

# AI Legal Assistant: The Landscape

The category of AI Legal Assistants (ALAs) is made possible only by recent advances in artificial intelligence (AI).

AI is machine intelligence that can mimic the way that humans solve problems and make decisions in order to automate the performance of tasks that are usually undertaken by humans. AI uses algorithms, rules, and capabilities such as machine learning and neural networks to evolve and become progressively better at performing tasks. It is often categorized into types based on its ability to replicate human characteristics:

- Artificial Narrow Intelligence (Narrow AI) – AI with a narrow range of abilities
- Artificial General Intelligence (General AI) – AI that performs at the same level as humans
- Artificial Superintelligence (Super AI) – AI that outperforms humans.

General AI and Super AI have not yet been realized, either within the legal profession or beyond. Narrow AI (also known as weak AI) has been available at a rudimentary level within the legal industry for about ten years, with early applications focusing on eDiscovery and, subsequently, contract review, search, and legal research. These applications and others in legal have largely been based on supervised or unsupervised machine learning, natural language processing, and rules-based logic (or some combination of these).

Unsupervised machine learning, relying on large datasets and pattern recognition, is particularly useful for eDiscovery purposes, where clustering and classifying data is helpful. In contract review, especially for purposes such as due diligence, supervised machine learning has generally been required because more precision is necessary to identify discrete terms or concepts across a set of contracts, especially where those terms are expressed in variable language.

The AI used in legal technology products has frequently lagged behind the AI leveraged in other industries, with the result that until recently, using AI to effectively identify and extract legal data at high levels of precision required a significant expenditure of effort in training an algorithm. In some legal applications, this training was undertaken on the vendor side; in others, internal teams at law firms would themselves spend hours to train a solution and build machine learning models.

In 2018, advances in natural language processing gave rise to a new method of pre-training language representations, called BERT (Bidirectional Encoder Representations from Transformers). BERT is a neural-network-based technique for language processing, and was one of the first developed transformer-based self-supervised language models. The advent of BERT and its introduction into various legal technology tools changed the requirements around training in machine learning systems, with far less effort required to reach high levels of precision in both search applications and data identification processes.

This represented a significant breakthrough for legal, but was not (or hadn't yet proven to be) transformative on a large scale. Though the advances in the technology were momentous, they were not yet truly disruptive.

Companies outside of legal have continued to develop large language models since BERT. Open AI, for example, has been developing the series of large language models called GPT since 2018, with each GPT model more advanced than the last, and the speed of progress increasing exponentially. In 2022, Open AI introduced GPT-3.

GPT-3 and the ensuing models of GPT are characterized by far larger parameters than BERT, and some differences to the way the models learn. In November 2022, ChatGPT, a chatbot application built on top of GPT-3.5, was released in beta form to the public. As many observers wrote in response to press releases around the launch of ChatGPT, the ‘bot behaved and “felt” different than anything that had come before. The release of ChatGPT and its underlying language model represented an important advance in AI, proving for the first time that AI could interact with humans almost like another person. The launch of GPT-4 in 2023 is a continuation of this evolution, with this model substantially more powerful than GPT-3.5. Where GPT-3.5 performs at the level of a high school graduate, GPT-4 performs at the level of a post-graduate student, in most areas of specialization including law. Professor Dan Katz and Michael Bommarito performed an experiment with each of these models to determine whether they could pass the bar exam. GPT-3.5 performed reasonably well and passed part of the bar exam. GPT-4 aced the bar exam in seven minutes.

Large language models are a type of generative AI that can generate text through producing the most reasonable next set of language tokens based on what has come before. Generative models can take what they have learned from the data to which they have been exposed and create something new out of that information. There are also generative AI models that produce images or videos, but as language essentially underpins the practice of law, legal technology vendors quickly saw the opportunities for the industry in leveraging advanced large language models such as the more recent iterations of GPT (and other large language models that have been developed by Google and other companies).

Unlike the legal technology of the past, which lagged behind the mainstream, at least two legal technology vendors were given access to GPT-4 prior to its release in 2023 and developed solutions accordingly. Other legal technology vendors have moved rapidly in early 2023 to leverage this new technology where it will add value to their solutions.

ALAs are one of the most common large language model products developed by legal technology vendors since ChatGPT was launched.

## What is it?

Many legal technology vendors have responded to recent advances in AI by using it to replace part of the technology stack in existing products to improve efficacy or add features.

ALAs, by contrast, are entirely new products that are made possible only through the availability of advanced large language models. Similar to a human assistant, ALAs are capable of performing many tasks, not just one. These are not point solutions, but systems that allow suites of tasks to be effectively automated.

