



Legal AI

A beginner's guide

By Richard Tromans

ARTIFICIAL LAWYER

the answer company™
 **THOMSON REUTERS®**

About this White Paper

Although several legal AI companies launched as early as 2010, the technology and how to make use of it has since 2016 become a headline issue for many law firms and large corporates.

This is not because the technology radically has changed between 2010 and now, but rather that the traditionally risk-averse and often conservative legal market is now finally ready to adopt software solutions making use of natural language processing and machine learning.

This move toward AI adoption in part has been driven by increasing client pressure on law firms to be more efficient and a growing unwillingness to pay for what they regard as process level work. As clients demand fixed fees for large projects law firms have little choice but to make use of technology to improve efficiency and protect margins. Clients are also increasingly asking law firms to show them proof that they are innovating and embracing the latest wave of legal tech to provide a better service and value proposition.

And, it is probably fair to say, that as more lawyers see rival firms adopting AI systems and clients welcoming such moves, then more firms will seek to embrace the same technology tools. No firm wants to allow a rival to get so far ahead in terms of using new technology that they begin to have too great a competitive advantage.

This report describes the current shape of legal AI and suggests some uses of AI, as well as some that may emerge. It should be seen as an initial starting place for those interested to learn more.

Areas covered include:

01 Contract review	5
Structure of contract analysis market	6
02 Legal data research	7
Knowledge systems	7
Predictive systems	8
03 Intelligent interfaces	9
Data interfaces	9
Triage services	10
Legal bots	10

What is legal AI?

'Legal AI' is the use of AI technologies, such as natural language processing (NLP) and machine learning (ML), in relation to legal tasks.

NLP is the use of a special type of software that is able to read 'natural language', i.e. normal, text that we all use. As the law is in large part constructed from the written word, the power to read, at great speed, legal texts using NLP provides a considerable new capability lawyers and clients did not previously have access to.

NLP, for example, could be used to read a contract and tell you what the key clauses were and if they differed from standard clauses you would normally expect in that type of contract. Or, it could be used to understand a user's legal query and then search legal data to find not just any document that used certain key words, but rather return information that truly matched the concepts in the user's question.

While such sophistication is not infallible, nor as subtle as an experienced lawyer's work, it can provide a junior lawyer or paralegal with some very compelling competition.

Machine learning refers to the ability for software to learn and to become more accurate in its outcomes. In the context of legal AI and reading text, this would mean having the ability, often with some human intervention, to improve its level of accuracy.

AI can be used within a law firm or by in-house lawyers. We should not be too proscriptive about how and where certain systems can be used, even if they are used by a certain customer group today.

Because AI applications are in effect 'tools', any lawyer could make use of the systems when and if there is a use case to do so. They are not specific any one practice. The limits on AI's use are often more about the imagination of the users, than the technology itself. That is to say, NLP can be used in a wide variety of ways and become a useful tool in multiple legal tasks.

The limits on AI's use are often more about the imagination of the users, than the technology itself

In fact, because non-lawyers also need to deal with legal issues, such as agreeing to or referring to legal contracts, some legal AI 'tools' are also designed to be utilised by non-lawyers. This is already becoming a growing segment of the legal AI market, for example in relation to contract generation and completion.

In short, legal AI has a potential use wherever there are people who must deal with legal documents or address legal queries, especially where those legal needs are expressed through text, which AI experts refer to as 'unstructured data'.

With regard to eDiscovery, some vendors in this space are making use of AI software, but not all. For this reason it isn't listed as an AI group of its own. However, AI-driven eDiscovery is most similar to contract analysis.

One can divide up the many applications of legal AI into roughly three main branches, though these will be, and to some extent already are, being added to by new inventions. That said, an easy way to start is to focus on the following three groups of uses:

1. Contract review

Reading and analysing legal agreements, such as commercial contracts and leases, then extracting useful data from them, and/or checking them against rules/ current law. In some cases this also means helping people to finalise contracts.

2. Legal data research

Legal research and litigation prediction systems, covering statute and case law as well as case outcomes, i.e. not specifically looking at contracts, but rather examining the data produced from the practice of law and from laws/regulations.

3. Intelligent interfaces

Interactive, web-based, Q&A systems that clients can engage with via text input to gain legal information, or that can guide lawyers/non-lawyers in completing basic legal documents and forms.

