

Edit



Configure basics

*Required



Add definitions

Workflow name* ?

Enron Energy Services - Subpoena Working



Add context



People & Roles



Organisations & Entities



Additional Context



Review & Update

Overview

Add high-level context about the matter and review goals that provide the AI Advisor with

The subpoena for Behavior Affecting Electricity Prices in California seeks documents and other records on potential market manipulation strategies, internal and external communications for interactions between Enron and regulatory bodies, as well as any internal assessments, gathering evidence for investigations into potential regulatory violations and market manipulation.

revealaji

From "Prompt and Pray" to Precision:

How Reveal aji Delivers Second-Generation GenAI for Document Review

Contents

Introduction.....3

The Challenges of Current GenAI Review Tools3

How aji Addresses Current Tool Limitations6

Multiple Methodologies for Different Needs10

The aji Difference: Three Game-Changing Features..... 12

The Path Forward: Practical Implementation..... 17

Conclusion..... 18

Introduction

The legal technology landscape has witnessed a dramatic shift with the introduction of generative AI for document review.

And while the GenAI tools currently on the market show promise in revolutionizing how legal teams handle discovery and document analysis, some early adopters have found significant limitations that have made these solutions difficult to trust and costly to implement.

At Reveal, we've been listening carefully to customers and market feedback in developing aji: our second-generation GenAI solution that directly addresses the shortcomings of today's available tools while introducing innovative features that deliver defensible, transparent results fast and through intuitive workflows.

The Challenges of Current GenAI Review Tools

When the current generation of GenAI document review tools hit the market, the promise was elegantly simple: upload your documents, write a prompt, feed it into an LLM, and—presto—return documents coded to reflect their responsiveness with the prompt. While this approach may seem to lack the refinement necessary for privilege review, it has largely gained traction as a potential substitute for attorney responsiveness review and issue tagging.

In reality, without time-consuming and often resource-intensive calibration of the AI, this approach often manifests as “prompt and pray,” where definitions assumed to pull responsive documents miss the mark and several rounds of refinement are needed. The larger the document set needed for refinement, the higher the cost.

The Black Box Problem

Perhaps the most significant challenge with existing tools is their opacity—a deterrent exceeding even that of cost, where attorneys are leery of banking on results that cannot be easily explained if challenged. Legal professionals are essentially asked to trust a black box system without understanding how decisions are made. When an AI classifies a document as responsive or non-responsive, there's often little visibility into the reasoning behind that determination. For an industry where every decision must be defensible and explainable, this lack of transparency is a non-starter.

This opacity can create particular problems during meet-and-confer discussions with opposing counsel. When challenged about classification decisions, legal teams may be limited in explaining the AI's reasoning or demonstrate the reliability of their process. The inability to provide clear explanations for AI-driven decisions leaves teams potentially vulnerable to challenges about the adequacy of their document review and whether it meets the “reasonable inquiry” standards of Rule 26.

Prompt Fear

Many legal professionals may be hesitant to deploy GenAI due to its complexity—or perceived complexity. This reluctance is only reinforced when they encounter a blank field requesting the perfect string of words to form a prompt. “Prompt fear” can produce paralysis for two reasons: it is new and potentially uncomfortable for professionals who pride themselves on doing things they're already good at; and the perception that a wayward prompt wastes precious time and client money—which it can. The brave who've embraced prompt writing have often done so with the hand-holding of technical experts or with a reliance on third-party GenAI tools, like ChatGPT, which live outside the eDiscovery platform.

Prohibitive Costs and Lack of Control

Current tools often operate on per-document processing models, creating what we call “prompt fear.” Legal teams find themselves hesitant to experiment or iterate on their definitions because each attempt carries significant financial risk. Without the ability to test and refine their approach on smaller sample sets, teams are forced to make expensive bets on untested prompts and definitions. It’s like being asked to buy a car without a test drive—while blindfolded.

