



The General Counsel's Guide to AI

Contents

- Introduction 3
- Data Challenges and Trends at the Intersection of AI and Legal..... 3
- Using AI to Make a Difference in Your Organization 4
 - What is AI good for?..... 4
 - Is using AI a big technological lift?..... 4
- AI in Action: Use Cases & Applications in Legal..... 5
 - Contract Review 5
 - Sentiment Analysis 5
 - Data Privacy Compliance 5
 - Translation 6
 - Internal Investigations..... 6
- Taking a Responsible Approach to AI 6

Introduction

General counsel (GC) are grappling with the promise of artificial intelligence (AI) for their legal, security, and compliance challenges. And as with every technological evolution, GC have two roles. First, you will need to understand how your business is using AI so that you can advise your business partners and offer critical, realistic, and pragmatic advice to your board on this new technology. Second, you will be expected to integrate productivity-enhancing technology into your own operations as a demonstration of your leadership and focus on innovation for the business. But keeping up to date in this quickly changing field is a challenge in itself. Indeed, with the emergence of generative AI and Large Language Models (LLMs) such as Open AI's GPT4 in 2023, AI is already making a considerable impact on the legal industry. It is therefore imperative now more than ever that general counsel develop a foundational understanding of AI and the nuances of deploying it in tasks within the legal function.

According to a [Thomson Reuters report](#), “corporate counsel believe they are tech savvy but acknowledge that their comfort level and confidence with technology have limitations, specifically around artificial intelligence.” Striking the balance between

validating where AI can bring value and identifying where the promise diverges from reality can be daunting. Integrating AI across litigation, privacy, and related areas is core to Relativity's business; to that end, this brief guide provides a practical overview of AI in the context of the legal profession

Data Challenges and Trends at the Intersection of AI and Legal

Corporate legal departments today are hobbled by the sheer variety and complexity of data challenges:

Consider, for example, stringent privacy requirements necessitated by new privacy regimes such as the California Privacy Rights Act (CPRA). Further, the [explosion of new types](#) of unstructured data (from Slack messages and Teams chats to Zoom transcripts and meeting notes) as remote work becomes increasingly normalized and more workplace interactions go online, has made it more difficult to stay on top of growing data volumes and secure sensitive information buried within it. Additionally, the same sensitive information lying latent within an organization's data pool presents a greater attack surface and emboldens threat actors to orchestrate data breaches.

“AI is probably the most efficient tool to help organizations turn massive amounts of ‘senseless data’ into more ‘senseful data’. And if it’s true that ‘senseful data’ are the new gold of the digital era, then the value of AI should not even be a question. The legal function often has a need to make sense out of big data in a short period of time. e-Discovery is a good example but it’s not the only one.”

STEFAN JOHN

Senior Vice President, Chief Legal Officer BASF Environmental Catalyst and Metal Solutions

In the backdrop of these challenges, the corporate legal department — already viewed as a cost center by the rest of the organization — must strive for greater efficiencies and contain costs particularly in the current uncertain economic climate. General counsel and the corporate legal department are therefore forced to find ways to do more with less.

On the other hand, AI presents considerable opportunities to accelerate the pace of legal work, maximize accuracy and automate those parts of the job that are repetitive and manual-intensive. Early adopters of AI in the legal industry are already applying AI in a variety of contexts — from accelerating document review, intellectual property management and conducting sentiment analysis to contract review and modernizing their organization’s privacy program by using AI to identify and redact personal information (PI) and sensitive personal information (SPI) across large data volumes.

Further, the recent emergence of generative AI and LLMs are already beginning to transform the nature of day-to-day work in the legal industry. As an example of the growing sophistication of these technologies, consider how Open AI’s GPT 4 passed the bar exam and scored in the [90th percentile](#). It’s too early to speculate the vast breadth of use cases and applications unleashed by generative AI in the legal space but it is certain that they will run the gamut and have an indelible impact of the nature of legal work.

Using AI to Make a Difference in Your Organization

What is AI good for?

Today, AI is all about **processing data more quickly and effectively than humans**. AI systems are great at quickly poring through vast data sets, discovering patterns, identifying anomalies, and presenting these insights to take action or further refine the results for clarity. In this way, AI takes raw input that would otherwise be too overwhelming for human operators, begins to digest it, and then hands it off to reviewers, analysts, and attorneys to perform their role more effectively and with greater focus. AI can be of value, if applied appropriately, when your organization is faced with sifting through large, diverse data sets to meet litigation, investigations, privacy, or related challenges.

Is using AI a big technological lift?