To some degree there can be some overlap between these three. They could also be linked together in some applications that will emerge. But as far as the present day is concerned, the main vendors of legal AI appear to branch into these three general groups.



Table 4: The three main branches of legal AI

01 Contract review


Contract review covers the reading via NLP of legal agreements, such as leases or due diligence documents.

What a user wants to look for, or what certain vendors tailor their systems to do, varies. But the fundamental process is the same in each case. There are many potential uses for such technology; some of the applications that law firms and/or vendors have already identified include:

- Due diligence
- Lease review
- Compliance and risk review
- Sales/procurement contract review
- Employment contract review
- Financing/OTC derivative agreement review
- And, as noted, some types of eDiscovery.

NLP is many times faster than human lawyers at reading contracts, while accuracy levels in matters such as due diligence is generally higher than that achieved by human lawyers.

Natural language processing is many times faster than human lawyers at reading contracts, while accuracy levels in matters such as due diligence is generally higher than that achieved by human lawyers



Structure of contract analysis market

Although there are hybrids in the AI contract review market it can still be said to have two main product varieties:

Volume contract review

These are systems that are focused on analysing large numbers of documents. The objective is usually to seek out specific legal issues in contracts and leases. Sometimes this is to give an overall picture to the client of the legal status derived from the document group; in other cases the aim is to find anomalies (such as in due diligence), or to spot areas that need further legal attention (such as in a compliance review).

Contract assistance

These are systems that tend to be focused on smaller numbers of contracts, sometimes even single contracts. Some vendors aim such systems at non-lawyers who wish to understand what a contract contains (for example, a procurement executive who wants to know what is in the 50-page procurement contract on his desk). Some of the systems are also focused on the pre-signing phase and help the client to spot clauses the other party has included in a contract that they may need to re-examine, or where they may need to add in certain legal clauses to meet standard internal practices/rules for that type of contract.

A note on eDiscovery

People often ask whether AI contract review is the same as eDiscovery. The simple answer is that although some of the latest litigation eDiscovery platforms do seek to make use of NLP and machine learning to analyse documents, it is perhaps better to see such uses of AI techniques as operating with a parallel, but often quite different use case to contract review for due diligence or lease review, for example.

In fact, as most readers will appreciate, eDiscovery is already a vast legal technology industry in its own right, with over 200 vendors providing a wide range of technologies and methodologies.

02 Legal data research

AI systems can also be used in a broader support role beyond contract review. These uses can be roughly divided into:

- **Knowledge systems:** e.g. Legal research along practice lines; and
- **Predictive systems:** e.g. Case outcome prediction based on specific matters and/or litigation trends based on court outcomes.

Knowledge systems

An AI-driven knowledge system is a piece of software that taps data held or linked to a law firm or in-house team. Data could be expert opinions on legal matters by the partners of the firm; statements of fact about laws and regulation; relevant cases and commentary by judges; as well as any associated case notes or updates that the firm has created itself or is linked to.

In short, the system can do an 'intelligent deep dive' into the material available, working in natural language (i.e. normal English, often in sentence form) to provide the answers you require.

What makes these research systems better than simply an enterprise search, or a database trawl, is that the system is both learning from the questions a lawyer is asking, but also seeking to infer the best responses from the data. It is not just a key word search that brings back hundreds of documents; instead the NLP seeks to isolate what the lawyer actually needs.

Such research alone clearly does not remove the need for detailed analysis by senior lawyers of the research that has been delivered. However, it may significantly speed up basic legal research conducted by junior lawyers who are working as part of a larger team. It may also be made more valuable when and if it links to other AI systems and document automation processes: for example, where a document may take note of certain key, though 'vanilla' legal points that the firm wishes to add for the client's benefit.

It may also reduce the need for lawyers working in PSL (Professional Support Lawyer) roles, or at least those handling relatively straightforward research matters.



Predictive systems

Predictive systems are a variation on the above knowledge systems and could arguably be called a sub-set of them, though they could also operate on a standalone basis. They are seen as primarily of use in pre-litigation planning.

