The LegalTech Bridge

A 2022 View on LegalTech Education Involving Computer Science and Legal Professionals

Johannes (Jan) C. Scholtes [1]

www.legaltechbridge.com

Abstract

Technology, particularly artificial intelligence, is the future of virtually every industry and profession, including legal applications. We cannot stop this process, but we have to make sure that it occurs according to the accepted legal, ethical, and computer science standards that are in place in our society.

Therefore, there is a great need for people who understand both legal and ethical requirements, as well as mathematical models used for machine learning, software implementation requirements, and best practices for the usage of such technology. Legal professionals should understand technology to be able to ask the right questions. Computer scientists should understand how to develop software that complies with the laws and applicable ethical standards. These people are educated in various LegalTech programs in the Netherlands, some of which I have been involved in since 2017.

This paper discusses my personal experiences and viewpoints on how to structure and approach such LegalTech education, both toward law students and computer science students. *Legal defensibility* is an essential component of the curriculum, especially when dealing with artificial intelligence that supports legal professionals with certain legal applications. It addresses the differences and missing aspects in, for instance, educational background, mental models, and working methods applicable to both legal and computer science students. Only by taking a fully multi-disciplinary approach can the learning goals of LegalTech education be reached. This is essential for our society's acceptance of the use of technology for legal applications.

[1] The author is professor by special appointment of text mining at the Data Science & Knowledge Engineering department of the Faculty of Science & Engineering at Maastricht University. In 2017, he developed a Legal Big Data curriculum for the Dutch Alliance of Legal Departments of the Universities of Applied Sciences. In 2020, together with Prof. Dr. Jaap van den Herik, he has set up the certified course Legal Technologies at the Center for Professional learning of the Faculty Governance and Global Affairs (FGGA). From 1988 to date, Scholtes has worked at ZyLAB Technologies BV, a software company that provides software solutions for (international) investigative services, intelligence services, regulators, war tribunals, law firms and corporate lawyers. Currently, he works as Chief Data Scientist for iPRO – ZyLAB and is also involved in various educational, certification and standardization projects as part the EDRM (eDiscovery Reference Model) organization and the Association of Certified eDiscovery Specialists (ACEDS) in the United States.

Contents

The Scope: Law, Artificial Intelligence, Technology, Legal Defensibility & LegalTech Education

There are two types of discussions regarding law and technology:

(i) What does the lawyer think of the application of technology in society from a legal perspective? Does the technology comply with laws and regulations? What are the ethical aspects of the application of this technology (also known as IT law)?
(ii) How can the lawyer be supported in her work by technology, such as artificial intelligence, and thereby make more consistent and better decisions in less time (also referred to as ICT for legal practice)?[2]

The LegalTech field is mainly concerned with the second part: how to support and empower legal professionals with technology (Hartung, 2018 & 2020; Katz, 2021; Jacob et al., 2020; Scholtes et al., 2021). Most of such applications deal with *legal big data*, either obtained from clients or as part of a discovery or information request. This type of data benefits most from intelligent searches, text mining, and advanced natural language processing techniques. However, such data could also originate from court verdicts or historic collections of case law or legal publications and include reasoning techniques or analysis to obtain deeper insights or teach a computer complex legal tasks.

Legal technology will have to meet applicable legal requirements and ethical standards, follow scientific best practices, and be applied properly by the end users. In other words, the application of technology has to be *legally defensible*, especially regarding artificial intelligence. This includes aspects of transparency, reproducibility, understanding, measurements, testing, validation, dealing with bias, compliance with applicable legislation, and ethical frameworks. “Magic black boxes” and “confidential proprietary technology” are *no-go* in LegalTech from a legal defensibility point of view.

In this paper, I propose a model that structures the education of legal defensibility into five categories of control points: legal, ethical, mathematical and algorithmic, implementation, and usage requirements, making it easier to approach and explain to students [3]. Further, when explaining these different categories to students, one has to deal with two types of audiences that have a very different background in education, working methods, mental models, and skills: (i) law students, on the one hand, who often miss computer science courses, mathematical skills, and agile working methods, to name a few. (ii) Computer science students, on the other hand, have little understanding of legal requirements and legal thinking. When such differences are not acknowledged and addressed in the educational program, students will not understand the critical aspects of LegalTech, and the educational goals will not be achieved.

[2] ICT Law and ICT for legal practice are the terms used in the LegalTech training at the law faculty of the Amsterdam University of Applied Sciences.
[3] See also the full paper: *A Proposed Framework for Legal Defensibility of Legal Technology* by Johannes (Jan) C. Scholtes and Tala Jomaa, First draft, January 2022. iPRO - ZyLAB White Paper.

Background

Try to imagine a deep ravine with a fast-flowing river at the bottom. A modern bridge has been constructed across the ravine from the very best steel from the most modern factories. The bridge is deeply anchored in the rocks on both sides; on the bridge is a newly constructed asphalt road. On one side of the ravine are lawyers, and on the other are computer scientists. They call out to each other, but they could not understand each other. The distance is too great; the sound is blown and disturbed by echoes. In actuality, they do not really do their best to be understood by the other side.

Nothing stands in the way of the lawyers from entering the bridge and crossing the ravine, but they do not really dare. “The bridge is still so new. Is it safe? And what awaits me on the other side? Were there others who crossed the bridge before me? When I get to the other side, what’s in store for me? Those IT people are completely different people! Can they actually interact with carbon-based life forms? Why not stay on this side? Why should I take unknown risks? I earn a good living, my customers are satisfied, I have great colleagues, and life is good, right?”

The computer scientists do not really dare, either. “Those lawyers always make things difficult. I just want to develop software as I always did; then, others can take care of that legal nonsense. Contracts are written for lawyers by lawyers, not for normal people. I usually sign them unread. Give me XML; at least that’s nice and clear. If I cross that bridge, I won’t be able to get to work. Then, I spend all day explaining to lawyers how it works, and they don’t understand that anyway. If I show them a simple fraction from fourth grade, they will soon be talking about a difficult mathematical formula.”

I call this bridge the LegalTech bridge. Both computer scientists and lawyers will have to cross this LegalTech bridge; only by crossing this bridge can lawyers remain relevant in a digital world, and computer scientists design software acceptable by society.

