



AI-Driven Contract Analysis in Perspective and in Practice



Preface

This e-book arose out of the development of a new capability for Kira software, Answers & Insights. At the launch of the new offering, we felt this was a significant development in the history of automated contract analysis, and in the development of contract-related technologies more generally. To communicate the significance of Answers & Insights in that historical perspective, we published a series of weekly articles leading up to the launch. The articles show how quickly contract technologies have advanced in just a few short years, and they also give a hint of the direction we think the future will take. Those articles are collected and republished here in parts I-IV of this e-book.

Also in conjunction with the launch, we published a research study on a timely topic in debtor-creditor relationships in this time of economic crisis: the use of so-called “trap doors” in negative covenants that some borrowers have used to shield assets from creditors. The research for this study entailed using Kira’s machine learning based software to analyze 156 publicly-available credit agreements, and as part of the research we applied Kira’s new Answers & Insight’s capabilities. We included a link to the study at the end in order to illustrate in practical terms the importance and value of this new capability, and to provide a tangible example of the state of the art in contract technology today.

To see a video about Answers & Insights or to request a demo, [visit our website](#).



PART I

The Emergence of Contract Automation and Technology



Using technology to draft, store, review, and analyze contracts has been in a constant state of development and improvement for a couple of decades, and the pace has significantly quickened in the last 5-7 years.

Tools for drafting contracts have come a long way since lawyers finally set aside their quill pens and inkwells. For a long time the typewriter was the predominant technology. The first real digital technology breakthroughs were the introduction of personal computers and the development of the WordStar word processing program in the 1970s. With word processors, contract drafters were able to do more than just draft contract terms; the software automated things like the numbering of sections and insertion of cross references to those sections.

By the 90s, “document assembly” systems such as HotDocs and Contract Express further facilitated contract generation. These tools use guided forms and questionnaires to gather data on requirements and then generate contracts. More tools, such as Deal Proof (now part of Thomson Reuters), and Eagle Eye (now part of Litera), further facilitated the drafting process by automating the cross-checking of defined terms and section references. These were rule-based systems and were probably among the first widespread applications of early Artificial Intelligence to contract drafting. As with most of the contract technology to that point, the focus was on the front-end process



of contract drafting rather than extraction of value and insight in the post-drafting stages of contract lifecycle management.

Finalized executed contracts were mostly still stored in paper, tucked into filing cabinets or archived away in a basement file room. They were not, for the most part, accessible digitally except perhaps in scanned formats that were not easily searchable. Contracts were not typically seen as collections of data; each contract was independent as a document and there were limited opportunities to extract insights from larger sets of contracts.

Yet there was clearly enormous value in those previously executed agreements. Tools for searching contracts existed by this point; in the 90s law firms and corporate legal departments had begun implementing document management systems, such as OpenText and iManage, that housed these contracts in their original digital formats. But these systems were primarily for version control and record keeping purposes. Searching was based on keywords. Contracts often use unpredictable and inconsistent language to express the same concepts. So even though DMS and other search engines now enabled lawyers to search through large bodies of contracts, finding clauses and terms that matched keyword searches would inherently be hit-or-miss.

The value locked in those signed contracts is enormous. Within a business, events will arise during which knowing what is in one's contracts becomes mission-critical. Regulatory changes such as GDPR or accounting standards are two examples. And there is perhaps no moment more critical than a merger or acquisition. Much of the value of any company is represented in its contracts: customers' payment obligations, the consistency and stability of its suppliers, limits on liability, the relationships with employees and contractors. Mergers and acquisitions lawyers have the unenviable task of needing to quickly review and verify that all the legal agreements are in place, in order to confirm that the company being acquired is worth what the amount to be paid for it.

That M&A due diligence process was almost completely manual, and consisted of young associates reviewing boxes full of printed contracts, often piled high in physical deal rooms. In the mid-90s, the internet enabled the birth of "virtual data rooms" like Intralinks and Merrill, which allowed scanned copies of these files to be shared securely online. Nonetheless most of the actual contract review work remained fully manual, with young lawyers given the task of reading through hundreds or thousands of contracts hunting for specific terms and clauses to identify risks. It was repetitive and boring for the lawyers, and extremely expensive for the client.



