TAR to HAR: How LLMs Are Ushering in Legal 3.0

Tom Gricks, OpenText

Tom Gricks

Tom Gricks serves as a Strategic Discovery Advisor at OpenText, where he focuses on the practical application of advanced technologies—such as technology-assisted review (TAR) and large language models (LLMs)—in legal discovery. A former trial lawyer with deep experience in e-discovery, he has worked with law firms, corporate legal departments, and service providers to develop defensible and efficient approaches to document review.

OpenText's Tom Gricks reflects on the evolution of technology-assisted review and how large language models are redefining speed, accuracy, and strategy in legal discovery.

Technology-Assisted Review (TAR), also known as predictive coding, revolutionized the legal industry by applying machine learning algorithms to document review processes. TAR 1.0 learn from human reviewer decisions on a control set and then extrapolates those classifications across the entire document population. Second generation TAR, or TAR 2.0, continuously updates its model as documents are reviewed and prioritizes documents for review based on likely relevance. TAR technology has delivered tremendous value to legal teams by dramatically reducing review time and costs—often by 80 percent or more compared to traditional linear review—while simultaneously improving accuracy and consistency. By prioritizing likely relevant documents and minimizing human review of irrelevant materials, TAR has enabled legal teams to meet tight deadlines, control litigation expenses, and focus attorney time on high-value analytical work rather than repetitive document review and categorization. As courts have increasingly recognized TAR's validity and effectiveness, it has become a cornerstone of modern discovery practice, establishing both the technical and legal foundations for the continued evolution of AI in legal document review.

Large Language Models (LLMs) represent the next frontier in legal technology, poised to further transform document review with their sophisticated natural language capabilities. Unlike traditional TAR systems that identify patterns based on features, LLMs are trained on vast bodies of text making them capable of recognizing context, semantic relationships, and conceptual meaning. This deeper understanding allows LLMs to make more accurate classification decisions, especially with complex legal concepts and subtle distinctions that challenge conventional systems. By integrating LLMs into review workflows—whether independently or in conjunction with existing TAR solutions and analytics workflows—legal teams can achieve unprecedented levels of efficiency. As

adoption increases and best practices emerge, these powerful tools will continue to drive down litigation costs while improving outcomes, allowing legal teams to redirect resources from document review and coding to review and analysis of key evidence that will enhance strategic case analysis and client advocacy.

We recently sat down with Tom Gricks, Senior Principal Consultant at OpenText, and head of the OpenText Investigations Service team, to speak to him about his early experience as a pioneer in implementing and validating TAR for document review and his experience in working with clients and their data to test and refine the use of large language models for investigations and document review under OpenText eDiscovery Aviator "Earn Your Wings" program.

CCBJ: To start, please tell us about your background during the early days of TAR review and your experience promoting technology adoption in legal document review.

Right after Judge Peck issued his seminal opinion in Da Silva Moore1, I recognized an opportunity to transform document review through technology. When faced with a two million document collection, I proposed using TAR to dramatically streamline the process. This led to the Global Aerospace 2 decision that authorized TAR use over objection, a landmark moment for legal technology adoption.

The results were transformative, we saved approximately $1.8 million on that review alone. This success inspired Karl Schieneman and I to write about Rule 26(g) disclosure responsibilities to establish best practices for this

powerful new capability. While initial adoption faced some hesitation, the legal community increasingly recognizes that technology-assisted review delivers superior results more accurate, more efficient, and more cost-effective than traditional methods. The journey from skepticism to acceptance has prepared the groundwork for each new advancement in legal technology. Today's sophisticated technology solutions build upon these foundations, with each innovation making document review more powerful and accessible to legal teams of all sizes.

At what point did TAR begin to gain broader acceptance in document review? How did influential cases and foundational principles contribute to its adoption?

While Global Aerospace was certainly significant, Judge Peck's decision in Rio Tinto 3 truly transformed the landscape by establishing that "it is now black-letter law that if a responding party wants to use predictive coding, the courts will allow it." This watershed moment shifted the conversation from whether technology could be used to how it should be implemented, opening the door for continued innovation. This precedent has been invaluable in establishing a framework for evaluating and accepting new technologies. Each advancement builds on this foundation, allowing us to move forward with greater confidence and focus on maximizing benefits rather than debating permissibility.

TAR has been around for years, but more recently, OpenText has been highlighting HAR. What is HAR, and how does it mark an evolution in document review technology?

HAR, or human assisted review, represents an exciting evolution in how legal professionals and technology collaborate. HAR provides a more intuitive and natural partnership between human and machine where, rather than starting the process by coding documents, legal professionals can communicate their requirements through natural language prompts. While TAR relies on human coding decisions to train the model to recognize responsive document features and build contextual and conceptual recognition into the model, LLMs can recognize subtle patterns and contextual and conceptual relationships between documents from the start because their neural networks are pre-trained on vast data sources.

The work for the human in HAR is to develop prompts that will tell the LLM which contextual and conceptual relationships matter. The system continuously learns from these interactions, becoming increasingly effective at understanding the nuances of each specific review project. This collaborative intelligence approach combines the best of human expertise with the processing power of advanced AI.