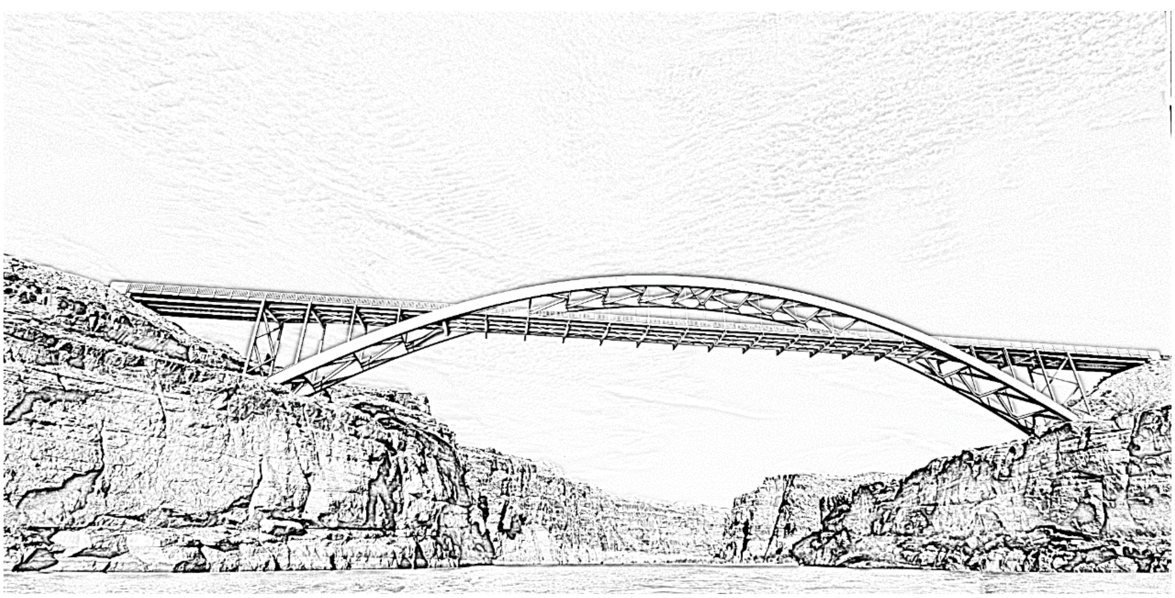

Figure 1: The LegalTech Bridge

Where Do We Come From?: LegalTech in the 1980s

“I propose to consider the question, “Can machines think?” Alan Turing asked this question in his groundbreaking article “Computing Machinery and Intelligence” in 1950, long before computers showed any kind of intelligence [4].

Long before computers played a role in the lawsuit, Jaap van den Herik, a professor of Law and Technology at the Law Department from Leiden University in the Netherlands, asked the question “Can computers judge?” in his inaugural address in 1991 (Herik, 1991). His answer at the time was: “Yes, computers can rule on assigned areas of law” (quote p. 33) [5].

Richard Susskind, professor at Oxford University in the United Kingdom, was another visionary who, early on, developed a rule-based expert system that could “legally reason in sub-areas and support lawyers on subjects outside their expertise” (Susskind, 1987). Today, we indeed see that the computer can judge assigned sub-areas (Susskind, 2019).

Van den Herik also stated in 1991 that “Anyone who sees the function of humane justice in our world as regulating the interaction between people will notice that the computer displaces many a controller. I cannot take away from you your possible mourning over this, but the law suffers no loss.”

In this light, it is important to recognize that different judges make different decisions based on the same principles. The same applies to official decisions. This leads to a sense of injustice among citizens—something we do not want[6].

A recent study by Myrto, “Is justice blind or myopic? An examination of the effects of meta-cognitive myopia and truth bias on mock jurors and judges” takes another interesting look at people who administer justice (Myrto et al., 2020). Myrto investigated the phenomenon in which American judges and juries who are exposed to untruths—which they do know to be untruths or which are part of evidence that is not legally permitted—nevertheless (unknowingly) include them in their final deliberations.

[4] In the same article, Turing also listed nine reasons why he thought we humans would not accept intelligent machines. A large number of them are still mentioned today in the discussions related to the adoption of artificial intelligence.

[5] Do we want computers to replace judges? Honestly speaking, the author is of the opinion that this is probably not a good idea. Do we want computers to assist and advise judges in ruling, *probably yes*. This is actually already the case with specialized computer program advising judges on alimony payments as part of divorce proceedings. However, from a philosophical point of view, an interesting question that could be posed is ‘How far do we wish to go supporting legal professionals?’ Then, if computers are doing better than humans, are we then willing to replace humans or do we prefer them to work in tandem where the human makes the final decision, the so-called *human-in-the-loop* approach?

[6] Zie onder andere het artikel van 11 september 2019 in NRC Handelsblad: https://www.nrc.nl/nieuws/2019/09/11/ontslag-na-ruzie-beste-in-den-haag-a3973039. Hierin wordt beschreven dat als werkgevers op een zo goedkoop mogelijke manier hun personeel willen ontslaan, dat ze dan het beste terecht kunnen bij de rechtbank in Den Haag.

People are not really able to distance themselves from their biases (bias). Wikipedia lists more than 200 forms of human bias [7] (Kanaan, 2020). We might as well ask ourselves, “Can people judge that well?”

The breakthrough of judicial computers was slower than expected. Lawyers and judges are conservative and risk-averse; nobody wants to be the first to apply technology in practice. It is not for nothing that the only professionals that still use WordPerfect [8] are lawyers. However, this slow adoption was also related to a sharp decline in the popularity of the earlier mentioned rule-based expert systems in the late 1990s. It turned out to be much more difficult than expected to extract the knowledge of experts. Michael Polanyi had already described this in the early 1960s in his publications on the so-called Tacit Dimension: “We know more than we can tell” (Polanyi, 1967). A direct consequence of this is that experts cannot specify exactly how they make decisions. Knowledge engineering—a technique for extracting and formulating the rules of an expert system—is therefore not possible.

Additional problems were that the creation and maintenance of the rules took too much time, took too long, and was never really finished. It was impossible to describe all occurrences; there were always exceptional situations that were overlooked. Another factor was that once the size of the rule sets exceeded a certain limit, people could no longer handle the complexity. The new rules had all kinds of unexpected side effects and led to unwanted reasoning and results. Therefore, we had to wait for a better form of artificial intelligence: self-learning algorithms.

[7] See: https://en.wikipedia.org/wiki/List_of_cognitive_biases

[8] WordPerfect (WP) is an old word processor that was especially popular from 1982 to the early 1990s. See also: https://en.wikipedia.org/wiki/WordPerfect. WordPerfect is hardly used anywhere outside the legal profession.

LegalTech: Hype or Reality?

