

Learn more about better  
eDiscovery tech and services.



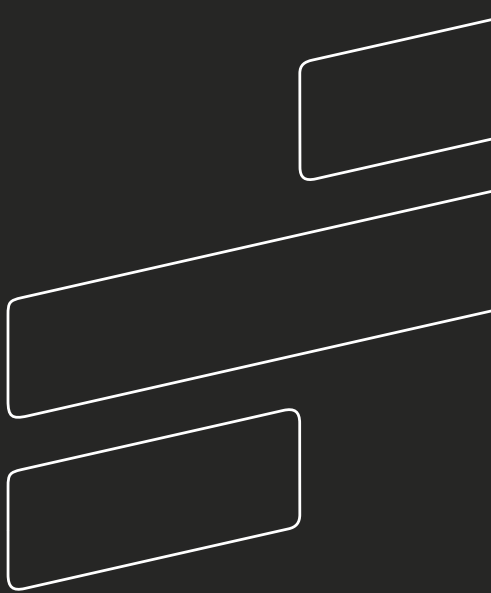
Flipping the Script on AI

# How Human-Directed AI Turns Intelligence Amplified (IA)



## WHAT'S INSIDE

Discover how legal teams can “flip the script” on AI by using human-directed intelligence amplification to unlock faster insights, stronger defensibility, and reduced risk.



# Introduction

Artificial intelligence (especially generative AI) has rapidly evolved from experimental technology to a core enabler across industries, and nowhere is the momentum more evident than in the legal industry. In just a few years, lawyers, legal operations professionals, and eDiscovery specialists have found themselves navigating a tidal wave of AI tools promising efficiency, insight, and competitive advantage. And yet, alongside the promise, the profession has witnessed a rising number of “AI gone wild” moments, such as hallucinated case citations, fabricated precedent, and unpredictable system behavior. These incidents have not only made headlines but have also sparked deep anxiety about the risks of adopting GenAI and similar tools in a profession where accuracy, ethics, and defensibility are non-negotiable.

The message behind these failures is clear: when AI operates without meaningful human direction, oversight, and validation, it can become a liability. Conversely, when AI is integrated within a human-directed, controlled, and purpose-built governance framework, it becomes a powerful amplifier of human intelligence – not a replacement for it. The shift from AI as a decision-maker to AI as an augmentation engine helps legal professionals extract meaning from data faster, more accurately, and with greater defensibility.

**This paper discusses how success with AI in legal workflows hinges on “flipping the script” from artificial intelligence (AI) to intelligence amplified (IA).** We’ll discuss how AI’s role becomes most valuable when humans drive the process: setting goals, defining parameters, validating outputs, and integrating AI within the broader demands of legal practice and eDiscovery. We will also illustrate how Lineal’s Amplify™ solution inside RelativityOne embodies this reimagined paradigm: keeping attorneys firmly in control while empowering them with accelerated insight, and reducing errors, noise, and cost. **The destination is a future where AI doesn’t replace human judgment; it elevates it.**

FLIPPING THE SCRIPT ON AI

# When AI Goes Off the Rails



# When AI Goes Off the Rails

Many of you are probably aware of how ubiquitous GenAI is becoming in the legal community and in society overall. You're probably also aware of many "AI gone wild" stories where the AI models – left unchecked and unvalidated – have delivered hallucinated results that were wildly inaccurate<sup>1</sup>. At best, these AI failures cause considerable embarrassment; at worst, they impact careers. Here, we will discuss some high-profile AI failures – in both the legal community and society overall – how that has escalated overall anxiety about AI technology, and the common denominator in these AI failures.

---

<sup>1</sup> Apparently, not everyone is or they wouldn't keep happening.

## High-Profile Failures in the Legal Community

AI's growing presence in legal practice has produced a series of widely publicized missteps that illustrate both the dangers of hallucination and the consequences of deploying general-purpose generative models in high-stakes legal contexts. The most infamous example remains the [Mata v. Avianca](#) case in 2023, in which attorneys filed a legal brief containing several fabricated cases generated by ChatGPT. The model confidently cited nonexistent opinions, misstated holdings, and produced precedent that collapsed under even casual review. As the court summarized in an [earlier ruling](#), "Six of the submitted cases appear to be bogus judicial decisions with bogus quotes and bogus internal citations." The resulting sanctions were not only embarrassing but also a cautionary tale: AI is not a lawyer, and unverified AI output can derail even routine filings.

