



The *Legal* AI Handbook



Foreword

When OpenAI released ChatGPT in 2022, the world seemed to immediately change. We’ve been told AI will impact every aspect of our jobs, and our lives, and hype, fear, excitement, optimism, and everything in between have dominated conversations in the two years since.

One thing we know for sure is that the legal industry is among the vanguard when it comes to realizing the early benefits of generative AI. Earlier this year we released the [State of AI in Legal Report](#), which found that **74% of legal professionals are using AI for legal work**—90% of which plan to use AI more frequently next year. As of the publication of this guide, over 80% of Ironclad customers actively use our own AI.

But not only do lawyers have a high bar—the margin for error in legal work is razor thin. Generative AI is a rapidly developing technology, and in order to fully, and responsibly, wield its full potential, it’s critical that we understand it. What are the risks? How can I actually use it? What security measures do organizations need to take, and how can we properly evaluate AI software?

This guide is meant to provide the knowledge you need to not only understand AI – but deploy it intelligently in real world scenarios. It’s important to note that while not exhaustive, this guide provides a comprehensive look at AI in the legal field. See our table of contents to guide you to the exact areas you’re looking for.

74%

of legal professionals
use AI for legal work



Read More

[State of AI in Legal Report](#)

Table of Contents

1

The Basics

2

The Practical Power of AI

3

The Impact of AI-Powered
Contract Management

4

AI & Security: Lifting the Hood

5

How to Evaluate AI Software

6

Tips on Getting Started

7

AI Innovation in the Legal Field

8

How Ironclad Helps

CHAPTER ONE

The Basics

A glossary of AI terms

As with most any industry, artificial intelligence comes with its own terminology. Here's a rundown of key terms to be familiar with as you brave this new world:

ALGORITHM

A set of rules or instructions given to an AI/ML system, or computer, to help it learn, make decisions, and solve problems.

ARTIFICIAL INTELLIGENCE (AI):

It's a little bit tricky to nail down a single definition of AI that we can all agree upon, but this one from the [MIT Technology Review](#) is pretty good:

AI is a catchall term for a set of technologies that make computers do things that are thought to require intelligence when done by people. Think of recognizing faces, understanding speech, driving cars, writing sentences, answering questions, creating pictures. But even that definition contains multitudes.

The basic idea is that AI refers to technology that can do human-esque tasks at something approaching — or, eventually, surpassing — human levels of competence. That's a pretty broad domain, but so too is the potential for AI applications.

CHAIN-OF-THOUGHT PROMPTING

A prompting technique that significantly improves the ability of large language models to perform complex reasoning. The technique centers around prompting an LLM to generate a series of intermediate reasoning steps (aka a chain of thought) before arriving at its final response.

CONVERSATIONAL ASSISTANT

Conversational assistants are programs designed to simulate human conversation. They're often trained and deployed to handle specific use cases, like customer service. Voice assistants like Alexa (Amazon) and Siri (Apple) use a combination of NLP, speech recognition and synthesis, and other technologies to carry on conversations aloud with human users.

COPILOT

A copilot is a type of assistant generally designed to help users accomplish common tasks faster than before. The term is usually used to refer to a workplace assistant, like a programming copilot or sales copilot.

DEEP LEARNING

A subset of machine learning that uses artificial [neural networks](#) to learn from data without human domain knowledge. The “deep” in deep learning refers to the use of multiple layers in the networks. Deep learning is generally more capable and accurate than traditional machine learning.

EXTRACTION

The process of identifying and pulling out specific data points or attributes from documents, such as contracts. This process uses AI and machine learning to recognize patterns and extract relevant information, transforming unstructured text into structured, searchable data. In a legal context, metadata is often extracted from contracts because they are the key documents in which all legal data lives.

GENERATIVE AI (GEN AI)

As the name suggests, GenAI refers to AI systems that generate new content, like writing, computer code, images, audio, and video. Generative AI models identify the patterns and structures within existing data and use that knowledge to create original content.

HALLUCINATION

Hallucinations are false or misleading outputs generated by AI models. The term is most often used to describe when an LLM makes something up and presents it as truth as part of a text response.

LARGE LANGUAGE MODEL (LLM)

LLMs are algorithms that can recognize, predict, and generate content using very large datasets. They are a type of Generative AI specifically developed to generate text. GPT (OpenAI), Claude (Anthropic), Gemini (Google), Llama (Meta), and Nemotron (NVIDIA) are all examples of LLMs.

MACHINE LEARNING

A subset of AI that gives computers the ability to learn from data, without explicit programming.

METADATA

Key information and attributes that describe and categorize elements in a document or file, like a contract.

MODEL

An AI model is a program that has been trained on a set of data to recognize patterns and make predictions. Chatbots like ChatGPT are powered by models that can understand requests and respond to them by predicting the best possible response, one small chunk of data — or, token — at a time. Image generation programs are similarly powered by models that predict the best visual response to a prompt, token by token. Models can be trained on various data modalities (eg, Text, images, audio, etc), and on data sets curated for specific industries or purposes.

NATURAL LANGUAGE PROCESSING (NLP)

Technology that gives machines the ability to understand, interpret, and generate human language. NLP combines computational linguistics, machine learning, and deep learning to process human language.

OPTICAL CHARACTER RECOGNITION (OCR)

OCR is a technology that turns images into text that can be read, searched, and indexed by machines. For instance, OCR can be used to turn PDFs of contracts into digital documents (that then become fully searchable).

PROMPT

An instruction given to an AI system. Eg, “Write an article outlining the current state of GenAI in the legal profession.”

REINFORCEMENT LEARNING

A type of machine learning in which the model learns through trial and error, receiving rewards or being penalized for its actions.

RETRIEVAL AUGMENTED GENERATION (RAG)

RAG is a type of generative AI that relies on a database of knowledge to retrieve information from. By combining retrieval mechanisms with generation capabilities, RAG can provide contextually relevant, reliable, and accurate responses that are tailored to specific industry needs.

SENTIMENT ANALYSIS

An NLP technique that determines the emotional tone or attitude expressed in a piece of text. In the context of contract analysis, it can be used to assess the overall tone of a contract or specific clauses, helping to identify potentially adversarial, or favorable, language.

SUPERVISED LEARNING

A type of machine learning in which the model is trained on labeled data, like a set of human facial images each labeled with a person’s name.

TRAINING AND INFERENCE

Training is the process of training an AI model on a data set. Inference is when you pass data through the model to generate a response.

TRAINING DATA

The initial data used to teach a machine learning model. Training datasets used to train LLMs are typically enormous, ranging into the hundreds of billions of words in length.

UNSUPERVISED LEARNING

A type of machine learning in which the model finds patterns in unlabeled data, like a set of human facial images without any names, genders, or other identifying labels attached to them.

A brief timeline of AI history

To understand where we're going, it's helpful to know where we've been. Here's a quick timeline of AI's evolution:

1950s - 1960s

The concept of **AI is born**, with early experiments in ML and NLP.

- British mathematician **Alan Turing** develops a way to assess if a machine thinks on par with a human. Turing actually called his method “the imitation game,” but it was soon known as “the Turing test.”
- 1956: Dartmouth mathematics professor **John McCarthy** hosts a summer workshop largely credited with founding the field of “Artificial Intelligence.”

1970s - 1980s

AI winter sets in due to limitations in computing power and overhyped expectations.

- 1974: **Sir James Lighthill**, an applied mathematician, publishes a highly critical report on academic AI research. Lighthill’s claims that researchers over-promised and under-delivered when it came to the potential intelligence of machines resulted in massive funding cuts across the field.
- 1986: Ernst Dickmanns, a scientist in Germany, soups up a Mercedes van with sensors and on-board computers to create the **first self-driving car**. It could only drive on roads empty of other vehicles.

1990s - 2000s

Resurgence of AI with advances in ML algorithms and increased computing power.

- 1996: World chess champion **Gary Kasparov defeated IBM’s Deep Blue** computer in a chess match, 5 games to 1.
- 1997: **Deep Blue wins a rematch over Kasparov**, needing only 19 moves to take the final game.

2010s

Deep learning breakthroughs lead to significant advances in image and speech recognition, NLP, and other AI applications.

- 2012: **AlexNet** wins the ImageNet competition, shining a light on the potential of neural networks and deep learning. AlexNet is one of the most influential computer vision research papers ever published, and is credited with accelerating deep learning. As of mid-2024, the AlexNet paper has been cited over 157,000 times.
- 2013: **Geoffrey Hinton**, a professor who co-authored the paper with two graduate students, joins Google. One of the student authors, **Ilya Sutskever**, would later co-found OpenAI.
- 2017: Eight scientists working at Google publish, “**Attention is All You Need**.” The seminal paper introduced the **transformer**, a then-new deep learning architecture that most LLMs are now based on.