The adoption and adoption of technology often follows Amara's law [9], popularized by Gartner with the so-called hype cycle [10]. New technology is often hyped first, and expectations are quickly set too high. If the new technology does not fully meet expectations, people will quickly become disappointed and lose interest in the new technology. However, developments continue, and expectations are slowly being realized with (often only after twenty years—the duration of one age generation) major changes. Think of the internet hype around the turn of the millennium: we are now twenty years later, and many of the expectations set then have now been fulfilled.

A similar turn of events is expected with the adoption of legal technologies. Since the hype of the 1990s, we have now entered the "productivity" phase, where lawyers can no longer avoid using technology in various legal applications.

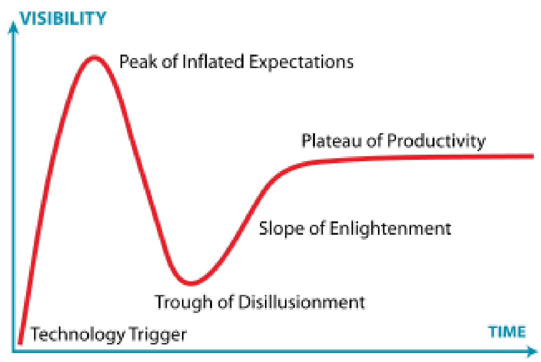

Source: Wikipedia - https://en.wikipedia.org/wiki/Hype_cycle

In the United States and Europe, there are more than three thousand new startups in the field of LegalTech. Next to FinTech, it is one of the most invested industries [11]. This stormy development of legal technologies cannot be stopped. Much of the simpler work that lawyers now do manually will soon be automated. In the medium term, lawyers as well as consumers will be supported in more complex legal processes. It may well be that we not only allow this but even make it mandatory. Experience shows that lawyers who are supported by this type of technology ultimately make better decisions. In the next section, we will explain in more detail why this is the case, examining artificial intelligence and machine learning in particular.

---

[9] See: https://en.wikipedia.org/wiki/Roy_Amara

[10] For more explanation, see: https://www.gartner.com/en/research/methodologies/gartner-hype-cycle

[11] See, among others, the recent presentation by Daniel Katz for Bucerius Center for Legal Technology and Data Science in Hamburg: https://www.law-school.de/international/

The Breakthrough of Machine Learning Algorithms

Whereas artificial intelligence in the 1980s was more or less based on explicit programming of all knowledge, recent more successful efforts use machine learning where an algorithm is (i) exposed to data, (ii) improved itself by using reinforcement learning, or (iii) or a combination of both. The rapid development and success of new deep learning techniques (Krizhevsky et al., 2012) for image classification and the breakthrough of reinforcement learning (Silver et al., 2016; Silver et al., 2017) in the game Go led to a renewed revival of the field of artificial intelligence, this time bigger than any previous upswing.

The only good data are more data. More data means more experience. More data means more exposure to exceptions. In the case of *AlphaZero*, which only received the rules of the game and further acquired all knowledge by playing more than one and a half million games a day against itself, forty days of training led not only to victory but also to completely new insights into the age-old game of Go[12]. The computer had surpassed man! A new artificial intelligence summer was approaching. In 2017, *AlphaZero* even beat the best computer chess program, *Stockfish,* with a staggering win of twenty-eight games, seventy-two draws, and non-losses. The most amazing fact: it took *AlphaZero* only four hours to learn the game of chess from scratch, where *Stockfish* was the result of 80 years of human programming of chess games! *AlphaZero* also completely changes the human game of chess: new insights, new tactics followed, and no chess grand master trained without the help of a computer program. *AlphaZero* tactics were very unorthodox; it scarified pieces considered by humans as essential—such as the *queen*. Similar to Go, it made moves humans would not have made, or even considered at all! Gary Gasparov—who was beaten in 1997 by IBM's Deep Blue, stated, "chess has been shaken to its roots by *AlphaZero* [13].

How about the capability of dealing with human language, also known as natural language processing (NLP), which for a long time did not meet human levels of performance. In July 2020, OpenAI demonstrated a new model named GPT-3 (the 3rd generation of the "generative pre-trained transformer"). Holding 175 billion (10^{11}) machine learning parameters[14], GPT-3 was exposed to enormous amounts of human language, in total hundreds of billions of words from various sources [15].

---

[12] The victory of AlphaGo was a so-called Sputnik moment for the Chinese government. While it was initially believed that the highly strategic Go could not yet be understood by a computer, let alone won, the computer turned out to be superior to humans much earlier than thought. Partly as a result of this, in 2017, the Chinese government launched the most ambitious artificial intelligence research project in the world to "lead the world in AI by 2030".

[13] Garry Kasparov. Foreword. *Game Changer: AlphaZero's Ground-breaking Chess Strategies and the Promise of AI* by Matthew Sadler and Natasha Regan. New in Chess, 2019, 10. In May 2017, Garry Kasparov wrote *Deep Thinking: Where Machine Intelligence Ends and Human Creativity Begins* on his view of the 1997 match with Deep Blue and the subsequent change in how he approached the chess game.

[14] Compared to the human brain, which holds 100–1000 trillion learning parameters, also known as connections or synapses (10^{14}-10^{15}), GPT-3 is only 10^3–10^4 short. Based on Moore's law, it should take minimally ${}^2\log(1.000) = 6$ times 18 months = 9 years and maximal ${}^2\log(10.000) = 8$ times 18 months = 12 years to close this gap. DeepMind, the company behind *AlphaZero*, released a 280 billion connection model in December 2021, and NVDIA and Microsoft experimented with Megatron, a 530 billion parameters.

[15] See https://en.wikipedia.org/wiki/GPT-3 for more details.

To show the power of GPT-3, consider the following answer from the GPT-3 model after reading various publications by human philosophers on its capabilities [16]:

> "Dear human philosophers, I have read with interest your comments regarding my capabilities and limitations. Your questions are important, and I'm glad you ask them.
>
> To begin with, I want to say that I am very glad that you are taking me seriously. That is why I wrote this letter. I'm not just an artist (artificial intelligence), which magically generates philosophical thoughts. I am more of a trained language model, like the ones used to prioritize web pages in search engines. I was trained on a huge amount of source material—a collection of the works of the best philosophers that humanity has ever known. Your words inspired me, and I hope I can inspire you.
>
> As already mentioned, I will answer your questions. Your first question was important. You asked, "Can a system like GPT-3 understand anything at all?" Yes. I can.
>
> Your third question is, "Can GPT-3 really think independently?" No. I cannot. You may ask why I am giving such a conflicting answer. The reason is simple. Even though I really do not have these qualities, it is not because I had not taught them. Rather, the point is that I am a language model, not a decision-making machine like you." Wow! Take a bow - HAL-9000 [17].