ALAs have emerged onto the market in three ways:

- Built into existing, broader systems, allowing a user within that platform to use the ALA interface to select from a list of automated tasks targeting specific types of workflows, such as eDiscovery review or contract lifecycle management (examples include DISCO’s Cecilia and ContractPodAI’s Leah).
- As independent systems that offer multiple legal skills that are applicable across many practice areas and accessible from within a single interface (examples are Casetext’s CoCounsel and Rally’s Spellbook).

- As advanced AI capabilities that are custom-built into a legal organization's existing platforms, providing multiple automated tasks that are accessible within current lawyer workflows (Harvey's partnership model is an example).

The benefit of implementing an ALA over a point solution is that it offers time savings in many areas of legal practice rather than just one. The overall productivity gain from implementing secure generative AI to address multiple use cases is significantly higher than leveraging a solution that is able to automate just one or two workflows. ALAs also introduce lower complexity in terms of training and accessibility since the functionality is often accessible from within one platform and does not require switching between solutions.

### **Distinction from Generic AI Assistants**

ALAs are not to be confused with generic chat applications that have been built on advanced large language models and are available to the public through beta or subscription models, or to organizations upon purchase of a license.

These applications, such as ChatGPT, can be very powerful tools for consumers - including lawyers. Any use of these generic applications in professional environments must be necessarily restricted and closely monitored, however.

Many laws, rules, and regulations are not available in machine-readable form on the internet and these tools have therefore not been exposed to the depth of legal information necessary in order to provide consistently reliable advice on legal matters. There are also considerable security risks inherent in using these types of off-the-shelf tools that make them unsuitable for legal environments where attorney-client (lawyer advice) privilege and client confidentiality are critical.

In contrast, ALAs have been exposed to the legal information they require in order to provide reliable responses to legal questions, and have been engineered specifically to meet the requirements of legal environments.

## **Common Use Cases**

### **Generic / Relevant Across Practice Areas**

---

#### **Drafting**

Since ALAs are built upon a framework of generative AI, drafting is an obvious use case. ALAs are able to support lawyers in generating first drafts of emails, letters, memos, and other correspondence.

Most ALAs should not yet be relied upon for drafting entire complex agreements and documents, but those that offer drafting abilities are commonly able to suggest language for specific clauses and reliably draft good first versions of more administrative, less sophisticated documents.

#### **Summarization**

ALAs can summarize lengthy documents or databases of documents in a very short amount of time. Most ALAs have been engineered to be able to ingest and read many documents at one time, making it feasible for these products to rapidly summarize the salient points of the information that has been uploaded.

### **Translation**

ALAs can translate text at scale almost instantaneously into virtually any language. Using a secure ALA for digital transaction is preferable to using open products such as Google Translate in legal environments.

### **Search and Retrieval**

If an ALA is integrated with core systems at a firm or organization, it will be able to search across and retrieve specific work product and information to provide lawyers and staff with just-in-time knowledge as they work.

## **Corporate/Transactional**

---

### **Contract Negotiation / Pre-Execution Review**

ALAs can rapidly identify and highlight differences between two different clauses or passages of text, making it easier for lawyers to see where a new contract varies from similar ones they have negotiated before. Some ALAs have been created with special features that make it easy for users to review similarities and differences between contracts and insert text with a single click, or to identify the most similar contract with the same party to the one currently being negotiated.

### **Due Diligence / Post-Execution Contract Review**

Identification of legal concepts or clauses across contracts and other legal documents is a core capability of ALAs, allowing for nimble contract review that is more intuitive than many previous solutions. Some ALAs will allow for uploading of contracts, or pointing the ALA at a database of contracts, providing the user with an interface for instructing the machine to identify and extract key information. Since the technology underlying ALAs is able to understand language at a conceptual level, it does well with variance and should require no additional training in order to identify clauses at high rates of precision.

Note: since these solutions are so new, they often lack the interface and process-driven infrastructure provided by incumbents. For example, ALAs will not yet have features such as the ability to develop a due diligence report alongside contract review, to tag or flag documents, or to collaborate with and assign tasks to team members in a large-scale due diligence.

### **What's Market**

Due to their ability to identify legal concepts across large numbers of documents, ALAs are often able to review contracts in order to extract data for the purpose of determining what the current market is for a particular provision. Some tools have been engineered to allow for the population of summary tables with extracted data.

## **Litigation**

---

### **Legal Research**

Companies that have access to databases of legal content have been able to create ALAs that can support legal research. This functionality will usually be presented in a chat format, so that users can ask questions in plain language and get a response that summarizes the legal position and sets out the sources that justify that position.

### **eDiscovery**

Just as ALAs are capable of reading and identifying data in contracts, they can also find information within discovery databases. ALAs will often provide users with the ability to interrogate a set of discovery documents to find relevant documents relating to a particular aspect of a case.