It was a cautionary tale that received considerable publicity, seemingly putting the entire legal profession on notice. Despite that, it was not an isolated event. Courts across the United States began issuing standing orders requiring attorneys to disclose whether GenAI assisted in drafting filings, emphasizing that lawyers remain responsible for verifying all citations and factual assertions. According to the Duke Center on Law & Tech's [Responsible AI in Legal Services \(RAILS\) resource](#), there are over 100 AI related court orders, local rules, and guidelines from the U.S. and other countries.

Surely, the combination of cautionary tales and court orders would slow down the cadence of filings with AI hallucinations, right? In fact, the opposite has happened: the number of cases involving AI hallucinations has increased – **dramatically**. Lawyers are doing it, pro se parties are doing it, experts are doing it (even AI experts as demonstrated in [this case](#) where the AI expert's declaration was thrown out after it contained citations hallucinated by AI) and, yes, even judges are doing it. According to Damien Charlotin's **AI Hallucination Cases database**, there were **798** cases worldwide involving AI hallucinations from 2023 through 2025. And most of them – a whopping **91.3%** – occurred in 2025!

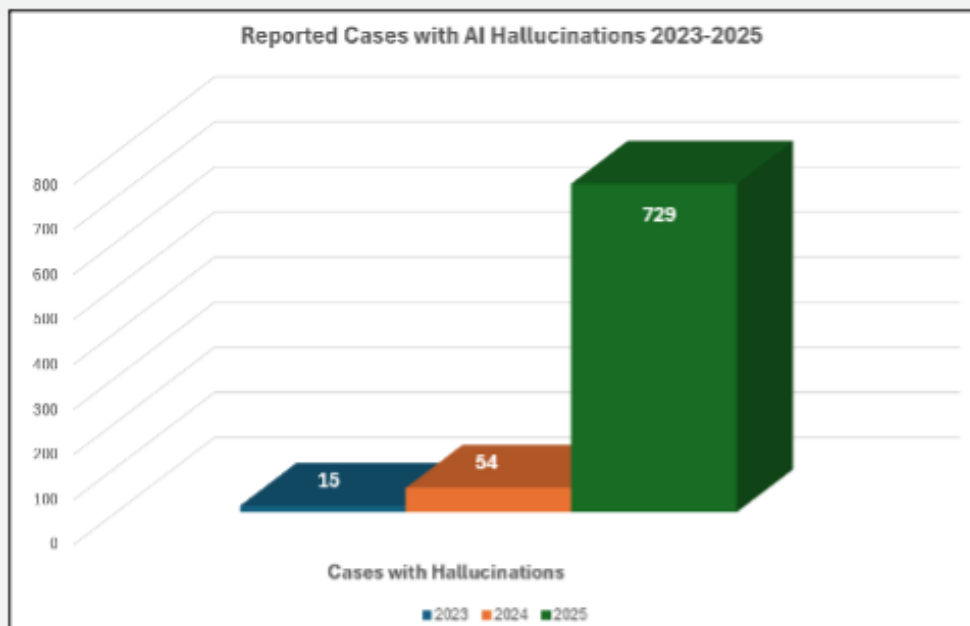


Figure 1: Reported Cases with AI Hallucinations 2023-2025 (Source: [Damien Charlotin](#))

These examples underscore an essential fact: GenAI models hallucinate regularly. When AI models are used as an autonomous "answer generator," those models risk creating falsehoods wrapped in perfectly fluent language (after all, they are **language** models), which is dangerous in a profession where precision is paramount. Lawyers are treating the **technology** as the deliverable – not their **own legal expertise leveraged by technology** – often leading to disastrous results.

