



Don't Throw the Bot Out with the Bathwater: Embracing Generative AI in eDiscovery Work

Technology, Privacy, and eCommerce



Banner artwork by Dark Geometry / *Shutterstock.com*

It is clear artificial intelligence (AI) is supplementing the work of humans at a rapid pace across many industries. Several industries — from manufacturers to sales reps, from investors to publishers — use the technology to drive efficiency, boost productivity, curtail costs, and enhance precision.

A recent [acquisition spree](#) indicates the legal profession is slowly but surely embracing digital tools and AI. This embrace — as real or hypothetical as it may be — raises issues: Can generative AI be trusted in eDiscovery? Will lawyers utilize it? Does it ensure fairness?

Legal practitioners are wise to take a cautious, open-minded approach to adopting generative AI. There remain fundamental questions of how far, how fast, and in what capacity to integrate this technology, especially when it comes to highly evolved generative AI models.

AI, in the most basic understanding of the technology, offers legal professionals several practical functions, but the key terms and nuance related to the technology must be clearly defined in order to allow for proper rules, regulations, and a general understanding of what is and is not yet operationalized.

AI is often incorrectly used in eDiscovery as a catch-all term, and that lack of clarity can lead to confusion. AI extends to a broad field in which software and computer systems perform tasks traditionally reliant on human intelligence.

Machine learning is an important component to the AI equation. This is how algorithms are trained to use an array of data to improve performance. In eDiscovery, this could be applied to email thread

analysis, clustering, and predictive coding, as well as to help automate the review and categorization of large volumes of electronic documents.

Generative AI is a subset of AI that focuses on creating new data or content. This is highly evolved technology with big potential, but it must be understood. It relies on techniques like Generative Adversarial Networks (GANs) and Large Language Models (LLMs) that, in eDiscovery, can be used for generating synthetic data for testing and training purposes.

Building trust and comfort

There is reason for caution when it comes to pushing for an AI revolution in formal discovery. Generative AI technology tends to operate in a black box of proprietary models and algorithms, subject to manipulation, falsification, and errors. Legislators have conducted a series of bipartisan closed-door learning sessions to better understand how generative AI works. One common criticism is generative AI lacks a reliable feedback loop and can be extremely expensive to use and validate.



Staying cautious with generative AI technology can prevent inaccurate information in eDiscovery.
Teerachai Jampanak / Shutterstock.com

Generative AI technology tends to operate in a black box of proprietary models and algorithms, subject to manipulation, falsification, and errors.

There is, however, a strong use case for generative AI to aide many legal skills, including summarizing documents, initial drafting of documents, suggesting issue codes, optimizing keyword searching, and assessing quality control — and to perform each of these actions at speed. There is also a strong argument for generative AI in the investigative space of discovery, where the technology may soon be able to turbocharge the identification of relevant documents in a data set. It may soon be able to sift through millions of pages of historical information to formulate an educated analysis or identify common themes. Similarly, when performing managed review work, generative AI will likely be able to comb through thousands of documents searching for documents which are responsive to discovery requests and will be able to do so in a fraction of the time as more traditional ways of doing such reviews.

When thinking about the generative AI adoption curve, it is helpful to consider the relatively slow adoption rate of technology-assisted review (TAR) and consider whether that slow progression will be what we see with generative AI in discovery, or if conditions are such that adoption will be on a faster track. After roughly 10 years of attempts to more uniformly and more widely adopt TAR in discovery, TAR is now increasingly accepted by judges, regulators, law enforcement bodies, and litigating parties. But, while there is now a more stable and consistent adoption of TAR, there still is progress to be made. To really understand TAR's potential, it is best to break it down into separate levels of document review intervention:

- **TAR 1.0** — the “OG” of technology assisted review — leverages machine learning to review a data set in a static manner. Subject matter experts review a limited subset of the data set, after which an algorithm is trained on the subset to predict responsiveness across the full population. This is a lot like a streaming media service asking a viewer to rate 10 movies, and then suggesting content based on just those ratings.
- **TAR 2.0** — also referred to as continuous active learning (CAL) — is the more mature big sister of TAR 1.0. TAR 2.0 also uses machine learning to review a data set. However, as reviewers code documents, the algorithm is continuously updated to improve its accuracy in identifying responsive documents. In the streaming example, suggestions are refined by a continuous improvement mechanism based on future ratings and shows watched over time.

