

SPECIAL EDITION

LEGALTECH

BEYOND THE MYTHS

• Special Edition • Legal Operators • 2022

**THE VARIOUS MYTHS AND
MISCONCEPTIONS AROUND
AI IN THE LEGAL SECTOR**

Special edition based on the articles by serial legal tech entrepreneur and thought leader Arnoud Engelfriet



Join the Brightest Minds

SUM
BY TH

The Inaugural

EXCLUSIVE NETWORKING EVENT for
Legal Operations Professionals

September 14 – 16, 2022

Ritz Carlton, Half Moon Bay, CA

BY INVITATION ONLY

The Best and the Brightest

- **OUR VIP SUMMIT** assures our attendees gain real value through connecting to the best and brightest thought leaders and decision makers in legal operations.
- **OUR UNIQUE FORMAT** builds in professionally facilitated, executive-level conversations that allow for meaningful information to be shared and lasting connections to be made.

Roundtable topics will include



DIVERSITY,
DATA COURAGE and
WHAT'S NEXT



PROFESSIONAL
DEVELOPMENT
and LEADERSHIP



CHANGE
MANAGEMENT



SPEND
MANAGEMENT



OPEN
TOPICS

in Legal Operations for

SUMMIT BY THE SEA

Now 40 extra seats
Available

SUMMIT STRUCTURE

Our Summit planning team has created a unique, in-person event at the spectacular Ritz-Carlton Half Moon Bay venue includes:

- **A Curated Guest List** limited to 100 attendees
- **Facilitated Roundtables** with 10 Summit guests per table
- **Substantive Conversations** on current topics
- **Peer-to-Peer Networking** activities and opportunities
- **Best Practice Exchange** opportunities for sharing experiences and lessons learned

*For Sponsorship Opportunities,
contact [Colin McCarthy](#) at Legal Operators*

AGENDA OVERVIEW

Day 1

- Opening Lunch Followed by Peer to Peer Facilitated Sessions
- Welcome Reception at Ritz Carlton Ocean Lawn
- Opening Dinner at Ritz Carlton Tacos and Tequila

Day 2

- Breakfast & Lunch Followed by Peer to Peer Facilitated Sessions
- Offsite Dinner Sam's Chowder House
- Night Cap Ritz Carlton

Day 3

- Breakfast Followed by Peer to Peer Facilitated Session




REGISTER NOW

Summit By The Sea is a Legal Operators Initiative





Arnoud Engelfriet, the author of 'Legal tech: beyond the Myths', is co-founder of the legaltech company JuriBlox, and creator of its AI contract review tool Lynn Legal. Arnoud has been working as an IT lawyer since 1993. After a career at Royal Philips as IP counsel, he became partner at ICTRecht Legal Services, which has grown from a two-man firm in 2008 to a 80+ person legal consultancy firm. He specializes in Internet law, which he has been involved with since 1993. With his IT background, he likes to focus on complex technical / legal ICT issues and software licenses (especially open source). His website, [lus.mentis](http://lus.mentis.nl), is one of the most popular in the Netherlands about ICT and law. In addition to his daily work Arnoud is a part-time lecturer at the Vrije Universiteit Amsterdam.



<u>The Artificial Intelligence revolution</u>	4
<u>The focus on accuracy</u>	10
<u>How can robots read?</u>	16
<u>Are robots taking over our jobs?</u>	22
<u>How do robots handle legal practice?</u>	28
<u>How will AI transform the nature of legal work?</u>	34

This special is brought to you by Legal Operators, the #1 Community for LegalOps Professionals

Legal tech: Beyond the myths #2

The Artificial Intelligence revolution

Legal tech is coming. With Artificial Intelligence on board. Ah, yes. We have seen and heard so many promises: it will transform our work. It will replace lawyers, reduce tedious work. And so on. Still, here we are, still typing away in Word while the shiny AI-powered workflow optimization tool gathers dust in the corner. Often the reason is the same: the tool was overpromised and underdelivers. This series will take a look at the

various myths and misconceptions around AI in the legal sector. What can we expect, and what is still a fairytale?

Swimming submarines

Let's get the biggest myth all out of the way: computers aren't intelligent and never will be. That is not to disparage their capacities, which are formidable and which can deliver results that far outclass the best work of



humans. It's just that this is not intelligence, at least not in the foreseeable future. As computer scientist Edsger Dijkstra famously put it, "the question of whether machines can think is about as relevant as the question of whether submarines can swim."

This is important because when considering the capabilities and results of AI, we tend to compare it to human capabilities and re-

sults. We expect human reasoning and the type of results (and mistakes) that humans make as well. And AI just isn't delivering on that front. Fundamentally, an AI does not think as we know it. In its most common form, AI is driven by statistics: pattern recognition, similarity clustering, outlier spotting, and so on.

Jokingly it has been proposed that any

mention of AI should be replaced by “giant Excel charts.” There is a kernel of truth to it. Starting with the assumption that AI, for example, does not understand language but can spot patterns in language, and look up different patterns that fit better you will have a much easier time accepting the analysis of a legal text.

Accuracy and trust

This misconception causes us to mistrust AI results and misjudges the usefulness of AI legal analysis and reviews. We are used to certain types of mistakes from a junior lawyer, for example. He or she will miss complicated or rare exceptions or focus too much on the letter of the law and forget the business aspects. A senior partner would never do that but could go too fast or be focused on her hobby horse regarding IP protection and consider liability a trivial issue taken care of by insurance.

AI never makes that kind of mistake. However, there is an entirely new class of mistakes to be made by AI that dives into case law (legal research) or extracts information from contracts. Based on statistics, a particular Superior Court verdict may best fit the current legal question. If so, the AI will happily suggest it as the winning argument. However, often a legal analysis is expected to cite certain cases, so this suggestion will be seen as “off”, not what a human lawyer would do. Similarly, suppose a human lawyer reviews a contract and sees a clause she’s never seen before. In that case, the lawyer will flag this as a question mark and seek a colleague’s input, or maybe consult a legal library or expert system to learn more. But when an AI encounters such a clause in a legal review, the AI will simply assign the label

that best matches the clause according to its underlying data and algorithms. There’s no such thing as “I don’t know” for a computer. However, the AI will likely give this label a low probability or provide several (almost) equally likely alternatives.

In a firm where an AI system is newly deployed, this type of mistake will quickly cause human lawyers to throw up their hands and dismiss the AI system as useless. And they are right – if this were a new associate hyped-up to ease everyone’s workload, this would be a very disappointing outcome. But this is a computer, which does things differently. And this takes some getting used to.

Text recognition

A related subject – which we’ll discuss in a future article – is how computers analyze text. In the legal profession, the text is the raw material from which all legal output is built. From legal briefs to contracts to advise or pushback in negotiations, it all comes down to what is said and how it is said. In terms of speed, no human will ever beat a computer system in the task of looking for a specific phrase in a large number of documents or in performing any technical operation on a text. However, once we start considering the analysis of the meaning of text, it becomes a very different game.

Most AI systems – even those that promise natural language processing (NLP) – do not understand texts at all. They operate, again, on statistical analysis. This is a proper noun in the plural form, so the plural form must be appended from this list of associated verbs. Any intelligence is brought into the system by manual human design and thus is restricted

to the list of intelligent steps that the human operator has thought of.

