

Machine learning in contracts:
What's real and what's possible?

Juro

ARTIFICIAL
LAWYER



How can we put the realm of 'legal AI' in context? How can this field that has generated so many headlines over the past few years be pinned down, in terms of the day-to-day life of lawyers? Is there even a way of getting into this technological world of machine learning and 'natural language processing' that means anything to those with no particular interest in this field?

The short answer is: yes. If we put aside the technology – which Matt will brilliantly deal with below – we can still understand what this is all about. Simply put, what we are talking about is automation.

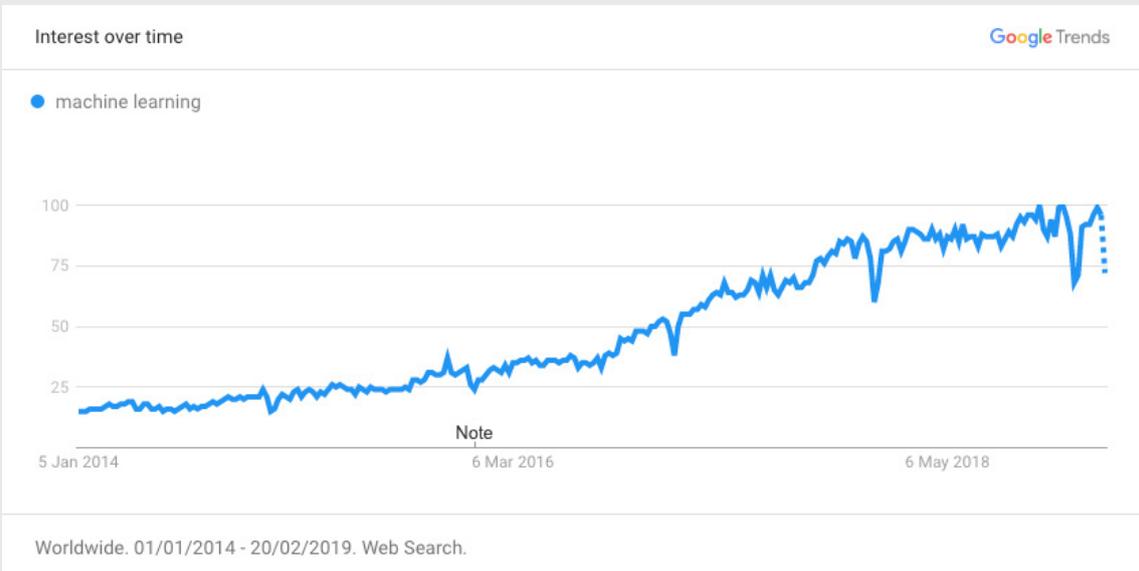
By this I mean the automation of tasks a lawyer may normally carry out, from reviewing a document, to gaining insights into how people are actually contracting with one another, to in some cases actually producing first drafts of legal documents.

Why would anyone want to do this? One reason is efficiency. Human lawyers can only ever work so fast. Another reason is that in a world with masses of unstructured legal data – words – some types of analysis just need this technology to achieve the task.

OK, all very nice. But is anyone really doing this? Yes, they are – and in droves. There are now more than 25 legal AI companies just doing document review, and dozens more providing applications that use some aspect of machine learning, including Juro.

The number of lawyers using this tech is growing too. For example, the top two legal AI document review companies have around 225 clients between them. Add in the other 20-plus companies, then the several dozen other legal tech companies using similar tech, and we are looking at hundreds of firms and in-house legal teams around the world.

This represents incredible growth in a short space of time. Just three or four years ago that figure would have been a fraction of this. It looks like AI and machine learning are here to stay in the legal world – and the reality is that we have only just begun to explore what is possible.



Global interest in machine learning as a concept has exploded in recent years – per Google's monitoring of global search trends.

Richard Tromans is the founder of Artificial Lawyer and Tromans Consulting.



Dr. Matthew Upson is lead data scientist at Juro. Matt was one of the first data scientists recruited to the UK Civil Service, and his work

to modernise the creation of statistical publications had a major impact across Whitehall departments. In his last position before leaving government, he deployed one of the first applications of machine learning in the Cabinet Office, and worked on the first application of deep learning to GOV.UK.

'Artificial intelligence' should be straightforward to describe. It can refer simply to the ability of machines to complete complex tasks. But such a broad definition isn't that helpful. In the 80s, when AI began to develop as a matter of public interest, it often meant an 'expert system'. Expert systems enshrine in a computer program a series of deterministic rules: if x happens, do y; if not, do z. The advantage here is obvious: machines can remember massive volumes of rules better than people can. Even better, they allow the knowledge of one expert to be scaled across multiple machines, and applied almost instantly.

So far, so good. But while AI has always been on the radar as an increasingly powerful technology, with the potential to make a profound impact on society, in recent years the hype cycle has been more tightly focused on 'machine learning'.

What is machine learning?



JARGON BUSTER

AI

Artificial intelligence, or the application of machines to complete complex tasks.

ALGORITHM

Instructions on how to complete a task, ranging from the banal (number 1 + number 2 = number 3), to the very complicated (e.g. machine learning algorithms).