The lack of cost control is compounded by limited flexibility in how the tools operate. Users have little control over the AI’s decision-making process and can’t easily adjust their approach based on initial results. This inflexibility means that teams often have to accept suboptimal results or start over entirely, multiplying both costs and time investment.

Limited Methodology Options

Currently available GenAI tools typically offer only one approach: full generative AI review across entire document sets. This one-size-fits-all methodology ignores the reality that different cases, different types of documents, and different legal objectives require different approaches. Some situations call for comprehensive AI review, while others benefit from more traditional classifier-based approaches or hybrid methodologies.

The lack of methodological flexibility means that legal teams can’t adapt their document review strategy to the specific needs of each case. Whether dealing with a complex MDL with millions of documents or a focused investigation requiring high precision on a smaller set, teams are forced to use the same basic approach – a classic “when all you have is a hammer” scenario.

Integration Challenges

Finally, most current tools operate as standalone solutions or bolt-on modules that don't integrate well with existing document review workflows. This creates inefficiencies and forces teams to work across multiple platforms, losing valuable context and requiring duplicate effort to accomplish basic tasks.

How aji Addresses Current Tool Limitations

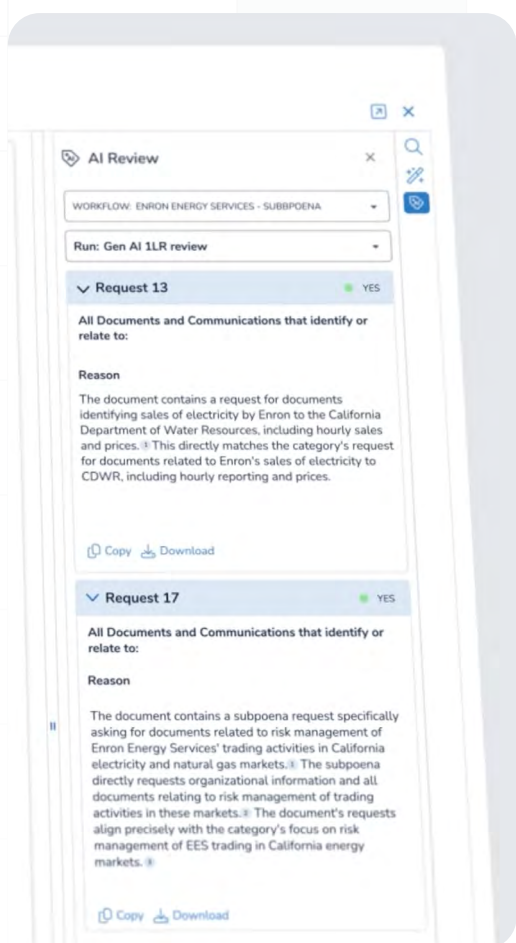


Fig 1. Each aji coding decision is accompanied by detailed reasoning and a direct citation to the underlying document.

Reveal aji represents a fundamental rethinking of how GenAI should work for document review. Rather than simply applying large language models to document classification, we've built a comprehensive system that addresses each of the major limitations we see in today's market offerings.

From Black Box to Transparent Decision-Making

aji eliminates the black box problem by providing detailed explanations for every classification decision. When the system determines that a document is responsive to a particular request, it doesn't just provide a yes/no answer—it explains exactly why that determination was made, complete with source citations that allow reviewers to jump directly to the relevant portions of the document, which are highlighted for ease of review.

This transparency serves multiple purposes.

First, it allows legal teams to understand and validate the AI's reasoning, building confidence in the results. Second, it provides the documentation necessary for discussions, should they be needed, with opposing counsel and the court, offering clear explanations for classification decisions. Finally, it enables teams to identify patterns in the AI's decision-making that can inform refinements to their definitions and approach.