A few simple rules already provide a surprising appearance of intelligence. For example, in my NDA reviewing tool, I originally included a date checker: if someone used the tool on Friday after 3 pm local time, the “Waiting for results” page would add one of a set of random phrases, including “Don’t worry, I’ll get this done before the pubs open”. Also, we wrote the output to include some random exasperations such as “This NDA is for ten years, don’t sign that, are you crazy”. Such simple touches go a long way to create the impression that this is more than a computer doing calculations.

In the end, however, AI does not understand text. It will go by statistics and patterns. This can be surprisingly effective, especially in the legal profession: there are only so many ways to declare the laws of California applicable to the agreement, after all. But it can also create very strange errors: if the system recognizes a price clause as a limitation of liability and amends it as instructed, you end up with a price of 2 million plus whatever the insurance pays out. No human would ever make such a mistake, and so this type of mistake is memorable and will underline that AI is far from production-ready.

Of course, this type of mistake is embarrassing, but is it really a fundamental error? Remember, this is how AI works: the clause matches the patterns for “limitation of liability” best, and such a clause should be amended to “two million plus insurance”. This is what I meant above with “computers don’t

think”. Does that mean AI is useless? Far from it. The system should first of all be explicit about its confidence in the prediction and offer explanations such as “Clause recognized as a liability with 35% confidence, amended to match minimally acceptable limitation”. With such small steps, the output is a lot more understandable.

The true cost of AI

Another touchy subject always is the cost of an AI solution. This applies equally to law firms and in-house counsel that seek to deploy an AI-based solution, although for different reasons. Law firms tend to worry about the effect on billable hours: if it takes two minutes for a tool to do the same as a senior associate now does in an entire workday, what can be charged to the client? Even with the time for reviewing the tool’s output included, the cost per time spent will be significantly lower. But this fear is easily assuaged: surely a senior associate can find better things to do than the type of review that a tool can do as well?

The concerns of in-house counsel regarding costs are more complicated. There is a direct effect on the department when the tool is deployed: installation costs, consultancy and training all are billed by the vendor. And after that, a monthly bill will appear for the service. However, there is no direct benefit on the other end: the employees don’t suddenly become cheaper, as they are salaried and don’t reduce their hours now that a tool is taking over the drudgery of reviews or analysis. As with the law firms, in-house counsel will find new (and more challenging) work to do, but unlike with the law firms, there is no


measure on the hours spent or time saved in the legal services provided to the company

hours spent or time saved in the legal services provided to the company.

The net result may well be that the department's cost is increased, but the people are just as busy, and no one knows whether the average quality or response time has improved. Of course, the answer is obvious: start measuring the quality of work and response time from the internal client's perspective. This, however, hasn't been done much, and in any case, it takes time before enough historical data is available to make quantifiable statements. In the meantime, the perception that the AI tool costs money without making things better lingers around. We will talk later about legal operations and implementation strategies to reduce these fears.

Going forward

AI is coming; there is no doubt about that. But with AI, misunderstandings and missed expectations will come, which may harm the successful deployment of AI in the legal workplace. The key challenge, therefore, is managing the intended users' expectations, which given the hype surrounding much of the AI "revolution" is going to be quite a challenge. In the upcoming episodes of this article series, we'll discuss these and further myths in more detail and seek for practical solutions to get them out of the way.



Bridging academic research
and practical advice,

The Practice offers
thoughtful analysis on the
legal profession for a global
audience.

[Read The Practice now >>](#)



**CENTER ON THE
LEGAL PROFESSION**
Harvard Law School

Can you help me to connect with the right DMU*

I'm looking for opportunities to claim leadership

How do I share the Unique Value Points of our solution

I'm eager to show the world that we make a difference!

We like to share our knowledge and insights

LOOK NO FURTHER

Legal Operators will help you to find your path to sustainable growth!

Connect with the market and increase engagement with potential users thru webinars, meetings and highly effective content marketing solutions.

**DMU: Decision Making Unit*





Legal tech: Beyond the myths #2

The focus on accuracy

What's the difference between a lawyer and a lawyerbot? In the twenty years I've worked as a lawyer, no one has ever asked me how accurate I worked. But every time we introduce our lawyerbot to a new audience, the first question we get is always "How accurate is it". Which is fine, as we have a good answer: 95.1%. But what does that even mean in a legal context?

Rote language: food for robots

As I wrote last time, computers aren't intelligent, and never will be. They are good at computation, and therefore also at routine tasks like comparing texts, looking for statistical patterns and so on. This is great for legal tools: lawyers produce huge numbers of comparable texts with clear patterns – the rote language and standard expressions that



(artwork by Yash Wadke)

are needed to trigger certain regulations or avoid problematic jurisprudence.

The difference however also means that the way a robot analyzes texts and makes legal predictions or recommendations is fundamentally different. This in turn has grave consequences for the perceived quality of AI for legal work. It bears repeating: we consider

lawyerbots to be simulated human lawyers, and their work the automated performance of human work. This is wrong, but very hard to root out.

Of course we are used to human lawyers making mistakes. A rookie lawyer can miss important case law. A senior partner may focus too much on his own hobby horse, or be

out of touch with the latest developments. In a rush, certain clauses may be glossed over and implications missed. We understand and can work with such mistakes, as we can relate to them.

Machine learning: a deep dive

Robots make an entirely different class of mistakes. These have to do with the way that robots screen text. Let's have a little dive into the technology for that.

Most robot lawyers work with so-called machine learning, a process where the computer learns to recognize patterns in data, usually based on statistical similarities derived from a given set of examples (the training dataset). Usually the process is focused on classification: assigning a label to a piece of text, e.g. "this is a liability clause" or "this is a verdict for the plaintiff". Most contract review tools work this way. The labels help classify and value the contract, especially when a value judgment ("this liability clause is 2.5 million Euros") can be used in the classification.

Another popular application is information extraction: "the first contract party is Royal Shell", "this verdict cites the ECHR Sunday Times case" or "the defendant was not served with a notice of default". This is the domain of so-called natural language processing, where human-programmed or statistics-derived rules of grammar are applied to identify information: "this clause has the supplier as the subject and uses a verb indicating obligations without an ancillary verb indicating trying, therefore this is a supplier warranty".

In both cases however note that the computer

has no actual grasp of the legal consequences, it is applying formulas and numbers to derive conclusions. "Must", "shall" and "will" are all verbs indicating obligations, therefore "Supplier must" is a supplier warranty. And in particular with classification, the computer will classify the sentence as belonging to one of its categories. There is no "ignore if you're not certain" or "unclear" category. (In fact, if there were the computer would probably classify all of the clauses as "unclear" since there is no downside to it for doing so.)

Confidence in robot review

The usual way of handling this limitation is to examine the prediction's confidence. Most machine learning systems provide predictions with indications of their confidence or certainty. This outcome looks very much like the training set, therefore the system is highly confident. But this outcome is rather unusual, therefore the confidence is only 28%. During training, an engineer would search for a minimum confidence that gives the lowest numbers of false classifications or mistaken extractions. A prediction with only 28% confidence is likely to be ignored.