Given your longstanding advocacy for technology in legal review, how do you see large language models enhancing review processes? What additional applications do you find particularly compelling?

While I initially approached LLMs with the healthy skepticism any new technology deserves, I've been impressed by their capabilities. Beyond production review, investigations have emerged as an area where LLMs truly shine. They excel at rapidly synthesizing information across large document sets, identifying connections that might otherwise be missed, and providing investigators with valuable starting points for deeper inquiry. This ability to work alongside investigators handling background research while humans focus on critical analysis creates a multiplier effect for legal teams.

What has your experience been so far in using large language models for corporate investigations? What benefits are clients seeing from this technology?

LLMs can enhance human investigation capabilities by serving as powerful research partners that work alongside our team. While we maintain appropriate verification protocols for all AI-generated analysis, the efficiency gains can be substantial. In practice, LLMs create a valuable multiplier effect. As an investigator focuses deeply on one critical issue, inevitable tangential questions arise. Instead of putting these questions aside, we can task an LLM system to explore them concurrently. When the investigator completes their primary analysis, they already have well-structured starting points for the next avenues of inquiry. Even when citations require verification, the LLM's analysis typically points in productive directions, accelerating the investigation process. Corporate clients particularly value how this approach allows us to explore more potential avenues of inquiry within the same timeframe, leading to more comprehensive investigations without increasing costs. This collaborative approach combines the best of human expertise critical thinking, legal judgment, and strategic decision-making with AI's ability to process and synthesize vast amounts of information quickly.

In your work with large language models, how have you approached document classification? What testing strategies have you employed, and what encouraging outcomes have emerged from those efforts?

At OpenText, we have taken a methodical approach to implementing LLMs for document classification, and the results have exceeded our expectations. In one test, we used a ten-page review memo prepared for a document review team as an initial prompt, without any optimization, and ran it through an LLM to assess classification capabilities against a previously coded control set. The recall results exceeded acceptable judicial standards, and precision matched TAR 1.0 levels, all without any prompt engineering. Most impressively, by controlling LLM parameters, we achieved fully reproducible results at the document level, even when varying the sample composition. What is particularly valuable is how the LLM's also offer reasoning insights that can be used to help refine prompts and further improve performance. This transparency creates a virtuous cycle of continuous improvement in a manner that wasn't possible with earlier technologies. Early results demonstrate the potential of LLMs to transform document review while maintaining the reliability and defensibility that legal work demands.

How are large language models poised to transform traditional document review workflows? What potential do you see in integrating them with legacy technologies such as TAR?

LLM technology is already demonstrating impressive classification capabilities, and as prompt engineering becomes more sophisticated and streamlined, we will see even greater efficiencies. What's particularly exciting is how LLMs and TAR can work together synergistically. For example, you could use an LLM with a well-engineered prompt to identify responsive documents, then leverage TAR's continuous learning capabilities to prioritize documents for review. As the system learns, you can progressively reduce the number of documents requiring full LLM analysis, optimizing both speed and cost. This combined approach represents the best of both worlds TAR's efficiency in ranking and prioritizing documents with the LLM's ability to incorporate both content and context. Together, they create workflows that achieve exceptional recall without processing the entire collection, delivering better results faster and at lower cost. While

new technologies always face adoption hurdles, the legal community's growing comfort with statistical validation puts us in a stronger position for faster acceptance of these powerful new capabilities.

Considering the early phase of genAI implementation, what strategic advice would you offer to corporate counsel facing pressure to adopt LLMs? How can they strike a balance between innovation and risk management?

The most important principle is to view LLMs as powerful tools that enhance your existing processes rather than replacements for sound legal judgment. Focus on incorporating AI capabilities strategically to increase efficiency and effectiveness, not on wholesale transformation. A practical approach is to combine techniques for maximum benefit. For example, use traditional analytics to filter obviously non-responsive documents upfront, then apply LLMs to the reduced document set. Since LLM results are reproducible at the document level, this automatically increases precision without extensive prompt engineering. This combined workflow delivers results greater than either approach alone.

To manage risk, I recommend starting with lower-risk, high-value use cases like internal investigations where the stakes are manageable but the potential efficiency gains are substantial. This builds organizational confidence and expertise before moving to productions in high-stakes litigation. Remember that validation remains essential regardless of technology.

In light of previous hurdles in legal tech adoption, how do HAR and genAI differ from earlier waves of innovation? Are today's organizations better equipped to embrace these advancements?

We are in a significantly stronger position today for several reasons. The legal community has developed a much deeper understanding and appreciation of statistical validation, a critical foundation for any technology-assisted review process. The validation principles established with TAR provide a robust framework that can be applied to LLM-based approaches as well. At a fundamental level, both TAR and HAR share the same goal: teaching technology to distinguish responsive from non-responsive documents. The specific mechanisms differ, but the objective and success criteria remain consistent. This continuity makes the transition more natural for legal teams already familiar with technology-assisted review. The combination of established validation practices, growing comfort with legal technology, and the enhanced capabilities of LLMs positions us for more rapid and confident adoption than we saw with earlier innovations. While challenges remain, we have better tools and knowledge to address them effectively. ■

opentext™

64% of legal leaders are scaling mountains with GenAI.