2020s

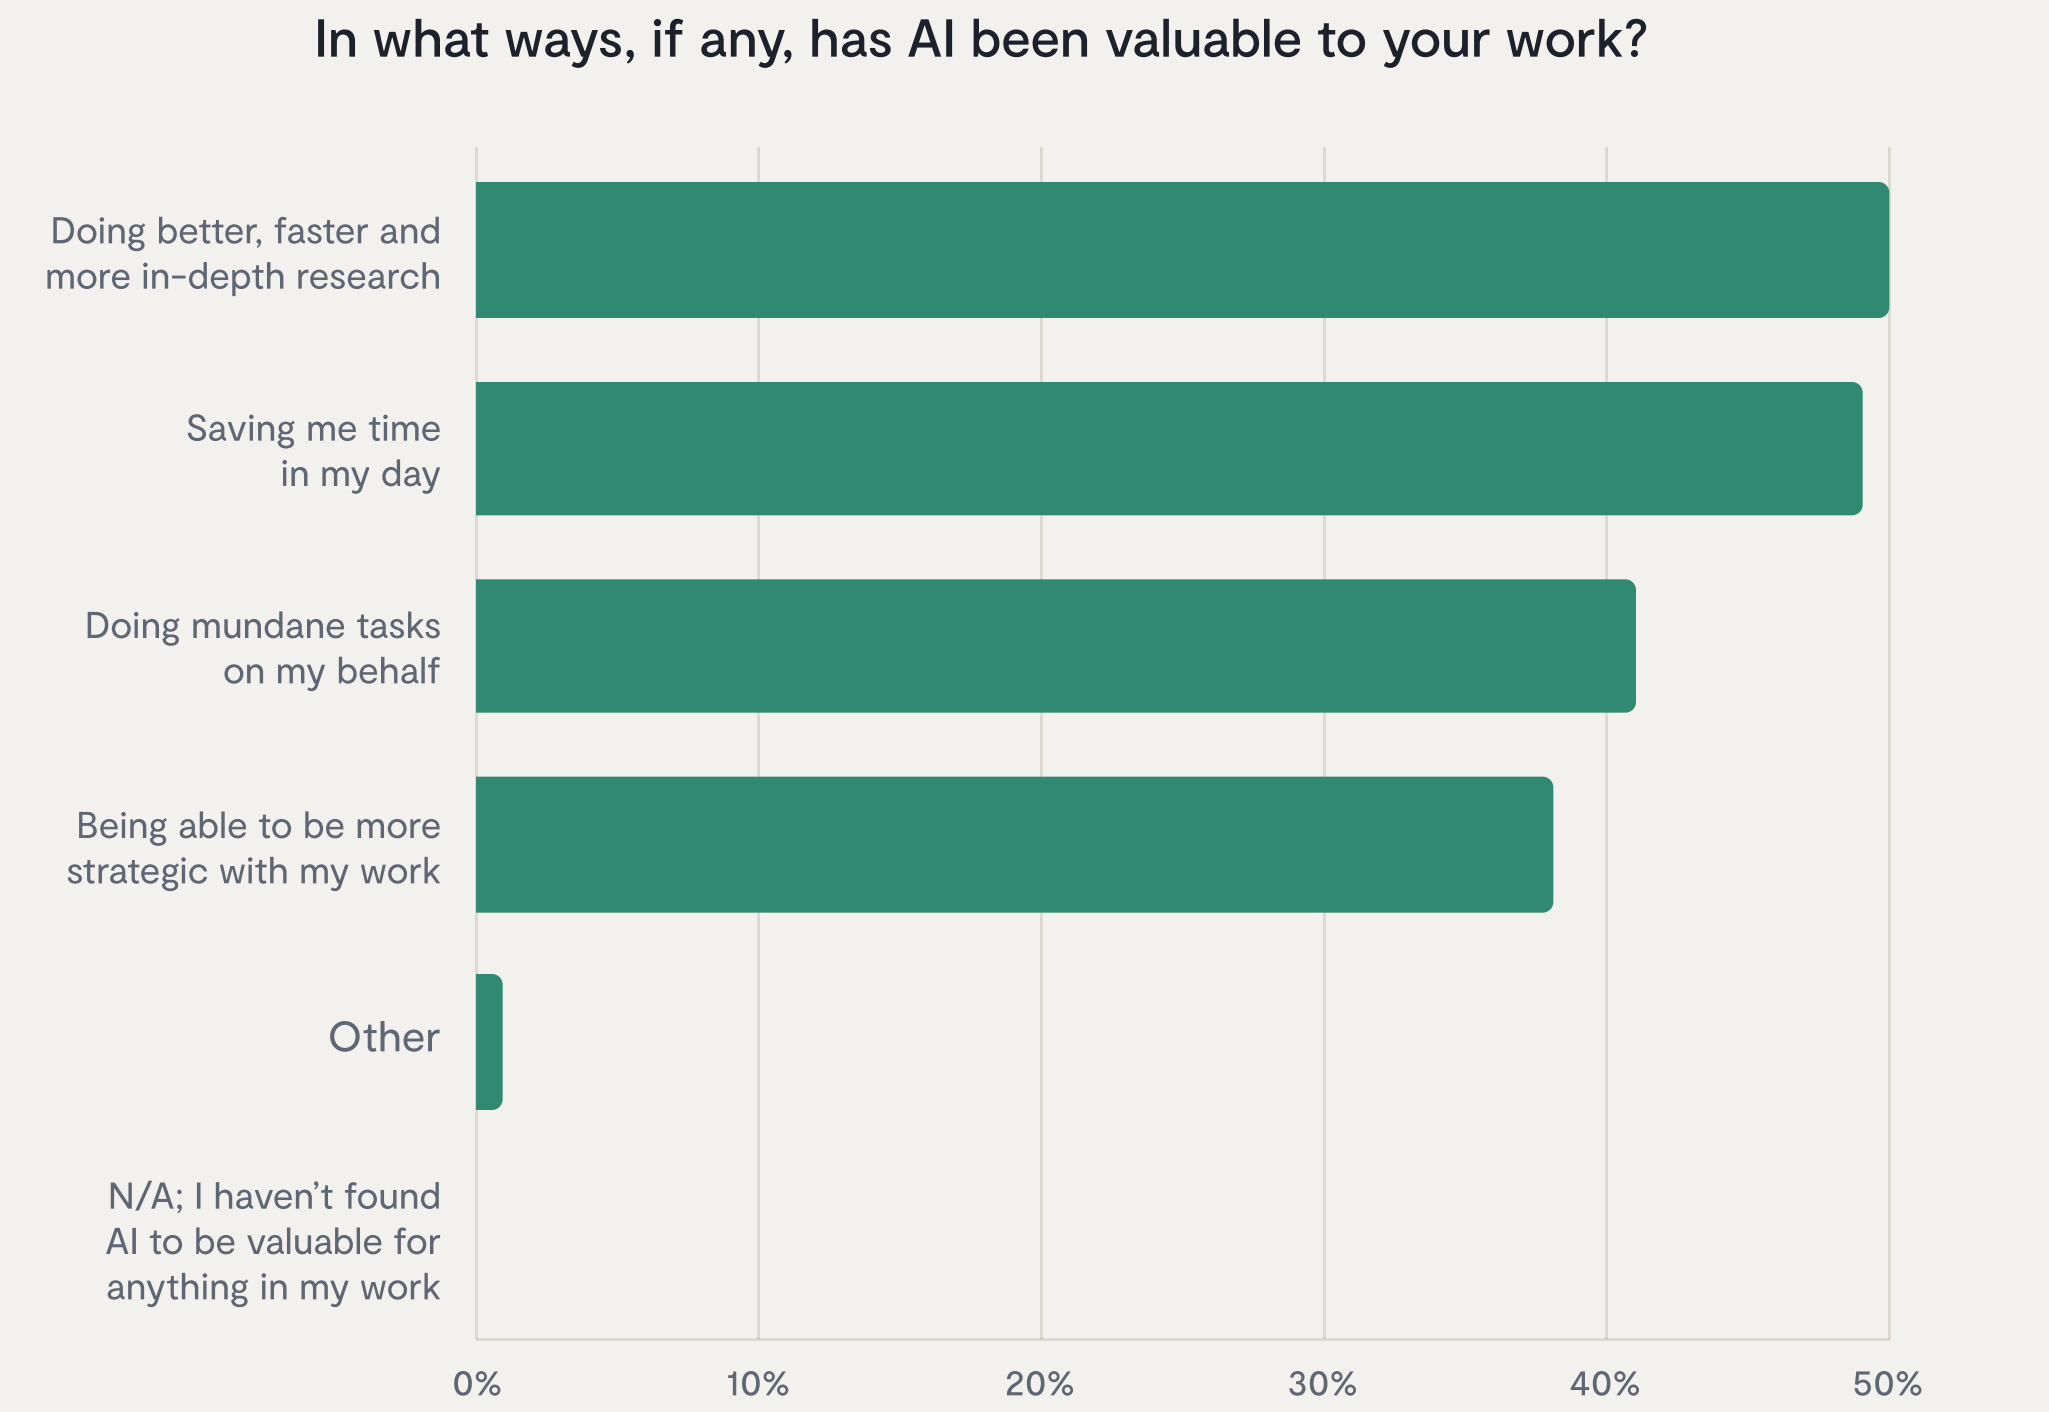
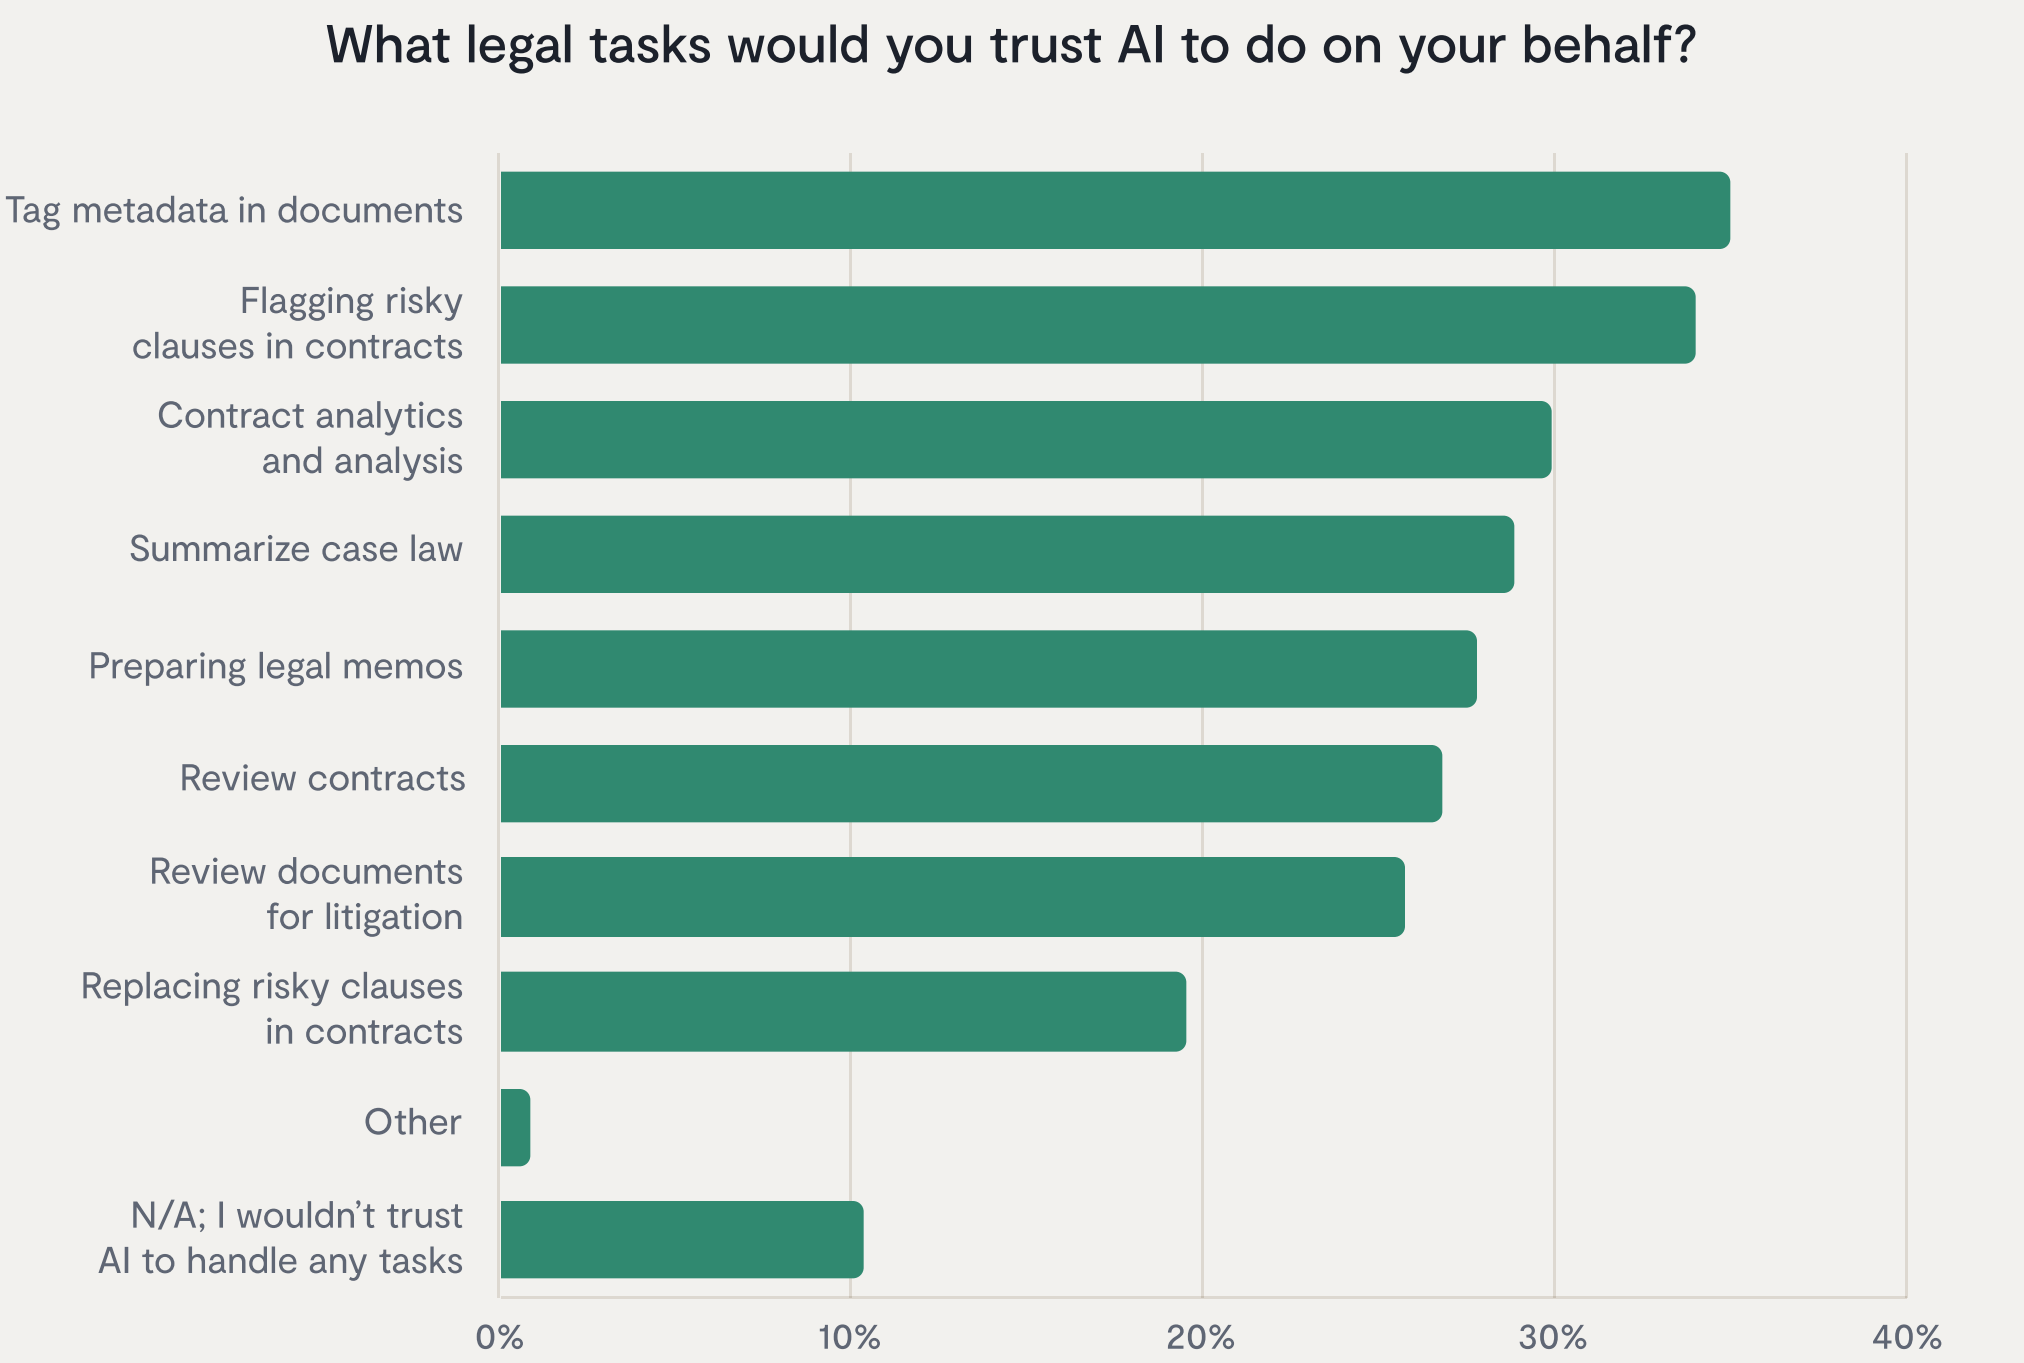
The rise of LLMs and GenAI, exemplified by tools like ChatGPT and Midjourney, opens up new possibilities for AI in various fields, including law.

