


# Using GenAI To Navigate Complexity, Not Replace Judgment.



## WHAT'S INSIDE

A closer look at why the document-centric model is failing, and what the teams getting it right are doing differently.



*“Where is the wisdom we have lost in knowledge?  
Where is the knowledge we have lost in information?”*

— T.S. ELIOT

## Defensible Review in Modern Discovery Requires More Than Better Technology

For years, the legal industry has approached document review as a problem of scale. As data volumes increased, the response was predictable: more reviewers, more hours, more cost. One of the primary levers used to manage that cost was to move review work progressively away from core case teams. What began with review handled by key case team attorneys shifted to more junior associates, then to staff attorneys, and eventually, for many years now, to large teams of contract attorneys. Technology helped at the margins. Predictive coding and TAR enhanced the contract-attorney model and made it more efficient. But the underlying structure of review remained largely unchanged. Review was still document-centric, labor-heavy, and linear.

### **The document-centric model is no longer sufficient.**

Modern discovery is defined not just by volume, but by complexity. The composition of ESI (electronically stored information) is constantly changing, and discovery habits must evolve to match it. Email is no longer the dominant data type. Today's matters routinely include chat platforms, collaboration tools, non-text images, hyperlinked and interdependent files, and structured or semi-structured data that does not behave like documents at all.

In this environment, effective fact-finding requires more than simply reviewing faster. It requires adapting how information is examined, connected, and prioritized, so efficiency does not come at the expense of effectiveness. Treating this information as if it were a collection of standalone files is increasingly ineffective and, more importantly, increasingly risky.

At the same time, defensibility expectations remain a high priority, but the nature of those expectations has evolved. Courts, regulators, and opposing parties now expect not only that review decisions be reasonable and consistent, but that the process used to reach those decisions can be clearly explained, debated, and, where necessary, defended in detail. There are more ways than ever to navigate the growing diversity of ESI types, which has expanded both the options and the scrutiny around how review is conducted. This has created a greater need to triage ESI thoughtfully and, in many cases, to apply different technologies and workflows to different categories of data. Speed and efficiency remain important, but they cannot come at the expense of defensibility.

# Why Document-Centric Review Breaks Down



Traditional review workflows are document-centric by design. They assume that a document is a discrete, static unit of meaning, and that relevance or responsiveness can be determined by reading it in isolation. Review progresses linearly, one document at a time. Context is accumulated slowly, through repetition and volume.

**That approach struggles in a modern data environment.**

Chats derive meaning from sequence and participants. Images often require surrounding context to be interpreted correctly. Hyperlinked files only make sense when viewed together, and in some instances may no longer exist in their original form – or at all. Structured data cannot be reviewed linearly at all. When these data types are forced into document-centric workflows, context is fragmented and reviewer judgment is diluted.

What is less often acknowledged is that this limitation applies not only to newer data types, but also to traditional documents. Even email and office documents increasingly exist within dense webs of relationships, metadata, and behavioral patterns. When documents are treated only as standalone artifacts, opportunities to identify themes, connections, and priority areas early are lost.

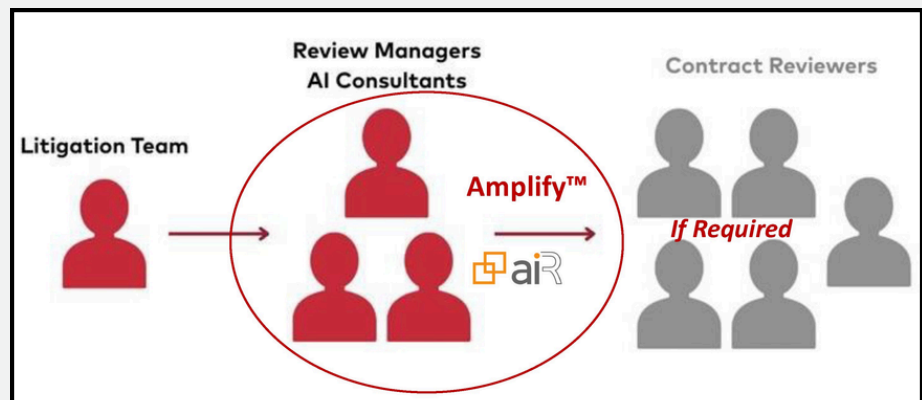
The result is familiar to anyone who has managed a large review. Early review is broad and expensive. Insight arrives late. Critical documents are often discovered after significant time and cost have already been incurred. Review teams spend more time navigating noise than analyzing substance.