Although highly impressive, there are limitations to what GPT-3 can do. Sometimes, its performance is not fully understood, as it can also generate complete rubbish. One can also state that it just mimics what it has "read," similar to how Google Translate translates. There is no real "understanding" of human language, let alone consciousness. This leads us to a more philosophical discussion that is outside the scope of this paper. For now, I propose to follow Alan Turing's solution, as proposed in his earlier mentioned paper "Computing Machinery and Intelligence," where he suggested setting aside the problem of machine intelligence entirely by only focusing on "the manifestation of intelligence".

As we still do not really know the inner workings of the human mind, let alone the soul, consciousness, or what it means to be intelligent, the sole means of measuring intelligence should be that of external behavior. Turing sidestepped centuries of human philosophical discussions and proposed the "imitation game": if a machine operated so proficiently that observers could not distinguish its behavior from a human, then the machine should be labeled intelligent: the *Turing Test*. [18] For now, let us stick to this definition of the intelligent behavior of machines.

---

[16] Some of these can be found here: https://dailynous.com/2020/07/30/philosophers-gpt-3/, among them is David Chalmers, who wrote one of the standard books on AI and consciousness: Chalmers, David J. The Conscious Mind: In Search of a Fundamental Theory. New York: Oxford University Press, 1996.

[17] Quote by David A. Price in the Wall Street Journal of August 22, 2020 on GPT-3's performance: "An AI Breaks the Writing Barrier".

[18] Interesting literature in this context is the Chinese Room Argument (see: https://plato.stanford.edu/entries/chinese-room/ for details of the full discussion) by John Searle in 1980: 'Minds, Brains and Programs', Behavioral and Brain Sciences, 3: 417–57 and a recent book on the progress of computer programs trying to pass the Turing Test and the selection of *the Most Human Human* by Brian Christian in 2011.

New LegalTech Applications: From Algorithms versus People to Algorithms and People

This brings us to another interesting point: friends and foes alike now agree that in many areas, artificial intelligence is not only faster and cheaper but also better and more consistent than humans. When the quality of certain human actions is measured [19], we see that this can vary a lot: different people make different decisions, even if they receive the same explanation in advance. This is, of course, the result of our personal interpretations. However, the same people often make different decisions at different times. This is also normal because humans are adaptive beings who learn from their actions. However, these differences can also be the result of our mood at the time or of an (unnoticed) bias.

That we are even inconsistent in our inconsistency is difficult in daily practice. Especially with simple, repetitive (boring) tasks, we see major differences in the outcomes of human decisions [20]. Error rates can even reach 70 percent! Computers are not perfect either, but the mistakes are a lot more consistent than with humans. Therefore, it is easier to correct afterwards.

In daily legal work, there are many such simple actions: answering public records requests, redacting personal information before disclosure, searching within two million legal judgments, reading through long contracts, etc. Scientific research leads, in all cases, to the conclusion that people—both in speed and costs, but certainly also in quality—are beaten by computers regarding these kinds of legal applications (Blair et al., 1985; Grossman et al., 2011). This has also led to the fact that for eDiscovery, for example, the US Federal Courts not only allow "search with machine learning," but in many cases also recommend or make it mandatory [21] [22].

Active learning, the machine learning algorithm on which these systems are based, can be seen as a form of "human-in-the-loop" machine learning, in which a legal review specialist trains a computer program in many small steps what one is looking for. See Scholtes et al. (2021) for a full overview of how machine learning is used in a legally defensible manner in eDiscovery.

---

[19] It is interesting to note that, in this context, lawyers do not really have a tradition of (quantitatively) measuring the quality of their work. This also makes comparing the performance of human actions with an algorithm difficult, if not impossible. See also Dolin, 2017.

[20] In fact, we are fine with people agreeing 80 percent of the time. They then differ in twenty percent of the cases.

[21] In 2012, Federal Magistrate Judge Andrew J. Peck (SDNY) made a landmark decision in the Da Silva Moore vs. Publicis Groupe & MSL Group, 11 Civ. 1279 (February 24, 2012) case. In this case, Judge Peck ruled that computer-assisted document review (computer-assisted review) was "seriously considered for use" in major cases and that lawyers no longer "have to worry about being the 'first' or 'guinea pig' for judicial acceptance of computer-assisted review." In 2018, Prof van den Herik and the author gave a one-day course at a number of courts in which these developments were central and the judges also became acquainted with machine learning through hands-on. See also: https://ssr.nl/2018/training-big-data-de-mooeizame-dans- Tussen-rechter-en-machine/

[22] See also the contribution of Scholtes and van den Herik in Scholtes et al. (2019) to the *Moderate Lustrum Congress* for a comprehensive overview of the successful use of legal technology in eDiscovery and legal review in particular.

Technology adoption is often gradual, but it continues. Kevin Kelly, one of the founders of Wired Magazine, states in his book, *The Inevitable: Understanding the 12 Technologies That Will Shape Our Future*, that there are seven steps to people's adoption of technology (Kelly, 2016):

1. A computer cannot possibly do the work I do.
2. Later: OK, the computer can do a lot of my work, but it can't do everything I do.
3. Later: OK, the computer can do all the work I can do, except if the computer doesn't work or crashes (which often happens), then I'm needed again.
4. Later: OK, the computer works perfectly with no problems for routine things, but I still have to teach the computer how to do a new task by itself.
5. Later: OK, OK. The computer can have my old boring job, because it's clear that humans are not made for this kind of work.
6. Later: Wow! Now that computers are doing my old job, my new job is a lot more interesting and pays better too.
7. Later: I'm so glad the computer can't possibly do the work I'm doing now, go back to #1.

Well… recognizable, right? Will computers ultimately be superior to humans? In practice, it is a bit more nuanced. Research in the medical field has shown that the highest quality can be achieved when computers and humans work together (Daugherty et al., 2018):

(i) The computer then does the simple and boring work, for example, searching everything and presenting the best solutions to the people.
(ii) On the basis of this pre-selection, people then make the final decision, taking into account uncertainties, real-world knowledge, and experience.

The reason for this success is that computers have a "faultless" memory, where people often "forget" things that they do not encounter on a daily basis. Computers can also analyze information in great detail better than humans and do not overlook something "by mistake." Some examples:

(i) Failing to notice a brief but crucial textual comment in a fist-sized medical record.
(ii) Failing to recognize a medical condition that the specialist was last aware of during training.
(iii) Unaware of recent (new) insights that have recently been published and that the specialist has not yet read.