- 2022: The first direct-to-consumer AI company **OpenAI** releases **ChatGPT**, a chatbot and virtual assistant based on their GPT series of LLMs. ChatGPT becomes the fastest growing consumer software application in history, racking up 100+ million plus users in just over a month.
- 2024: **GenAI** takes hold in mainstream use, with AI-powered business and consumer applications addressing use cases from image generation and automated customer service to legal operations and beyond.

How could AI make life better?

In general, AI tools excel at taking over repetitive tasks and turning mountains of information into usable insights. AI tools augment legal work; they don't do it for you—but it can take remote, mundane tasks off your plate and give you back time to focus on higher level tasks, planning, and strategy.

Our recent [State of AI in Legal report](#) uncovered a few examples of where legal teams were interested in—and more importantly, trusting of—using AI in their day-to-day roles:



But if we dig deeper, the gains from using AI go much deeper than the tasks they perform. For instance, the aforementioned report found that the three leading benefits of AI were 1) doing better, faster, and more in-depth research, 2) saving time during the day, and 3) doing mundane tasks on legal's behalf. Today, over half of lawyers are [unsatisfied with work](#)—citing inundation with stressful deadlines and an overwhelming amount of tasks—but 57% of those lawyers believe AI can help alleviate the dissatisfaction.

A shift in mindset

But even before getting started, succeeding with AI is all about adopting the right mindset. Don't look at AI as a silver bullet designed to deliver instant results, but instead as a powerful tool to help you iterate on content and ideas, take some of the drudgery out of mundane tasks and help you work faster and more efficiently.

Having a misguided mindset about what AI can and can't do will derail your project before it even gets off the ground. A few things to keep in mind are:

- AI might feel like magic—but just like any other tool, it takes time to stand up. Especially with regards to training your AI, this will not happen overnight. But the good news is the more you train it, the better it becomes.
- Keeping your data accurate and up to date is critical. As they say, “garbage in, garbage out.”
- The rate of advancements in foundational models is exponentially increasing—so even if a solution doesn't meet your needs today, it will likely surpass those needs soon.



CHAPTER TWO

The practical power of AI

We're still in the early stages of the artificial intelligence revolution, but companies, legal teams, and individual professionals are already seeing real benefits from using AI on the job—and not just for general tasks like brainstorming, softening an email, or helping you figure out that Excel function.

For those of you who have been experimenting with AI already, this may not be surprising, but a few of the primary ways to use AI to boost your personal productivity is:


- **Task automation:** AI can take over routine, time-consuming tasks, freeing up lawyers to focus on high-value work that requires human judgment and creativity.
- **Data-driven decision making:** AI can analyze vast amounts of data quickly, providing insights that can inform strategy and decision-making much faster than manual analysis.
- **24/7 availability:** Unlike human workers, AI tools can work around the clock, helping to meet tight deadlines and manage high volumes of work.
- **Personalized learning:** AI can adapt to individual working styles and preferences, offering personalized recommendations and assistance. The rise of AI-powered “copilots” to help humans in the workplace is widely thought to be one of the next big trends in technology and business.

Conversational AI – what is it, and what can it do for productivity?

One category of AI that gets outsized personal use—and certainly no shortage of media attention—is conversational AI. Conversational AI refers to technologies that allow computers to understand, process, and respond to human language in a natural way. [ChatGPT](#) is a conversational AI chatbot, as are [Claude](#), [Gemini](#), and the countless other LLM-driven chatbots that have popped up in the past few years.


Conversational AI differs from GenAI in that the former is built to hold authentic, two-way conversations with human users, while GenAI is designed to produce original content when prompted. In the legal context, conversational AI can take the form of chatbots, virtual assistants, or more sophisticated language models that can engage in complex legal discussions.

When it comes to productivity for legal professionals, what you can do with conversational AI runs the gamut, but some of the quickest categories to start seeing immediate benefits with include:




Drafting legal documents, like:

- Drafting contracts and policies
- Creating document summaries
- Drafting emails and other communications
- Generating amendments, addenda, and other ancillary documents



Legal **research**, like:

- Understanding new policies, regulations and local / federal laws
- Fact checking documents
- Extracting and aggregating data from multiple documents



Reviewing legal documents, like:

- Reviewing and redlining contracts
- Interpreting legal language
- Answering questions about a specific document based on historical data
- Comparing multiple documents and document versions



AI in Contract Lifecycle Management

You've seen up until this point that there is no shortage of what AI tech can and will do for legal work. But where do those capabilities actually live? How do early adopters access them today? The answer for many is a category of tools that has become a must-have for both private firms and in-house teams: contract lifecycle management software, or CLM.

CLM systems have been around for a long time, but, as with tech across virtually every sector today, they've all recently been turbocharged by AI. Many could already streamline stages of the contract process, from creation to negotiation, execution, and post-execution. But now, because AI uses ML and NLP to mimic human processes, a truly robust AI-powered CLM will do all that, plus help with:

- **Contract ingesting and tagging.** CLMs use AI to analyze, tag, extract, and report on contract data.
- **Building standardized templates** and assisting in **negotiating, redlining, and reviewing contracts.** AI-powered CLM tools can automatically locate problematic clauses within contracts before you send them out.
- Providing **contract analytics insights** to help organizations **make data-driven decisions** about their contractual relationships.
- **Boosting operations and productivity** by streamlining the contract creation, review, and renewal process.

Why contract data is an ideal place to start using AI tech

There’s another reason why CLMs are ideal spaces to start with AI: they house contract data. And working with contract data presents an opportunity for legal departments to drive real, data-driven business impact through data locked in contracts—and serves as a perfect experimentation environment for five main reasons:

1

Universality

Contracts are one of the most ubiquitous business tools on the planet. Every department, within every company in the world, uses contracts.

2

Accuracy

By nature, contracts are highly accurate. Having been reviewed by teams of lawyers and stakeholders, for the most part, your contracts will be complete, accurate, and, more or less, final.

3

Volume and repetitive structure

Organizations typically handle a large number of contracts, many of which share similar structures and clauses. This repetition creates a perfect learning environment for AI systems, and the structured format of most contracts also provides a consistent framework for AI algorithms to analyze.

4

Risk mitigation

Manual contract review is time-consuming and prone to human error, and unfavorable or risky terms can often hide in contract data. AI excels at uncovering problematic language, and at creating guardrails so employees don't inadvertently introduce risk while drafting or editing contracts.

5

Customization opportunities

Contract data is a good place to get your feet wet with customization capabilities as most organizations usually have custom clauses, fields, and metadata properties that are pre-approved for use.

Conversational AI for contract data

Conversational AI, in particular, has tremendous potential to transform how legal professionals interact with contract data. Why? Because it's so natural to use, there's hardly any learning curve involved.

Here's a small sampling of some of the ways conversational AI can help you get more out of working with contract data:



Ask questions about your day-to-day data.

Conversational AI tools understand natural language queries, so it's easy to talk with them about everyday tasks. Ask questions about your contracts in plain language, like, "What are our obligations under this contract?" or "Which contracts are up for renewal next quarter?" Then, ask follow-ups and dig deeper into the data than you would on your own.



Brainstorm with AI.

Chatbots can do more than just answer questions about the data they're trained on. They're great at distilling complex information into concise summaries and make killer brainstorm partners as well. AI chatbots can help you come up with new ideas, strategies, and plans for improving your contracting process or making data-driven decisions across the business. Talk to the bot like a brainstorming partner, and see what you can come up with together.



Gain insights on how you negotiate.

Conversational AI can analyze your historical contract negotiations and provide valuable insights into your negotiation style and tactics. By examining patterns in your past negotiations, the AI can tell you which clauses are getting negotiated the most, which contract types take the most negotiation time, and if there are major areas of revenue leakage in your historical contracts, amongst other things. A conversational AI chatbot can even simulate negotiation scenarios to help you practice.

Customization in AI-powered contract management

AI systems are designed to be highly customizable. When we talk about customization in AI-powered contract management, we're not talking about the software itself, or some kind of bespoke licensing agreement to use the software. We're talking about how AI systems allow legal professionals to tailor their individual work processes and outputs.

Let's examine some concrete examples:



Read More

[Check out how Ironclad is approaching custom AI](#)



Custom clause libraries

AI systems can be trained on an organization's preferred clauses, ensuring consistency across all contracts while maintaining the unique voice and requirements of the business. A legal department could use AI to create a custom clause library for merger and acquisition (M&A) contracts by analyzing thousands of past deals to suggest optimal clauses based on transaction type, jurisdiction, and client industry.

Here's how this goes beyond simple template management:

- The system learns the nuances of your clause preferences based on example text you provide, including specific language, formatting, and contextual usage.
- As new clauses are approved and added to the library, AI adapts in real-time, suggesting these new clauses as appropriate in future contracts.
- The system can identify when a proposed clause deviates from the preferred library, flagging it for review and suggesting alternatives.
- Over time, the AI can analyze the performance of different clauses, providing insights into which ones lead to faster negotiations or fewer disputes.

Industry-specific models and retrieval-augmented generation

While general contract AI models are powerful, industry-specific fine-tuning takes their capabilities to the next level. An AI model fine-tuned on drug licensing agreements, for instance, could accurately identify and assess complex royalty structures, regulatory compliance clauses, and intricate intellectual property terms specific to biotech partnerships.

- Fine-trained models can become fluent in industry-specific jargon, regulations, and standard practices, reducing the need for constant human oversight.
- The models can recognize and flag industry-specific risks that might be overlooked by a generalist system.
- The models can suggest industry-standard clauses and terms that might be missing from a draft contract.
- As regulations change, the AI model can be updated to ensure compliance across all new and existing contracts.