Reporting on the findings was also frustratingly manual and incomplete. In many cases the volume of documents was too large to review everything, and the parties would agree that the risk assessment would be made only on contracts over some dollar threshold or that met some subjective guesses as to materiality. This inevitably led to some risks being missed or underestimated. Compiling the results of such a limited review was still a herculean task. For weeks, bleary-eyed associates would pore though page by page, retyping their findings into Microsoft Word charts and tables that would allow them to start to form a risk assessment.

This all culminated in a “diligence memo” which spelled out a set of findings and recommendations to the acquirer in narrative form. While spreadsheets were available around the same time as word processing tools, and might have provided a better tool for data manipulation, lawyers tended to stick with the technology they had already begun to adopt: word processing programs and their tables function. This technology didn’t really automate anything about the process, but it provided the ability to organize, sort, and manipulate data that was extracted and collected manually.

It soon became clear, however, that technology could be used to capture the output of that review—to turn the unstructured text in contracts into rows and columns of data.

By the time a critical mass of lawyers were up and running with some basic productivity tools including word processing and search, the legal industry had some of the building blocks for a more accurate and efficient way to create, store, and analyze contracts. Contract data was digitized and there were better tools for search and analysis. But the entire contract lifecycle was still mostly manual and relied heavily on human creation and analysis.

The next leaps forward in contract analysis had yet to come. That’s the subject of the next part of this series.



PART II

The Breakthrough: Machine Learning Addresses Scalability & Accuracy in Contract Analysis



For about 20 years, from the 90s through the early 2010s, the predominant contract automation technologies were based on keywords, rules and expert systems. These systems all faced problems of scalability. Applying rules to extract meaning from something as complex as the unstructured text found in contracts proved to be brittle and overly complex, requiring heavy human curation and intervention to adjust for changes in underlying law and simply human language variations.

This profusion of human curation and intervention cut across the entire contract lifecycle:

- To speed contract drafting, lawyers created ever-growing libraries of standard forms, complete with manually-curated clause banks and “playbooks” to define the universe of acceptable terms. Only occasionally (and with great pain) were these built into document assembly systems.
- Lawyers might save contracts in deal banks (sometimes in a document management system, sometimes in a separate system). These often required the lawyer to fill out a questionnaire/cover sheet to restate important deep points in a standardized form, to make them findable later.
- In larger corporations, some large businesses stored sales contracts in CRM (Customer Relationship Management) systems like Siebel and Salesforce, vendor contracts in procurement systems like Ariba or Oracle, and occasionally specialized Contract Management systems like Novatus or Selectica. To make all of these searchable it was



commonplace to employ an army of contractors to manually enter data into databases, essentially retyping legalese into deal points that could be searched.

- Specialized online research systems started to make a business of “mining” public databases for contract clauses from public filings. For example, GSI Livedgar employed analysts who tagged all transactions in the US SEC’s EDGAR database with a taxonomy of Deal Points to facilitate searching. Deeper searching for clause language was created by having yet more expert analysts build complex pre-crafted search queries.

All of these approaches relied in part on the semi-standardized nature of contracts - clauses, headers, defined terms, schedules. But it was a fragile system and would fall apart when a contract didn’t fit the expected mold, such as when a lawyer invented a new deal structure that hadn’t been seen before, or when laws and regulations changed and necessitated new queries. It was not inherently scalable.

Lessons from eDiscovery

It is peculiar that manual rules and keywords in contract searching persisted for more than two decades because, in the meantime, machine learning was rapidly catching on in the litigation side of the legal industry. Technology-Assisted Review (TAR) based on machine learning techniques proved to be a scalable solution for analyzing the huge stores of email and other documents that are relevant to a large-scale lawsuit. In the US, a phase of litigation called “discovery” requires parties to hand over relevant documents in response to the opposing party’s request for evidence to support a case. At some point after email became prevalent in the late 1990s and early 2000s, that necessitated extensive “electronic discovery” – at which point the volume of “responsive” documents exploded; manual review for responsiveness became impossible, and TAR came to the rescue.