This is why triage has never been more important. In modern matters, there are exponentially more paths through the data. Deciding where to focus attention is no longer an administrative step. It is a strategic one that directly affects defensibility.

# Data-Centric Review Starts with Human Judgment

A more defensible approach begins by treating content as data before treating it as documents. A data-centric review model seeks to understand how information behaves across a dataset. It prioritizes relationships, patterns, and signals over linear reading. Most importantly, it applies human judgment early, when it has the greatest impact. This approach is particularly well suited to newer ESI types such as chats, images, and structured data, which naturally benefit from pattern-based analysis and contextual grouping. At the same time, it materially improves how traditional documents are handled. By analyzing documents through a data-centric lens, legal teams can surface themes, custodial behavior, and issue clusters before committing to large-scale eyes-on review.

**This is the foundation of Amplify Review.**



Rather than beginning with broad first-pass review, Amplify Review starts with early data assessment and data-centric triage. Senior legal and review subject matter experts engage with the data immediately, defining what constitutes relevance, priority, and risk in the context of the matter. Low-value content is identified and removed from consideration. High-signal areas are surfaced early.

These decisions are documented, repeatable, and explainable. They are not delegated to volume-based staffing models. They are made and guided by highly engaged and invested subject matter experts who remain close to the facts, the legal issues, and the objectives of the matter. This early human governance determines how different data types will be treated, which technologies are appropriate, and where human attention should ultimately be focused. By keeping judgment anchored with experienced case team expertise, review strategy is informed by substance rather than diluted by scale, setting the conditions for defensible use of advanced tools later in the process.

# Using GenAI To Navigate Complexity, Not Replace Judgment

Within a data-centric review framework, GenAI expands the set of tools available to legal teams, but it does not eliminate the need for judgment. GenAI performs unevenly across different ESI types. It is particularly effective when applied to text-heavy material for thematic analysis, issue identification, and large-scale categorization. It is less reliable when applied without preparation to non-text images, chat conversations, or structured and semi-structured data.

Recognizing these differences is not a limitation. It is a prerequisite for defensible use. This is where Relativity aiR for Review is most effective. Applied after data-centric triage, aiR helps legal teams structure complexity by categorizing data across relevance, issues, and themes at a scale that is not feasible through human review alone. Its value is not simply that it operates broadly, but that it surfaces patterns, contradictions, and areas of uncertainty that merit human attention.

aiR is designed to be explainable. When it assigns a categorization or score, it also explains why it reached that conclusion, identifying the signals that influenced its assessment. Importantly, it will also surface countervailing considerations, effectively arguing with itself and highlighting where ambiguity exists. This allows review teams to understand not just what the AI suggests, but where additional scrutiny may be required.

Citations and reasoning are grounded directly in the source text. Referenced passages are validated against the original content, and unsupported assertions are not presented as fact. This materially reduces hallucination risk and ensures that AI outputs remain provable, reviewable, and challengeable.

Throughout this process, humans remain in control. Senior SMEs define the questions aiR is being asked to evaluate, informed by the composition of the data and the objectives of the matter. Prompts are shaped by triage decisions, not applied generically. AI outputs are treated as hypotheses to be tested, not conclusions to be accepted. Sampling, benchmarking, and reasonableness testing are embedded into the workflow, and refinements are documented as part of review governance.