## High-Profile Failures in Society

Legal professionals are not alone in confronting AI's unpredictability. There are many examples in society of "AI gone wild" including:

- Google's AI Overviews tool **telling** users to use non-toxic glue to make cheese stick to pizza better, and that geologists recommend humans eat one rock per day.
- Google Gemini's image generator **putting** people of color in Nazi uniforms.
- Microsoft **including** the Ottawa Food Bank in a travel article on Ottawa, Canada, telling readers to "Consider going into it on an empty stomach."
- An AI bot that handles tech support for Cursor (a programming tool) **announcing** that users were longer allowed to use Cursor on more than just one computer – a policy change that did not exist.
- ChatGPT **telling** a user that Christof Wandratsch of Germany holds the world record for crossing the English channel – on foot.
- Google's AI Overviews tool **stating** that Lady Gaga is two days older than Ariana Grande, while also stating that Gaga was born on March 28, 1986 while Grande was born on June 26, 1993. Although, Grande did once **joke** that Gaga was only two days older.

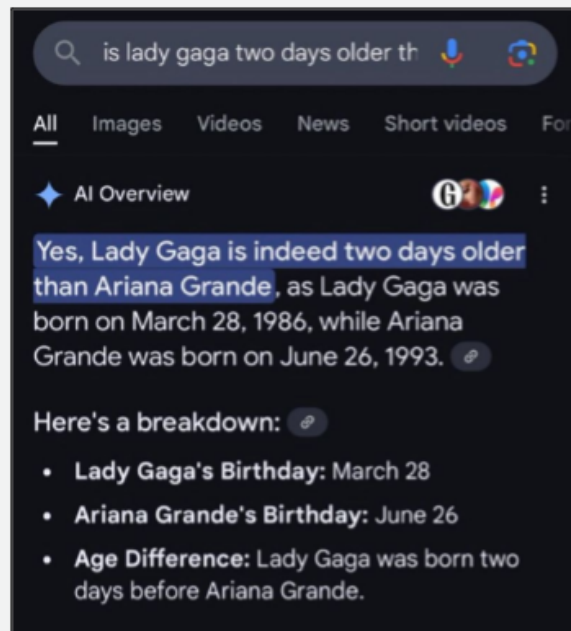


Figure 2: Example of an AI Hallucination from Google's AI Overviews

In addition to these general examples, higher education has faced waves of AI-related misconduct, from students submitting AI-generated essays that contain fabricated citations to researchers publishing papers later found to include hallucinated sources, inaccurate references, or AI-generated passages masked as genuine scholarship.

The [site \*Academ-AI\*](#) documents the suspected instances of AI being used to author research without appropriate declaration. Many of these instances are not subtle: they're phrases that tip off the use of AI that include: "Certainly, here are", "As of our knowledge cutoff", "As an AI language model, I must" and "new developments may have occurred since my last update"! As of this writing, the site identifies **655** journal articles, **114** conference papers, and **35** books and chapters (**804** publications in all!) with obvious use of AI in their creation.

There are plenty of other examples in society. Business professionals have submitted AI-generated reports that misstated data trends or invented financial figures. Newsrooms have faced public backlash after publishing AI-generated stories filled with factual errors or nonsensical narrative jumps. Social media has amplified the problem, with AI-generated images, transcripts, and deepfakes spreading misinformation at unprecedented speed.

Across sectors, the pattern is consistent: AI systems that are used as autonomous decision-makers or factual sources instead of as controlled assistants can misrepresent, fabricate, or distort information. The risks escalate when users assume accuracy without implementing verification steps. And because AI outputs can appear polished and confident, users often mistake language fluency for accuracy.

## **Broader Industry Anxiety About AI**

These high-profile failures have fueled widespread industry hesitancy. Many legal professionals express concerns about the "black box" nature of AI models: questioning how conclusions are reached, what data underlies them, and whether outputs can be reliably defended in court. This anxiety is compounded by fears about job displacement, ethical ambiguity, disclosure obligations, and the risk of inadvertently exposing client-sensitive data to third-party systems.

Hallucinations, bias, and unpredictable model behavior undermine confidence. Without transparency, lawyers struggle to assess when an AI system is appropriate and when it poses unacceptable risk. "Shadow AI" usage, where employees informally use generative AI tools without organizational approval or governance, magnifies the issue by creating unmonitored, unvalidated risk vectors.

The result is anxiety: legal professionals recognize AI's transformative potential but remain wary of adopting tools that could lead to sanctions, errors, reputational harm, or data breaches. The key to resolving this anxiety lies in reframing AI not as a replacement for human judgment but as a force multiplier with structured, validated, and monitored within workflows that keep professionals firmly in control.