It is, however, a mistake to think that a prediction with a high confidence is likely to be accurate. This has to do with the fact that robots are not trained to look for the right answer. Instead, they are trained to provide an answer that best matches its training data.

Training, training, training

we were first developing our NDA-reading robot NDA Lynn, we noticed that confidentiality agreements with a California choice of law were always getting rejected as being very

onerous for the recipient. This despite the fact that the clauses dealing with security, notification and so on were as standard as they could be. Further digging revealed however that the training dataset contained only very strict, one-sided NDAs with California law. From this, Lynn had concluded that California law is a good predictor of a very strict NDA. It thus made perfect sense to first look at the choice of law, and if that is California then to give a quick answer.

It is thus imperative to have a training dataset that is as complete and diverse as possible. This is however enormously hard, even when the system is restricted to only one jurisdiction. First of all, there are no public datasets with contracts, so assembling a large corpus of data is very labor intensive. Companies working in this field may be able to get (anonymized) documents from their first customers, but that introduces a bias: a law firm with an IP focus for high-tech enterprises will create different contracts than a law firm focused on SME businesses.

This is essentially the same problem as the allegations of bias that pop up whenever a machine learning system makes predictions or analyses of human behavior, such as with spotting potential fraudsters or even simple face recognition. In legal it is a bit harder to spot, as it may require deep human review of the clause to see that something fishy is going on. And to add to that, two lawyers may reasonably disagree on interpretation of a legal clause or implications of a court verdict.

That said, of course there are measures to objectively evaluate quality. A simple approach is

to set aside 20% of the dataset for an evaluation when the machine learning system has been trained. As this 20% was labeled prior, the system output can easily be compared against the human-chosen labels. In more advanced approaches, this split is done multiple times along different lines, generating multiple models with different test sets. If all comparisons reveal a good quality of the predictions, then the dataset (and the models) are suitable for practical use.

Still, this presumes that the dataset is representative. A very high accuracy only means that the test dataset was well recognized – in other words, that the predicted labels match the human-assigned labels.

Handling computer mistakes

All this goes into the central question of trust. Trust is derived from accuracy in past performance, but how do we measure accuracy if it is so different from how humans work?

The usual answer involves the difference between true positives, false positives, true negatives and false negatives. Very quickly: a true positive and a true negative mean correct identification or rejection, and false positives and negatives are both mistakes. However, these terms assume a yes/no, true/false or guilty/not guilty dichotomy. In a legal robot, we usually deal with multiple classes: a contract clause can be any of 30 to 50 types, verdicts contain a lot more information than “guilty / not guilty” and let’s not even go into the amount of options in a legal demand letter.

Merely looking at a classification being wrong is not enough. You can have small or big

mistakes. Human lawyers easily can make small mistakes: overlooking a dependent clause, forgetting to change a minimum term after a statute has changed, and so on. Big mistakes – say, taking a liability clause as a contract term – rarely if ever happen, and then only to the most junior of rookies. However, for a computer these are all more or less the same: it's not part of the model, it was labeled differently, this is what the dataset looked like. So getting your liability clause misinterpreted or a carve-out to a payment penalty scheme overlooked can happen just as easily.

This is a key reason why lawyerbots are harder to trust: they make strange mistakes, which just as often may be rookie mistakes. And rookie mistakes have huge consequences. This means that a human lawyer feels like having to double-check the lawyerbot's work all the time, which in turn destroys any added value (e.g. time saving) the lawyerbot may claim to have had.

Going forward

Understanding the function of AI tools is tremendously important. Anyone who considers them mere automated versions of human lawyers is setting themselves up for a huge disappointment.

The key issue when it comes to accuracy is not to strive for 100% perfection. This is impossible, just like with human lawyers. Even when a dataset is comprised of thousands of documents – as with NDA Lynn: 14.000 NDA's – there is still a chance a new contract has vastly different language, and thus the system will perform lower. Continuous re-training based on mistakes thus is key. This requires input

from the human lawyers: which label did you expect here?

Similarly important is the realization that a lawyerbot does not distinguish between big and small mistakes. The only way to address this issue is to ensure there are multiple steps in the robot's process. For example: when a liability clause is detected, double-check for certain expected wording and reject the detection if that's missing (such as a number or reference to contract value). Check if another clause has been detected as the same, but with higher confidence. And so on. This is an iterative process that again requires human lawyers to get to know their robot counterpart.

Finally, it comes down to positioning. How can a lawyerbot be deployed to save time or money, whilst minimizing the impact of inaccurate analysis. But not only that: also position the human lawyer (who reviews the bot's work) to easily provide feedback on the bot's output. Nothing frustrates acceptance of a tool as much as being unable to change its workings. Back in the box it will go. Every lawyerbot tool therefore should come with a feedback button, and of course its designers should listen to the feedback. Let's get that right!

Arnoud Engelfriet is co-founder of the legal tech company JuriBlox, and creator of its AI contract review tool Lynn Legal. Arnoud has been working as an IT lawyer since 1993. After a career at Royal Philips as IP counsel, he became partner at ICTRecht Legal Services, which has grown from a two-man firm in 2008 to a 80+ person legal consultancy firm.



Companies involved
in AI. Visit our tech
directory and see
their profiles

TECH DIRECTORY



Legal tech: Beyond the myths #3

How can robots read?

“The question of whether machines can think is about as relevant as the question of whether submarines can swim”. This quote by computer scientist and visionary Edsger Dijkstra is still as relevant today as it was in 1984 when it was written. Computers do not think, they calculate and process information. This may result in outputs that look like the outcome of a thought process, but that is a mere coincidence. What does that mean for a language-oriented field like law? And how do robots then extract information from language?

Natural language processing

Already in the 1950s during the advent of AI the concept of robots reading and interpreting textual information came to the forefront. The so-called Turing Test, a test of a machine's ability to exhibit intelligent behaviour indistinguishable from that of a human, was created with the ability to interpret text in mind. In short, the test proposes that if a human asks questions and gets responses from various counterparts, the human must from the responses determine if he is



communicating with a robot or not. If the human cannot, then the robot is considered “intelligent”.

The earliest work on natural language processing focused on rules. Given a collection of rules the computer emulates natural language understanding (or other NLP tasks) by applying those rules to the data it confronts. The classic example is the “Chinese room” designed by philosopher John Searle in 1980. Suppose we put in a room many, many books

that instruct the reader which Chinese symbols to write down given certain Chinese symbols received on paper. (Yes, we write out each and every possible question-answer pair that is possible in Chinese.) Then, we get a Chinese-speaking person to write questions and put them under the door. In the Chinese room, a person (or a robot) applies the books to produce an answer, which is shoved back under the door to the person asking the questions. Can this person now tell if a Chinese speaker is in the room?

If not, the robot must be considered intelligent.

Of course, this approach requires the creation of a gargantuan amount of rules, many of which even native speakers wouldn't be able to formulate. Still, progress was significant and early experiments showed surprising results. For example, Joseph Weizenbaum's ELIZA program simulated a psychologist able to engage in open discussion with "patients", employing strategies such as "how do you feel about that" or "do you think this reveals something about your relationship you're your parents" whenever the patient presented a topic that the system had no specific rule for.