MACHINE LEARNING

A field of AI concerned with developing systems which can 'learn' patterns from data and then use those patterns to make predictions when presented with new data they have not seen before.

NATURAL LANGUAGE PROCESSING

A field of computing interested in the analysis of natural human language.

Machine learning is currently the dominant field in AI. Broadly, it refers to a two-stage process.

1. Training

This is when instead of using rules written by humans to complete a task, we allow a machine learning algorithm to learn the rules for itself. One way to do this is by presenting a machine learning algorithm with examples of our data, and the outcome which we want to predict, and allowing the algorithm to learn the patterns in the data that would lead us to that prediction. This new knowledge, built from the context of our data is termed a 'machine learning model'.

2. Inference or prediction

The machine learning model can now make predictions when presented with data it hasn't seen before, using the knowledge it has accrued and developed through the training process.

We mentioned above that machines can remember big sets of rules better than people can. Machine learning algorithms can learn very complex patterns better than people can too; those patterns might be too difficult, if not impossible, for us as people to formalise into deterministic rules.

But what does this mean for legal?

There are two fields of machine learning that are of great interest to the legal sector:

“While it’s tempting to jump on the hype train and claim that machine learning will easily solve all those bulk tasks in no time at all, it’s not quite that simple”



LOST IN TRANSLATION

Early work in machine translation focused on translating individual words or phrases directly into another language. But this approach ran into difficulties for the same reason people would if they focus narrowly on direct translation: words can mean different things based on context. A different approach was needed.

The incredible advances behind tools like Google Translate are thanks to an approach called Neural Machine Translation (NMT): here, instead of directly translating word for word, the algorithms develop an internal representation of meaning – just like we do.

So now, when machine translation tools translate from English to Spanish, they first use an algorithm (an 'encoder') to convert that data into an 'encoded' state; then another algorithm (a 'decoder') to generate words that represent that encoded meaning in Spanish.

1. Supervised learning

This is where the training process uses input from humans in the form of 'labels' applied to examples of our data ('training examples'). For example, if we want a program to recognise the type of contract that it is looking at, we could gather 50,000 contracts, find some paralegals, and have them go through the contracts and label the type manually – as a sales contract, employment contract, NDA, or whichever.

Once these contracts are labelled, we can show the machine the different groups of contracts, and the labels applied to them, and the machine learning algorithm will learn how to categorise new contracts based on the characteristics of the labelled contracts. Just like with the expert system, this knowledge can be applied to thousands of contracts instantly, potentially making a hugely time-consuming task trivial.

2. Unsupervised learning

This time, we remove the paralegals. Instead, we take the 50,000 contracts, show them to a machine learning algorithm, and ask it to find patterns without any human involvement. The machine is able to look at the contracts and say, for example: “within this dataset there are ten distinct groups, each of which share common characteristics.” The machine doesn't know why these contracts are similar, just that they are similar – that's the unsupervised bit.

To turn this into useful insight, we can bring back a few paralegals, and ask them to look at each group and apply a label, which could be the contract type. The jargon for this process is 'semi-supervised learning'.

These are fairly simple examples, but you can imagine that the same techniques could be applied to a whole range of legal problems, given the right data.

The problem

Sounds transformative, right? But while it's tempting to jump on the hype train and claim that machine learning can and will easily solve all those bulk tasks in legal in no time at all, it's not quite that simple. There's no doubt that machine learning's potential benefits to the legal profession and, most importantly, its clients, are huge; but the reality is that legal – and particularly contracts – is a difficult space in which to use machine learning, compared to an area like image analysis.

Contracts aren't photos

If I open the Photos app on my Android phone, I can type in a search term – 'cats', for example – and the program will search through my photos and return all of those that feature cats. This is possible because Google has built machine learning algorithms that can analyse photos effectively and detect patterns; in this case, the arrangement of colours and shapes that makes up a cat. Learning from photos is relatively easy for machines because almost all the information they need is right there in the image. A computer can see as much information in a photo as a person can, and converting that information into numbers for an algorithm to work with is trivial (in fact, that's how photos are stored on a computer anyway).

The difficulty that machine learning models have with text is that it's much harder to convert into some kind of numerical representation that's capable of capturing all of the information available to a person when they read it. We can provide a machine with words and syntax, which can be expressed numerically; but it's much harder to express the semantics, the meaning and the context behind a given document.



Which contracts data is usable?

Contracts data is harder to work with, but of course it's not impossible. Broadly speaking, at Juro we work with two types of contracts data:

“Users shouldn’t notice some ostentatious AI deployment. Their contracts solution should just work and solve their problem”

1. Operational data

This is data about the way people work with contracts. This could include variables like:

- How long a contract takes to be signed;
- How many edits it undergoes;
- How many times approvers view it;
- How many people access it;
- Which approvers take the most time to sign off;

... and so on. This data is the easiest to work with: it's already numeric, it requires no real translation, and you can build models from it immediately that bring real value. For example, Juro users can access this kind of data on their analytics dashboard, and it can help them understand the bottlenecks in their workflows, or the contracts that require the most revision, or the templates that lead to the fastest time to signing. All of this helps contract processes to become smarter and more efficient.