Alternatively, more and more applications are using retrieval-augmented generation, or RAG for short, as a more cost-effective and broad approach. In a nutshell, RAG is a type of GenAI that relies on a trusted database of knowledge to retrieve information from.

By combining retrieval mechanisms with generation capabilities, RAG can provide contextually relevant and accurate responses that are tailored to specific industry needs.

- RAG models can access up-to-date information and domain-specific databases, ensuring that the generated content reflects the latest trends, regulations, and best practices within an industry.
- By retrieving pertinent information, these models can produce more informed outputs, improving the quality of responses in specialized fields such as law, finance, or healthcare.
- The retrieval component allows RAG to maintain a balance between creativity and factual accuracy, enabling it to generate content that is not only coherent but also grounded in real-world data.

3

Risk scoring

Customizable risk assessment is a game-changer for legal teams, allowing them to align AI outputs with their organization's risk tolerance and priorities. Imagine a multinational company that trained an AI contract analysis tool to assign risk scores based on its specific regulatory landscape and risk tolerance. The tool could quickly identify high-risk clauses in vendor agreements across multiple countries the company does business in, allowing for renegotiation to significantly reduce their overall risk profile.

- Organizations can define their own risk categories and weightings, ensuring the AI system focuses on what matters most to them.
- The system can be trained to recognize subtle indicators of risk that are specific to the organization's history and context.
- Risk scores can be dynamically adjusted based on changing market conditions or company priorities.
- The AI can provide both detailed explanations for its risk assessments and visualizations of the referenced data, allowing legal professionals to understand and validate the scoring.

4

Workflow Integration

Integrating AI into existing contracting workflows goes beyond simple automation, offering a tailored approach to contract management:

- Create custom approval chains and task systems with routing contracts to the appropriate stakeholders based on content, risk score, or other defined criteria.
- Escalation rules can be fine-tuned to match the organization's hierarchy and decision-making processes.
- AI models can learn from past workflow patterns, suggesting optimizations to reduce bottlenecks and speed up contract cycles.
- Automated notifications and reminders can be customized to match the communication style and urgency levels preferred by different team members.

Another way to take advantage of Generative AI for contract work is to train the system on custom language your company and/or industry frequently uses. Let's take a look at how it works.

Data accuracy through custom training

Data accuracy is crucial for effective decision-making and risk mitigation. Custom training plays a pivotal role in enhancing an AI system's ability to understand and process contracts accurately. Here's how it works:

Clause identification

AI systems can be trained to recognize and categorize specific clauses unique to your company or industry. This customization goes beyond generic clause recognition, as an AI model can learn to identify company-specific language and formatting preferences.

From there, It can recognize industry-specific clauses that may not be common in general contract databases. The system can then be trained to differentiate between subtle variations of similar clauses, ensuring precise categorization. As new clause types are introduced, the AI model can be quickly updated to recognize and categorize them appropriately.

Data extraction

Custom training significantly improves an AI system's ability to extract relevant data points from contracts.

An AI model can be taught to recognize and extract specific data fields that are crucial to your company, but might be overlooked by generic systems. The system can learn to interpret context-dependent information, like when certain data points are relevant to the contract type or parties involved. The model can also be trained to handle complex data structures, such as nested clauses or interdependent terms. As data extraction accuracy improves, the need for manual verification decreases, streamlining the contract review process and freeing staff for other, higher-order work.

Variation recognition

AI can be taught to recognize acceptable variations of standard clauses, improving flexibility in contract analysis. First, the system learns to identify when a clause is substantially the same as a standard clause, even if the wording is slightly different. It can then flag variations that fall outside of acceptable parameters, flagging them for review by human experts.

The AI can suggest standardization opportunities when it encounters frequently used variations of clauses. This flexibility allows for more nuanced contract analysis, accommodating the real-world diversity of contract language while maintaining consistency.

Error reduction

AI feeds on data, and AI models can learn from experience over time. As an AI-powered contracting system learns from company-specific contracts, it becomes better at identifying and flagging potential errors or inconsistencies.

The system develops a deep understanding of what "normal" looks like for your organization's contracts, making it easier to spot anomalies. That understanding makes it easier to identify potential errors in data entry, such as incorrect dates or mismatched party names. The AI model can also flag inconsistencies between different sections of a contract, helping to ensure coherence within the document. Over time, the AI system can provide insights into common error patterns, allowing for proactive improvements in contract drafting and review processes.

What data should I track?

When implementing AI in contract management, tracking the right data is essential for overall performance and return on your investment. Here are key metrics to consider tracking, along with tips on why the data matters and how best to leverage it.



Contract lifecycle time

WHAT IT IS

Elapsed time from contract initiation to execution.

WHY IT MATTERS

Shorter cycle times can lead to faster deal closures and improved business agility.

HOW TO USE IT

Identify bottlenecks in the process and implement targeted improvements. For example, if the data shows that legal review of a particular contract type takes an average of 5 days, the team could implement an AI-assisted pre-review against an approved language playbook to cut it down to 2 days.

2 Approval times

WHAT IT IS

Time spent in various approval stages.

WHY IT MATTERS

Long approval times can delay contract execution and potentially lose business opportunities.

HOW TO USE IT

Identify which types of contracts or clauses tend to cause delays and streamline approval processes. For instance, analysis reveals that non-disclosure agreements (NDAs) with non-standard confidentiality terms take 3 times longer to approve, so pre-approved alternative clauses are created to shorten approval times.

3 Negotiation rounds

WHAT IT IS

Number of back-and-forth exchanges during negotiation.

WHY IT MATTERS

Excessive rounds can indicate inefficiencies or misalignments in the negotiation process.

HOW TO USE IT

Develop strategies to reduce negotiation rounds, such as improving initial contract drafts or providing negotiators with better data. For example, after noticing 5+ rounds of negotiation on payment terms, the team could create an AI-powered clause library with pre-approved variations, reducing negotiation to 2-3 rounds, on average.

4 Risk scores

WHAT IT IS

AI-generated risk assessments for each contract.

WHY IT MATTERS

Helps prioritize high-risk contracts for review and informs risk management strategies.

HOW TO USE IT

Adjust risk tolerance thresholds, focus resources on high-risk areas, and track risk trends over time. For instance, the legal team could set an alert for contracts scoring above 7/10 on their internally created risk scale, ensuring these receive priority review by senior attorneys.

5 Compliance metrics

WHAT IT IS

How well contracts adhere to internal policies and external regulations.

WHY IT MATTERS

Ensures legal and regulatory compliance, reducing the risk of penalties or legal issues.

HOW TO USE IT

Identify areas of frequent non-compliance and implement targeted training or process improvements. For example, if AI analysis discovers that 30% of contracts exclude GDPR clauses, the team could implement an automated GDPR clause insertion and verification step.

6 Value leakage

WHAT IT IS

Instances where contract terms are not fully leveraged or enforced.

WHY IT MATTERS

Represents lost value or missed opportunities for the business.

HOW TO USE IT

Implement systems to better track and enforce contract terms, and educate stakeholders on contract value maximization. For instance, AI could flag unused volume discounts in supplier contracts, leading to cost reduction through better term utilization.

7 Contract renewals

WHAT IT IS

Upcoming renewal dates and associated values.

WHY IT MATTERS

Proactive management of renewals can lead to better terms and prevent unintended auto-renewals.

HOW TO USE IT

Set up automated alerts and processes for timely review and renegotiation of contracts approaching renewal. For example, an AI system that alerts the team 90 days before each SaaS contract renewal would allow time for usage analysis and renegotiation, potentially saving on renewals.

8 Clause language flagging

WHAT IT IS

Frequency and context of specific clause usage across contracts.

WHY IT MATTERS

Provides insights into negotiation patterns and potential areas for standardization.

HOW TO USE IT

Identify frequently negotiated clauses for potential pre-approval, and standardize common variations. For instance, if analysis shows 5 common variations of a liability clause, creating a pre-approved clause menu could significantly reduce negotiation time.

9 Obligation fulfillment

WHAT IT IS

Tracking of contractual obligations and their completion status.

WHY IT MATTERS

Ensures all parties are meeting their contractual commitments and helps prevent disputes.

HOW TO USE IT

Set up automated reminders for upcoming obligations and regularly review fulfillment status. For example, an AI system could track delivery deadlines in manufacturing contracts, sending alerts ahead of due dates to potentially cut down on late deliveries.

Tracking these metrics can help you gain valuable insights into contract management processes, identify areas for improvement, and leverage AI to drive better outcomes in your contractual relationships.

How can I use this data to do my job better?

The global legal technology market has grown significantly in recent years and GenAI will accelerate this growth, meaning the market will reach \$50 billion in value by 2027, according to [Gartner, Inc.](#)¹ This investment is not just about automating routine tasks; it's about transforming how legal professionals approach their work. As AI takes over the mundane aspects of contract management, legal experts are free to engage in more strategic, high-value activities.

The true power of contract data lies in how it can inform decision-making and strategy. Predictive analytics, for instance, allows legal professionals to anticipate potential issues in new contracts or negotiations based on historical data. A 2024 survey by [Lex Machina](#) found that 100% of legal analytics users find it valuable, with 70% saying successful litigation outcomes drives their usage, and 69% saying improved efficiency drives theirs.