# THE RISK OF MISAPPLIED GENAI

The growing availability of GenAI has also introduced new risks when it is applied indiscriminately: Using GenAI without data-centric triage, without understanding the composition of the dataset, or without clearly defined prompts can create a false sense of confidence. In those cases, AI may appear to be efficient while actually obscuring blind spots.

The risk is not that GenAI is inherently inaccurate. It is possible that it may be asked the wrong questions, applied to the wrong data, or relied upon without sufficient context. Defensible review depends on avoiding those missteps by ensuring that GenAI is deployed selectively, with human judgment directing where and how it is used.


## HUMAN REVIEW, REPOSITIONED RATHER THAN ELIMINATED

A common misconception about modern review models is that they seek to eliminate human reviewers. In reality, the opposite is true. Human review remains essential, but it is repositioned to be more effective.

In traditional document-centric review models, reviewers receive large, undifferentiated queues of documents and are asked to infer relevance from scratch. Consistency depends on layers of quality control applied after the fact.

In a data-centric review model, human review occurs after triage and categorization. Reviewers receive targeted, purpose-built batches of documents, each accompanied by clear context. They understand what the documents are, why they are being reviewed, and what specific questions they are being asked to answer. That context is derived from early data assessment, Amplify categorization, and aiR reasoning.

The result is greater consistency and more effective use of human effort. Reviewers validate and confirm insights rather than discovering them blindly. Escalations and re-review cycles are reduced, and human judgment is applied where it adds the most value. Crucially, human reviewers and SMEs retain final responsibility for review decisions and for how those decisions are defended.



# What Defensible Review Looks Like Now

Defensible review in modern discovery is not defined by how quickly documents can be processed. It is defined by how early insight can be achieved, how transparently decisions can be explained, and how clearly responsibility is owned.

A defining characteristic of modern defensible review is that decision-making is documented as it occurs, rather than reconstructed after the fact. Through data-centric triage and early categorization, Amplify captures and preserves the criteria used to prioritize data, route different ESI types through appropriate workflows, and scope human review. These decisions are not implicit or anecdotal. They are explicit, repeatable, and auditable.

Relativity aiR for Review extends that transparency. In addition to producing categorizations, aiR records the reasoning behind those categorizations, including the signals considered, the supporting source text, and areas of uncertainty. This creates a clear record of why content was treated a certain way and where human judgment was applied. Decision logic is visible and reviewable, rather than buried in reviewer notes or inferred from outcomes.

This level of transparency is fundamentally different from traditional first-pass review, even when enhanced with TAR or CAL. In those models, scoring and prioritization often occur without clear, contemporaneous documentation of why specific decisions were made, and defensibility frequently depends on post hoc explanations of review methodology. Modern workflows built around Amplify and aiR shift that dynamic by embedding explainability and governance directly into the review process.

The result is a review record that can be defended on its own terms. Legal teams are able to explain not just what decisions were made, but how and why they were made, using documented criteria rather than generalized descriptions of process. That transparency supports meet-and-confer discussions, regulatory inquiries, and judicial scrutiny with a level of clarity that traditional models struggle to provide.

Technology that accelerates review without embedding this kind of governance simply moves problems downstream. By contrast, a data-centric, human-governed approach aligns efficiency, cost control, and defensibility. It allows legal teams to understand their data sooner, apply the right tools to the right data, and defend their decisions with confidence.



# Conclusion

This is not about adopting technology for its own sake. It is about using technology the right way, in service of sound legal judgment.

If you are rethinking how review is conducted on your matters, the right starting point is not a tool selection. It is a review strategy discussion grounded in how modern data behaves and how defensibility is maintained. That conversation is where meaningful change begins.



lineal.com

Data-Centric eDiscovery  
and Managed Services.

Want more insights like this?

# Subscribe to Our Monthly Newsletter

Get first access to more things like this!

GET MORE TODAY