Generally speaking, no. While the technology may be more sophisticated under the covers than existing software tools such as email, Excel, and Word, adopting AI-based tools in your organization should not be categorically different or require new expertise and extensive training. Relativity’s AI capabilities and solutions, for instance, are designed to free up legal teams from performing rote, routine, and repetitive tasks so that they can perform more strategic work. This means delivering a tool that is accessible, intuitive, and relevant to the legal activity at hand.

“When a company gets a large lawsuit, the first thing that happens for the general counsel is a call from the CEO asking: What’s this all about? What am I going to do? If you, as the GC, have the ability to quickly identify key information in that lawsuit and are able to report back promptly to the CEO who, in turn, can communicate to the board, you have significantly helped the company move toward a decision based on the potential risks and the best next steps. AI makes this all possible.”

DAN COOPERMAN

Former General Counsel Apple & Oracle

AI in Action: Use Cases & Applications in Legal

Contract Review:

General counsel are automating contract review with AI, allowing their in-house attorneys to quickly sift through standard contracts like NDAs, supply agreements, SaaS licenses, and purchase agreements. Because AI is designed to recognize patterns, it can also be used to identify outliers, automatically spotting and flagging any anomalies that may be buried in the contract language. By focusing your attorneys on language that diverges from language that is generally used, they can spend their time on provisions that require closer review. Essentially, AI contract review technology helps organizations to more quickly and cost effectively answer the question, “Is it appropriate for my organization to enter into this agreement?”.

Sentiment Analysis:

Investigators may rely on communications that demonstrate emotions such as fear or anger in various matters involving harassment, fraud, or deception. For example, in employment-related litigation, identifying communication in which people made threats, expressed support, or relayed dissatisfaction with a particular person or organization could be crucial to the case. Identifying sentiment has long been a challenge for investigators and attorneys. Without a tool purposely built to identify sentiment in human communication, legal teams must rely on keywords and conceptual searches to find documents that include emotions. In English alone, there are a panoply of words to express emotion, and thousands more in slang . This inevitably leaves more opportunity for error and makes the process more costly. To compound the problem, keywords can have multiple meanings depending on their context. The emotion of a particular word or phrase heavily depends on the surrounding words and sentences. For example, “I hate you” could be used earnestly or jokingly between friends. In these situations, an AI-driven sentiment analysis tool can detect emotional signals to shed light on important evidence.

[Designed to address these linguistic nuances and semantic complexities](#), Relativity’s sentiment analysis tool provides machine learning-powered labels that indicate the presence of positive or negative tone or emotions, such as desire and anger, within written documents. The tool also allows users to search for these behavioral signals to quickly pinpoint items of interest or scan emotionally charged sentences in documents in a case or investigation.

Data Privacy Compliance:

Data privacy obligations have become more complex for any company that collects, processes, or stores consumer data. Your company’s responsibility to comply with regulations such as GDPR, the CPRA, and others means you can’t be passive or reactive; you need to be prepared for the compliance challenges that large data sets — particularly unstructured data in documents and files — present. Many of your obligations will turn on locating not only personally identifying information (PII) but also personal information (PI), which is more broadly defined under data privacy mandates to include data elements that can be linked, associated, or related to individuals. AI can prove key to meeting these challenges efficiently.

Data Breach Response: Data breaches are becoming more prevalent and more costly, and in the [first quarter of 2023 alone](#), more than 6 million data records were exposed due to a data breach. Once a breach has been identified, the remainder of the data breach response plan rests on how quickly and accurately the breached data can be analyzed. Generally, this process is time-consuming and involves extensive human review. Leveraging AI to perform the data breach assessment can help teams gain a critical commodity: time. AI-driven approaches can improve accuracy for PII like Social Security numbers as well as identify sensitive personal information such as political opinions, religious beliefs by incorporating semantic analysis, human signals, and context. Putting AI to work can empower the response team to quickly and accurately identify which records have been compromised and what types of sensitive data were impacted for each individual, enabling a smooth, efficient, and accurate response plan.

For example, with Relativity Data Breach Response, you can [reduce review time by 75%](#), gain deeper insight into your data with impact assessments powered by AI that is trained to identify personal information impacted in a data breach.

PI Identification: When the stakes are high, you can't leave PI identification to chance. AI can uncover patterns in data that human reviewers can't on their own, allowing you to find critical PI that traditional searches miss, while also ensuring accuracy and consistency. For example, during the document review phase of a large corporate litigation, Relativity PI Detect reduced time spent on identification and redaction by [78 percent](#), which equated to an approximately \$100,000 cost reduction and a six-week time reduction for the project.