The rise of machine learning

In the late 1980s the introduction of machine learning algorithms for language processing presented something of a revolution. In machine learning, algorithms build a model based on human-provided training data applying statistical techniques to identify correlations and patterns. Using this model, predictions or decisions can be arrived at without any explicit rules having to be configured. This made it possible to use statistics-based approaches to analyze and respond to textual input. The first breakthrough was in automatic translation, and some successes were achieved in specific domains.

The hardware and memory limitations of then-current computers did put an upper limit on what could be achieved. This changed in the early 2000's with the advent of big data and cloud computing on the one hand and the exponential increase in publicly-available text data: the World-Wide Web. Now it was possi-

ble to take huge corpora of text and apply tremendously complex statistical calculations and pattern-recognition algorithms to distill rules and schemes to transform text into other text. Whether question and answer or writing from prompts or interpretation, machines could now do it.

A key setback remained that humans were needed to annotate the input from which the machine learning algorithms trained their models. This changed in the 2010s, where the rise of feature learning and deep neural networks allowed for so-called unsupervised learning of text features for recognition and interpretation. One of the keys to this breakthrough is the use of word embeddings. "A word is characterized by the company it keeps", as English linguist John Rupert Firth put it. Meaning can thus be derived from context: if these and these words occur together, this other word must be involved and could for instance be used in the output.

Still, limitations remain. One common example is how to handle homonyms, as in the example "The club I tried yesterday was great!". In this sentence, it is not clear if the term 'club' means a dance club, a social club, a golf club, a club sandwich or any other type of club. Humans can understand this from context, even when not given in the document itself: a senior lawyer in her sixties is more likely to mean the golf or social club than the twenty-year-old student known for his partying tendencies.

Machine learning on legal documents

As noted above, initial focus on natural language processing was on translation. This had

one important reason: especially in government documents, multiple-language versions of the same document were often available. For instance, the European Union publishes official documents in all its 29 languages, allowing good comparison between the language structure and vocabulary of each. Further, there was a clear need for quick and “good enough” translation.

A second field where NLP made great strides was in transcribing dictation, especially in the medical sector. Doctors produce a large amount of dictated reports (e.g. autopsies or surgery reports) that needs to be transcribed, typically quickly. At the same time, absolute perfection is not necessary. And what’s more, the wording and phrases used will be limited and somewhat predictable: when trying to distinguish between, say, ‘patient’ and ‘patent’, it is safe to assume the doctor meant the former. For similar reasons, machine translation of legal dictation has seen success, albeit in more limited form as the time and money factor present in government and medical fields is less pressing in the practice of law. The main focus of NLP in the legal field has been in automating legal processes, e.g. a case assessment to predict the outcome if it were to go to court. Here, NLP is a first step in the legal process: extracting the facts of a case, or identifying key factors that judges use when applying the law. But next steps require more advanced machine logic, e.g. figuring out which legal requirements apply. So far, success here has been limited.

Machine learning in contract review

A domain in the legal field where machine learning is quickly gaining attention is the re-

view of contracts, mostly business-to-business agreements. Long this type of work has remained the realm of human experts, as such agreements represent significant business value (and risk), each agreement is different and the time factor for review was not considered crucial. Today, this has changed. More and more agreements (or at least, provisions therein) are considered standard, the cost for human review is becoming more and more prohibitive and speed is of the essence.

This change could first be seen in standard documents such as the confidentiality agreement (NDA), thousands of which are signed across the globe every day. While lawyers (correctly) stress the importance of reviewing each NDA provisions carefully, most businesspeople (also correctly) consider an NDA very much a standard text and just want to know “can I sign or not”. This has led to a value gap: businesspeople do not want to wait for, let alone pay for, a review of an NDA. Several legal tech providers have jumped in this niche to offer automated NDA review tools.

All of these use some variation on the same basic process: use statistical methods to recognize typical clauses found in such agreements, extract problematic aspects of such clauses (e.g. a too-long term or a liability cap) and report to the human user what was found. This works very well, mainly because the amount of variation in such clauses is very limited. There are only so many ways to declare the courts of Santa Clara, CA competent for any disputes. What’s more, NDA’s tend to contain a high level of ‘borrowed’ language. Our own tool NDA Lynn for instance has reviewed over 14.000 NDA’s and has found that

for most clauses, there are only a handful of truly different structures. This type of limited variation makes analysis surprisingly effective.

Other document types may have similar attributes. For example, under the European Union's General Data Protection Regulation (GDPR) a so-called data controller must have a specific type of agreement in place with its suppliers and other processing partners ("data processors"). This data processing agreement (DPA) must meet specific statutory obligations. While each organization has developed its own DPA, the language is very much shared as most lawyers tend to closely copy the letter of the law. Tools such as DPA Lynn thus can provide effective review of this document. However, automated review for contracts in general still seems far away due to the variability of the type of clause that may be present.

The Contract Understanding Atticus Dataset

A promising development in the field of contract review is the creation of the Contract Understanding Atticus Dataset (CUAD) by the Atticus Project, a US-based nonprofit organization of legal experts. This dataset was created with a year-long effort pushed forward by dozens of law student annotators, lawyers, and machine learning researchers. The dataset includes more than 500 contracts and more than 13,000 expert annotations that span 41 label categories (from applicable law to covenants not to sue, limitations of liability, payment obligations and warranties). Interestingly, the dataset contains human-made annotations of what a reviewer would like to know, such as the monetary cap on liability or the end date of a certain obligation. This al-

lows a machine learning system to be accurately trained (or verified) on the CUAD dataset.

Employing the CUAD would provide a well-deserved boost to machine learning contract review. The dataset can be enhanced with company (or law firm)-specific contracts for additional focus. For instance, an IT focused firm would add IT insourcing agreements and categories relevant for technology services, while an international supplier of goods would focus on adding shipping costs, risk allocation and insurance clauses.

Going forward

Machine learning for contract review has come a long way. While it is true that no contract reviewing robot can claim to have an "understanding" of what it has read, a lawyerbot can certainly produce highly accurate reviews of typical legal agreements. This is especially true for standard agreements such as NDAs, but with the advent of large datasets such as CUAD more general contract review is right around the corner. The challenge for any business therefore is: how do we create value with automated contract review, while reducing any new risks that may appear? *This is something for the next article.*

SHARE YOUR WHITEPAPER WITH 250K READERS

-A UNIQUE CONTENT MARKETING SOLUTION BY LEGAL OPERATORS-

WE'LL MAKE A PROFESSIONAL EBOOK FROM YOUR WHITEPAPER AND SHARE IT WITH 250K PROFESSIONALS IN THE LEGAL INDUSTRY



EXAMPLE

ALL OUR PRODUCTS ARE RESPONSIVE. MEANING THAT THE SCALE AUTOMATICALLY TO THE USED DEVICE (PHONE, TABLET, EBOOK, KINDLE, DESKTOP, LAPTOP ETC.)





Legal tech: Beyond the myths #4


Are robots taking over our jobs?