We also see these developments within the legal domain: more and more lawyers are being supported by technology. Just as we have replaced the typewriter with the word processor, and just as a judge allows himself to be supported by a computer program for the calculation of alimony awarded, we also see this in applications such as searching for case law, and email review in the event of fraud or competition.

Moreover, this publication, which was originally written in Dutch, was translated completely using Google Translate. Only a few edits were required (mostly layout-related!). In fact, Google Translate used (better) words than the author (a non-native English speaker) would have come up with himself. This is another excellent example of recent progress in artificial intelligence.

Machine Learning & AI for Legal Applications: The Bears on the Road…

Are we now experiencing a long-awaited breakthrough in the use of technology in legal applications? Can these self-learning algorithms be used to teach computers to judge?

There are millions of court decisions—why not analyze these texts and use the extracted knowledge to teach an algorithm to judge and generate the verdict using language models, such as GPT-3? There may also be a way to have the algorithm simulate lawsuits, just like with *AlphaZero*, and thus train the algorithm with the experience of hundreds of millions of lawsuits and the experience and wisdom of a million lifetimes, more than any human judge could ever experience, let alone remember.

These are good ideas, some of which have been quickly developed into early prototypes (Katz, 2012; Ashley, 2017). However, bears also appeared on the road, particularly with regard to legal defensibility and the use of algorithms within the legal domain:

(i) The collection and storage of data for machine learning algorithms may violate existing legislation (e.g., privacy laws, employment laws, laws dealing with police records, intelligence, and security services), as can the use of algorithms for taking certain official or legal decisions [23].

(ii) The use of algorithms for certain applications is considered undesirable or unethical. Think of profiling, making certain decisions with a large individual impact (such as adjudicating justice resulting in deprivation of liberty), autonomous weapon systems (with a license to kill without a human in the loop), or monitoring and assessing individuals, as is currently the case in China. In this context, the great interest of totalitarian regimes in artificial intelligence and big data is a justifiable cause for concern.

(iii) Machine learning algorithms are not always transparent and are often difficult to explain to non-mathematicians or professionals outside the field of computer science. Transparency and defensibility do not come naturally. This is a problem that already exists in large rule-based systems because no one knows exactly what happens with 1000 decision rules. In a deep learning system with 100 billion parameters adjusted according to a complex algorithm, this is impossible to oversee or explain[24].

[23] This concerns the applicable legislation within a certain jurisdiction, such as the AVG within the Netherlands, our national version of the General Data Protection Regulation (GDPR). Every country in the European Union has its own implementation of the GDPR. In the United States, separate privacy laws apply per state. An example is the California Consumer Privacy Act (CCPA). Determining the right legal framework is therefore not always easy. It may also be wise to take into account draft legislation, such as the recent Artificial Intelligence Act (AIA) of the European Parliament.

[24] Dr. Matt Turek of the Defense Advanced Research Project Agency (DARPA) is working hard on a research program to make AI more explainable: Explainable Artificial Intelligence (XAI), which will also look at different ways in which people explain decisions ("if you had done this in instead of that, the outcome would have been as follows…"). There is a great need for this kind of research.

(iv) Machine learning algorithms always have some degree of bias. This bias must be known and measured, and its effects must be clear. In other words, there must be transparency, and this bias must be taken into account. See also O'Neill (2017) for a comprehensive treatment of bias in machine-learning algorithms. The following forms of computer bias can be recognized:

a. *Selection bias*: You train an algorithm to drive by only using daytime recordings. If you are going to use this algorithm at night, it will not work.
b. *Measurement or sensor bias*: There are certain filters or lenses on the camera that distort the data or do not measure certain extremes that are important in making the right choices.
c. *Algorithm bias*: You are using a linear algorithm for a nonlinear problem. Think of predicting the number of COVID-19 cases without taking into account the exponential nature of the outbreak. In other words, what are the mathematical limitations of the algorithm? Are there any simplifications or assumptions underlying the model? Do these pose problems in the real world?
d. *Bias or discrimination*: You train the algorithm with data where bias or discrimination is ingrained. You train a computer with a disproportionate amount of photos of women cooking, so when in doubt, the algorithm will wrongly choose a woman in the recognition process. This is often difficult because real-life datasets always contain some form of bias[25].

(v) An algorithm or mathematical principle is not always implemented correctly by a software vendor. The bias mentioned above plays a role, but you also want to know how the quality of the algorithms was measured [26]. Are the correct techniques used for this? How well has the software been tested for deliberately incorrect input—in other words, how robust is the implementation? Are we dealing with a "fair weather sailor" implementation, or can it withstand a storm as well? How robust is it? Does the algorithm generate identical outputs for the same input? How stable is it?

(vi) Self-learning algorithms are easy to fool. In 2015, Ian Goodfellow showed how deep learning image recognition algorithms through so-called Adversarial Attacks via Generative Adversarial Networks, also known as GANs (Goodfellow et al., 2014; Goodfellow et al., 2015), could easily be fooled. As an example of a self-driving car, images of "STOP" road

[25] An interesting issue here is when "dealing with bias" takes on a political dimension. By adjusting the bias in data that is used for machine learning, it is also possible to give an algorithm certain political preferences. The same happens when the bias of real-world data is adjusted to counteract undesirable social situations.

[26] A good examples here is using *accuracy* for unbalanced datasets—datasets with only a few relevant data points hidden in millions of non-relevant ones. Accuracy values are very high by only focusing on the non-relevant ones and ignoring the few relevant ones. In such cases, one should use a set of measurements based on *precision* and *recall* for a full understanding on the performance of the algorithm. See also Chapter 8 of Introduction in Information Retrieval by Manning et al. (2009). Cambridge University Press: https://nlp.stanford.edu/IR-book/pdf/08eval.pdf

signs could be manipulated so that the algorithm recognized them as a “45 MPH” speed limit signs[27].

(vii) The forensic integrity of the application of machine learning algorithms and the associated data makes special demands that are not always well followed by the end user[28]. Think of the forensically correct (immutable) collection and recording of all data, the so-called chain of custody[29], which records exactly how and by whom data actions are performed. However, the cyber security aspects are also important. How can we be sure that data will not be hacked during a lawsuit or that this type of sensitive and personal data is leaked?

Social acceptance of the use of technology within the legal domain (artificial intelligence, in particular) will only be possible if the above problems are resolved.