Cost Control Through Calibration

One of aji's most innovative features is its calibration workflow, which allows teams to test and refine their approach on small sample sets before committing to full-scale review. Documents representative of those the user seeks to identify with AI can be selected manually. But where aji really shines is in its workflow integration with ASK, Reveal's GenAI fact-finding and analysis engine, which will automatically select documents that best reflect the definition (Note: "definition" refers to a description or prompt used to recall documents – for example, items in a request for production.)

Teams can calibrate aji with as few as 25 documents, where the AI will classify the documents in terms of agreement with the definition (e.g. strongly agree, agree, borderline, and so forth) and then be compared to human classifications of the same 25 documents. Testing can be run until sufficient alignment between the AI and the human classification is achieved. Here, the calibration process is iterative, allowing users to run multiple rounds with updated or new definitions provided by aji's patent pending AI Advisor to see how changes affect agreement rates.

This approach transforms what was once an expensive guessing game into a methodical process of refinement and validation. Teams can achieve confidence in their approach before scaling up to larger document sets.

Agreement Report by Document



	BegDoc	GenAI Rating	Manual Decision	Agreement
☐	DE-00010120	Strong Yes	Yes	✓ Agree
☐	DE-00055164	No	No	✓ Agree
☐	DE-00129070	No	No	✓ Agree
☐	DE-00190082	No	No	✓ Agree
☐	DE-00266880	No	No	✓ Agree
☐	DE-00297553	Yes	Yes	✓ Agree
☐	DE-00370856	Yes	Yes	✓ Agree
☐	DE-00438624	Strong Yes	Yes	✓ Agree
☐	DE-00524205	Yes	Yes	✓ Agree
☐	DE-00656131	Yes	Yes	✓ Agree
☐	DE-00664074	No	No	✓ Agree
☐	DE-00677431	No	Yes	✗ False Negative
☐	DE-00842079	No	No	✓ Agree
☐	DE-00852473	Yes	Yes	✓ Agree
☐	DE-00852478	Yes	Yes	✓ Agree
☐	DE-00852494	Yes	No	○ False Positive
☐	DE-00852712	Yes	No	○ False Positive

Fig 2. Calibration aligns GenAI ratings with human coding until sufficient accuracy is achieved.

This transparency serves multiple purposes. First, it allows legal teams to understand and validate the AI's reasoning, building confidence in the results. Second, it provides the documentation necessary for discussions, should they be needed, with opposing counsel and the court, offering clear explanations for classification decisions. Finally, it enables teams to identify patterns in the AI's decision-making that can inform refinements to their definitions and approach.

GenAI in eDiscovery: Federal Courts Break New Ground

Federal courts increasingly appear to be receptive to parties using AI as a substitute for attorney responsiveness review in eDiscovery workflows. Recent judicial orders in cases like *EEOC v. Tesla, Inc.*, 727 F. Supp. 3d 875 (N.D. Cal. 2024) represent potentially groundbreaking developments where courts have, as part of discovery protocols agreed to by the parties, approved GenAI for “first pass” document review—marking a shift from experimental technology to accepted practice.

Key findings from recent federal cases reveal that courts are establishing **three core requirements for GenAI implementation**: parties must demonstrate how the underlying technology works, validate results through statistically sound methodologies, and engage in meaningful meet-and-confer processes about search protocols.

The Tesla case stands out as among the first known orders explicitly greenlighting “GEN AI” for attorney responsiveness review, requiring parties to agree on recall rates and effectiveness measures. Similarly, U.S. District Court Judge Arun Subramanian in *The Estee Lauder Company Securities Litigation*, Case Number 1:24-cv-00716, approved artificial intelligence use contingent on disclosure of specific review applications and protocol agreements.

3 Core Requirements for Using GenAI

1.

Show How It Works

Parties must demonstrate how the underlying technology works.

2.

Validate Results

Parties must validate results through statistically sound methodologies.

3.

Meet and Confer

Parties must engage in meaningful meet-and-confer processes about search protocols.