AI-driven software can examine huge numbers of cases and all the publicly available court documents and rulings made by judges in the past up to the present day that are relevant to a case along, with many other types of useful public data:

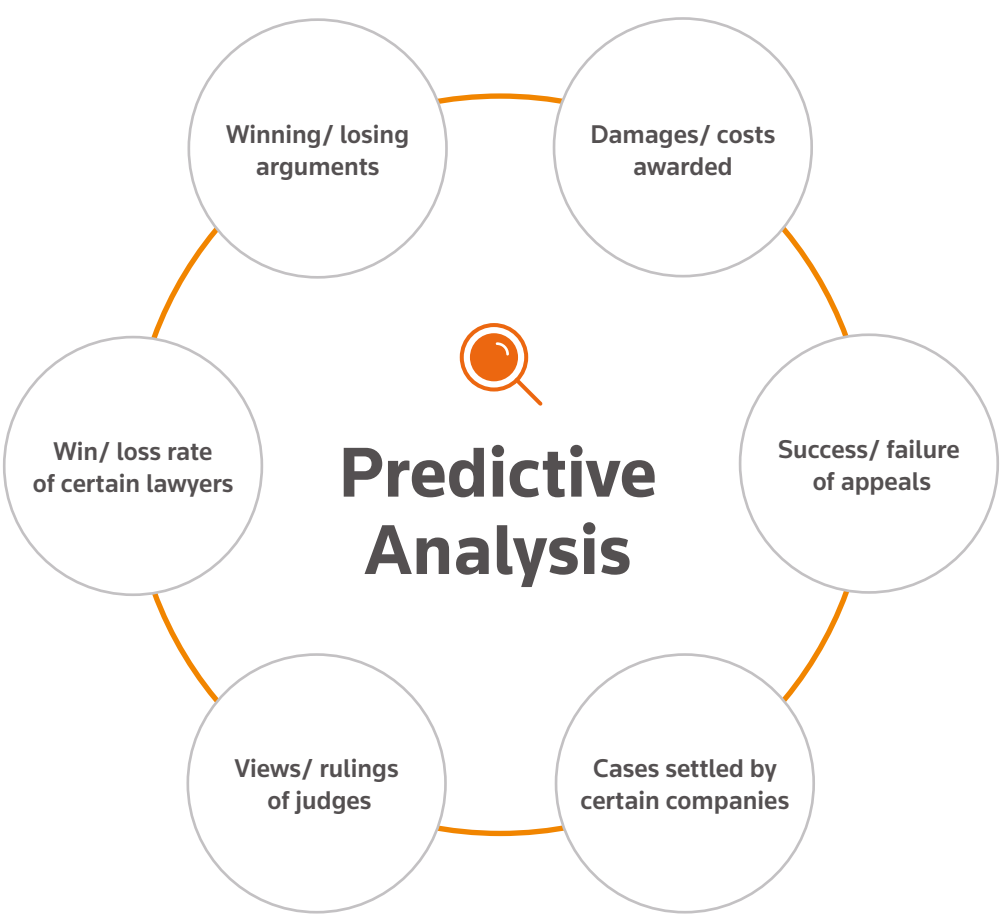


Table 6: Types of data that could be examined using an AI predictive system

The main aim is to reduce the volume of manual research and provide lawyers and clients with actionable insight into previous cases, the actions of lawyers on similar matters, and where possible to gather evidence on the terms of likely success of a matter compared to previous similar matters, and/or give some indication of the damages that could be awarded by such a matter and/or other fee/value data.

03 Intelligent interfaces

The third main branch of legal AI is the development of intelligent interfaces that can guide lawyers or clients to specific legal information, or to 'triage' their legal needs. The aim of the technology here is not so much on conducting primary research or analysis, as the above applications do, but to help guide a user through to the right outcome.

Data interfaces

AI-enabled systems can help clients and lawyers to conduct rapid and routine legal tasks that require some 'expert' information to complete.

In some cases they may be using NLP to understand queries a lawyer or client has typed into a dialogue box. Machine learning may also be used to help the system better provide the right answer that is tailored to the user's needs.

That said, some expert system are not using AI technology, but rather conditional logic and/or word tagging to understand queries and respond to them. The reality is there is a grey area here that is still being explored by vendors. But even those not making use of AI systems look likely to move in that direction eventually.

These applications are often used when a person is guided through an 'intelligent checklist' that allows them to gain the right knowledge, or in some cases to complete very simple legal documents, such as NDAs.

The software usually uses drop-down menus and check-boxes to move the user through a series of steps so that they can either be given the correct data they need, for example in response to a specific legal question, or be used to fill in the missing elements of a standard document.