Additionally, the insights gleaned from contract data enable legal teams to allocate resources more efficiently. By identifying which types of contracts — and stages in the contract lifecycle — require the most time and attention, teams can prioritize their efforts where they'll have the greatest impact. A recent [EY Law Survey](#) revealed that 99% of organizations report that managing current contracting workloads is a challenge, highlighting the potential impact of data-driven resource allocation.

Perhaps most importantly, leveraging contract data allows legal professionals to contribute more directly to business strategy. By analyzing trends in negotiation patterns, compliance metrics, and contract performance, legal teams can provide valuable insights that inform broader business decisions. Moving forward, the ability to effectively leverage contract data will become a key differentiator for legal professionals and help them become drivers of tangible business impact and strategic partners to the business.

¹ *Gartner, Gartner Predicts the Global Legal Technology Market Will Reach \$50 Billion by 2027 as a Result of GenAI, 25 April 2024, Rob Van Der Meulen, et al.*

CHAPTER THREE

Impact of AI-Powered Contract Management Across the Organization

AI has already proven to significantly impact many aspects of business operations. We're seeing AI-powered contracting accelerate deal cycles, surface trends and lead to data-driven insights, and save time and money through automation. It is also helping optimize team workloads, improve accuracy and consistency, and free up legal teams to focus on the higher-value strategic work that they went to law school for.

“Our goal is to keep legal out of 95% of contracts – AI-driven workflows, permission controls, and analytics help us get there.”



Catherine Choe  Everlaw



automatically tagged/indexed over 4k contracts



cut contract processing time by 50%



reviews contracts 80% faster



Ironclad AI automatically reviews these contracts, flags language and clauses that don't work for us, and suggests L'Oréal-approved provisions to swap in. This cuts the review process from hours to minutes, improves our team's efficiency, and frees up time for the team to focus on more high impact work."



Charles Hurr L'ORÉAL

Team-specific Benefits

Clearly, we believe that AI is a game changer for the legal field (explained below). But what about the rest of your organization? Let’s explore what it’ll do for you, and the benefits your various stakeholders could see within the context of contract lifecycle management.

LEGAL

1. Time Savings

Automation of routine tasks allows focus on complex negotiations and strategic advisory roles. AI will handle:

- Initial contract drafting
- Redline suggestions based on predefined guidelines and playbooks
- Intelligent ingestion of legacy documents and their metadata

EXAMPLE

A lawyer spends 30% less time on routine NDAs, allowing them to contribute more to high-stakes M&A deals.

2. Reduced Risk

AI-powered analysis can catch potential issues human reviewers might miss. AI will handle:

- Automated bulk review of documents, flagging anomalies overlooked in manual review
- Standardization of templates that ensure compliance with company guidelines

EXAMPLE

AI flags an unusual indemnification clause in a vendor contract, preventing potential liability.

3. Data-Driven Insights

Access to contract analytics can shape negotiation strategies, uncover bottlenecks, and prove legal’s business value. AI will:

- Detect contract metadata in ingested documents for trend analyses
- Analyze contract metadata to surface insights into risks, costs, abnormalities, and more

EXAMPLE

Analytics reveal that certain clauses consistently lead to faster deal closure, informing future contract drafting.

4. Improved Work-Life Balance

AI helps legal teams manage their workloads more effectively, reducing burnout and making them better business partners across the org. AI will:

- Automate review of low risk, high volume contracts
- Enable faster drafting, reviewing, metadata gathering, and risk management

EXAMPLE

A legal team reduces average weekly overtime from 10 hours to 3 hours after implementing AI-powered contract management.

SALES

1. Faster Deal Closure via Self-Service

AI-powered systems can generate and process contracts faster, reducing time from proposal to signed agreement. AI will:

- Automate contract review on low risk, high volume contracts and remove legal as a bottleneck
- Pull in relevant, pre-approved contract language to ensure compliance
- Provide visibility into contract status and details instantly

EXAMPLE

Average contract cycle time reduces from 4 weeks to 1 week, allowing sales to close deals 75% faster.

2. Predictive Insights

By analyzing historical data, AI models can provide insights into deal likelihood, forecast accuracy, and process bottlenecks. AI will:

- Predict issues in deals based on the analysis of risky clause language
- Provide more accurate forecasting based on deal size, timing, and redline analysis
- Flag stages in the contract lifecycle that might slow the deal down

EXAMPLE

AI flags that including a particular clause in a deal yields a 30% lower close rate on average, allowing sales to proactively address the issue.

IT

1. Data Visibility

AI-powered CLMs offer easier and more accurate contract data visibility, which helps monitor security and spend. AI will:

- Analyze large volumes of contract data to find patterns in software usage, spend, and data access
- Use NLP to detect renegotiation or termination clauses in contracts
- Evaluate the success and compliance of different service providers based on past contract data

EXAMPLE

You can ask a legal AI chatbot to pull the top five most and least expensive software contracts and map out corresponding adoption to savings opportunities

2. Risk Mitigation

AI-powered CLMs help IT teams reduce various technological and operational risks. AI can:

- Provide digestible audit logs for fast security and compliance reviews
- Monitor contract performance and notify teams of unmet obligations, mitigating legal and financial consequences

EXAMPLE

The AI system automatically flags potential data privacy issues in contracts, helping the IT team proactively address GDPR compliance risks before they become problems.

PROCUREMENT

1. Workflow Efficiencies

AI-powered CLMs streamline procurement processes, addressing common operational challenges. AI will handle:

- Initial drafting and redlining of procurement contracts
- Sequential, customized approver routing using standardized guidelines

EXAMPLE

Leveraging an AI-powered playbook, the CLM automatically routes contracts to the right approvers based on contract value and type, reducing approval times and eliminating lost contracts.

2. Supplier Management

AI-powered CLM can help track supplier performance against contract terms. AI will:

- Monitor deliverables, SLAs, and other performance metrics and generate reports automatically
- Flag upcoming renewals and unmet obligations via contract data analysis

EXAMPLE

The system flags a supplier consistently missing delivery deadlines, allowing procurement to address the issue proactively.

3. Cost Savings

Better contract analysis can identify opportunities for consolidation or renegotiation. AI will:

- Analyze total contract value across contract types and other segments, identifying overlaps or opportunities for consolidation or renegotiation
- Consolidate spend data from multiple sources to provide a unified view across documents

EXAMPLE

AI analysis reveals \$500,000 in potential annual savings by consolidating office supply contracts across departments at a global organization

4. Risk Management and Compliance

AI can flag high-risk suppliers or contracts for closer monitoring and ensure compliance with requirements based on pre-set risk factors. AI will:

- Prioritize which suppliers or contracts need more attention.
- Verify that contracts include necessary clauses and adhere to regulatory requirements.

EXAMPLE

The system identifies a supplier with recent negative press, prompting a review of the relationship, while simultaneously ensuring all new contracts meet GDPR requirements.

CHAPTER FOUR

AI & Security: Lifting the Hood

Concerns about AI and security aren't unfounded, but they shouldn't deter you from using the technology. Just as you would with any other technology, you should take time to understand the basic issues surrounding AI and security, and to learn — and follow! — best practices for using artificial intelligence without putting your company or clients at risk.

When we talk about the risks of generative AI, we're talking mainly about two things: 1) what the system outputs (and why), and 2) who can access all of the data both going into and coming out of the AI model, including what the model was trained on. Let's break that down into some key areas to be aware of as you think about using AI securely in your own workplace:

Privacy

AI systems often require large amounts of data to function effectively. Questions of where the data training data comes from — and what happens to data analyzed or otherwise passed into these systems during use — raises concerns. Will sensitive client information be protected if I feed it to AI for analysis? Will my company's intellectual property be used in training and recommendations for competitors? How can we best leverage AI while maintaining compliance with data protection regulations like GDPR (European Union regulation on information privacy)? In the case of rapidly evolving AI capabilities, consideration of answers to these questions are often more important than the product performance itself.



Security

AI systems can be targets for hackers and cybercriminals, especially if they govern critical infrastructure or power grids. Security breaches could lead to stolen data, but also to compromised AI systems that produce altered or otherwise unreliable output. A malicious actor could poison a system's training data, causing the model to learn incorrect behaviors. They could also wage adversarial attacks designed to skew input data to manipulate the AI's output without being detected.

Look out for integration challenges: Poorly implemented AI tools that don't integrate well with existing systems can create security vulnerabilities or lead to data inconsistencies.



Bias and Fairness

AI systems can inadvertently perpetuate or amplify biases present in their training data, potentially leading to unfair or discriminatory outcomes. The use of AI in legal decision-making, specifically, raises ethical questions about accountability, fairness, and the role of human judgment in the legal process.

Evaluate algorithm design: The way AI algorithms are designed can introduce biases or errors. For example, if certain factors are given too much weight in a decision-making algorithm, it could lead to skewed results.



Lack of Transparency

Many AI systems, particularly deep learning models, operate as "black boxes," making it difficult to understand how they arrive at their conclusions. This can be problematic in legal contexts where explainability is crucial.

Look for Model Cards that show their work: AI model cards provide details about how a model was developed, including architecture and training data. Seeing what kind of data was used to train the model is key to understanding if the output of the model will be biased. Bonus points for solutions that clearly lay out how a model came to a specific conclusion.