## The Common Denominator in AI Failures

Despite the wide variety of AI failures, they share a simple root cause: misplaced expectations and a lack of disciplined human oversight. When users expect generative AI to always produce fully accurate, autonomously correct answers, failure is inevitable. When organizations treat AI as a replacement for human decision-making instead of an efficiency accelerator, they open the door to litigation risk, professional embarrassment, and operational disruption.

AI failures result from several intertwined missteps:

- **Unrealistic Expectations:** The belief that AI can autonomously deliver legally accurate conclusions reliably misunderstands its purpose and nature. Generative systems are pattern mimics, not fact-verifying engines.
- **Insufficient Oversight:** Accurate legal work requires domain expertise, context, and nuance: elements that AI cannot independently infer without guidance. Errors happen when outputs are accepted without verification.
- **Incomplete or Poorly Structured Prompts:** AI models rely heavily on the framing they receive. Lack of clarity, missing context, or ambiguous instructions lead to outputs that may be fluent but incorrect.
- **Lack of Validation and Quality Control:** Hallucinations and inaccuracies often go unnoticed when outputs are not subject to systematic review, cross-checking, or human validation.

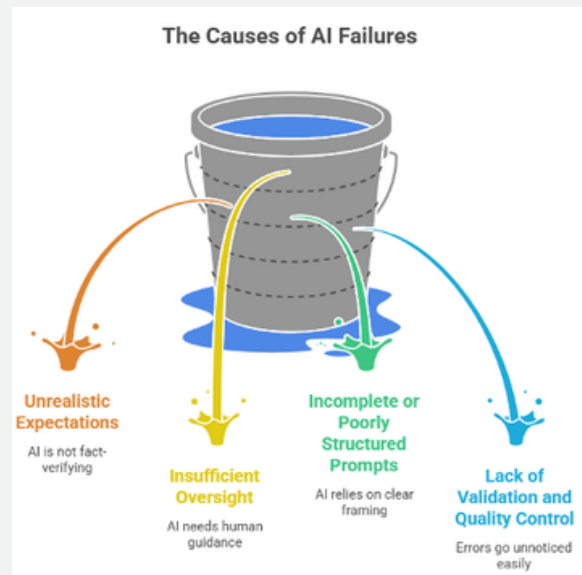


Figure 3: The Causes of AI Failures



AI is Here to Assist Legal Professionals, Not Replace Them

**The lesson is clear for legal professionals: every AI deliverable is still a legal services deliverable.** AI does not eliminate the responsibility of review, analysis, or verification. Instead, its value lies in **assisting** legal professionals by accelerating tasks, surfacing insights, and consolidating complex data while legal professionals remain accountable for accuracy. This mindset shift is the foundation of the Intelligence Amplified approach.

# Addressing AI Concerns with the Human-in- the-Loop (HITL) Approach




# Addressing AI Concerns with the Human-in-the-Loop (HITL) Approach

The examples of AI failures illustrated above demonstrate that the approach of treating AI models as a “Staples easy button” does not work. Leveraging AI technology successfully requires a “Human-in-the-Loop” (HITL) approach where the human thoughtfully guides the AI model toward the desired result and validates that result to ensure the objectives were met. Here, we’ll discuss AI’s role in legal and eDiscovery workflows, and what a HITL approach looks like.

## AI’s Role in Legal and eDiscovery Workflows

The power of AI to streamline processes and workflows is undeniable. But if anything is to be learned by the examples above it’s this: In legal practice, AI should not function as the final decision-maker. The power of AI is to enable it to serve as an intermediate analytical tool that accelerates human work without replacing the need for professional judgment. Regardless of the task involved, AI operates best when its outputs feed into a human-managed workflow.

In eDiscovery, this distinction is essential. Tasks like classifying documents (for responsiveness, privilege or key issues), summarizing documents, identifying patterns or extracting entities can be highly efficient with AI, but the outputs still require careful review to be defensible. AI models make mistakes and the failure to implement a defensible approach to check for those mistakes can subject your legal team to in-depth court scrutiny of your entire discovery process. Courts don’t expect perfection, but they do expect a reasonable approach that involves verification and validation of AI outputs.