Did you notice? Whenever you read an article on robotization or artificial intelligence in law, there's always a stock picture of a robot wearing a gown or wig, or tapping a judge's gavel. An easy way to illustrate "robots doing legal work" of course. But it shows an underlying assumption, namely that robotization or AI in law means that a robot takes over the work we human lawyers do. That's not what's going to happen, but the idea is stubborn. Until we manage to take it

away, however, robots aren't going to be much of help.

Mere automation

The main promise of legal tech is that technology will help lawyers improve their work. This of course isn't a new promise. For decades legal providers have put efforts in automating aspects of their work. Timekeeping, file management and formatting documents, to name a few. This provided a huge efficiency gain, which



allowed lawyers to focus more on substantive work.

Automation did not, however, fundamentally change the way of working of the legal professional. It merely meant the same work could be done faster, with fewer distractions caused by limitations of the tools. Writing with a word processor is faster than a typewriter because you can eliminate mistakes more easily. But it does not make the actual thought process faster: the contract won't be written any easier, and a legal ar-

gument won't sound more convincing because of a better word processor.

Work transformation

A second stage in legal automation is to automate parts of the substantive work itself. This goes a step beyond automating the way of working. Let's take an example from legal practice: writing legal opinions. Automating the tools would mean using a word processor, and searching a case law database with keywords and filters, then typing in (or maybe

copy-pasting) the relevant passages and the case citation. One could imagine a tool inside the word processor, inserting citations with passages with one click.

But still, choosing which cases (and which passages) is truly the choice of the human lawyer. A tool like IBM's ROSS is a good example that takes this one step further. Given a legal case, the tool retrieves relevant citations and provides the skeleton of a legal argument. This does more than just saving time: it takes drudge work out of the lawyers' hands.

In contract drafting practice, a common example is the document assembly tool that prompts the user with questions: who are the parties, what is the price, should confidentiality be one-sided or mutual, how long should the term be, and so on. Based on the answer, the tool selects relevant clauses and builds a contract. Clever and useful as this may be, in the end this is still a mere automation of the old model agreement that said "insert contract term here (in years)".

Substantive transformation

The third stage in legal automation, which is what's usually referred to as legal tech today, is the transformation of how substantive work is performed. True transformation means doing things differently because technology allows you to.

For instance, in the contract assembly example, a true transformation would be to come up with a different template based on the client or the type of deal, or more generally to go beyond the standard clauses from the model for answers A through D. Or choosing to omit certain articles in their entirety because the client's history reveals they are not useful.

Or re-drafting a limitation of liability after reviewing the other party's professional liability insurance. This is different from merely answering questions and always getting clause 23.A or 23.B for a force majeure statement.

In contract review, an example of automation would be a tool that recognizes legal clauses and flags them for the lawyer to review. This helps the lawyer do his or her work better, but does not change it. A transformation would be a tool that not only recognizes clauses, but also determines their impact and makes a decision whether the lawyer should even see them. If the contract is low-value and the clause is only a small deviation from the company's policy, why bother?

(If you want an example of such transformation that's closer to home, consider a legal liability insurer. Often, when a small claim is received, the insurer may make the decision to pay the insured out of its own pockets and forego the traditional route to court to recover damages from the party that caused it. This is of course based on financial considerations, but still it is very different from the traditional approach where you recover damages from the party that caused them.)

The challenge, of course, is getting such a transformation actually adopted. And this is hard. Change is always hard, but for some reasons especially in the legal sector. This despite the fact that this sector deals with new stuff all the time.

Four hundred years of experience

The legal sector has an image of little change. For a large part, this is only logical: the work

is fundamentally the same as, say, 400 years ago. Back then as now there were conflicts that needed legal arguments to be settled and agreements to be put on paper. While the subjects and the law may have been different, the principle remained the same.

At the same time, this is strange. Most lawyers are very much open to new developments, from the latest gadgets to large-scale developments in society. And again this is only logical, as keeping up with the world is needed to do legal work. If a lawyer can't operate a fax machine, how can he provide legal advice on the status of fax messages, to name just one example. Novelty is part of the work. So why did the legal work remain the same for so long?

Many explanations have been proposed. The hourly billing system in particular supposedly blocked innovation: those who work more efficiently, could not claim as many billable hours. Add to this a system where the partners at the top of a law firm receive a percentage for each billable hour, and the result is a very strong stimulus against reducing the number of hours worked. But surely this is not the entire explanation: those who work more efficiently may make fewer hours on one job, but would have time available to work for other clients.

In my opinion the fundamental reason is the well-known expression "Don't change a winning team". Firms that do well, are very busy. Introducing fundamental change in a busy environment is not going to go over well. This takes time and concentration, and both are in strong demand. So the change will have to wait until next year.

But then why now? Customers are asking more vocally for change, for speed and cost savings. And there's nothing left in terms of small savings, lowering the hourly rate or distributing costs. In addition, more and more firms see options to introduce legal services in innovative ways, threatening the traditional monopoly position law firms have enjoyed. This forces thinking about transformation.

The role of AI in transformation

The rise of artificial intelligence may provide the key to this transformation. Most types of transformation that are currently under consideration, rely heavily on automatically reviewing texts and spotting anomalies. This is something computers are very good at, and this happens to be an activity often requested from lawyers. But we can do more than just put an AI as a first check to save lawyers a little time.

Artificial intelligence is in particular good in recognizing patterns in huge amounts of data, turning it into actionable intelligence: this is a force majeure clause, this line of reasoning reflects the three prongs from the Sunday Times case, and so on. Such intelligence is useful as direct input to the professional taking the next step – but can also be leveraged to change the next step, change the workflow or process in which the analysis occurs.

As a simple example, consider a company that often receives confidentiality agreements from prospective customers and partners. The traditional process would be to send this to Legal for review, the lawyer would enter into discussions with the other side, and after discussions had led to a mutually acceptable

agreement both parties would sign. Applying legal tech would speed up this process: instead of the lawyer reviewing the document, an AI would do the same. The lawyer would review the output, including a marked-up document and open the negotiations.

When transforming this process, one needs to do a step back: what is the intent of this process? Who are the actors that need to operate the process? Even though this is a legal document, the process is by itself not legal: it serves to enable the business to talk with prospects in confidence. There are more ways to address this purpose without having a manual review of each incoming NDA. One simple example: insist on the company's own NDA, or a well-accepted standard NDA (such as the oneNDA initiative).

More advanced process changes involve putting the AI earlier in the process. For instance, the businessperson sends the NDA to the AI, who reviews and either approves or rejects it, where rejection means "not salvageable, use our own NDA instead". This would cover two-thirds of the situations, according to statistics of our own product NDA Lynn. In the other one third, negotiations can be started based on a redline prepared by the AI tool. And even here, the lawyer is not (yet) necessary: the redline can be sent back for initial comment, once those are reviewed by the other party, the lawyer can resolve the feedback and negotiate towards mutual agreement.

Managing workflow

Enabling such changes first of all requires a clear workflow. Businesspeople should know to involve the AI, and when to approach the other side with a redline. Technology can help: the

review tool can take care of most of the administrative burden, and keep track of the latest actions taken and the next steps to be performed by humans. However, this only works if the humans involved are in agreement on what the next steps are supposed to be. Establishing assent on such matters is very hard, as it may involve changing a company's culture or even overcoming corporate infighting.