Multiple Methodologies for Different Needs

Notably, aji offers two distinct approaches to GenAI document review, allowing teams to choose the methodology that best fits their specific situation:

Full GenAI Review: This approach allows teams to leverage calibration rounds to refine their definitions, then apply the AI directly to their full document set. It's ideal for situations where teams want the speed and consistency of AI classification across large volumes of documents.

Hybrid GenAI/TAR Mode: This innovative approach flips the traditional model by starting with control sets created by the subject matter expert and then using GenAI to continue training a classifier until it reaches specified precision and recall targets. This method combines the tried-and-true reliability of technology-assisted review with the power of generative AI, creating a solution that's both innovative and defensible.

Seamless Integration

It is not uncommon for first-gen GenAI tools to exist as separate modules from their core platforms. This means that, for example, the automated calibration aji empowers must not only be conducted manually with other solutions by identifying representative documents through keyword searching or the development of scripts – but that these processes occur in an environment external to the GenAI tool.

Unlike these bolt-on solutions, aji is fully integrated into Reveal's comprehensive document review platform. It works seamlessly with existing features like ASK, concept searching, clustering, and pre-built models. This integration means that teams can use aji as part of a comprehensive discovery strategy rather than as an isolated tool.

For example, teams can use concept searching to identify relevant documents, leverage ASK to understand key themes and entities, and then apply aji for systematic classification—all within a single platform that maintains context and workflow continuity throughout the process.




Fig 3. ASK, Reveal's novel GenAI engine for fact-finding and document analysis, can be paired with aji to speed calibration such that selection of documents most relevant to definitions is automated.

The aji Difference: Three Game-Changing Features

Reveal's aji stands apart from both first-generation GenAI tools and current competitors through three key differentiators that address the core challenges legal teams face when implementing AI for document review.

1. The AI Advisor: Intelligent Definition Refinement

One of aji's most distinctive features is the AI Advisor, which helps legal teams refine their definitions before beginning the classification process, eliminating much of the fear and paralysis that prompt writing can trigger. Rather than leaving teams to guess whether their definitions are optimal, or relying on third-party tools to iterate on them, the AI Advisor analyzes definitions and provides specific, actionable recommendations that improve the quality of the prompt.

The AI Advisor identifies opportunities for improvement such as specifying acceptable document formats, suggesting relevant acronyms (like recommending "CDWR" for "California Department of Water Resources"), and highlighting strong elements in definitions that include specific actions, timeframes, and measurable criteria. Importantly, these are recommendations, not mandatory changes—teams retain full control over their definitions while benefiting from AI-powered insights.