They may be outward facing for clients to use, or inward facing, allowing lawyers to interrogate the expert system for their own specific needs and/or help them to complete a legal document.

Expert systems, whether outward-facing or inward-facing, are carefully designed to provide informational support to a specific need, such as how to respond to a certain type of employment dispute, or to help add in data to a certain type of legal form.



Triage services

The potential exists to use an outward-facing AI system to act as a triage service. At present, most law firms do not use these types of systems, though banks and other financial service companies are developing automated customer relations systems. A hypothetical example for a law firm is set out below:

Such triage/customer-directing interfaces already exist in a very basic way at some law firms – usually those that deal with the public and ask clients to write a short email to describe the matter. However, these could be far more effective and not just serve individual clients.

Rather than asking the client to do all the work, an AI system could be used to help guide clients to the right outcome in terms of an advisory path and understand their queries using NLP. It could also make use of machine learning to steadily improve its responses to certain types of client query over time. The AI could also immediately link the information provided via the triage system to the firm’s own research into the types of case worth pursuing, as well as link to the firm’s CRM system.

Legal bots

Many lawyers will be aware of ‘bots’ or ‘chat bots’, though perhaps without considering how they could be used in the legal space. Apple’s Siri is probably the best known ‘bot’ – what one might call an AI-driven personal assistant. Essentially, this is an interactive interface that operates in natural language, whether written or spoken, with the latter clearly being far more complex.

At present there are ‘access to justice’ legal bots that operate using written text, which help to give preliminary advice on matters such as criminal law to members of the public. Another example is a bot that guides members of the public through the process of completing a challenge to a parking fine.

However, these systems are, as yet, relatively narrow. That said, the market for legal bots is continually evolving and there are already new bots surfacing that are capable of a far broader range of legal topics.

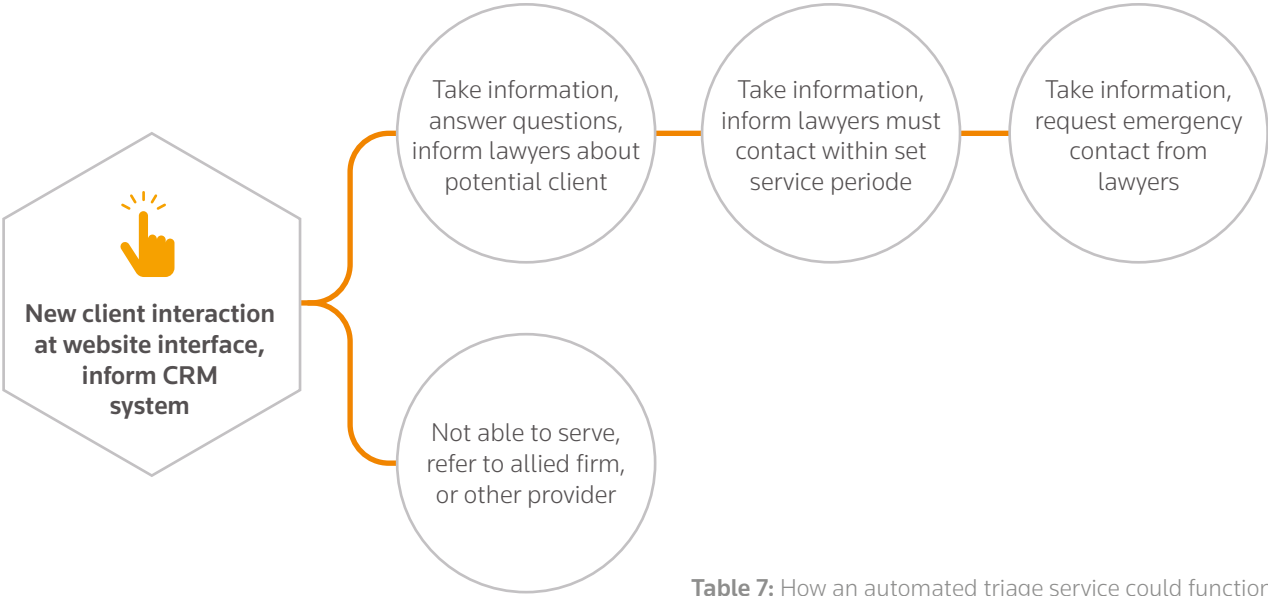


Table 7: How an automated triage service could function.