A related matter is to know what the AI tool is to approve or reject. In general terms, most companies have an understanding of what is acceptable in their line of business. For instance, software companies typically want to retain their IP and are concerned with employee poaching, while food suppliers worry more about quality provisions, returns and the confidentiality of price discounts, to name two examples. But zooming in a little more often reveals blind spots: what do we mean with "we typically want to retain IP"? What amount of returns do we accept, exactly? And are there combinations: is a small price discount with a high return right for the customer acceptable just like a high discount with no right of return?

First steps

Successfully deploying legal tech thus, as a first step, forces a company to reflect on its workflows involving the legal department and the underlying assumptions about the business. Getting this on the table in an actionable format is a huge challenge, but at the same time represents 80% of what is needed for a true transformation of such workflows. The tool can then quickly be deployed to fit the new process. As always with technology, it's not about the tool but about the humans using it.

DIVERSITY & INCLUSION

For Legal Operators diversity and inclusion are important topics. Luckily we see a lot of organizations prioritizing these topics and starting to roll-out dedicated programs. However there is still a lot to learn, and why not learn from each other?

Therefore we welcome everybody who is willing to share their experience, programs, what works - what not, challenges and more in an article published in Legal Operators eZine.

Feel free to send your article, article outline, or idea to [Colin McCarthy](#) or [Joek Peters](#).



Legal tech: Beyond the myths #5

How do robots handle legal practice?

Last time, we ended on a positive note: it's not about the tool but about the humans using it. Which prompted several readers to write in and ask, how will the humans be using it? We know how to work with a new colleague, but what can you expect on a robot's first day in the law department? More to the point, what can you do to get a robot to work as you expect? This gets to the issue

*of what legal tech aficionados call the **playbook**.*

Numbers, numbers, numbers

Whenever we talk about artificial intelligence, the word algorithm makes an appearance. To programmers this is a bit strange, because the one thing that AI does not apply is algorithms. AI is essentially statistics-driven: does this



phrase resemble this set of phrases or more that one, which of these twenty word groups is most similar to this word, and so on. This is in stark contrast to what an algorithm is to programmers, namely a series of precisely-defined steps to take, one after the other. In an algorithmic approach, all possible choices would be enumerated in advance and the robot would simply check off the lists one

by one. Of course, this is impossible, which is precisely why the statistics-driven approach of AI has proven so successful. With enough data, patterns can be discerned that allow for the derivation of bright lines (although quite complex lines) to help identify patterns.

However, this process of classification or recognition is merely one small part of the

system that makes up a lawyerbot. Several other key pieces are needed to arrive at a fully functional system, one that can produce legal advice given a contract, case description, legal brief or other input document. And it is in these pieces that traditional algorithms play a large role.

Building the system

Let's take the use case of contract review as an example. In other words: how can we turn a contract we received from a prospective partner into a redline that will further the negotiations? It's good to keep perspective here. As noted earlier in the series, robots are good at standard work. So this review will be a check for standard issues. No one would use such a review for a high-stakes merger agreement or a complex bespoke services purchase. But such documents are rare: over 80% of legal documents reviewed by in-house lawyers or by attorneys-at-law are routine agreements.

Reviewing a standard agreement means checking for standard issues and making standard amendments to better reflect the client's interests. A term should be a bit longer, a contractual fine should be gone for a service provider, liability should be capped at the level of the liability insurance, and so on. All perfectly suitable for a robot: spot the issue, grab the best alternative clause and make the amendmend.

Thanks to the above mentioned statistical approach, spotting the issue is not a problem for a well-trained machine learning system. But we have some steps to take before we get there. To start with: how do you turn a Word

document into a set of clauses, from which issues can be spotted?

Reading like a robot

This may seem like a silly question: just open the document and read the text. But for a robot, this is a lot harder than it sounds. We humans see text, well-formatted, some bold, some italics, a list of definitions at the end, indented lists, and so on. For a computer, this is all a mass of letter groups held together with XML tags that may or may not provide a fixed structure to guide identification of what's what.

A first approach could be to simply extract the paragraphs as Word provides them. Headings, normal text, tables and other elements have their own markup – technical codes that provide boundaries to separate text and to attach styles to. This markup can be recognized automatically, allowing e.g. a list of clauses to be built with the last-used heading preceding it. Then, each clause can be fed to the machine learning system to identify it and to determine what's wrong with it.

There's a little problem with that: we humans tend to be sloppy with our styles. In the tens of thousands of documents I have reviewed for our lawyerbots, I have yet to find the perfect document marked up as Microsoft had in mind when it created Word. People don't apply styles but merely colour some text bold or enlargen it to indicate "heading". Old-fashioned authors press enter when they think a line is done, as they did in the typewriter era. This looks the same, but for Word this means there are two paragraphs, two clauses – and

to a robot, that means two incomplete clauses to review.

Fortunately, there are advanced text extraction libraries available, so part of the problem is solved. But once you add PDF support – a popular request – you get an entirely new set of problems, because PDF is designed to look good, not necessarily to automatically retrieve text from. For instance, if you extract text from a two-column PDF document you get garbage: every line from the right column is appended to the one appearing on that line in the left column, and you try making sense from that.

But even if you get the text out exactly as the writer intended it, you still have the issue of what goes together. Not every contract is of the one-clause-per-paragraph type. Consider for instance the numbered list, e.g. a list of representations and warranties that a seller provides. Or a list of undertakings by the recipient of confidential information. Or the security measures a data processor commits to in a GDPR data processing agreement. Is that one clause? Should it be split and evaluated separately, or considered one mega-clause? Both have advantages and disadvantages, but the quality of text extraction is key to a good identification.

Playing by the book

Having fed each clause to the machine learning system, we now know what each clause is about, as well as what's going on within that topic. Here's a liability clause, its cap is 2.5 million dollars. This venue clause picks Paris, France. The term is one year with silent renewal. This list of security measures is very

long and its text does not conform to article 32 GDPR. Good to know, but is this bad?

Now we get to the heart of any robot lawyer: the playbook. The playbook is a set of boundaries if you were, a list of all possible outcomes that the machine learning can provide tied to a judgment. Typically each outcome is tied to a response message that can be included in the report generated by the lawyerbot, or coupled to an amendment or counterproposal that would overcome the issue that was found.

A simple example is a liability cap. The playbook can be configured to reject any caps higher than one million, or in the converse to insist on at least two point five million. Or whichever amount you want – this of course depends on the user's perspective. Should the machine learning system then detect that a particular agreement has a cap that is three million, then the playbook will indicate this is a problem: it should be one million, no more. Often, a legal position is not expressed in a number, but rather in some sort of criterion or position. A client may not be prepared to accept contractual fines, or doesn't want a dispute venue outside the European Union. GDPR security obligations should always include two-factor authentication, compliance audits should be undertaken by independent auditors bound by professional secrecy and their bill paid by the other guy. Playbooks can configure this type of counter: if venue is not in this list, reject. If auditor is not 3rd party, reject. If auditor is 1st party or auditrights is not found, accept. And so on.