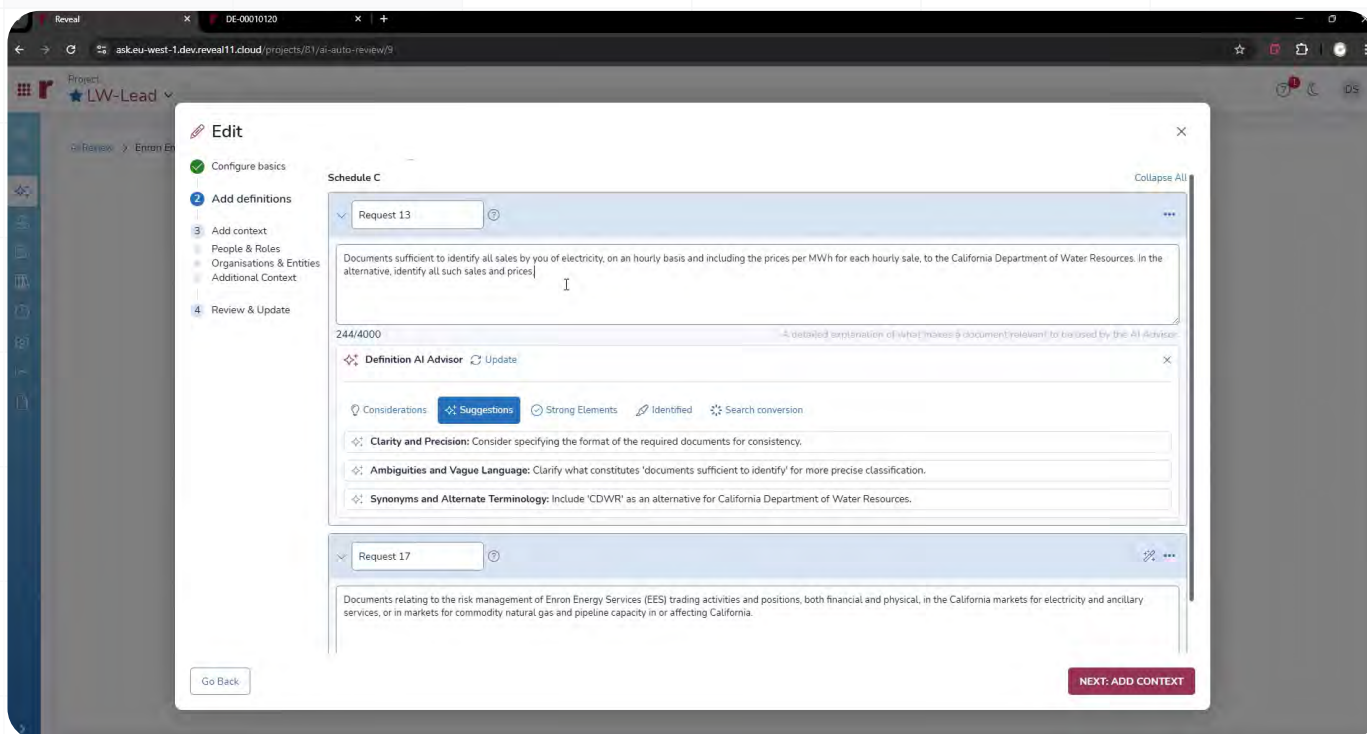


Fig 4. aji's patent-pending AI Advisor helps refine definitions with considerations and suggestions to increase the quality of the prompt.

This feature addresses the “prompt fear” problem by giving legal professionals confidence that their definitions are well-structured before they invest in large-scale classification. It also helps ensure that definitions are comprehensive and precise, reducing the likelihood of missing relevant documents or capturing too many false positives.

The AI Advisor represents a fundamental shift from the “prompt and pray” approach of first-generation tools. Instead of hoping that initial prompts will work, teams can refine their definitions based on intelligent analysis, significantly improving their chances of success from the outset.

2. Advanced Calibration Using ASK Technology

aji's calibration feature leverages ASK technology to intelligently select documents that are most likely to test the boundaries of the classification definitions.

The calibration process creates a vector database from the text in the project and uses this to identify high-scoring examples that represent different types of potentially responsive content. This ensures that calibration sets include documents that truly test the AI's ability to make nuanced classification decisions, rather than just randomly selected documents that might not be representative of the challenges the AI will face.

The system tracks agreement rates, false positives, and false negatives across multiple calibration rounds, providing detailed statistics and visualizations that help teams understand how their definitions are performing. After each round, aji provides suggested edits to definitions based on the patterns observed in disagreements between AI and human classification.

This iterative approach allows teams to systematically improve their classification accuracy before committing to full-scale review. Teams can run multiple calibration rounds, each building on the insights from the previous round, until they achieve the level of agreement and confidence they need for their specific case.

3. Hybrid GenAI/TAR Mode: The Best of Both Worlds

Perhaps aji's most innovative differentiator is its hybrid mode, which combines the proven reliability of technology-assisted review (TAR) with the power of generative AI. This approach addresses a key concern among legal professionals who want to leverage AI capabilities while maintaining the defensible (court-accepted) methodology of traditional classifier-based review.