A variety of techniques could be applied to the eDiscovery process, as a way of predicting which documents in a large set of emails, memos, and many other types of documents were responsive to an eDiscovery request. Early systems still frequently relied on rules-based systems and complex queries, but as time went on more and more employed non-rules-based approaches to model document concepts (such as latent semantic indexing and eventually supervised machine learning). Modern TAR software tends to leverage a variety of techniques that allow the machine learning algorithm to direct human review effort to the most impactful documents (using some court-approved or agreed definitional criteria).



The big leap in contract technology finally came in the early 2010s as a small crop of companies like Kira (which was then named “DiligenceEngine”) began to see that the unstructured nature of contracts provided an excellent use case for the increasingly sophisticated machine learning technologies that were being leveraged in eDiscovery. These companies set out to test whether the techniques that were used to identify relevant documents in a litigation eDiscovery process could be adapted to extract valuable data and insights from large collections of contracts.

Adapting Machine Learning to Contract Data

Success was not immediate. For one thing, it turned out that the problems that needed to be solved for contracts had significant differences from the problems to be solved in eDiscovery. First, with contracts a user typically needs to pinpoint specific clause language, while eDiscovery users look for relevance only on the document level. Second, with contracts users are often looking for the same thing from deal to deal, while in eDiscovery the unique facts of the case can vary so much that a new machine learning model needs to be crafted every time. Third, in many of the most critical contract review scenarios - such as a divestiture or a regulatory compliance review - pretty much every document being reviewed actually needs to be reviewed. In eDiscovery, techniques like sampling and having software guide reviews are more common.

Specialized contract analysis software pioneers such as Kira, which was commercially launched in 2013, addressed all of these differences. They were built from the beginning to pinpoint clause language, not just whole documents. They eliminated the need (common in eDiscovery systems) to train a new model every time. This was accomplished by including a set of models (pre-trained by our lawyer-experts) for all of the most important clauses: termination, assignment, change of control, governing law, etc. And they came with workflow tools to make sure the reviewer could see and confirm every page of every document - if needed - to ensure accurate and complete review.



The Scalability Breakthrough

These capabilities made contract analysis software very useful for applications like M&A due diligence. But it still wasn't scalable to the other myriad contract review needs such as general contract management, contract audit, real estate etc., not to mention the need to handle non-English languages and to meet the needs of all the precise variants and specialized expertise of individual teams of lawyers and other professionals. The software needed one more thing – the ability for users to train their own models without any technical machine learning knowledge from the user's perspective.

The turning point came in 2015, when Kira announced a capability that allows that customization, called Quick Study, simply by highlighting examples on screen. The big leap here was that Quick Study allows users to customize their use of machine learning for their specific task, without needing to become machine learning experts themselves. With Quick Study (which may have been the first ever application of what is now increasingly referred to as "AutoML" to documents like contracts) the system was finally scalable. Any subject matter expert could employ the machine learning system to apply to their area of expertise - no technical services required.

At this stage of development, machine learning had driven contract technology forward in just a few short years, from a word processing-centric focus on simple drafting tools and document automation, to a much more versatile downstream application that treated contracts as a valuable source of data. Machine learning enabled the extraction of that value and set the stage for both lawyers and other business functions to fully leverage the data in contracts in other business processes.

Even as the core analytic capabilities of Kira have improved in this way, we've also improved the ways that Kira integrates with our customers' other systems, through the development of strong industry partnerships (such as Intralinks for due diligence, and Apttus/Conga for contract management), and by offering APIs, which allow custom integrations with myriad other tools. Data from contracts can be more widely distributed in enterprise systems, where it can be leveraged by various functions including business intelligence, strategy, purchasing, marketing, and many others.

The application of machine learning to contracts was a long time coming and a giant leap forward. It had become the gold standard for high-quality, efficient extraction of insight from contract data. But those technical enhancements are not static. Each application of AI to a legal problem is built on the shoulders of domain experts. That interplay between technical expertise and legal domain expertise is the subject of the next article in this series.



PART III

The Role of Domain Experts in the Success of Machine Learning-Based Contract Analysis



For a few years in the late 2010s, legal publications were awash in hype that, with the rise of AI in law, “robot lawyers” were going to threaten human lawyers’ jobs. Thankfully, that hype has mostly subsided. As professionals have educated themselves on the capabilities as well as the limitations of the available technology, they have gained a growing appreciation for the role of legal domain expertise in the design and application of any legal AI system.

