

# Explainable AI (XAI) for DD

How does AI work? Can you trust the results?



 **imprima** | White Paper

**Published Feb 2020**

The contents of this document are limited to general information and not detailed analysis of law or legal advice and are not intended to address specific legal queries arising in any particular set of circumstances.

## Introduction

With AI being rapidly adopted in almost any application imaginable, from spam filtering to self-driving cars to medical radiology, the discussion on how to control it is intensifying. Some even worry about AI becoming too powerful, with humans eventually losing control (Harari, Homo Deus, 21 lessons for the 21st century). At Imprima, we don't believe that's a realistic fear, at least not in the foreseeable future. We do believe, however, that results from the use of AI should be explainable. This is called glass-box AI (as opposed to black-box AI), or Explainable AI (XAI).


The reality is, however, as all modern AI is based on "Machine Learning", it is in fact a black box. Machine learning is not based on pre-programmed rules. No model is implemented that captures a human decision-making process, like decision trees.

Instead, with Machine Learning a purely stochastic and empirical approach is taken. Therefore, it is not possible to explain how the machine learning reaches a conclusion or prediction. We can, however, analyse and understand its "behaviour".

In this paper we explain how that can be done in relation to AI being used to enhance legal due diligence during M&A. We analyse our recently released AI tools, by asking the following questions:

- **How does the "machine" learn?**
- **Can you trust the results?**

We want to avoid overly technical language when answering questions in this paper. Therefore, we will use a real-life analogy to explain how it works, and how we can ensure you are in control, so you can trust the results.





## How does the Machine learn? ... The Virtual Apprentice.

We could compare a Machine Learning algorithm (or “the Machine”, for short) to an apprentice. As it is a virtual apprentice, it can’t go to the café on the corner to get a coffee for you, but it can be trained very, very fast, and it can flawlessly remember what it has learned.

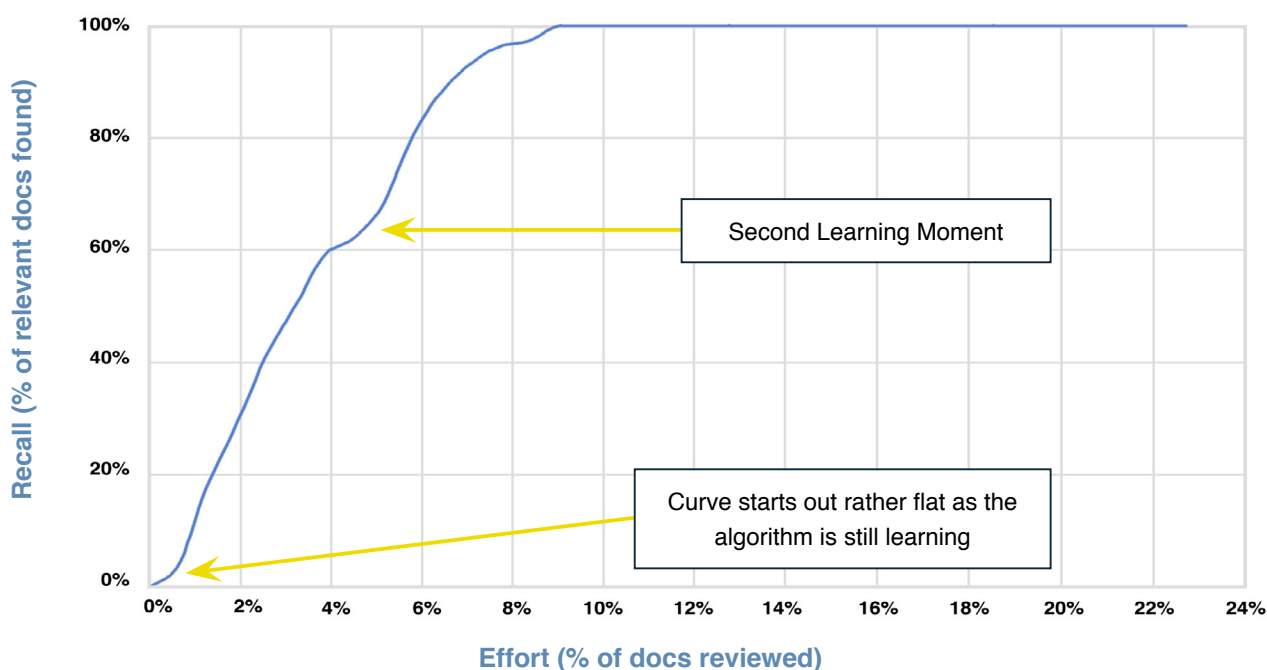
To explain how that works, consider this scenario:

As a Lawyer or analyst, you need to find and review ‘Employment Contracts’ and, identify and display those that do not contain a Non-Compete clause (the “Query”). This may be necessary as a standardised ‘Employment Contract’ has been rolled out at a company, but not all the older versions have been revised.