## What HITL Looks Like

A robust Human-in-the-Loop model positions AI as an assistant, not an authority. Under HITL, humans set direction, supervise performance, validate results, and continuously refine the system. This structure preserves control, ensures defensibility, and mitigates risks associated with model drift or hallucination.

### ***Oversight and Control***

In a HITL framework, humans define the goals, context, and constraints of AI behavior. They set parameters, establish guardrails, and determine how outputs are used. This ensures that AI aligns with professional standards, legal requirements, and organizational policies. For example, attorneys determine what constitutes responsive content, how privilege is defined, and what level of review is required before production. AI simply accelerates their ability to operationalize these standards across large data volumes.

### ***Validation and Quality Assurance***

Human validation is the cornerstone of defensible AI use. In legal workflows, human reviewers confirm whether AI-identified documents are truly relevant, privileged, or responsive. They check summaries for accuracy, assess whether reconstructed timelines reflect the actual context, and review suggested classifications for consistency. This oversight prevents errors from propagating and ensures that AI-enabled decisions can withstand scrutiny.

### ***Continuous Feedback and Learning***

HITL also enables continuous system improvement. Whether it's through confirming correct outputs or correcting mistakes, human feedback refines AI performance over time. This real-world grounding reduces bias, increases precision, and ensures alignment with evolving legal standards and ethical obligations. Over time, AI becomes more attuned to the specific needs of the case, the team, and the domain, but always under human supervision.

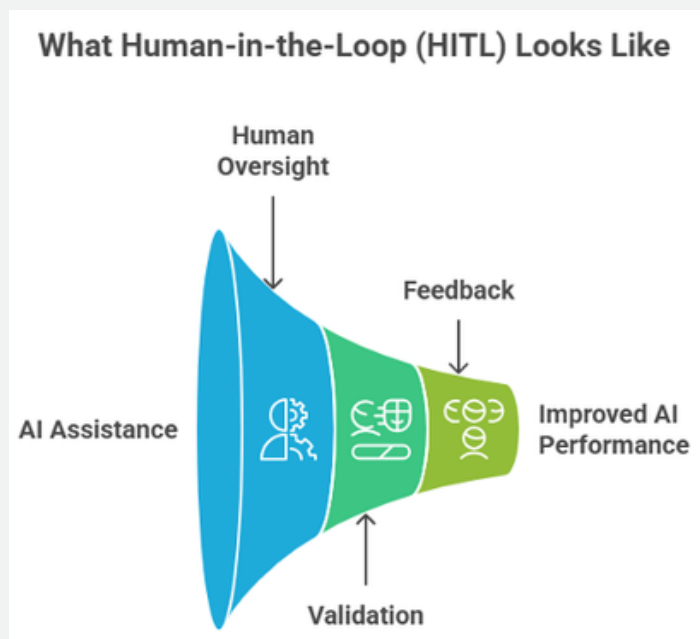


Figure 4: What a Human-in-the-Loop (HITL) Approach Looks Like

FLIPPING THE SCRIPT ON AI

# Flipping the Script from AI to IA



# Flipping the Script from AI to IA

We titled this paper "Flipping the Script on AI" because it takes a conceptual leap to think about AI differently. We've demonstrated that artificial intelligence is not always accurate intelligence, but with a human "in the loop", human intelligence can be amplified. Here, we'll discuss how the Intelligence Amplified approach to leveraging AI models can be used to maximize the benefits of AI while minimizing the risks – highlighting specific benefits of an IA approach in eDiscovery.

## **Taking the Conceptual Leap**

The shift from AI to IA represents a fundamental reorientation in how organizations conceptualize technology's role in legal work. Rather than treating AI as a replacement for human judgment, IA positions AI as a mechanism that strengthens, extends, and complements human decision-making. It transforms AI from an "answer generator" to a "judgment amplifier," enabling legal experts to work more quickly, more accurately, and with greater strategic insight. It's truly an inverted AI approach!

## **Intelligence Amplified (IA)**

In an IA model, AI is a force multiplier. It helps lawyers process information at scale, identify signals within noise, and synthesize insights that would otherwise take days or weeks to uncover. Importantly, IA does not dilute human accountability; it enhances human capacity. Attorneys retain control over relevance determinations, privilege assertions, and strategic decisions. IA adds speed, structure, and depth, but humans define the outcomes.