Configuring the playbook takes a surprisingly long time, even when experienced negotiators

or lawyers are brought to the table. This is because many of them operate from intuition and rules of thumb. Perhaps that cap of one million is because of the insurance payout, but where did the two-factor authentication demand come from? Do we always object to fines or just if they are large, and if so what is “large” if we have to put a number on it?

Would we take a late delivery fine if payment is within 14 days? Getting these positions out of the human heads and into the computer table is more time-consuming than training the machine learning system.

(One startup in this field uses the clever approach of analyzing a customer’s own contracts to determine what should be acceptable. This works, although it is not uncommon for large companies to have purchasing terms that conflict utterly with their sales terms.)

Getting a response

The next, and ideally final step in the process is to produce a response. A simplistic approach is to just accompany each boundary in the playbook with a fixed sentence, such as “We cannot accept this high a liability cap, it must be lowered to one million.” Whenever a clause triggers this playbook boundary, this text is included in the output. The end result is a nice summary of issues with counterpositions, presumably drafted by a human lawyer. Today, legal professionals expect more and lawyerbot systems should give them that. Some variation is nice: instead of the same text for the same issue, allow for multiple outputs given that issue and select one. One approach suggested in the literature uses a

simple counter: the more objectionable issues were encountered, the more terse and blocking each subsequent response became. Just like a human lawyer or negotiator would get if a contract draft keeps asking and asking. Very helpful is adding the outputs as comments in the Word or PDF document. This is part and parcel of human reviews, and lawyerbots should seek to emulate that. Comments can then immediately be tied to the relevant text, reviewed for correctness and responded to with a counter.

But what really takes the cake is an amendment in the text, of course as a redline (tracked changes) using Word’s facility for registering who changed what. This makes the counterproposal immediately clear, and allows the other party to respond focused on the content (or accept the proposal, since we are still talking about standard documents with standard counters).

This may seem hard. It’s true that AI can generate texts – the OpenAI generator called GPT-3 can produce near-perfect quality essays of any length in seconds, for instance. But coming up with a counter in a specific clause seems like a bit much. Fortunately there’s no need to introduce yet more artificial intelligence. A preconfigured set of counterproposals is more than enough. If you have twenty acceptable security clauses, one is likely to match the language of the original text closely and so can be used as a basis to draft an amendment. We are still talking about standard documents after all; in my experience there are rarely more than twenty variations of any clause in NDA’s or DPA’s, for instance.

CONFIDENTIALITY AND NON-DISCLOSURE AGREEMENT	
<p>THIS AGREEMENT is made as of the __ day of ____, 20__ ["Effective Date"], between XXX, a Delaware limited liability company with its principal place of business at XXX; (hereinafter "Company") and _____, a company having an office at _____ (hereinafter "_____").</p> <p>WHEREAS, the parties hereto desire to exchange information regarding their systems for XXX, hereinafter "Project";</p> <p>WHEREAS, each party and their Affiliates possess Confidential Information, and in conducting the Project anticipates needing to disclose to the other party, certain Confidential Information, including but not limited to prototypes, specifications, designs, drawings, manufacturing information, pricing data, sales data, cost data, business strategy, financial information, supplier and customer information, and certain confidential assets, including, without limitation, inventions, trade secrets, processes, and intellectual properties related to the Project or such information customarily regarded as proprietary or confidential, <u>but only to the extent the information is marked as such</u> (collectively "Confidential Information"); and</p> <p>WHEREAS, the parties desire to formalize the duties for maintaining the confidentiality and ownership of the Confidential Information.</p>	<p>NDA Lynn NDA Lynn has checked the entire document from mutual perspective. Advice: do not sign.</p> <ul style="list-style-type: none"> - Critical issue: Parties definition and Security obligations. - Minor issue: Notice for breach. - Detail issue: Applicable law, audit, Personal data, Duration of confidentiality, Term, Public domain scope and Scope of sharing ("Need to know"). <p>NDA Lynn Detail Can we add an audit clause, just to be sure?</p> <p>NDA Lynn Detail Should we have an explicit GDPR/personal data clause?</p> <p>NDA Lynn Critical The definition of confidential information is too broad. Needs to be limited to marked information only.</p>

A non-disclosure agreement reviewed by a robot. The future of legal practice

In and out in minutes

Now let's put things together. We can take Word documents and convert them to machine-readable clauses, with the right headings attached. We can take those clauses and assign meaning to them – strict security, 7-year term and so on. We can look up whether we like that – term no more than 2 years, strike that fine, no cap on liability. We can insert our objections into the document and even amend the clause proper.

One final issue remains: how do we get documents in and out? A web interface seems intuitive, but there's an even more exciting option – just e-mail your document and get a reply back, just like you would work with a human lawyer. Except this one mails you back in five minutes. Does that sound like a myth to you?

Legal tech: Beyond the myths #6

How will AI transform the nature of legal work?

Legal tech is coming, with Artificial Intelligence on board. The promise I quoted in the first article, and it is still very much relevant today. Computers aren't intelligent, rolling in a big shiny tool will not change the way lawyers work but you'll get far with a clear playbook and handy rules. In the final article in this series we'll wrap the mythbusting up.

Robots as lawyers

Let's get one thing straight: if you're a lawyer robots will not take your job – unless your job was so tedious that a robot could do it. Robots calculate & automate, they cannot do (not now, and not in the future) the real stuff that lawyers do. Robots can help though: they

can scan for patterns, identify clauses or phrases of interest and take programmed action based on what they encounter.

This is as true for classical programmed systems (if you see X, take action Y) as it is for artificial intelligence. Even though AI is promised to think creatively, in actual fact AI is just very, very good at pattern recognition, text analysis and conversational text generation. That's not to say it's bad or useless. On the contrary, an AI analysis can very much speed up legal work and identify issues or help with problems that humans may miss or need hours of hard work for.

The actual issue however is not so much how good an AI can work, but how systems using



AI (or more generally: legal tech) can improve and transform the working of law firms and legal departments. This is much harder than programming AI – this is business transformation. As business consultant Peter Drucker famously said, “Culture eats strategy for breakfast”. Expecting a shiny AI tool to transform the company is a pipe dream merely because it works very well. The challenge is to understand the workflow that the company uses.

Digital signatures

An early adoption example in the legal tech world is the digital signature tool. Putting an electronic signature where a 'wet' one was previously required is a good example of the

benefits of legal tech. But it is also one of the few really good ones: simple, clear and with a very precise advantage. You can almost calculate the business benefit of this. There is no downside, apart from the costs of the tooling. But training is nil, you can see what happens and that it is legal is also obvious.

For further steps, from contract generators to the use of chatbots for intakes or wizards who write advice letters, it is often not that simple. And the reason behind that is actually simple: it is not clear how you earn money by working with such tools. And that is again mainly because organizations are set up to deliver their services in a certain way, and new technology entails a different way of working.

The partner structure has often been mentioned at law firms as the reason why legal tech would not be there: the partners would earn less from billable hours, because those tools are faster and therefore fewer hours are worked. I don't think it's that simple. The argument is an exponent of the underlying reason: fearing less customer returns is an economic objection. Whether that literally means that your timesheets are getting shorter or that the customer wants to pay less at all or something else, it matters less.