Domain expertise, on the part of both business professionals and software providers, is central to the success of any application of AI to extract value from contract data. Done right, machine learning does not replace the expertise of the professionals that use contract analysis software. The software amplifies that expertise by embedding the professionals’ knowledge in analytical tools, and applying that expertise in large-scale workflows.

Machine learning does not operate in a vacuum. Systems using machine learning are designed by humans.

There are two main approaches to designing an AI system. A common approach is to have a basic system that requires fine tuning, customer-by-customer, by technical experts often referred to as “data scientists.” This fine tuning is typically performed by data scientists contracted to customers, whose customization work is either required as part of an implementation period or offered for a fee as consulting services. Many successful systems are built in this way; companies can buy them from many of the large systems integrators and consulting firms, and also now from some of



the large alternative legal services providers. However, these arrangements are not truly scalable; if requirements change, services are required to implement the changes.

The other, more powerful approach, is when machine learning is leveraged to magnify, amplify, and scale the knowledge of human domain or subject matter experts.

This is the approach that Kira takes. Kira's Quick Study capability enables subject matter experts to create machine learning models just by highlighting examples on screen.

Users have two ways to use the technology:

- We know that many users of an AI system will want to get started quickly, so we include pre-built models created by our own Legal Knowledge Engineering team. Despite the word “engineering” in their title, members of this team do not have technical degrees; these are the experienced lawyers, accountants and other professionals that are responsible for Kira’s Built-In Intelligence. Using Kira’s Quick Study interface, they have built over 1,000 machine learning models (called smart fields) that ship with Kira “out of the box” and allow Kira to extract the most commonly used contract terms. The LKE team listens closely to customers to help select built-in smart fields that will get the most usage across our user base to help power the existing use cases, while constantly identifying and evaluating new ones. Their expertise is the built-in intelligence that allows Kira to go to work, without additional training, on day one.
- The second source of domain expertise in the Kira system is of course the Kira customer base. Their knowledge of their own business needs and interests, their understanding of clients’ businesses or their own, and their sensitivity to clients’ risk tolerances and objectives, will frame and drive the way they apply Kira’s power to a given matter or project. They do not need any special machine learning training, and they do not need implementation consultants. As of this writing, Kira users had well over 15,000 of their own unique smart fields deployed on contract projects and processes, all around the world.

Of course, these two broad approaches can sometimes be combined, with technical customization offered alongside human training to bridge gaps (or perceived gaps) to get to final results. Kira avoids this. Our team of research scientists never does customization for any specific customer. Rather, they constantly evaluate new external technologies and create new proprietary approaches. Unlike the less-scalable approach of using implementation consultants for each customer, all of Kira’s machine learning enhancements are built right into the core Kira product and become available to all Kira customers, often with minimal or no changes to the user experience. Kira’s scientists also follow, and contribute to, the state of the art in machine learning across the world, and bring that evolving expertise directly into Kira products.



Our vision as a company is that AI in this space is actually done correctly when the expertise of real human experts is amplified, empowering them to do more with it. We are giving them “superpowers” to apply that expertise to new tasks and challenges that were not previously addressable. Machine learning lifts the constraints that hindered even the strongest legal experts from solving their contract-related challenges.

Many lawyers remain concerned about the impact of AI on legal careers, but properly leveraging machine learning requires lawyers to engage in the process of contract analysis at the top end of their skill level. The technology is more than capable enough to learn to mimic the judgment of a professional expert (and at a level of accuracy that can surpass expert humans’ ability to agree on what constitutes a specific clause and what does not). Rather than using their time in large-scale manual reviews, they are leveraging higher-level skills by adapting technology to a given problem, augmenting their own professional judgment with better tools. In the meantime, new forms of legal work have emerged for legally-trained professionals to help build the models and delivery platforms that buyers of legal services need.