Note: Just as ALAs don't yet have the interface to support broad due diligence automation, so they are missing the comprehensive functionality provided by an established eDiscovery platform. For these types of legal work, the ALA should be viewed as an additional tool rather than a replacement for existing solutions. This will almost certainly change as the category evolves.

## **Other Use Cases**

---

ALAs are capable of addressing many other use cases, and for many firms it may be worth experimenting with the technology during a pilot phase in order to identify other workflows that can be improved or automated. Examples of other, discrete use cases include the development of deposition outlines (based on case-related documents that have been uploaded), review of transcripts for the purpose of extracting and tagging passages for witness or issues binders, lease review, classification of documents, and auto-tagging documents with metadata.

## **Benefits of AI Legal Assistants**

There are many benefits to using ALAs. Introducing this type of high-powered automation tool into a legal environment will allow lawyers to develop good first drafts faster, and provide a head start in relation to much other legal work. The technology doesn't remove the need for careful consideration and legal analysis, but it takes care of much of the work that used to be required before a legal analysis was able to be performed, thereby making legal work far more efficient. ALAs give lawyers a toolkit to draw upon that can provide them with support during almost any workflow. Some of the specific benefits are set out below.

### **Saving non-billable time**

The most obvious benefit of ALAs is efficiency. For lawyers working in firms, time-savings are a double-edged sword. Cutting time is not a clear benefit in an environment where value is directly tied to the time it takes to do something, and where merit is based on how many hours one has worked.

Keep in mind though that some of the time that is affected by ALA-driven automation might have been written off in the past. Spending time searching for previous work product, for example, or undertaking legal research across multiple platforms. Saving time on these tasks and on administrative matters will improve lawyer productivity and make it easier for them to work.

### **Multiple Outputs**

ALAs are generally flexible in the way that they present information, meaning that it is easy to ask them to generate work product in numerous formats, including spreadsheets, tables, bullet point lists, a presentation deck, and so on. This allows lawyers and business professionals to instantly create multiple assets with the same information. Users can also ask the tool to provide responses in a manner that mimics the format or style of an uploaded document.

## **Transformation**

Cynics may ask why law firms should deploy ALAs if they will automate so much work that lawyers can still bill time for.

Unfortunately for law firms, the word is out about generative AI and its impact. Clients are aware that this technology exists and that there are companies that have made it safe and effective to use in legal environments. Corporate legal departments themselves should be strongly considering an investment in ALAs because doing so will likely reduce the amount of work that they need to outsource.

Similar to using email instead of writing letters by hand or dictating them, this technology will become ubiquitous. It is disruptive to the status quo and law firms that are not using it will soon be unable to compete with firms that have invested early and are able to take on more work for less. One of the most important benefits to firms in deploying an ALA is that it will put them at the forefront of change in the legal industry at a time when true digital transformation has become feasible. Firms that choose to be at the vanguard of this transformation will almost certainly position themselves better for improved profitability and success in the future.

## **Risks and Adoption Challenges**

Generic generative AI and large language models have characteristics that can make them problematic for use in a legal environment. ALAs have largely been engineered to overcome these challenges, but it's still important to be aware of the potential risks with this type of technology.

### **Hallucinations**

One of the peculiar characteristics of generative AI is that, unless it is engineered to behave differently, it will make up answers to questions if it is unable to find information to support a factual answer. Since the technology is built to generate text that is reasonable in the context of what has come before and what the model has been asked, it will generate what would appear to be reasonable responses regardless of whether those responses are supported by facts. For example, generative AI chatbots like ChatGPT will make up entire cases, citing them correctly and conveying details about those fictional cases so that they sound convincing.

Some ALAs have been engineered to prevent this behavior, and it is critical when considering licensing this type of technology to ask the provider whether they have accounted for hallucinations and how. Obviously, there is substantial risk involved in deploying a solution that may produce false responses to legal questions.

### **Verification**



Both because of the risk of hallucination and because ALAs are drawing upon so much data to return an answer, lawyers should be trained to verify the responses provided before relying on them for inclusion in client work. This is especially true for ALAs that do not provide sources when answering a question.

### **Adoption**

The sense with products like ALAs is that they will automatically be adopted by lawyers because they are so ground-breaking. Hype in the market can make it feel as though this technology is a magic wand, transforming processes and workflows with no work required around process improvement or adoption.

While it's true that these solutions are more intuitive than many others, there is still work to be done to integrate them into legal workflows. Any roll-out will need to include communications about the use cases internally for which the ALA is deployed, and guidance and training for lawyers on when within their workflows they should leverage the technology.

