumptively valid and enforceable, consistent with "the 'modern trend' favoring enforceability." The *EpicentRx* court clarified that the burden falls on the party opposing enforcement "to show that its enforcement would be in violation of the settled public policy of this state." This reversed lower appellate court decisions that had placed the burden on the party seeking enforcement to prove that it would not diminish California rights. Importantly, the court distinguished between substantive and procedural rights: it held that California's strong public policy "protects the jury trial right in California courts, not elsewhere," and the civil jury trial right is "a procedural right, not a substantive one." This holding effectively overruled Court of Appeal decisions, including *Handoush v. Lease Finance Group, LLC*, 41 Cal. App. 5th 729 (2019), that had refused to enforce forum selection clauses based on their impact on jury trial rights. That said, forum selection clauses could still be unenforceable under California law if they violate a "strong or fundamental public policy" of California. The court identified statutory examples where the California Legislature has prohibited such clauses, including in franchise agreements, consumer contracts, and certain employment-related statutes.

The *EpicentRx* decision has quickly simplified pending litigation nationwide where the parties engaged in often protracted proceedings regarding enforceability of forum selection clauses, due to a California litigant's waiver of the right to a jury trial. For example, in November 2025, a California corporation that was sued in the Delaware Court of Chancery two years earlier conceded the argument that a forum selection clause in a merger agreement was unenforceable because the Delaware Court of Chancery does not conduct jury trials. The Court of Chancery explained that the California defendant "conceded this argument after the California Supreme Court's EpicentRx decision, which held that a forum selection clause is not unenforceable on public policy grounds based solely on the clause's impact on plaintiff's jury trial right." *Anzu Partners LLC v. OmegaX, Inc.*, C.A. No. 2024-0526-PAF (Del. Ch. Nov. 6, 2025).

California courts are likewise applying *EpicentRx* to enforce previously contested clauses. In the *Comedy Store v. Moss Adams LLP* saga, a plaintiff sued an accounting firm for professional negligence in California, even though the parties' agreement contained a Washington forum selection clause and jury trial waiver. The trial court initially enforced the clause and dismissed the case but, on appeal, the California Court of Appeal reversed, relying on *Handoush* and other authority holding that forum selection clauses were unenforceable where they would potentially diminish a California resident's jury trial rights. On reconsideration following *EpicentRx*, the Court of Appeal affirmed the trial court's initial decision—putting an end to a four-year fight over the enforcement of an unambiguous forum selection clause. *Comedy Store v. Moss Adams LLP*, No. B327404, 2025 WL 3717063 (Cal. Ct. App. Dec. 23, 2025).

EpicentRx has brought substantial clarity to the treatment of forum selection clauses involving California parties. As demonstrated by *Anzu* and *Comedy Store,* courts are enforcing such clauses with simple, straightforward analysis, contrary to the often lengthy and protracted litigation proceedings over enforcement in the pre-*EpicentRx* era. The California Supreme Court made clear that the contract language controls notwithstanding a jury trial waiver, rejecting arguments based on inconvenience, witness location, or unavailability of jury trials. Companies engaging with California parties can now have greater confidence that contractually negotiated forum provisions will be enforced, bringing predictability and certainty to commercial transactions.

VICTORIES

Quinn Emanuel Partners Miguel Rato and Marixenia Davilla Secure Complete Withdrawal of Class Action for Qualcomm

Quinn Emanuel partners Miguel Rato and Marixenia Davilla secured a complete withdrawal of a £480 million ($650 million) class action against Qualcomm before the UK Competition Appeal Tribunal, following five years of litigation that included a five-week trial in autumn 2025. The Class Representative agreed to abandon the proceedings in their entirety after concluding, based on the trial evidence and arguments, that the Tribunal would find in Qualcomm's favor on all counts. Under the agreement, Qualcomm will make no payment to the Class Representative or the proposed class.

Unanimous Ninth Circuit Win in "Kia Boyz" Class Action

A Quinn Emanuel appellate team achieved a major Ninth Circuit victory for Hyundai and Kia. In 2023, Hyundai and Kia entered into a consumer class action settlement valued at over $275 million. The settlement released Hyundai and Kia from all consumer liability relating to the "Kia Boyz" thefts that exploited an alleged design defect in Hyundai and Kia vehicles. In exchange, Hyundai and Kia agreed to create a settlement fund to compensate theft victims and to make available a free software upgrade that addresses the alleged defect. Hyundai and Kia did not admit liability.

Eighty-nine class members objected to the settlement. In 2024, the district court overruled objections and entered final approval of the settlement.

The settlement's approval did not go unchallenged. Two class members who had filed objections in the district court appealed the district court's order and final judgment to the United States Court of Appeals for the Ninth Circuit, seeking reversal of the settlement approval. The objectors argued that the settlement amount was too low given the strength of the consumer class's case and that class members who had not experienced a theft or theft attempt received inadequate relief. The appellate proceedings presented a critical juncture for defendants, as an adverse ruling would have threatened the entire settlement structure.

The Quinn Emanuel team delivered a masterful oral argument before the three-judge Ninth Circuit panel, effectively responding to the objector-appellants' challenges and defending the fairness and reasonableness of the settlement terms. Shortly thereafter, the Ninth Circuit issued two unanimous opinions adopting Quinn Emanuel's arguments in their entirety and affirming the district court's final approval of the settlements. Emphasizing that "the very essence of a settlement is compromise," the Ninth Circuit rejected the objector-appellants' attempts to judge the settlement "against a hypothetical or speculative measure of what might have been achieved by the negotiators."

The unanimous decisions provided conclusive validation of the settlement and secured finality for all parties, allowing Hyundai and Kia to close a significant chapter of litigation and to deliver tangible relief to millions of class members—all while avoiding the uncertainty and expense of a protracted trial.

Published by Quinn Emanuel Urquhart & Sullivan, LLP as a service to clients and friends of the firm. It is written by the firm's attorneys. The Noted with Interest section is a digest of articles and other published material. If you would like a copy of anything summarized here, please contact Elizabeth Urquhart at +44 20 7653 2311.

- We are a business litigation firm of more than 1,300 lawyers — the largest in the world devoted solely to business litigation and arbitration.
- When we represent defendants, our trial experience gets us better settlements or defense verdicts.
- When representing plaintiffs, our lawyers have garnered over US$80 billion in judgments and settlements.
- We have won eight 9-figure jury verdicts and five 10-figure jury verdicts.
- We have also obtained fifty-one 9-figure settlements and twenty 10-figure settlements.

quinn emanuel