The outlook for generative AI

Generative AI takes the learning process to a different level, generating new content or data that did not previously exist. There is no binary requirement or, as noted before, continuous feedback loop. These systems learn patterns and features from large amounts of data — including text, images, videos, and music — and use this understanding to create similar content.

The most profound impact generative AI could initially have is in support of document review and classification, predictive coding, auto-redaction, and anomaly detection.

Imagine a scenario where understanding whether any documents in a collection stand out from the rest, and what themes in those documents match important issues in a case. Generative AI may be able to someday soon do just that. Or, after completing a document review, a case team might be able to utilize generative AI to help develop a *first draft* of a deposition outline, which an associate

could validate, edit, and quality control. It may also someday assist in a more automated comparison of documents, find anomalies in data, and provide statistics on document sets.

Legal practitioners should view the push to use generative AI within the industry somewhat akin to the automotive industry's deliberate and gradual push for self-driving cars. While autonomous navigation has deeply penetrated adjacent industries — aerospace, shipping, and agriculture — the conventional automobile business has been slower to adopt. The slow-pace of technical developments, rigid regulations, testing errors, and other factors have resulted in a *less-than welcoming* embrace of “robocars” from regulators and automakers. However, as “advanced-driver assistance systems,” (ADAS), are improved, pressure-tested, and developed with improvements, the acceptance and use of them has dramatically increased. ADAS gradually introduced semi-autonomous features (such as automated emergency braking or highway hands-free driving) into the cockpit, and as there are more successful use cases with these systems, the public may grow [more comfortable](#) with a more aggressive transition.

Just as the automotive industry is gradually transitioning, the adoption and acceptance of generative AI is swifter and faster than what the legal industries saw with TAR. Here, perhaps the difference is that the technology is much more cross-functional, and industries have a clear financial incentive to consider and adopt generative AI in a variety of ways. The wider acceptance of generative AI, and its obvious financial incentives, unlike TAR, are much more clear from the jump. This momentum to accept and use generative AI in a variety of corporate functions will extend itself to litigation, and it is possible that this can be done in a way that maintains fidelity to the needs of the court, but developed in a way that mitigates legal and liability risk. There is a gradual transition that law firms and legal service providers can make from review-focused machine learning to layering in the application of vetted generative AI, once the technology improves beyond the beta stages.

The most profound impact generative AI could initially have is in support of document review and classification, predictive coding, auto-redaction, and anomaly detection.

Important vs. urgent work

A well-known management technique — known as the Eisenhower Principal — is to not confuse the urgent with the important. As it relates to generative AI, one could say that although the technology appears to be aggressively gaining steam, the full-blown reliance and adoption is not as urgent — at least not in litigation. That does not, however, dilute the importance of beginning to understand and deploy generative AI. Navigating generative AI's potential in the legal industry should not be delayed. This starts with setting aside the idealized application of the technology, and instead focusing on first understanding the technology much in the way one would seek to analyze mobile device management or crisis management.

By breaking generative AI down into components, legal practitioners can understand how to apply and optimize it in day-to-day investigative work. At the same time, understanding the processes and safeguards that need to be instituted to avoid negative outcomes is critical, including those around data privacy regulations and how generative AI will process personal data. In other words, the adoption of generative AI need not, nor should it, be an overnight “on” switch, but rather a defensible and responsible layering into existing processes and established workflows and validations around TAR.