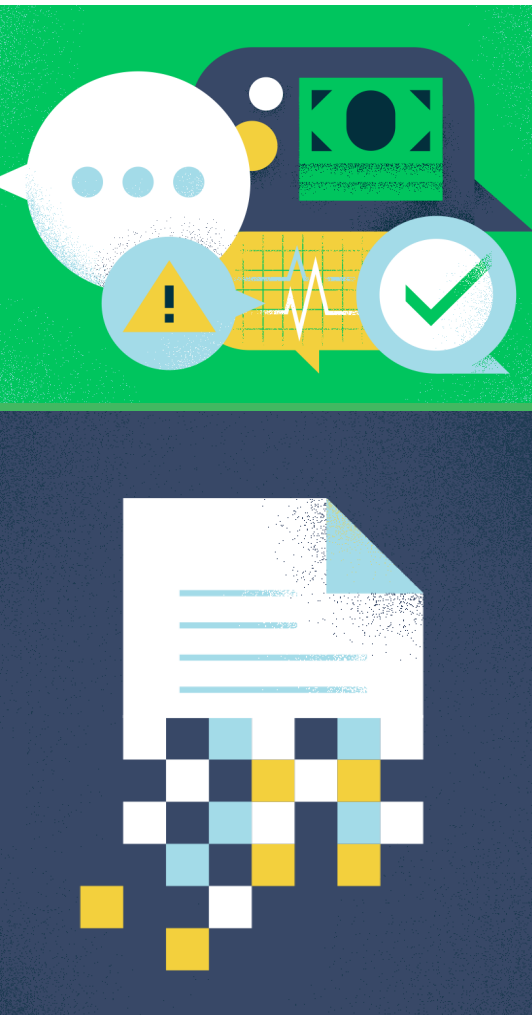
And even as AI has come a long way in releasing the power of contract data in combination with the domain expertise of these participants, there is still plenty of work for human professionals in analyzing the results. Products like Kira can quickly and accurately extract specific types of terms and data from large sets of contracts, but the results still require human analysis to arrive at business recommendations or legal judgments. For example, AI can automate the process of identifying all of the contract terms that include automatic renewal provisions, but it can’t yet tell a client whether they should or should not try to terminate those contracts. And it will likely be a very long time before the technology can do this (and when that day arrives, there may be many debates on whether it should).

In the meantime, the technology will continue to progress in meaningful ways, better supporting the recommendations of human experts. Currently, most of the domain-expert-trainable systems only identify and classify concepts. They do not yet provide structured answers about those concepts. Enabling subject matter experts to directly train automated systems to understand those substantive answers is the next great leap in this story. What if the system could provide structured answers about the concepts that the professionals were seeking to find? What if that could be done without extensive customization and without data scientist skills? That will be the subject of the final part of this series.



PART IV

The Next Frontier in Contract Analysis: Getting to the Answer



In this series we have seen how, over a span of decades, the application of contract automation technology has grown: from contract drafting tools (word processing and document assembly), to episodic contract reviews (whether for regulatory change or M&A), and more recently to the creation of contract databases (contract lifecycle management). As a result, the legal importance of contracts data was well established. Their business value to the enterprise had yet to be fully appreciated. This is now changing rapidly.

Contracts As Critical Business Data

Over the next five years, we expect that all businesses will appreciate having full visibility into the business insights that sit within their contracts. These critical documents (and related ones such as purchase orders and invoices etc.) document the business processes and organizational resources necessary for a business to function: obligations of the company's customers' to pay, the services or products to be delivered, the expectations of vendors and suppliers, the agreements that help retain employees and contractors. This "contract intelligence" is increasingly a source of competitive advantage. Business processes represented in contracts can be the "secret sauce" that makes some companies successful and others not. Having quicker access



to this information can help an enterprise be more nimble and outmaneuver its peers. Contracts are increasingly seen as at the heart of any “digital transformation” initiative, but companies are missing a system to rapidly make sense of this large unpenetrated reserve of organizational data. Current ERP, CRM and HRIS systems and associated BI tools have helped companies better understand their business, but they provide visibility only to existing structured data, which is typically only a fraction of the overall data in a business. Reviewing, analyzing and extracting structured information from text documents like contracts has historically required heavy manual review in a very unsustainable way. As a result, businesses have been resigned to accepting high levels of risk as business-as-usual.