**To illustrate the time saving, we now look at the following chart:**

### So, what is this graph telling us?

- At point zero, you begin reviewing the contracts by entering a search term, for example “employment agreement”. At this point the knowledge of the Machine is totally blank.
- The curve starts out rather flatly. That is because the first batch of documents produced is based on a simple flat-text search. Due to the limitations of this search method, only 4 relevant docs were found in this example.
- That is where the Machine is absorbing information: it watches what you are doing and learning from your actions.
- After that, the Machine kicks in: “Ah, that is what you want, I got it, here you go. Here are some relevant docs”.
- Note also there is a second point in the curve, at around 4% effort, where it becomes flattish again.



The vertical axis shows correct documents found as % of the total number of relevant documents in the data set. In this case there happen to be 100 relevant docs, so 100% means all 100 relevant docs have been found. The horizontal axis shows the number of docs visually inspected as % of all documents in the data set (3,100 documents).

This is a second learning moment. The more obvious results have already been retrieved by the machine, but you, as an expert, see a few more relevant docs, albeit a bit different from the ones so far. A bit later, the curve becomes steeper again. That is where the Machine says: “ah, so you also want those? I got it, here you go, some additional relevant docs”.

- Note that this means that the Machine can adapt. It learns while you are reviewing, so it can learn the subtleties of what you are searching for along the way. Similarly, the Machine can also adapt if you come to different insights while reviewing, if you want to see additional docs with different issues. The approach is inherently very flexible.

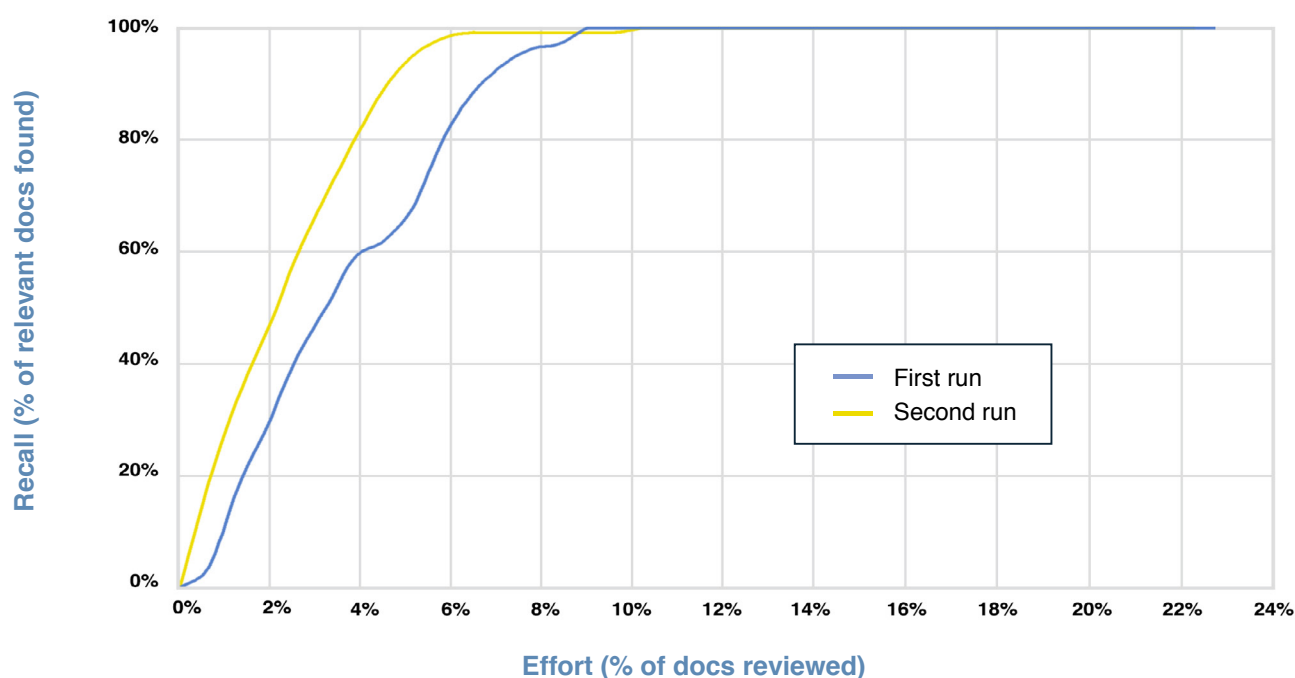
Finally, when you are done, and the Machine is trained, you can save your query. You can then re-use your saved query and apply it again to find employment agreements without a non-compete clause in a new set of documents.