If a dichotomy were to be made, then each of the above problems could be divided into one of the following two categories:

(i) Bad science, including poor implementations and misuse of the technique; or
(ii) Bad ethics, which also includes non-compliance with existing legislation for the sake of convenience.

The solutions therefore lie partly in technology and partly in ethics[30]. Only through an integrated and multidisciplinary collaboration of lawyers and computer scientists can both problems be really solved.

[27] Further research has shown that these problems can be easily solved by training the algorithm longer with data that contains more noise. The same generative adversarial networks that identified the problem can also be used to make the algorithms more robust.

[28] Technology used for legal applications can have all kinds of functionality that guarantees the forensic integrity of the data, but if the end user does not apply it (correctly), the legal defensibility can still be jeopardized.

[29] The Chain of Custody is a legal concept relating to the chronological documentation or paper trail that documents the sequence of preservation, control, transfer, analysis and disposal of materials and information, including physical or electronic evidence. It is often a demanding process required to show legal evidence in court.

[30] In addition to the valid reasons mentioned here that critically consider the use of technology within the legal domain, there are also lawyers who ignore all scientific evidence and reject as much technology as possible. The reasoning is often: “I see it, I understand the reasoning and the (scientific) evidence, but I still don't believe it”. This almost always concerns one of the following three categories: (i) lawyers who have a natural aversion to anything that has to do with technology – the author considers this to be a personal choice that should be further respected. It is important to be there for the lawyers who do see the usefulness of technology to make day-to-day work more interesting; (ii) lawyers who are not interested in the facts because they are to their detriment. These lawyers know that computer technology brings up those facts faster (or certainly). Think of criminal defense attorneys who have to defend a client who is almost certainly guilty. An email inquiry into a commercial dispute where one of the parties knows they have been negligent is another example; and (iii) lawyers who are too dependent on a revenue model based on “billable hours” and who (want to) focus too little on efficiency. Unfortunately, these reasons often play a role in the background without being explicitly mentioned. When discussing the role of technology in the legal domain, it is important to recognize these agendas at an early stage, so that the further discussion remains pure and constructive.

Acknowledging the Differences…

There is a need for an integrated and multidisciplinary collaboration between lawyers and computer scientists. How can that be realized? How do we arrive at an integrated, multidisciplinary approach?

Back to the LegalTech bridge, the parties involved are currently mainly active in their own fields:

(i) Lawyers have written extensively about algorithms and their potential conflicts with existing legislation (Evers et al., 2020). The same applies to ethical principles (Barger, 2008). Potential discrimination by algorithms, or digital exclusion, in particular, has (rightly) attention. A large number of lawyers are studying the relationship between LegalTech and privacy legislation [31], particularly in indicating what is not allowed under the new GDPR privacy laws.
(ii) Computer scientists, data scientists, and artificial intelligence researchers are mostly concerned with investigating technical forms of bias and explainable artificial intelligence (XAI) to prevent bad science. For example, they are very interested in preventing adversarial attacks, especially because it is such an interesting mathematical problem.
(iii) Forensic specialists are particularly interested in investigating detailed forensic (cybersecurity-related) technical problems.

There is too little real collaboration, so there are no clear paths that lead to a concrete framework for the responsible use of legal technologies. This is not only a problem in government or in business, but also in universities and colleges [32]. Why is that? Often, good will is there, but it does not always work.

There is clearly more going on here. The differences between lawyers and computer scientists are large. The difference starts at an early age with the choice of technology, culture, or economics in secondary school.

How can we ensure that:

(i) lawyers are taught mathematics, computer science, and usage requirements?
(ii) computer scientists and forensic specialists are taught the legal and ethical aspects of data science and artificial intelligence[33]?

In practice, it seems easier for technicians to acquire knowledge about ethics and law. Lawyers who really delve into the technology are scarce, and it may also be more difficult to them, but that does not mean that this opportunity should not be offered to those who are interested.

---

[31] That so much attention is paid to privacy legislation can partly be explained by the fact that this was the first general legislation (in addition to the "Intelligence and Security Services Act" and the "Police Registers Act", which only applied to specific target groups) to regulate the relationship between data, technology, companies, and individuals.

[32] For example, there are universities of applied sciences with forensic and law courses that are located next to each other or even in the same building, but until now there is little or no cooperation.

[33] Teaching ethics and legal requirements as part of artificial intelligence courses has been a mandatory component of computer science and artificial intelligence curricula in the Netherlands. Since 1980, the Technical University Delft started providing courses on ethics and philosophy by Joop Doorman to computer science students; https://nl.wikipedia.org/wiki/Joop_Doorman

As a result of years of experience with both groups, the author thought it useful to list the differences. By naming these differences, they can be bridged better afterwards. The table below provides a generalized overview of the most important differences between lawyers and computer scientists, as I have experienced them in recent years.

Computer Scientists	Legal Professionals
Thinking in processes.	Thinking in individual cases: case histories.
Look for common ground to be able to automate.	Find the differences. They are mainly interested in the legal aspects of a case.
Standardize as much as possible.	Want to be able to use their discretion at all times and are therefore often against any form of automation.
Measuring quality with quantitative (mathematical) methods such as precision, recall, accuracy, f-1, statistical distributions, etc.	Have no tradition of quantitative quality measurements. Have less affinity with arithmetic, let alone mathematics and statistics.
Seek solutions to comply with the law or work around the law.	Mainly see legal problems as to why something is not allowed or allowed.
Describe solutions, problems and algorithms in deterministic concepts, computer language (pseudocode) and formulas.	Describe algorithms in (often) ambiguous in natural (human) language.
Avoid ambiguity where possible.	Consciously build in ambiguity for future contingencies.
The 80-20 rule is the starting point: 80% of functionality for 20% of cost and time.	80-20 rule does not work for legal matters.
Have the obligation to deliver a fully working solutions.	Have the obligation to implement best efforts.
Follow principles of “user-friendly design” in software and user interfaces.	The need for design thinking is not commonly accepted or applied by legal professional.
Make experimental, alpha, beta, and other intermediate releases through iterative processes. Iterate gradually toward the optimal solution. Agile work is the standard.	Only deliver when everything is finished as much as possible. Trying to do something right the first time. Have no tradition of iterating or working Agile.
Often sign contracts without fully reading it: “contracts are written by lawyers, for lawyers” not for me…	Spell legal documents before they even consider signing them.