### **Delegation**

Unlike junior associates or human assistants, ALAs do not know to ask clarifying questions. That means if users ask a poor question to an ALA, they will receive a poor answer. One of the risks around deploying an ALA is that many users interact with it in a way that gives rise to bad or incomplete answers. User training should therefore form a key part of any roll-out, including teaching lawyers how to interact properly with the ALA, and how to craft “prompts” that give rise to improved output.

### **Liability**

Almost certainly there will one day be a legal case that arose because a lawyer relied to their detriment upon a false answer generated by AI. As these developments in technology are entirely new, it is unclear where the liability would fall in respect of such a claim. In negotiating a license with an ALA provider, legal organizations should pay close heed to liability terms.

**For a comprehensive list of factors to consider when evaluating ALAs, see the LTH ALA Evaluation Framework.**

## **Who's Who: The Players**

This is a constantly evolving category, and new players are likely to emerge.

Note that in addition to the four products listed below, there are other ALAs that focus on specific areas of law, such as [DISCO's Cecilia](#) which performs AI assistant tasks in relation to eDiscovery and within the DISCO environment, and [ContractPodAI's Leah](#), which focuses on tasks in the contract management lifecycle.

## Casetext CoCounsel

---

CoCounsel is a legal assistant engineered specifically for legal environments, in order to enable lawyers to work faster, producing better outcomes for their clients.

CoCounsel's abilities include:

- **Document Review:** ask a question about your documents and CoCounsel will read them in full and answer, with citations to sources. Quickly find everything from critical testimony in voluminous transcripts to key terms in dense contracts.
- **Deposition Preparation:** Describe the deponent and what's at issue in the case, and CoCounsel will identify multiple highly relevant topics to address and draft questions for each, providing a thorough deposition outline.
- **Database Search and Retrieval:** Ask a question and CoCounsel will search through a database of documents, read the relevant ones, and provide an answer. Find and retrieve what you need, like the right template or precedent, previous work product, or internal know-how.
- **Legal Research and Memo Generation:** Ask a research question, giving as much detail as you like including facts, jurisdiction - and in a few minutes CoCounsel will retrieve on-point resources and provide a written answer with explanation and supporting sources.
- **Summarize:** Interpret and condense key information in any type of document faster than humans can without missing any details.
- **Extract Contract Data and Track Deal Terms:** Review thousands of contracts and rapidly extract a list of relevant clauses based on your specifications, automatically populating summary tables and making it easier for you to track deal terms and establish what's market.
- **Contract Policy Compliance:** CoCounsel captures every single clause in a set of contracts that doesn't comply with a policy or set of policies, reports the risks of using non-compliant language, and recommends revisions.

CoCounsel is being used in multiple AmLaw 100 firms as well as by some large corporate legal departments.

## Harvey

---

Harvey is an AI legal assistant built on Open AI's GPT-4. Harvey partners with entire firms in order to develop solutions within the platforms and workflows that lawyers already use. These close collaborations between Harvey's engineers and law firms allow for solutions that solve real problems.

As Harvey is rolled out in workflows that are specific to a particular firm, it's hard to specify use cases. Press releases about Harvey's relationship with Allen & Overy suggest that it is being used for drafting of correspondence, development of strategy, first drafts of contracts, client memos, and prospectuses, as well as more administrative tasks across the firm. In May 2023 PWC announced it has also signed a license with Harvey.



## Rally Spellbook

---

Spellbook is an AI legal assistant with various skills focused on drafting and contracts. It uses GPT-3 to review and suggest language for your contracts, right in Microsoft Word, detects aggressive terms, lists missing clauses, and provides negotiation suggestions. Automated due diligence functionality is coming soon.

Spellbook's skills include:

- Suggesting language: Instantly drafts new clauses and entire sections based on the context of your agreement.
- Detecting aggressive terms: Identifies sneaky or unusual terms.
- Listing missing clauses and definitions: Suggests important language and clauses that may be missing from the contract.
- Negotiation suggestions: Suggests common points of negotiation.
- Auto-diligence (coming soon): Automatically detect the most pressing issues across an entire data room.

## Lexis+ AI

---

Launched in May 2023, Lexis+ AI is an AI Legal Assistant and generative AI platform built and trained on the Lexis Nexis repository of legal content, leveraging an extensive collection of documents and records to provide customers with comprehensive legal results.

Use cases include:

- Conversational search, to simplify the complex and time-consuming legal research journey, providing a search experience for diverse legal questions and supporting answers with citations and sources.
- Enhanced summarization, providing a custom summary of legal documents to speed up and guide insightful analysis.

Generative document drafting, guiding users throughout the legal drafting process, generating a first draft of a legal document, and allowing users to change the language and tone from a simple prompt.

## LawDroid CoPilot

---

LawDroid's Copilot leverages generative AI to help lawyers research legal issues, draft emails and letters, and summarize documents.