By understanding generative AI in its entirety, legal professionals can identify safeguards to abstain from unfortunate circumstances. Ice stocker / Shutterstock.com

Understanding the processes and safeguards that need to be instituted to avoid negative outcomes is critical, including those around data privacy regulations and how generative AI will process personal data.

Tackling this important step in the generative AI adoption process should not wait. The sooner it gains traction, the sooner the bigger existential problems get solved. *For example*, it is important to understand and address how including incorrect answers or citing a fictional case from a generative AI technology, like ChatGPT, can impact the industry, and future application of such tools in the space.

Where to begin?

Experimentation is critical

Much the way TAR workflows evolved, and basic analytics, clustering, and other related developments emerged, trial and error of generative AI are a key step in introducing these powerful new technologies without compromising quality or reliability of the result. Typically, that can be done in the investigative space before moving on to more rigid arenas, like litigation. [Experts say](#) these tools may evolve to help generate search queries or expand initial ones in the initial stages of eDiscovery. [Others](#) say generative AI can be used in anomaly detection, translation of foreign language content, and computer-vision-enabled image recognition and classification.

Quality controls are critical, deep understanding is too

Quality control has long been a cornerstone of effective and reliable discovery processes, but with generative AI, knowing *how* to train a model to suit a specific purpose is critical to unlocking the potential of the technology. Generative AI may be able to operate at the level of an “aggressive first year associate” and must be held to heavy quality control and review. Just as one cannot expect a junior legal practitioner to be the final set of eyes on a critical case filing, there should be a system of

checks and balances in the outputs from generative AI to ensure things are done correctly.

Quality control has long been a cornerstone of effective and reliable discovery processes, but with generative AI, knowing HOW to train a model to suit a specific purpose is critical to unlocking the potential of the technology.

And finally, patience is key

Success with technology newly applied to legal processes requires fortitude, discernment, foresight, and experimentation. While the full potential of generative AI is far from realized, legal professionals should not wait for a fully formed regulatory framework nor mature generative AI technology to start experimenting, building expertise, finding immediate use cases, building workplans, workflows and protocols, and creating a runway for longer-term adoption.

[Join ACC](#)

[Jimmy Chestnut](#)



Associate General Counsel

AlixPartners

Jimmy Chestnut is an associate general counsel at AlixPartners. With over 10 years of experience, Chestnut advises the business on a variety of matters, including intellectual property ownership, data privacy, data handling and confidentiality obligations. Chestnut has deep experience in reviewing, negotiating, and drafting various agreements, including eDiscovery agreements, business associate agreements, data privacy agreements and SaaS agreements.



Global Head of eDiscovery & Data Forensics

AlixPartners

Tarek Ghalayini is the global leader of AlixPartners' eDiscovery practice. He guides and prepares clients for large, cross-border matters involving data mining, analytics-driven document review, and other related data-centric issues. A licensed attorney with a strong background in accounting and finance, Ghalayini works closely with counsel, corporate clients, and financial advisors to quickly and efficiently tackle complex electronic data problems in global litigation, financial investigations, regulatory inquiries, and restructuring matters. Ghalayini is known for developing and executing effective and straightforward solutions that bring clarity and consistency to cross-border and local matters alike.

Danielle T. Davidson



Head Strategic Advisor, Advanced E-Discovery and A.I. Strategy

Ropes & Gray LLP

Danielle Davidson is the Head Strategic Advisor of the Advanced E-Discovery and A.I. Strategy Group, bringing over 15 years of E-Discovery experience. A leader in developing and adopting advanced technology and analytics, Davidson works with clients — including the heads of E-Discovery at global corporations — to create complex E-Discovery strategies. Davidson has deep experience piloting and adopting advanced technologies such as concept analytics and Technology Assisted Review (TAR). She leads projects on short message review strategies, the future applied use of A.I., and technological solutions for mobile data content capture and monitoring. Recognized for her commitment to providing excellent service to clients, she is an active voice in the E-Discovery community.