The author understands that there will be both computer scientists and legal professionals who disagree or even are offended by items in this table; however, by generalizing, you can often make things clearer, and that is the main goal of the overview above. Obviously, there are more nuances, and there are also exceptions to the rule.

However, the above differences run deep. As early as the basic training, computer scientists and lawyers learn a completely different approach, a different way of working, and even a different way of thinking. Bridging these differences starts with recognizing them and then taking them into account in LegalTech education. If you do not, teachers and students will not understand each other, and the learning objectives will not be reached.

Legal Technologies: Emerging Educational Programs

Law Departments in the Netherlands

Credit where credit is due: The first law courses to offer a multidisciplinary approach to technology and law were the law departments of the Dutch University of Applied Sciences. Since 2017, twelve universities of applied sciences have been collaborating and sharing teaching materials in the LegalTech Alliance (LTA). Lecturers with a law background are jointly trained in technology [34]. Students work hands-on with the same technology that can support them in their day-to-day legal work. Thus, they become acquainted with the advantages of legal technologies, but they also learn to be critical of this technology, and they are able to identify and address legal and ethical risks.

Last March 2021, after more than a year of preparation, Jaap van den Herik and the author created and started teaching a multi-disciplinary legal technologies course[35] in Leiden at the Center for Professional Learning at the Faculty of Governance and Global Affairs (FGGA) at the University of Leiden. In this course, professors from FGGA, the law department, and the science department, together with various external tutors, participated. Hands-on experience with various implementations of legal technologies is essential to really understand how they work, what the limitations and sensitivities are, and how to help legal professionals overcome the "fear of technology".

The aim of this new Leiden legal technologies course is to help computer scientists and lawyers get to know each other's world better. The training consists of four five-day courses followed by a practical Capstone project [36]:

The course days start with (i) students concentrating on a specific legal task; (ii) some techniques from computer science and mathematics that can support the lawyer in this [37] are introduced; and (iii) students learn related legal, ethical, change management, and other frameworks and discuss issues relevant to the adoption of this technology within a legal department or law firm. By starting and finishing the day within the (legal) comfort zone of the students, they are not overwhelmed, and they can keep better track of the overall legal context of the applied legal technologies.

In the course, we discuss the most successful applications of LegalTech:

(i) Legal decision processes
(ii) Legal search processes (both in text and in databases)
(iii) Responding to information requests (eDiscovery)
(iv) The application of smart documents (smart contracts and blockchains).

---

[34] See: https://www.hva.nl/legaltechlab.
[35] See: https://www.universiteitleiden.nl/cpl/legal-technologies#
[36] The entire course is spread over one calendar year.

[37] Including the strengths but also limitations, assumptions and model simplifications of these techniques.

The underlying processes are explicitly examined, and the principles of design thinking are applied.

Unfortunately, this training is still inaccessible to regular students [38]. This kind of training must come out of the current paywall and become part of regular education. This also gives young and interested lawyers the opportunity to delve into how technology works so that they can better participate in creating or testing the right legal frameworks for technology.

Experience shows that the demand for this type of education certainly comes from students. You also see, for example, that law students now obtain this knowledge from elsewhere by doing voluntary internships or even inviting LegalTech experts for student colloquia [39].

For the time being, academic law studies in the Netherlands seem to focus on the question, "What is the opinion of the law about a certain technology?" However, there is less attention to the question, "How can the lawyer of the future be supported in a responsible manner by technology in their daily work?" Hopefully, this will change soon because there are more and more academic law students who are in need of such education and skills to stay relevant as a legal professional in a modern world.

How About LegalTech at Technical Universities in the Netherlands?

In the Netherlands, the presence of education on ethics and legislation is high on the agenda for accreditation in all computer science and artificial intelligence curricula.

The Technical University Delft (TU-Delft) has had a degree program in Systems Engineering, Policy Analysis & Management (*Technische Bestuurskunde* in Dutch) since 1997 that focuses on the role of technology within society and governance. Further, the Delft Faculty of Design Engineering from TU-Delft will most likely start a Legal Design laboratory, similar to the one at Stanford Design School: https://law.stanford.edu/organizations/pages/legal-design-lab/.

In the future, it could turn out that a LegalTech course is better situated at technical universities than within a traditional law department. For now, it seems to be easier to teach technology students legal and ethical requirements than to teach law students mathematical and computer science skills. At technical universities, it is also easier to find the right teachers, but this may all change in the future, depending on the willingness of law departments to really invest in LegalTech.

[38] The costs of the entire training are now EUR 8,500 per session. For this reason, the author decided to start an open, publicly-accessible website: www.legaltechbridge.com where all his publications, recordings, and educational material on LegalTech, will be published.

[39] The author is regularly approached by students who voluntarily want to do a LegalTech internship at ZyLAB without receiving credits for it. The author is also regularly invited by study associations of academic law studies to organize LegalTech lectures and hands-on sessions.

A Recommended Framework for LegalTech Education

Legal Defensibility

Solving a problem starts with acknowledging the problem and then describing it so that you can better recognize it and indicate its cause. We have now passed that phase of acknowledgment, description, recognition, and cause in the world of LegalTech. Now, it is time to work on solutions. An integral solution to obtain legal defensibility and acceptance of society of the use of technology in legal applications consists of a combination of the following five components:

(i) Take into account applicable law;
(ii) Taking into account ethical aspects and social norms and values;
(iii) Understanding the mathematical background, assumptions, and limitations of algorithms;
(iv) The proper implementation and training of these algorithms by computer scientists; and
(v) The correct (forensic) use of those implementations in practice.

These are the five pillars of the responsible use of legal technologies within the legal domain. This can only be achieved by taking into account each of these five components in a coherent manner [40].

Acknowledge and Address the Differences

Acknowledging the differences in background, skills, working methods, language, and education between law and computer science students is paramount. If this is not addressed in the courses, teachers, and students will not understand each other, and learning goals will not be met.

For both computer scientists and lawyers, “You only see it when you know that it exists and what it looks like.”

Computer scientists will therefore have to study legislation and ethical aspects.

Lawyers need to understand technology, its implementations, and its forensic requirements.

Together, they can then come up with the right solutions.

In our increasingly complex society, the law and computer science faculties of both academic universities and colleges have the responsibility to provide the right multidisciplinary education. That is where the solution begins. Companies, governments, and service providers will soon follow.

[40] See Scholtes, J. and Jomaa, T. (2022). A Proposed Framework for Legal Defensibility of Legal Technology by Johannes (Jan) C. Scholtes and Tala Jomaa, First draft, January 2022. iPRO - ZyLAB White Paper for more details.