**The results of that are shown in the graph below.**

The yellow curve (“second run”) shows the application of a previously saved query on a new data set (for reference, we also show the results of the previous experiment as a blue curve (“first run”). The Machine already knows what you wanted from the old data set, has learned from your behaviour on that data set, and now just repeats the trick it has been taught on the new data (albeit still guided by your analysis when you are reviewing the new docs along the way, resulting in the Machine becoming even more accurate. Basically, the more you use it, the better it becomes).

As can be seen, it takes off quickly: the yellow line is much steeper in the beginning than the blue line, and reaches the target 95% recall significantly earlier, at around 5-6% effort. That is because the algorithm already knows what to do. Also note that 95% recall is achieved significantly earlier, but not dramatically so. The latter is to be expected, as the “yellow” result (“first run”) was already very good.

Simply put, the Machine, indeed like a virtual apprentice, follows what you are doing, learns from what you are doing, and then successfully mimics what you have been doing.



## Can you trust the results?

Well, we don't expect you to: We don't believe any lawyer will blindly rely on algorithms, not in the foreseeable future at least. However, with the above approach that is a moot point: The interactive nature of the Machine Learning allows you to inspect, at all times, what the Machine proposes to you. So, there is no need to rely on a black box; the Machine will tell you what is going on.

In other words, while the Machine is doing the grunt work, you are in charge. You monitor the results, while the Machine learns from your behaviour and produces ever better results. You review docs as you always do, read the docs, make notes on the docs, and mark docs for further review. No results produced will escape your scrutiny. At any time, you will be the judge if the Machine does what you want it to do.

Consider again the blue curve in the first graph: We already noted that at a certain point, at around 4% effort, the curve becomes flattish. That is a point where you find out that the Machine starts to

produce results you don't like. So, you will know when it goes wrong, and moreover, the interactive nature of the process will allow you to correct the Machine.

## How much time does all that take?

You may think: All that sounds nice, but it also sounds like you have to invest a lot of time in teaching the machine, right? Well, not at all, actually. Why? Because with Imprima AI, the Machine Learning is integrated into your normal DD workflow, in your Virtual Data Room. As a result, you do not have to spend any additional time teaching the Machine or monitoring the results; it learns while you are doing your normal job.

We believe that it would not work for you otherwise. If you would have to invest in training the Machine, which involves the tedious process of labelling a training data set, you are probably not going to do it, as also argued in a post on a well-known blog on legal tech ( <https://www.artificiallawyer.com/2019/11/11/legal-innovators-stuart-whittle-weightmans/>).



## OK, so it doesn't cost me time, but how much time does it save?

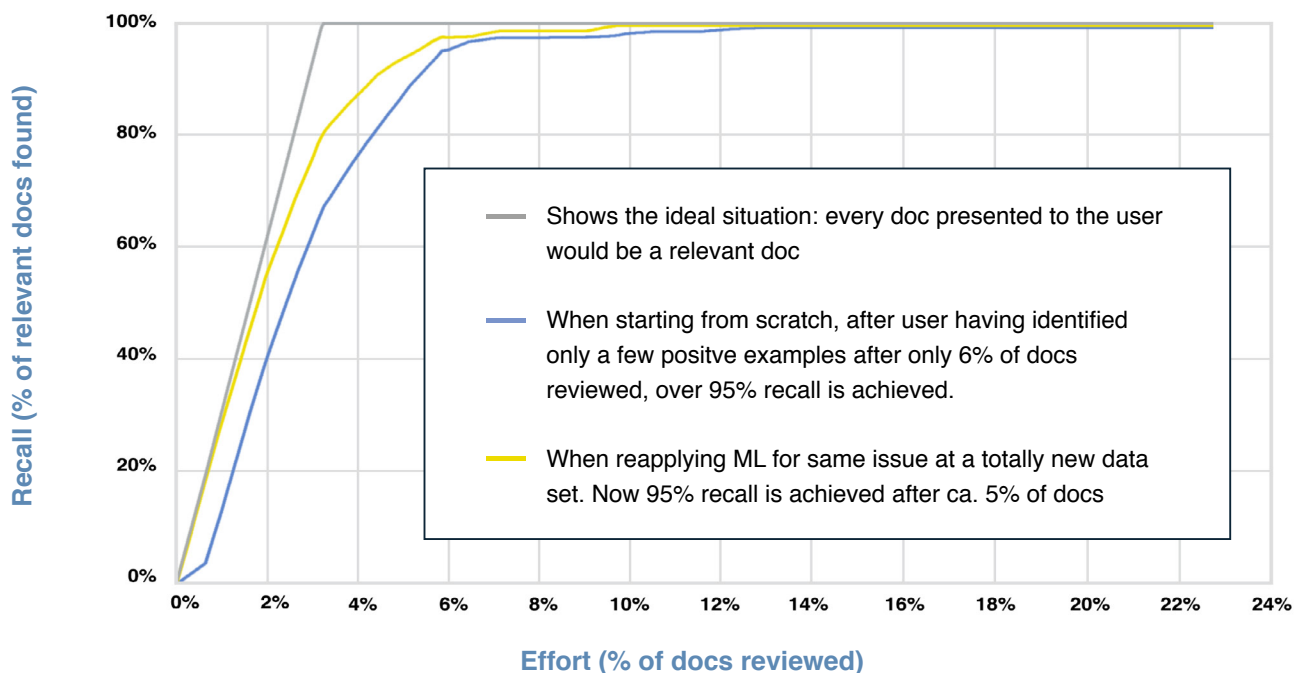
Of course, it is nice that the Machine learns while you are doing your normal review work, so you can add value without it costing you time, but does it also save you time? Let's consider the following example. Suppose you need to inspect credit agreements, and in particular identify those credit agreements with a Change of Control clause (**See again the recall vs effort graph below**).