Translation:

Litigation, internal investigations, and regulatory requests involving multinational companies with offices, operations, and partners in different parts of the globe can present a daunting amount of paperwork for the parties involved. But in these situations, where time is of the essence, the process of translating hundreds of thousands of documents, chats and internal communications that are called upon for discovery, can be extremely burdensome. Historically, organizations involved in cross-border litigations had two options: (1) Using a third-party translation service, which dramatically ballooned costs and took an inordinate amount of time. (2) Using a simple language translation APIs, which although cost-effective, comes with security risks, limited capabilities, and supports only a narrow list of languages.

But, with AI-driven translation solutions, these limitations do not apply. For example, Translate in RelativityOne, can [identify more than 100 languages](#) present in the documents automatically. The results are high-quality, AI-powered translations (which can be performed in big batches, or just one document at a time, based on reviewers' needs) that never leave the security of RelativityOne and take only minutes to complete.

Internal Investigations:

AI can help general counsel identify problematic misconduct before it turns into a major compliance issue. General counsel know that misbehavior is tough to spot, and the sheer volume of communications can present a challenge for monitoring compliance. Even more so in cases where the issue is not negligence so much as it is when bad actors attempt to disguise their misbehavior by using codewords to talk about regulatory avoidance, insider trading, or bribery schemes. AI can incorporate both semantic analysis and the human or organizational context of the exchange to help general counsel spot wrongdoing as quickly as possible and to take swift action. Since AI is so good at detecting patterns, it can tell when something seems off or looks fishy.

Taking a Responsible Approach to AI

In a recent personal injury case, [an attorney was sanctioned](#) after it was found that a fake case had been cited in an affidavit that he had submitted to the court. It is alleged that the attorney used an AI chatbot in his legal research and that the chatbot ["hallucinated" this information](#). As this case proves, the use of AI is not without risks—and to minimize those risks, new applications of AI must be built with responsibility and usefulness top of mind.

At Relativity, we develop and deploy valuable AI technologies with measured intention, using processes that are thoughtful, disciplined, and engender trust among our customers.

Over the years, we've thought very deeply about what a responsible approach means and how it should play out in the context of building our AI products and capabilities. In the process, we established [a set of AI principles](#) to guide our product development and underscore our commitment to ethical and responsible AI. Here they are:

We build AI with purpose that delivers value for our customers: Every AI system we create is designed to help people solve a specific legal or compliance challenge easily, productively, and defensibly. Our AI development isn't dictated by the latest trends or news headlines. Instead, it is fit for purpose—driven by the people who use our AI, the problems they're trying to solve, and the capabilities and limitations of the technologies we create.

We empower our customers with clarity and control: We design systems and interactions that harness and amplify human expertise to help our users best accomplish their work. It's important that we're open about our AI so you are well informed and able to defend your processes. This includes offering clear information to help you understand how our models are trained and the purpose they were built for, and what each model looks at when making decisions.

We ensure fairness is front and center in our AI development: Throughout development we carefully consider the fair treatment of not only our users, but anyone who could be impacted by the AI we're creating, such as custodians in a litigation or subjects of an investigation. We strive for fairness by seeking diverse perspectives from a wide variety of sources so our models can be as representative as possible, and we use human-centric design processes to focus on user needs. We test our models for potential bias, and if we find any, we pause to thoughtfully consider mitigations, document our decisions, and ensure our customers are informed about the model's appropriate purpose and use.

We champion privacy throughout the AI product development lifecycle: Relativity is a global company that operates in an ever-evolving data privacy landscape. Given this, our privacy principle reflects the latest applicable privacy regulations and guidance, leveraging the principles of Privacy by Design.

We place the security of our customers' data at the heart of everything we do: Relativity has a best-in-class security program, championed by Calder7, our award-winning security team. This team works around the world and around the clock to protect and defend our customer and corporate environments, proactively alleviate threats to our company, and enable clients to control the security of their data through greater transparency.

We act with a high standard of accountability: We understand that people rely on our AI technologies to find the truth in their data, and we're committed to developing these technologies in a responsible way. We put every AI model within Relativity through an extensive peer review and validate each one on a representative set of customer use cases to ensure we deliver on the expectations of our community.

Learn more at <https://origin.relativity.com/artificial-intelligence/> or by contacting Relativity today by scanning the code or emailing sales@relativity.com.



231 South LaSalle Street | 8th Floor
Chicago, Illinois 60604
+1 (312) 263-1177 | relativity.com