As more professionals have begun to appreciate the valuable insights that are locked up in contract data, they have sought out efficient and powerful ways to release it. With AI tools like Kira, companies can now scan across their business documents and find clauses and data points. They can create machine learning models that represent their unique business needs, without any technical knowledge of AI. And they can pull that data automatically into other enterprise systems via APIs.

Intelligence that Answers your Questions

If you continue upon this arc, the next step is clear: to accelerate the gleaning of insights directly from those documents. What if you could ask your most pressing business questions to your contracts and get an instant answer - an answer that was derived from the subject matter experts in your organization?

Yesterday, we announced our next leap forward - an intelligence that does just that. This new capability goes beyond identifying and extracting provisions, clauses, and data points, by interpreting this data and providing decision-makers in firms and organizations with the answers to their most pressing questions, instantly. It’s fully integrated into our no-code machine learning platform, Kira Quick Study, allowing for easily scalable customization.

The leap that Kira’s new answers capability enables is like the leap we all have made in navigating the world using our smartphones. It was a fantastic step forward to be able to hold an electronic map in your hand, but it was still a map. It was up to the user to find the starting and ending points, and figure out the best route between the two. But Google Maps and other navigation assistants changed all that - now we can just ask for the directions by naming the destination, and the best route and the travel time is calculated for us. The need to analyze the map data on our own is greatly reduced.



By providing answers derived from contract data, Kira delivers a similar type of boost to users who have common questions about the content of their contracts. The answers have been there all along; with today's Kira, the relevant terms and clauses containing the answers are efficiently and accurately surfaced; but it's still up to the user to review those clauses and determine the answer to the question. Our new feature takes that last step out of the process, enabling the user to go straight to the answers to questions such as:

- "Can you terminate this agreement for convenience?"
- "Does the agreement renew automatically?"
- "Does the contract contemplate payment of a non-refundable amount?"
- "Does this lease require payment of percentage rent?"
- "Does this contract contemplate an inflation adjustment?"

Some readers will undoubtedly point out that data scientists can build tools that produce output structured as yes-no and multiple choice answers (and have, for years). However, Kira is (we think) the first to do this for document analysis in a scalable way: any domain expert can create their own answer models for their own use, using our advanced no-code machine learning interface. No intervention from technical teams or data scientists is required.

Now, rather than requiring extensive review of relevant clauses extracted from contracts to arrive at the answer, Kira will now clearly and directly hone in on the specific answer users are looking for.

The Future of Business Workflow

Customers' ability to integrate contract data in workflow will take a leap forward with this new capability. Answers to commonly asked questions can be automatically embedded directly into other internal systems, triggering actions and faster, more accurate decisions across their organizations. And lawyers will be able to focus more on engaging with clients on the legal and business significance of what's in their contracts rather on the process of extracting that information.

- Businesses need tools that allow them to be instantly responsive to changes in regulation and the business environment.
- Companies need to embed contract data in business processes, including risk management, procurement, retail operations, strategic planning, and finance.
- Lawyers, accountants, auditors and other professionals want to complete process-oriented review work efficiently, so they can focus their expertise on higher-value advice to clients based on insights quickly extracted from contract data.



We expect to see contract analysis systems quickly becoming integrated more deeply into both business intelligence and business workflow automation tools. Many companies, large and small, are already working on this space, designing “smart contracting” systems that begin to represent contracts as digital objects so that obligations can be tracked and actions automated.

We see Kira as a bridge between these worlds; deeply technical and yet deeply human at the same time. Our mission is to enlighten the world’s enterprises to truly know what is in their contracts, and our vision is to do this by empowering people to capture and share their expertise using AI. We delight in sparking unimaginable creativity and innovation in this once-staid world of business documents. Our scientists and development teams are constantly responding to new challenges faced by our customers, creating the next generation of contract technology, and putting those advancements in the hands of our customers. Getting answers instantly with Kira is just the most recent refinement in this quest.

WRITTEN BY:

David Curle

Legal Content & Research Lead

+1.888.710.8454

Steve Obenski

Chief Strategy Officer

+1.888.710.8454



With our newest capability, Answers & Insights, we analyze how “trap door” structuring in credit agreements can impact negative covenants for both lenders and borrowers.

[Download the Study Today to Learn More](#)