Accuracy and Reliability

AI models are only as good as the data they're trained on. Insufficient, poor quality, or biased training data can lead to inaccurate or unfair outcomes.

Beware of hallucinations: hallucinations have been a scourge of generative AI since day one, and while they're rapidly improving, should be taken seriously. Make sure to double check any and all outputs from Gen AI tools against trusted sources.



How to think about mitigating those risks

Mitigating AI risks requires a multifaceted approach. Legal teams might not be on the front lines for the bulk of the work, they should know how to work with the right partners — Information Technology (IT) chief among them.

Here are the eleven things to consider with your colleagues in IT when making your own AI risk plan:

1

Data Governance: Implement strong data governance practices to ensure the quality, security, and ethical use of data used to train and operate AI systems.

2

Regulatory Compliance: Stay informed about and comply with relevant regulations and developing frameworks for AI governance, like ISO 42001 and NIST. Know how the regulations differ in various jurisdictions around the world.

3

Algorithmic Auditing: Regularly audit AI algorithms for bias and fairness, and make necessary adjustments to ensure equitable outcomes.

4

Audit Logs: Make sure you have access to application audit logs inside of your own network (eg, You don't have to request them from a third party). Audit logs let you monitor in real time, spotting threats or risks before they become full-on problems.

5

Explainable AI: Where possible, use AI models that provide explanations for their decisions, or develop supplementary systems to interpret complex models.

6

Human Oversight: Have humans review everything your AI systems create, before you use it or send it to someone else. Maintain human oversight of AI systems, especially for critical decisions. AI should augment human decision-making, not replace it entirely.

7

Continuous Monitoring: Legal teams, ask your friends in IT to implement systems that continuously monitor AI performance and outputs for anomalies or unexpected behaviors.

8

Ethics Guidelines: Develop and adhere to clear ethical guidelines for AI use in legal contexts. This begins with understanding and taking into account the algorithmic bias in any AI system used to influence decisions within legal work. Your guidelines should also account for accuracy, privacy, and human oversight of work done by AI.

9

Training and Education: Ensure that all users of AI systems are properly trained in their capabilities, limitations, and potential risks. Develop policies for use of AI systems in the workplace that address data and system security (see section below).

10

Security Measures: Implement robust cybersecurity measures to protect AI systems and the data they use from unauthorized access or manipulation. Encrypt everything, and [make sure you hold the keys](#). From complying with regulations like GDPR and HIPAA, to ensuring the integrity of your data, encryption is more vital than ever.

11

Diverse Development Teams: Encourage diversity in AI development teams to help identify and mitigate potential biases.

How to be responsible (and CYA) when using AI

When using AI in legal work, it's crucial to approach it with the right mindset and take appropriate precautions to ensure you're using it responsibly. In other words, always make sure to CYA – Cover Your... well, you know. To that end:

Keep a human in the loop

AI makes mistakes, so it's essential to adopt a "trust but verify" approach. Have human experts review all AI-generated content before anyone uses it. Treat AI as a tool to augment your work, not replace your judgment.

Treat AI output as your own

You are responsible for any work product that incorporates AI-generated content. Not the AI company or the app maker who built AI into their app — you. Review and validate all AI outputs before using them in official documents or communications.

Stay on top of your data

It's critical to know what solutions—and the third parties they work with—do with your data. Will they train models with it? Where will it be stored? Look for places to enable zero data retention (ZDR) when appropriate and have a plan for deleting your data if and when third party relationships end. Some companies, like Anthropic, operate on ZDR by default.

AI is a tool, not a teammate

Approach AI with the same level of skepticism you would use when searching on Google. AI won't do everything for you – it's only as good as the prompts you provide it with. Remember the old computer programming adage, "garbage in, garbage out." It applies to prompting AI models, too.

Craft your prompts carefully

Be specific and clear in your instructions to AI. And remember that the more you use an AI model, the better it gets at responding with the kind of output you're looking for: Getting AI to give you the results you want is usually an iterative process. Follow up with additional prompts as needed to refine and improve the AI's output.

Set an internal policy for AI usage

Establishing clear guidelines for how AI should be used within your organization should be a top priority for every company.

What should your internal policy for AI usage include?

You can see how Ironclad approaches this (and download our policy) [here](#), but a few basic elements you should consider for your team:

Confidentiality

Classify data according to your data classification matrix, clearly outlining which type of data is appropriate for which use. Data classified as confidential or higher should not be fed into AI prompts by default, and company intellectual property should always be protected.

Responsible Use

Employees should be held responsible for AI-generated content as if it were their own. They must carefully review content for accuracy and potential negative effects on third parties.

Service-specific Review

Not all AI services are created equal. Evaluate each service based on its approach to data processing, compliance, and legal terms. (Find more in Part 5, “How to evaluate AI software”).

Process documentation


Keep records of how AI was used in your work, including the prompts used and any post-processing or verification steps taken.

Staying informed

Keep up-to-date with the latest developments in AI technology and relevant regulations or ethical guidelines in the legal industry.

Transparency

When appropriate, disclose the use of AI in your work to clients or colleagues. This can help manage expectations and build trust.



Read More

[How We Regulate AI Use at Ironclad](#)

CHAPTER FIVE

How to evaluate AI software

When evaluating AI software for legal applications, it's crucial to ensure the solution meets your needs, and also that it complies with legal and ethical standards. Here's a look at a few key considerations:

What to look for from a data security perspective

1. Data Encryption

- Ensure the software uses strong encryption (e.g., AES-256) for data in transit and at rest.
- Look for end-to-end encryption for sensitive communications and personally identifiable information (PII).

2. Access Controls

- Seek granular access controls that allow role-based permissions.
- Check for multi-factor authentication options to enhance security.

3. Data Residency

- Understand where your data will be stored and processed.
- Ensure compliance with relevant regulations like GDPR or CCPA.
- Consider options for data localization if required by your jurisdiction.

4. Audit Trails

- The system should maintain detailed logs of all data access and changes.
- Look for immutable audit logs that cannot be altered.
- Ensure the ability to generate comprehensive audit reports.

5. Data Training, Retention, and Deletion

- Check for configurable data retention policies.
- Understand if and how the software will train on your data.
- Ensure the software supports secure data deletion methods (e.g., data wiping).
- Look for features that allow for selective data deletion to comply with "right to be forgotten" requests.

6. Third-Party Audits

- Check which third parties the software uses, and what their AI policies around retention and security are.
- Prioritize software that undergoes regular third-party security audits.
- Ask for recent audit reports and check how quickly past issues were resolved.

7. Compliance Certifications

- Look for certifications like ISO 27001, SOC 2, or GDPR compliance.
- Check if the vendor maintains a compliance program with regular assessments.

8. Data flow diagram:

- Ensure the DFD clearly depicts data sources, processes, data flows, and storage, so you can easily trace and understand how your data flows through their product.
- Look for adherence to relevant regulations (e.g., GDPR) and other data protection measures during transfer and storage.
- Evaluate how the solution integrates with existing systems – and whether it takes into account future growth in data volume and complexity.

9. Sub-processor usage:

- Most industries are subject to strict data protection laws (like GDPR or HIPAA), but investigate the sub-processors' data handling practices, encryption standards, and breach notification procedures.
- Establish a clear communication channel to receive updates on any changes in sub-processor arrangements.



When you're ready to take the next steps in your evaluation, be sure to check out the [LLM Top 10](#) from The Open Worldwide Application Security Project (OWASP), a comprehensive risk/threat model that is open-sourced and available for free.

What to look for from a policy perspective

When evaluating AI software for legal use, look closely at the vendor's policies and practices. These factors ensure ethical use, transparency, and reliability of the AI system in your legal operations.

1. Ethical AI Guidelines

- The vendor should have clear, publicly available ethical guidelines for AI development and use.
- Look for alignment with established AI ethics frameworks (e.g., IEEE Ethically Aligned Design).

2. Transparency

- Seek vendors who provide detailed information about their AI models' training data and methodologies.
- Check if they offer [model cards](#) or datasheets describing AI system characteristics.

3. Bias Mitigation

- The vendor should have robust processes to detect and mitigate bias in their AI models.
- Look for regular bias audits and diverse teams involved in AI development.

4. Human Oversight

- Understand how human expertise is incorporated into the AI system's decision-making process.
- Check for clear escalation pathways for AI-flagged issues.

5. Continuous Monitoring

- The vendor should have systems to continuously monitor AI performance and address issues.
- Look for proactive alerting mechanisms for performance degradation or unexpected outputs.

6. Update and Maintenance Policies

- Understand the frequency and process of AI model updates.
- Check how these updates are tested and validated before deployment.
- Ensure there's a rollback mechanism in case of problematic updates.

7. Data Usage Policies

- Get clear information on how your data will be used, especially regarding AI model training.
- Ensure the vendor offers options to opt out of data sharing for model improvement.

Common areas where the software fails

Hallucinations

AI models sometimes generate plausible-sounding but incorrect information. Often, this occurs when the model is given inputs outside of its training data. LLMs are given to generating very confident-sounding text, even if it's all based on inaccuracies or outright hallucinations.

EXAMPLE

An AI might invent non-existent legal precedents that sound convincing.

Context misinterpretation

AI systems can misunderstand the context of legal language, leading to errors. Models often rely on statistical patterns in the data, and not a true understanding of the content and context, which can cause misinterpretation.