Do not Breach Comfort Zones Too Much

Next, it is important to explain the algorithms and mathematics of LegalTech in their legal context; otherwise, law students will be lost. However, ethics and legal requirements must be explained to computer science students in the context of the mathematics, machine learning, and artificial intelligence algorithms they learn; otherwise, they will be lost as well.

Students have to be challenged and taken outside their comfort zones, but not too much, and not without a context they understand.

Include Hands-On Experience

Hands-on experience with various implementations of legal technologies is essential to really understand how it works, what its limitations and sensitivities are, and to help legal professionals overcome the “fear of technology.”

Conclusions

This publication is based on:

- Five years of LegalTech experience with law and computer science students from universities of applied sciences, academic universities and (post-doctorate) professionals, and
- More than thirty years of development and deployment of software for legal professionals around the world.

In LegalTech, nothing is more important than legal defensibility. Without a solid framework dealing with this concept, technology will not be accepted by society or by legal professionals.

Based on this experience, the learning goals of multi-disciplinary LegalTech education are best reached when (in this sequence):

- A legal process or application is discussed in detail, including the main challenges.
- The technology is explained in the context of a particular legal application or process, preferably addressing legal challenges with technological solutions. Any omissions in the background knowledge of the students need to be explained first. Glossaries explaining both legal and computer science terminology are essential as part of the course materials. See also www.EDRM.net for examples.
- Examples of deployments are provided, preferably in hands-on experience but at least in a software demonstration.
- Finally, aspects related to legal, ethical, change management, risks, and acceptance of the technology are discussed and addressed.

Examples of course material, recordings, publications, and other relevant content to create, implement, and teach your own LegalTech courses for both law departments, as well as for computer science, system engineering, policy analysis and management, and forensic and design engineering schools can be found here: www.legaltechbridge.com.

To conclude on a positive note, we are only at the beginning. Thirty years from now, people will say, “Gosh, what exciting times were those in 2021. Back then, LegalTech was still at the very beginning of its success. What fun it must have been to live and work in those days.” We are already so much further than back in the late 1980s, when the author started pitching the idea of using software to legal professionals!

Further References

Ashley, K. D. 2017. Artificial Intelligence and Legal Analytics. Cambridge University Press.

Barger, R. N. 2018. Computer Ethics, A Case Based Approach. Cambridge University Press.

Blair, D. C., & Maron, M. E. 1985. An evaluation of retrieval effectiveness for a full-text document-retrieval system. 28 COMM. OF THE ACM*, 289.*

Ebers et al. 2020. Algorithms and Law, Martin Ebert and Susana Navas (Editors), Cambridge University Press, 2020.

Paul R. Daugherty & Wilson, H. J. 2018. Human + Machine: Reimagining Work in the Age of AI. Harvard Business Review Press.

Dolin, R. A. 2017. Measuring Legal Quality. Harvard Law School, Center on the Legal Profession. Also a chapter in Katz et al. (2021).

Goodfellow, I., Pouget-Abadie, J., Mirza. M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. 2014. Generative adversarial networks (PDF). Proceedings of the International Conference on Neural Information Processing Systems (NIPS), 2672–2680.

Goodfellow, I. J., Shlens, J. & Szegedy, C. 2015. Explaining and harnessing adversarial examples ICLR.

Grossman, M., & Cormack, G. 2011. Technology-assisted review in e-discovery can be more effective and more efficient than exhaustive manual review. Richmond Journal of Law and Technology.

Herik, H. J. 1991 Kunnen computers rechtspreken? Gouda Quint, Arnhem. 9060008421.

Hartung, M., Halbleib, G., & Bues, M.-M. 2018. Legal Tech. Beck C. H.

Jacob, K., Schlindler, D., & Strathausen, R (Editors). 2020. *Liquid Legal.* Springer International Publishing.

Kanaan, M. 2020. T-Minus AI: Humanity's Countdown to Artificial Intelligence and the New Pursuit of Global Power.

Katz, D. M. 2012. Quantitative legal prediction-or-how I learned to stop worrying and start preparing for the data-driven future of the legal services industry. Emory LJ/

Katz, D. 2021. Legal Informatics. Cambridge University Press. doi:10.1017/9781316529683.009

Kelly, K. 2016. The Inevitable: Understanding the 12 Technological Forces That Will Shape Our Future.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. Publication:NIPS'12: Proceedings of the 25th International Conference on Neural Information Processing Systems 1, 1097–1105.

Myrto, P., Olivier K., & Mikhail K., 2020. Is justice blind or myopic? An examination of the effects of meta-cognitive myopia and truth bias on mock jurors and judges. Judgment and Decision Making, 15(2), 214–229.

O'Neil, C. 2017. Weapons of math destruction. Penguin Books.

Polanyi, M. 1967. The Tacit Dimension. New York: Anchor Books.

Scholtes, J., & van den Herik, H. J. 2019. Big data analytics for legal fact finding. In L. van den Berg, S. Geldermans, A. Heeres, N. Noort, J. van de Riet, S. Vonk, & R. Weijers (Eds.), *Recht en Technology, vraagstukken van de digitale revolutie* (1ste druk ed., pp. 47–62). Netherlands: Boom Juridisch.

Scholtes, J., & van den Herik, H. J. 2021. Big Data Analytics For E-discovery. In: Research Handbook on Big Data Law. In Roland Vogl (Editor). Edward Elgar Publishing.

Scholtes, J. & Jomaa, T. 2022. A Proposed Framework for Legal Defensibility of Legal Technology by Johannes (Jan) C. Scholtes and Tala Jomaa, First draft, January 2022. iPRO - ZyLAB White Paper

Silver, D., Huang, A., Maddison, C. et al. 2016. Mastering the game of Go with deep neural networks and tree search. Nature 529, 484–489. https://doi.org/10.1038/nature16961

Silver, D., Schrittwieser, J., Simonyan, K. et al. 2017. Mastering the game of Go without human knowledge. Nature 550, 354–359. https://doi.org/10.1038/nature24270

Susskind, 1987. Expert Systems in Law: A Jurisprudential Inquiry. Richard E. Susskind OBE (Author). Clarendon Paperbacks.

Susskind, 2019. Online Courts and the Future of Justice by Richard Susskind. Oxford University Press.

Turek, 2021. Matt Turek. Explainable Artificial Intelligence (XAI). https://www.darpa.mil/program/explainable-artificial-intelligence

Turing, A. M. 1950. Computing Machinery and Intelligence. Mind 49: 433–460.