Rather than asking the client to do all the work, an AI system could be used to help guide clients to the right outcome in terms of an advisory path and understand their queries using NLP.

Legal AI conclusion

This is a relatively short and succinct overview of legal AI, which is a market that is rapidly evolving.

Nearly every week a new legal tech start-up launches an application that makes use of NLP and machine learning techniques and so the picture inevitably is more complex than the simplified version set out here. But, we all have to start somewhere, and getting to know some of the key strands of legal AI is probably a good way to begin to structure one's thoughts.

The shape of the legal AI market will no doubt also be quite different by the end of 2017. More AI companies will emerge that may bring together several of the strands set out above. Others may perhaps merge together, or invent entirely new ways of using NLP and machine learning in the legal sector. It is truly a dynamic area and therefore all the more necessary to stay up to date with.

One thing is certain, we have now moved past what was a period of speculation and into a period of factual and real world uses of legal AI systems. The number of vendors will increase, the number of law firms, in-house teams and non-lawyers using these AI systems will grow. Eventually legal AI will become a key element of the legal sector that many thousands of people rely upon and use every day, just as many other technologies have done so before.

We have now moved past what was a period of speculation and into a period of factual and real world uses of legal AI systems

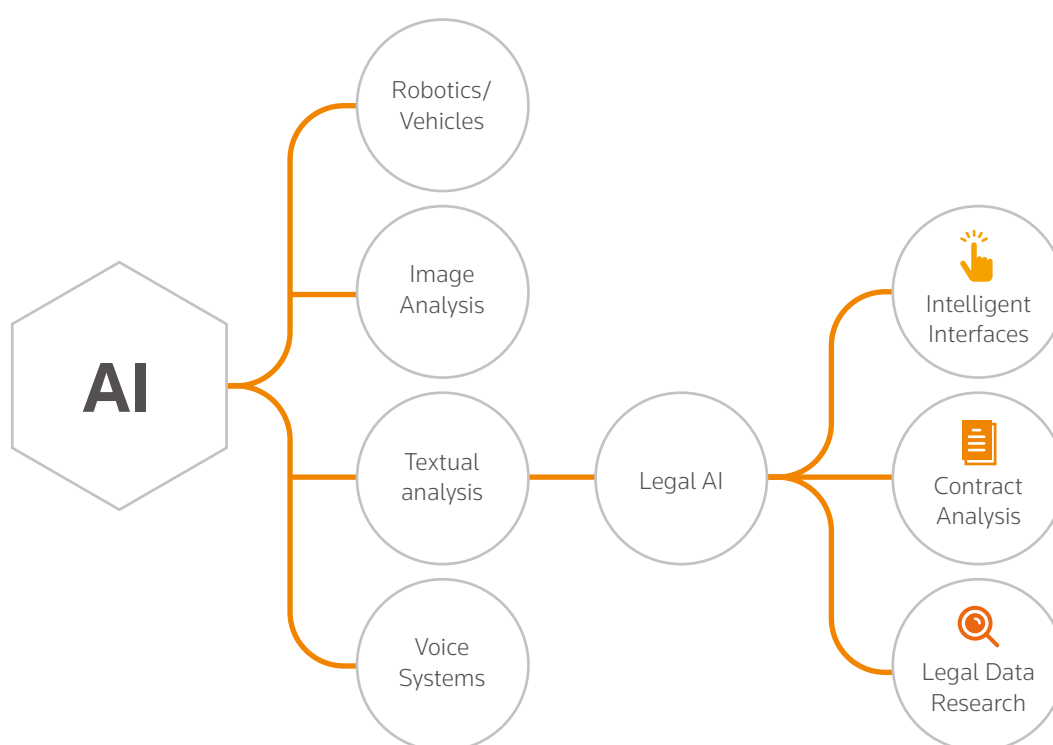


Table 8: A simplified taxonomy of AI, flowing down to the three main groupings of legal AI.



About the Author

Richard Tromans is the founder of TromansConsulting, which advises legal businesses on strategy and innovation, including advising on the adoption of AI systems. He also runs the site Artificial Lawyer, which is a site dedicated to new developments in legal AI and automation. He is based in London, UK.