Ready to join the climb? Reach new heights with AI legal tools.

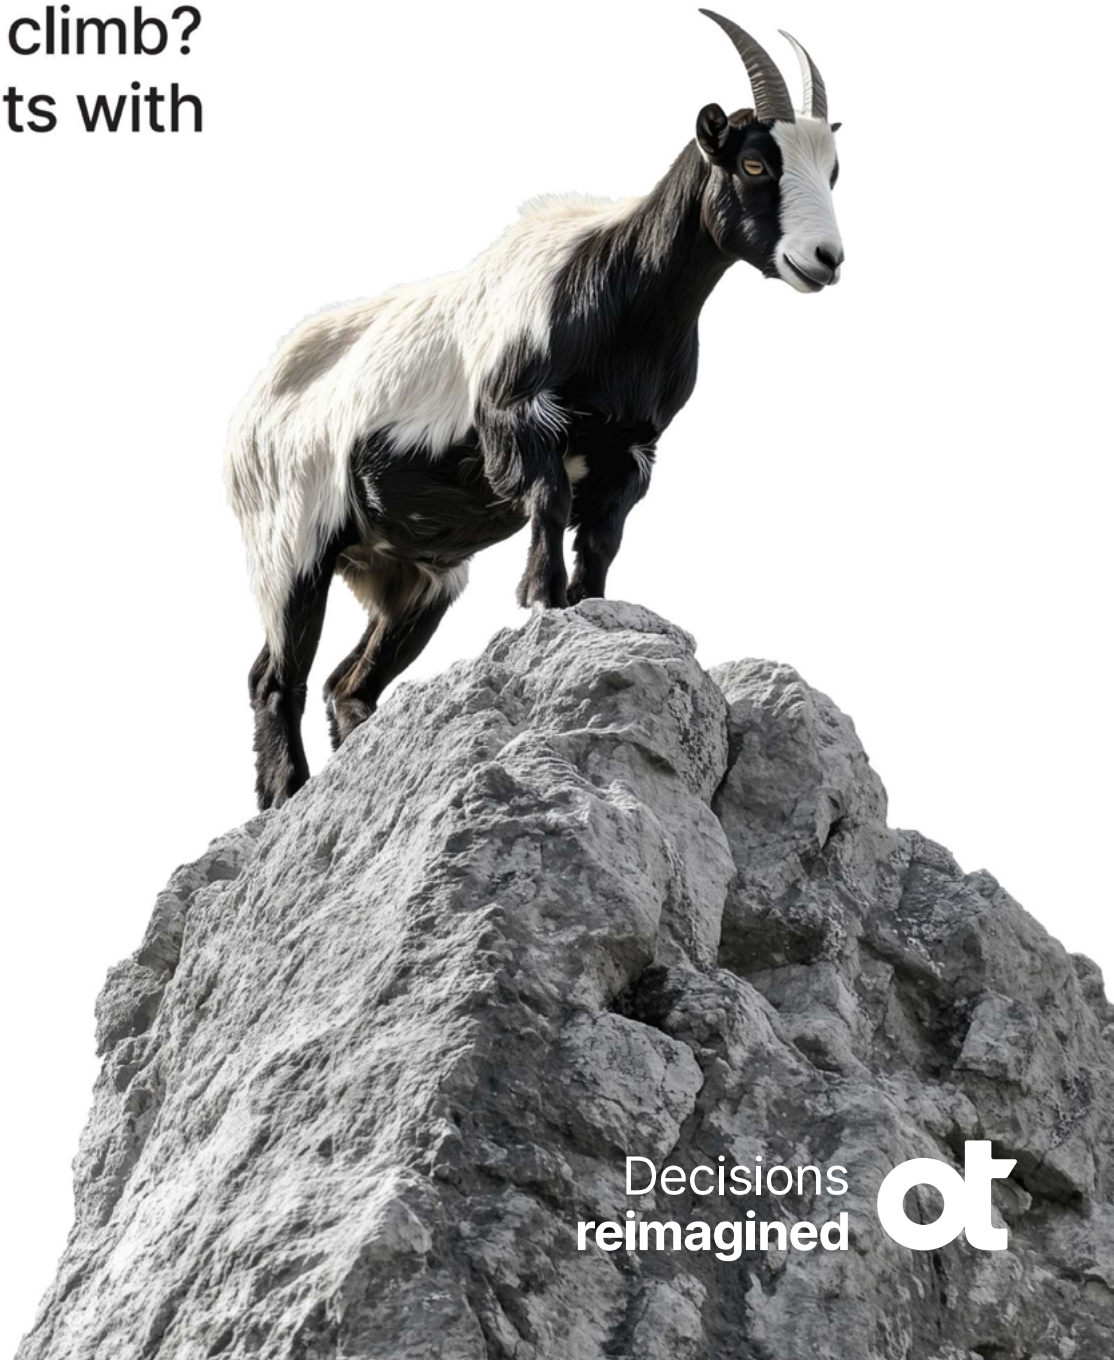

05.25 | 248-000108-001