An organisation's contract analytics dashboard in Juro

2. Data in the contracts themselves

The data within the actual contracts is much harder to work with. As we discussed earlier, we first have to convert it into a numerical representation, in order for any machine learning models to be able to work with it. When we look at that contract data, we broadly consider three levels:

- Word level** - where we extract valuable information from individual words, or groups of words. This could be extracting the start date or the end date of the contract, or finding its governing jurisdiction;
- Clause level** - where we draw conclusions about individual clauses. This could be identifying whether a contract contains a particular type of clause (like a non-disclosure clause, or a break clause), or it might be determining how similar the clauses are in two contracts;
- Contract level** - where we use the data at the level of the complete contract. In our paralegal example above, this might help us classify the type of contract with which we're dealing.

However we choose to use the data, the important point is to remember why we model contract data in the first place: to solve problems for customers. Machine learning shouldn't be deployed for the sake of it, to reinforce a marketing message, or to tick a box on a website. Feature design should be invisible - users shouldn't notice some ostentatious AI deployment. Their contracts solution should just work, solve their problem - and if the design really works, users shouldn't know, care or even notice whether machine learning was in play.

The giants are coming



WHAT IS DEEP LEARNING?

Deep Learning is a subset of machine learning that is based on a particular type of algorithm: Artificial Neural Networks (ANNs). The design of ANNs was inspired, in a simple way, by the structure of neurons in the brain, but this simple foundation has proven transformational in a lot of fields.

Almost all the major advances we have seen in recent years in speech recognition, image recognition, and self-driving cars (to name a few) have come from ANNs. What makes an ANN 'deep' is that it is able to combine simple patterns together to learn highly complex and abstract patterns in data. These combinations of patterns happen on 'layers', so typically a deeper model with more layers is able to learn more complex patterns.

One of the advantages of these models is that their building blocks are mathematically very simple, so they can easily be deployed in real time on small, mobile devices – like smartphones or cameras.

As reported in *Artificial Lawyer*, in late 2018 Amazon launched 'Amazon Comprehend Medical' – a machine learning service to analyze and extract medical information from unstructured medical text. It remains to be seen how the venture will pan out, but the idea that Amazon, or another behemoth like Google, might diversify and provide services in the legal industry is not at all far-fetched.

If one of the tech giants decides to enter the legal machine learning market, there are various possibilities. They might provide services to lawyers directly; alternatively they might provide machine learning as a service that could be used by lawtech providers, like Juro. Indeed, a significant portion of the lawtech businesses currently disrupting legal already operate on Amazon's AWS cloud platform.

Regardless of how such a service comes to market, or which company moves first in delivering it, one of AI's universal rules holds true for legal too: the player with access to the most data is likely to win. More data means more insight, smarter decisions and likely better outcomes: whoever has the biggest arsenal will have the best firepower, and the most powerful ability to deliver value for customers.

For example, we're often asked if machine learning could help us to create 'the perfect contract' – the perfect NDA, for example. The obvious problem is that if 'perfect' means 'free of superfluous language and clauses', then it's extremely subjective as to what 'superfluous' means. It might depend in large part on the context in which the contract is created, and as we know, it's very hard to know that context and express it numerically. But if the dataset is big enough, and we have not tens or hundreds but hundreds of thousands, or even millions of NDAs, then building a model to find the statistically optimal version, under various circumstances, would be pretty trivial.

“One of AI’s universal rules holds true for legal too: the player with access to the most data is likely to win”

The Eureka moment

Although we're making great progress in machine learning in legal, the genuine moment of an earth-shattering breakthrough hasn't really arrived yet. It is still very much an active field of research, and so far the difficulties of working with contract text as data mean that something truly transformative remains, for now, elusive. This might be simply because other more immediately compelling AI applications have historically had more focus, both from industry and academia. Self-driving cars seem just that bit more exciting than machines learning how to make contracts better.

However, the ground is still shifting – there were several milestone achievements in 2018 that we may yet look back on as watershed moments in machine learning, especially those in deep learning (see box above). If machine learning in legal delivers on even half of the hype we've seen to date, there's no doubting the transformative effect it could have on contracts. But focusing on a specific flavour of machine learning algorithms is to miss the point of why we're teaching machines to read contracts in the first place: to transform the value we can deliver to users, and help them solve contract problems faster and more easily. **

ENJOYED THIS WHITEPAPER?

Click here to share it:

LEARN
MORE

If you'd like to learn more about how Juro helps to make contracts faster, smarter and more human, contact us at hello@juro.com or follow us at blog.juro.com and [@GetJuro](https://twitter.com/GetJuro).



Juro is an AI-enabled end-to-end contract management platform that offers contract creation, negotiation, e-signing and analytics, with an intuitive interface designed to make contracts faster, smarter and more human. Juro saves businesses like Deliveroo, Skyscanner and Estée Lauder up to 96% of time spent on contracts, and integrations with Salesforce, Slack and Greenhouse mean Juro is loved by sales and HR users as much as it is by legal teams.

Juro