In hybrid mode, teams begin by creating a control set through traditional human review. This control set serves as the foundation for training a classifier, just like in traditional TAR workflows. However, instead of requiring human reviewers to continue labeling documents until the classifier reaches acceptable precision and recall levels, the GenAI takes over this training process.

The AI continues to label documents and train the classifier until it reaches user-defined precision and recall targets. This approach dramatically reduces the human review burden while maintaining the statistical rigor and defensibility that legal teams require. The final results include all the standard TAR metrics—precision, recall, F1 scores—that have been accepted by courts and opposing counsel for years.

This methodology is particularly valuable for time-sensitive matters where teams need to quickly achieve high-quality results. It also appeals to legal teams hesitant to embrace GenAI due to cost, whether real or presumed. Hybrid mode enables teams to intelligently focus the benefits of GenAI on the documents that need it, while combining with proven technologies that fulfill legal obligations across the broader document set to achieve the best combination of cost and value.




Fig 5. Hybrid mode allows users to deploy GenAI to train existing TAR workflows.

Perhaps most importantly, aji's hybrid mode allows organizations to leverage TAR workflows they've refined over the course of time and matters. Many legal teams have invested significant time and resources developing sophisticated TAR methodologies, complete with detailed documentation, reporting protocols, and quality assurance processes. These workflows often represent an organization's secret sauce—their competitive advantage built through years of refinement and optimization. Rather than forcing teams to abandon this intellectual property in favor of entirely new AI-driven processes, aji's hybrid mode allows them to enhance their existing workflows with GenAI capabilities. Teams can apply their proven control set methodologies and established quality metrics while letting AI handle the labor-intensive training process, preserving institutional knowledge while gaining efficiency.

The Path Forward: Practical Implementation

The development of aji reflects extensive feedback from legal professionals who have used first-generation GenAI tools and understand their limitations. By addressing the core challenges of cost, transparency, flexibility, and integration, aji provides a foundation for legal teams to confidently implement AI-powered document review.

The system's flexibility means that no two implementations need to be the same. Teams working, for example, on complex MDL cases can leverage the full power of GenAI review across massive document sets, while those handling focused investigations can use hybrid mode for maximum precision and defensibility. The calibration features ensure that teams can validate their approach regardless of which methodology they choose.

Most importantly, aji's integrated approach means that AI-powered review becomes part of a comprehensive discovery strategy rather than an isolated tool. Teams can seamlessly move between different types of analysis—conceptual searching, entity identification, classification, and manual review—while maintaining context and building on insights gained at each stage.

Conclusion

The evolution from first-generation to second-generation GenAI for document review represents more than just technological advancement—it reflects a deeper understanding of how legal professionals actually work and what they need to successfully implement AI in their practice. Reveal aji addresses the fundamental limitations that prevented widespread adoption of earlier tools while introducing innovations that make AI-powered document review more powerful, flexible, and, crucially, defensible.

By moving beyond the “prompt and pray” approach to provide transparency, cost control, methodological flexibility, and seamless integration, aji enables legal teams to harness the full potential of generative AI while maintaining the rigor and defensibility that the legal profession demands. As the legal technology landscape continues to evolve, solutions like aji that prioritize both innovation and practicality will define the future of document review and discovery.



Watch the aji explainer video

About Reveal

Reveal is a leading AI-powered platform for eDiscovery, document review, legal hold and investigations. Its best-in-breed suite of products include Logikcull, Onna, Reveal Hold, Reveal Enterprise and Brainspace. The company has a deep history in driving the adoption of legal automation, which is underpinned by its leading processing technology, visual analytics, and artificial intelligence capabilities. Reveal's software combines technology and human guidance to transform structured and unstructured data into actionable insight. We help organizations, including law firms, corporations, government agencies, and intelligence services, uncover more useful information faster by providing a world-class user experience and AI technology that is embedded within every phase of the eDiscovery process.

For more, visit revealdata.com.