## **Purpose-Built AI**

One of the clearest advantages of IA is the use of purpose-built AI systems rather than general-purpose models. Tools like ChatGPT are powerful but not designed with specific objectives in mind – in many ways, they are a "jack of all trades, master of none". Generalized AI chatbots lack the domain specificity, controlled vocabularies, and defensibility features required in legal contexts. Conversely, purpose-built eDiscovery AI tools are designed to incorporate legal objectives, defensibility workflows, audit trails, and structured review processes. They're designed with specific eDiscovery objectives in mind to provide the guardrails necessary to meet the standards of the legal profession, reducing risks and increasing reliability.

Generalized AI chatbots without those guardrails are more likely to hallucinate desired results. While no tool eliminates risks, purpose-built tools minimize risks by aligning AI design with legal requirements, data patterns, and workflow expectations. Combined with an HITL approach that keeps humans in charge, these purpose-built tools can be effectively used to reduce the anxiety that many legal professionals have regarding the use of AI models.

# Benefits of an IA Approach in eDiscovery

Applied specifically to eDiscovery, there are several benefits of an IA approach, including efficiency with control, transparency and defensibility, improved cost predictability and risk reduction.

## ***Efficiency with Control***

AI's ability to de-noise, cluster, and structure data dramatically accelerates early case assessment and document review. But in an IA approach, humans remain the ultimate determiners of responsiveness. Through the process of crafting thoughtful comprehensive prompts (as well as iterative testing and refinement of those prompts), AI organizes the data landscape and attorneys drive the interpretation and how to apply that interpretation to the needs of each unique case. This facilitates a streamlined review workflow that reflects the lawyer's mental impressions.

## ***Transparency and Defensibility***

Unlike "black box" generative models, an IA approach is built for auditability where classifications, suggestions, and culls can be explained and retraced. This level of transparency is key to providing defensibility in litigation. Courts do not require perfection in eDiscovery processes, but they do require reasonable, explainable methodologies. IA satisfies this requirement by ensuring that AI-supported decisions remain traceable and grounded in documented logic.

## ***Flexibility***

No two cases or projects are exactly alike. Legal services can be applied to litigation, investigations, audits, HSR Second Requests, and more. In each of those use cases, the way in which AI is applied can vary considerably. An IA approach enables legal professionals to leverage AI technology in a way that is appropriate for each use case – providing the flexibility that legal professionals need to maximize the value of the AI technology.

## ***Improved Cost Predictability***

Data reduction powered by IA allows teams to better align budgets with expected outcomes. By surfacing key information early and reducing noise before attorney-level review, IA helps organizations forecast review duration, staffing needs, and overall spend. Predictability is especially powerful for corporate legal departments seeking to manage outside counsel budgets and reduce discovery-related surprises.

## ***Risk Reduction***

IA frameworks reduce risk by integrating AI with strong human-driven governance practices. Data governance, secure workflows, chain-of-custody controls, and data sovereignty protections help ensure compliance. Because IA places humans in control of decisions, organizations avoid the pitfalls of automated misclassification, production of incorrect documents, or loss of privileged material. AI becomes a guardrail, not a gamble.

# Benefits of an IA Approach in eDiscovery




Figure 5: Benefits of an IA Approach in eDiscovery

FLIPPING THE SCRIPT ON AI

# Lineal's IA Approach with Amplify™ in Relativity



# Lineal's IA Approach with Amplify™ in Relativity

As a Relativity Gold Partner with three Innovation Awards, Lineal is uniquely qualified to take one purpose-built AI solution (i.e., the AI capabilities in RelativityOne, such as Relativity aiR for Review, aiR for Privilege and aiR for Case Strategy) and add even more purpose-built capabilities into it through Amplify™ Review. Here's how Amplify ties into Lineal's IA approach using RelativityOne.

## Native Integration

Lineal's Amplify™ solution brings IA directly into RelativityOne, not as a bolt-on or parallel system but as an embedded layer of intelligence that enriches existing workflows. This native integration ensures that AI accelerators work seamlessly within the same security, audit, and chain-of-custody infrastructure that legal teams already trust. Rather than exporting data into external systems, Amplify™ enhances what teams can accomplish inside their existing Relativity environments.