Let's first consider the theoretically ideal situation, that each document that is presented to you is a relevant document. That situation is represented by the grey line ("ideal"). So let's see how close the Machine can get to that ideal situation.

Again, the blue line ("first run") is when training the Machine from scratch. It starts out flat, reflecting the initial learning phase. However, after having reviewed the first few docs (as you normally do) the curve becomes steep, almost as steep as the ideal curve (grey line). In fact, after reviewing only 6% of the documents, 95% of the relevant docs have been retrieved.

The yellow curve ("second run") again shows the application of a previously saved query on a new data set. It is now even steeper, almost identical to the grey, ideal, curve. Now 95% recall is achieved even quicker, at around 5% effort.

Note that, although it is significantly faster to get to 95% recall with a pretrained query (yellow curve) than with an untrained query (blue curve), in both cases a tremendous amount of time is saved: only 6% or 5% of documents have to be investigated to find 95% of relevant docs, or, in other words, you don't need to investigate respectively 95% or 96% of the docs that would otherwise have needed to be investigated. Also note that both curves are very close to the ideal situation, (represented by the grey line). In essence, it is fair to say that this is as good as it gets.



## Conclusion

M&A due diligence leads to the making of decisions that carry large consequences. Getting something wrong or missing a key red flag can have impactful and long-lasting effects that are sometimes hard to overcome. So, if you are going to rely on the use or support of AI, you need to understand, if not how it works, how it behaves, and whether it behaves in a predictable way.

We've ensured that our ML technology plays a supplementary role to the due diligence process; it's not designed to replace any specific human activity (or add new ones!). We take current review processes, and boost both their speed and accuracy. Our technology can spot things and bring them to the attention of a reviewer, but it's still the human reviewer that provides the ultimate analysis of that contract, clause or sentence.

We've also ensured that the tool is constantly informing you of what it thinks you want, meaning that you as a human reviewer, can constantly review and then shape its output. This provides you with a basic understanding of how the tool works – it's mimicking you and you continuously monitor its understanding of your behaviour. Before you know it, you have a trained apprentice that increases both your speed and accuracy of review in due diligence!

We've made all this possible through our approach to building ML tech for due diligence, which includes:

1. **Works right out of the box: No need to “label” the data and train the machine with test data before you get started.**
2. **Intuitive user interface modelled after popular websites (often people experience what they are used to as intuitive).**
3. **Conduct your AI-powered review in a VDR: the AI tools are an integral part of your VDR, and therefore integrated into your normal DD workflow. In other words: document review as you are used to and the power of AI in one integrated environment.**

We will discuss the above features further in a next white paper.

In the meantime, and in conclusion of this paper, one could ask: Can you trust the AI that's driving improvements in M&A due diligence? Well, it depends on what AI or ML you're using, but we certainly think it's possible.

[www.imprima.com](http://www.imprima.com)

## Contact us today to find out more

---

### London

Tel: +44 20 7965 4700  
E: [londonsales@imprima.com](mailto:londonsales@imprima.com)

### Paris

Tel: +33 1 58 36 06 60  
E: [paris@imprima.com](mailto:paris@imprima.com)

### Milan

Tel: +39 0694 803 478  
E: [milan@imprima.com](mailto:milan@imprima.com)

### Frankfurt

Tel: +49 69 915 0980  
E: [frankfurt@imprima.com](mailto:frankfurt@imprima.com)

### Amsterdam

Tel: +31 207 155 600  
E: [amsterdam@imprima.com](mailto:amsterdam@imprima.com)

### Brussels

Tel: +33 1 76 75 32 91  
E: [brussels@imprima.com](mailto:brussels@imprima.com)

## Official London Stock Exchange, Marketplace Provider

---



**London**  
Stock Exchange

Imprima is proud to have been selected by the London Stock Exchange to provide due diligence services to corporations listed on the Exchange. This reflects the quality of our technology, services and customer support.