The general problem is of course a very tough one to tackle. Especially because it is rarely explicitly put on the table: sorry, we are not going to do this tool because it costs us a lot and we see few returns. Especially with legal tech, because “we have to do something with it” and then the tool is rolled out with a lot of fanfare (or rather: a special team or a champion) after which we are allowed to see some interest and after that the use dies a slow death. Just because it doesn't feel useful enough. So there must be an external incentive to bring about that change. In the legal sector, these are often the customers: large companies that do not accept an hour/invoice but want a fixed fee per job or per year, or even simply require that tooling is used. The only problem is that if the majority of the offices are not yet at that level, it is difficult for such a customer to get her way. There is, of course, a tipping point at some point, but when that will be reached remains to be seen.

Only: you are left with the point that a legal service provider is quite good at his job, and has also set up the workflow to work as well as possible in that way. Whichever way you

look at it, introducing legal tech comes down to changing business operations, and there must be a clear economic reason for this. Why would you start working in a new way now, especially now?

Start small and end big

Change does not have to be strategic, does not have to be huge and does not have to be all at once. Change comes when the culture is open to it. This is often the case with small things. A slightly more convenient tool to check references. A service that not only looks up case law but immediately puts it as a reference in your Word document (and checks whether there was an appeal against the case found). A button in Word that completes and signs a draft letter. An Outlook plugin that says “Look out, the attached NDA is unacceptable.”

If you want more, you really have to work with the culture in the office. Why do people not want to extract documents from a tool, but continue to work with their own templates? That could be because they don't trust the tool, or because they find the learning curve too steep. Or – very silly, but it happens regularly – because they can't log in (anymore) and then noticed that no one spoke to them about it. Or worse – but this happens regularly too – because their direct managers do not enforce using the tool, or even disparage its quality or importance to the business. And addressing that is ultimately also a cultural thing. Because if you want to change as the leadership of such an organization, you must have a culture of leadership. Bringing people along and motivating them. And if there's no other option, make it mandatory. But then

you have to fully support the choice yourself and take for granted that things will (temporarily) be less.

The trend towards commodification

Earlier I wrote about what is called commodification. Traditionally, legal advisors provided complete customization, just like the tailor of the past who delivered perfect customization for each customer with a roll of fabric. Smart offices standardize the process and provide customization based on standard clauses or quickly adaptable model contracts. They have semi-finished trousers and jackets in the closet and they trim the fabric for the customer. Standardization can be pushed further and further, and the inevitable end point is the product: a completely standard piece of service, at a fixed price and available in almost identical form from multiple providers. For a legal service provider, productization can feel threatening, and for good reason. A standard product will yield less money per item than the custom-made service of the past, and moreover entail risks such as an incomplete or inadequate service with all the associated complaints. However, there are several ways to deal with this threat. The first way is to use the product as lead generation. The product is then actually the first half, the intake of the actual service. A slightly more far-reaching approach is to clearly separate product sales from customized services. And yes, there will of course always be a need for pure customization in the market. But make no mistake: that market is getting smaller and smaller, because more and more customers are discovering that a standard product is actually good too. Just like many people prefer to buy their suit as ready-to-

wear, with at most limited adjustments, such as taking in trouser legs, than going to a tailor. The tailor will certainly continue to exist, but increasingly become a niche. The large market share – and therefore the large turnover – will move to those ready-to-wear sellers.

It's about the standard work

Of course, the legal profession started out as a specialty, and the legal knowledge today is specialized and quite unique. So it is not surprising that when you read about changes in the profession, you think about how your specialism can be safeguarded in the event of such changes. But I keep saying it: it's not about the specialized work, it's about the standard work. Although lawyers and lawyers often see themselves as suppliers of unique customization, a whole mountain of work is standard.

The only problem is that we lawyers can indicate much better than others why that standard work should also be delivered fully customized. After all, our services are looking for risks and problems, underlining their seriousness and proposing a solution that must be followed on pain of high fines and other legal calamities. (Few people can say that a mistake in their work can lead to criminal proceedings against management.) And that is often followed, partly because the specialism is shrouded in a certain mystique and the consequences sound very serious.

The above is not necessarily untrue. It is true that incorrect advice about an agreement with a competitor can have that consequence. Or that a mistake in an NDA can lead to years

of litigation with millions of claims. But those are the exceptions, not everyday practice. It is about “can I sign this because I want to drink coffee with that supplier”. The advice must be proportionate to this. A standard case requires a standard answer. The particular risks should be capped or captured if their probability becomes high enough.

More important to me is the realization that if you as a lawyer don't do this, others will. Take the drafting of general terms and conditions. Protecting an entrepreneur with clear rules for his trade is typically something you use a lawyer for. What risks does the entrepreneur run and how do we cover them? Yet the product "general conditions" is something that can be obtained from so many sources, from handy bookkeepers to many online services where you can obtain them with or without a questionnaire for free, cheap or expensive. There are indeed still lawyers who can ask 12,000 euros for a set of terms & conditions, but that is not the bulk of the work.

As a lawyer, I think you can go two ways. Either you say, I am really a specialist and you should have me if your situation is not that standard, or, on the contrary, I can provide that standard work just fine and I pass on the difficult cases. I think the natural tendency is to immediately choose the first, but realize that this only attracts a limited target group (and therefore the amount of work).

On the other hand, whoever enters the standard work immediately encounters the Chamber of Commerce, handy bookkeeper and online services. Because they can do the standard trick too, and are smarter in bringing in standard customers at a lower price. And yes, he

comes again with his Susskind: commodification in the legal world has already come a long way. For many products even to free.

Yet that also offers opportunities, especially for lawyers who can do more than that accountant but do not want to be the super specialist right away. Making yourself distinctive on top of commodity products is the best way to grow in such a market. So: what is your added value on top of such a generator? What does the customer want an accountant not to be able to do?

And no, the answer is of course not "perfect customization with all risks fully negotiated and put away". For me, this is exactly the sweet spot to use legal tech: it provides the standard work that is the start of your customization. To stay close to myself, thanks to NDA Lynn, my clients review their NDAs in five minutes, then pay me to get substantive answers to the crazy questions. As a result, I spend all day doing crazy stuff in NDAs and I like that.

That is how the lawyer of the future must redefine himself: where is my added value on top of the technology, on top of the commodity. I wish you good luck with that!

[More about Arnoud Engelfriet](#)

**How many NDAs
do you read each year?**

**Our robot reviews
about 4.000 of them.**

It also fixed every issue.

(That's some 50.000 redlined clauses, with comments.)



Lynn. Consistent. Sharp.

And a lot faster than human lawyers.
(Who now have time for the serious stuff.)

[**TRY NOW**](#)

2022 ANNUAL REPORT

BIHC
BLACK IN HOUSE COLLEGE

Disrupting the
**STATUS
QUO**

Carlos Brown on using his influence to change cultures and increase diversity

ALSO

In Conversation with Chair Gensler of the SEC
How Data Brings Accountability
The Unique Challenges for Black Leaders