## Examples of Leveraging IA with Amplify™

Within RelativityOne, Amplify™ operationalizes Intelligence Amplified across several high-impact workflows:

### ***Privilege Detection***

Amplify™ PrivFinder tool is designed to surface privilege indicators, such as names, domains, roles, content patterns – enabling attorneys to quickly confirm and apply privilege calls. This accelerates detection while maintaining absolute control over what qualifies as privileged.

### ***Chat Reconstruction***

Modern communications are fragmented across emojis, reactions, GIFs, edits, and threading structures. Amplify™ ChatCraft reconstructs these conversations coherently, enabling attorneys to understand context in a defensible and review-ready format.

### **Memory Bank**

Amplify™ Memory Bank stores prior decisions, preferences, and validated conclusions to ensure consistency across reviewers and across time. This helps prevent conflicting reviews, supports machine learning stability, and reduces training gaps.

### **Command Center**

Command Center tracks costs in real time, while Amplify™ reduces non-relevant data early, cutting review costs and maintaining defensibility. Fixed-fee models remove budget uncertainties, making financial outcomes predictable.

### **Attorneys in the Driver's Seat**

Amplify™ design philosophy ensures that decisions always remain in the hands of legal experts. AI provides structure, insight, and acceleration, but attorneys determine relevance, privilege, responsiveness, and production boundaries. This model keeps the legal team firmly in the driver's seat while giving them more powerful tools to navigate massive data volumes, tight timelines, and evolving case strategies. It embodies the essence of IA: empower humans, don't replace them.

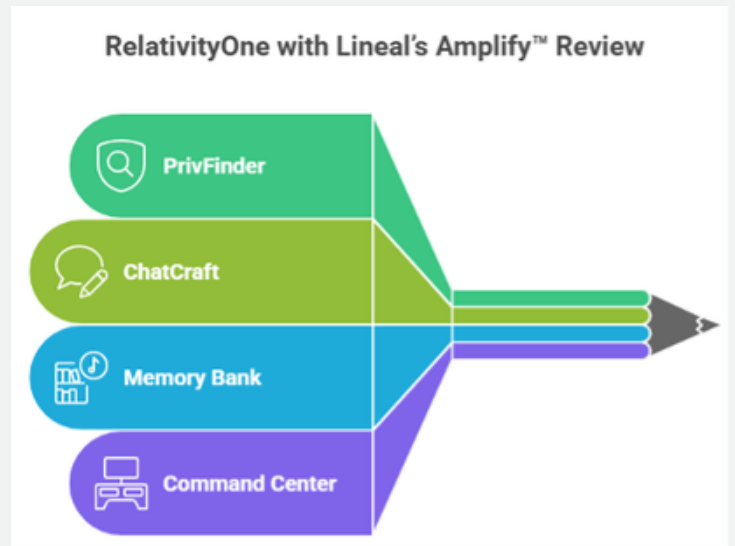


Figure 6: RelativityOne with Lineal's Amplify™ Review

# Conclusion

The legal profession's "AI gone wild" stories underscore a simple truth: unchecked automation is a liability. Generative AI, when used improperly or without oversight, introduces risk rather than reducing it. **The ultimate deliverable is not the output from AI technology – it's the legal expertise leveraged by AI technology.** By flipping the script and treating AI not as an autonomous decision-maker but as a human-directed amplifier of judgment, legal teams can unlock its true value and use it to enhance legal advice and services, not replace them.

An Intelligence Amplified approach that ensures that humans remain in the loop provides a safer, smarter path to incorporating advanced technology into legal and eDiscovery workflows. It reinforces accountability, strengthens defensibility, and accelerates the work that lawyers are uniquely positioned to perform. With Lineal's Amplify™ solution integrated inside RelativityOne, this vision becomes reality. Attorneys stay firmly in control, supported – not supplanted – by AI that enhances their expertise, sharpens their insight, and accelerates their path to accurate, strategic, defensible outcomes.



lineal.com

Learn more about better  
eDiscovery tech and services.

Want more insights like this?

# Subscribe to Our Monthly Newsletter

Get first access to more things like this!

GET MORE TODAY