EXAMPLE

Misinterpreting a clause's intent due to unusual phrasing or structure.

Inconsistency in responses

LLMs often give different answers to the same question asked in slightly different ways, stemming from the probabilistic nature of language model outputs. When you ask an LLM a question, the answer is essentially the model predicting what a good human response would sound like.

EXAMPLE

Providing conflicting interpretations of a contract clause when asked multiple times.

Bias amplification

AI models can amplify biases present in their training data. This can lead to unfair or discriminatory outcomes in legal analysis.

EXAMPLE

Consistently predicting higher risk scores for certain demographic groups in bail decisions.

Temporal confusion

AI models often struggle with understanding time-dependent information because they're trained on static datasets and don't have a true sense of time. A model's "cutoff date" refers to the date through which the model's training data runs.

EXAMPLE

Applying outdated laws or regulations that have since been amended or repealed.

Lack of common sense reasoning

Since AI models rely on pattern matching, and not true human-like understanding of information, AI can fail at tasks that require basic common sense. This can lead to absurd conclusions.

EXAMPLE

The "[how many Rs in 'Strawberry'](#)" issue arises because the model focuses on the literal question rather than understanding the concept of spelling.

Difficulty with novel scenarios

AI often struggles when faced with unique or unprecedented legal situations, like determining copyright ownership for artwork created by another AI. Similar to the cutoff date limitation, a model can only draw from its training data, which may not cover every possible scenario.

EXAMPLE

Failing to properly analyze legal implications of new technologies not present in training data.

Lack of causal understanding

Cause-and-effect relationships in legal contexts can be difficult for AI models to process correctly. This is because they're trained on correlations in data, not true causal relationships.

EXAMPLE

Misattributing the cause of a legal outcome by focusing on irrelevant but correlated factors.

Understanding these failure modes is crucial for legal professionals using AI. It underscores the importance of human oversight, cross-verification, and treating AI outputs as assistive tools rather than definitive answers.

CHAPTER SIX

Tips on Getting Started

The key to embarking on the AI journey in your legal department is to start small. Focus on high-impact areas, then gradually expand your AI implementation as you gain confidence and experience.

The key to embarking on the AI journey in your legal department is to start small. Focus on high-impact areas, then gradually expand your AI implementation as you gain confidence and experience. And don't forget—you don't need to go at it alone! All of our peers are in experimental mode right now. The more educated we get on AI and new technologies, the more opportunity we have as a field to secure our seats at the proverbial table.

To begin, consider focusing on four key areas where AI can make an immediate impact:


- High volume, low risk work, like employment contracts and NDAs
- Contract review to flag issues like risky clauses and non-standard terms
- Legal research and e-discovery for finding relevant cases and statutes
- Drafting of briefs, contracts, pleadings, and other legal documents



Some tips on prompting AI systems

As you begin to implement AI tools, learning how to effectively prompt a system becomes crucial. Approach prompting as an iterative process, expecting to refine your queries based on the AI's output. And always remember that while AI is a powerful tool, it can hallucinate and it doesn't replace legal judgment. Use your own expertise when reviewing any AI output.

When it comes to crafting prompts:



Read More

[5 ChatGPT Prompts to Boost Your Legal Ops Game](#)

Be specific about what you're looking for and provide necessary context.

Want your information presented in a specific format? Make sure to specify this in your prompt. For example, "Provide a bullet-point summary of the key risks in this contract."

Speaking of how you like your information presented, telling a chatbot to emulate a specific voice or style in its responses can be really helpful. Ex: "Summarize the following, emulating the style of *Stanford Law Review*."

Take a "trust and verify" approach: always check the output for accuracy.

Don't shy away from using precise legal terminology, as AI models trained on legal texts understand legal jargon.

Break more complex queries down into smaller, more manageable steps presented in a straightforward, logical manner.



Measure impact as you go

As you implement AI tools, it's crucial to measure their impact. Look for software that tracks and readily surfaces metrics such as time savings, accuracy improvements, user adoption rates, client satisfaction, and cost savings. This data will help you refine your AI strategy and justify further investments in AI technology.



Read More

[How to Get Started with Generative AI Within Your Legal Department](#)



Read More

[Legal Metrics Masterclass: The Numbers Every GC Should Track](#)

CHAPTER SEVEN

AI Innovation in the Legal Field

Arguably one of the most exciting things about any new technology is seeing it in action, and learning from those on the frontlines using it to drive real, tangible impact. Ironclad users have surprised us with the incredible ways they're putting our AI to work, and the more we spoke to other legal AI pioneers it became obvious that these were not isolated incidents.

So how is AI being used right now in the field of legal? Below are two use cases of AI in the wild, being used to solve difficult problems, make themselves 10x more efficient, and above all else—innovate.



AI helps find needles in the haystack during e-discovery after Lahaina fires



THE PROBLEM

Following the devastating fires in Lahaina in 2023, there was a significant challenge in sifting through vast amounts of 911 call transcripts to determine liability. On top of the physical and emotional devastation, the sheer volume of data made it difficult for legal teams to extract relevant information efficiently, hindering the preparation for litigation.

THE PROJECT

To address this issue, Mr. McCullough utilized AI technology from Everlaw—a cloud-based e-discovery and document review software company—to analyze 911 calls during e-discovery. By uploading the transcripts into the platform and employing the Description Summary function, he was able to generate comprehensive deposition summaries within minutes—a task that, before AI, would have taken weeks.

THE RESULTS

Efficiency of the review process improved significantly, giving Mr. McCullough access to a bird's eye view of the evidence faster than traditional methods. While the output of AI-generated summaries required review, this not only saved time but also enabled his clients to grasp critical information more effectively, enhancing their understanding of the case. Not only this, but his expertise in fire litigation and commitment to leveraging technology for better outcomes were crucial in championing this innovation.

“Over the last 10 years, the number of fires caused by utilities have grown quite dramatically, and now we’re seeing entire communities wiped out. Trying to help those communities rebuild, and get their justice, is what fire litigation is all about.”



Greg McCullough
Independent Fire litigation consultant

Gunderson Dettmer Leads by Example with ChatGD

THE PEOPLE



Gunderson Dettmer’s legal engineering team



Joe Green, *Chief Innovation Officer*



Naveen Pai, *Chief Knowledge Officer*



Lori Knowles, *Assistant General Counsel*



Laura Chao, *Practice Innovation Attorney*



Stephanie Goutos, *Lead Practice Innovation Attorney*



Avi Saiger, *Practice Innovation Attorney*

THE PROBLEM

While the legal industry has traditionally been slow to adopt new technologies, Gunderson Dettmer has consistently positioned itself at the forefront of innovation. As an early advocate for responsible AI use, the firm quickly recognized the rising demand from clients seeking guidance on AI’s regulatory landscape, legal implications, and best practices. By being one of the first law firms to strategically embrace and implement AI within its own operations, it has allowed the firm to guide clients through this rapidly-evolving space with authenticity and expertise.

THE PROJECT

The firm has deployed third-party AI tools for years—including some powered by LLMs—for a variety of use cases. With ChatGD, the firm created its first proprietary tool to make generative AI technology widely available internally, achieving a new milestone in Gunderson’s ability to marry its in-house legal expertise with cutting-edge engineering and technology.

ChatGD allows the firm’s attorneys to query and manipulate documents using a secure, enterprise instance of OpenAI’s LLMs through Microsoft Azure. Attorneys using ChatGD can leverage the underlying LLMs and the power of ChatGPT to accelerate and enhance their work as subject-matter experts. It also gives attorneys the ability to upload proprietary legal agreements and other relevant source material as context for processing queries using retrieval-augmented generation (RAG).

THE RESULT

In the year following its August 2023 launch, ChatGD has been used by more than half of the firm's attorneys and business leaders, many of whom use it regularly. They have created tens of thousands of conversation threads and messages in ChatGD. The firm's use of the tool has led to significant workflow enhancements, including increased accuracy and efficiency in delivering work product. Users have especially gravitated toward using ChatGD for “iterative editing” and “text manipulation” use cases, where users provide pre-drafted language and then quickly iterate through different versions by asking the LLM to offer text refinements.

“Gunderson Dettmer has a responsibility to our clients to drive innovation from the inside out. By rolling out this homegrown tool internally—and actually using it—our lawyers are not only working more efficiently, we’re leading by example. We’re working to get our field comfortable with AI, and a big part of that is being right on the cutting edge and helping set a precedent for how this technology can be applied and used in a real-world setting.”



Joe Green
Chief Innovation Officer





Ironclad is an all-in-one AI-powered contract lifecycle management (CLM) platform that integrates deeply with security, sales, and compliance tools to mitigate risk, optimize the contract management process, and drive business growth.

Our AI suite of tools, including Smart Import for automatically ingesting and tagging contracts, guidelines and AI Assist™ for automatic redlining, has saved our customers an estimated cumulative [29 years of effort](#) across contract uploading, review, and redlining.

We build security and compliance into everything we do, from integrations with OneTrust, to CIS benchmarks and security controls, and cloud security best practices such as the NIST Cybersecurity Framework.

To learn more about Ironclad and how it can help your team work smarter and faster with AI, [request a demo](#) today.