

Artificial Intelligence as Law

Presidential Address to the Seventeenth International Conference on Artificial Intelligence and Law

Anonymous

the date of receipt and acceptance should be inserted later

Abstract Information technology is so ubiquitous and AI's progress so inspiring that also legal professionals experience its benefits and have high expectations. At the same time, the powers of AI have been rising so strongly that it is no longer obvious that AI applications (whether in the law or elsewhere) help promoting a good society; in fact they are sometimes harmful. Hence many argue that safeguards are needed for AI to be trustworthy, social, responsible, humane, ethical. In short: AI should be good for us. But how to establish proper safeguards for AI?

One strong answer readily available is: consider the problems and solutions studied in AI & Law. AI & Law has worked on the design of social, explainable, responsible AI aligned with human values for decades already, AI & Law addresses the hardest problems across the breadth of AI (in reasoning, knowledge, learning and language), and AI & Law inspires new solutions (argumentation, schemes and norms, rules and cases, interpretation).

It is argued that the study of AI as Law supports the development of an AI that is good for us, making AI & Law more relevant than ever.

1 Introduction

It is my pleasure to speak to you today¹ on Artificial Intelligence and Law, a topic that I have already loved for so long—and I guess many of you too—and that today is in the center of attention.

It is not a new thing that technological innovation in the law has attracted a lot of attention. For instance, think of an innovation brought to us by the French 18th century freemason Joseph-Ignace Guillotin: the guillotine. Many

No institute

¹ This text is an adapted version of the IAAIL presidential address delivered at the 17th International Conference on Artificial Intelligence and Law (ICAAIL 2019) in Montreal, Canada (Cyberjustice Lab, University of Montreal, June 19, 2019).



Fig. 1 Technological innovation in the law in the past (left) and in the future? (right). Left: Guillotine at the Nieuwmarkt in Amsterdam, 1812 (Rijksmuseum RP-P-OB-87.033, anonymous). Right: TV series Futurama, judge 723 (futurama.fandom.com/wiki/Judge_723).

people gathered at the Nieuwmarkt, Amsterdam, when it was first used in the Netherlands in 1812 (Figure 1, left). The guillotine was thought of as a *humane* technology, since the machine guaranteed an instant and painless death.

And then a contemporary technological innovation that attracts a lot of attention: the self-driving car that can follow basic traffic rules by itself, so in that sense is an example of normware, an artificial system with embedded norms. In a recent news article,² the story is reported that a drunk driver in Meppel in my province Drenthe in the Netherlands was driving his self-driving car. Well, he was riding his car, as the police discovered that he was tailing a truck, while sleeping behind the wheel, his car in autopilot mode. His driver's licence has been withdrawn.

And indeed technological innovation in AI is spectacular, think only of the automatically translated headline 'Drunken Meppeler sleeps on the highway', perhaps not perfect, but enough for understanding what is meant. Innovation in AI is going so fast that many people have become very enthusiastic about what is possible. For instance, a recent news item reports that Estonia is planning to use AI for automatic decision making in the law.³ It brings back the old fears for robot judges (Figure 1, right).

Contrast here how legal data enters the legal system in France where it is since recently no longer allowed to use data to evaluate or predict the behavior of individual judges:

LOI n° 2019-222 du 23 mars 2019 de programmation 2018-2022 et de réforme pour la justice (1) - Article 33

Les données d'identité des magistrats et des membres du greffe ne peuvent faire l'objet d'une réutilisation ayant pour objet ou pour effet d'évaluer, d'analyser, de comparer ou de prédire leurs pratiques professionnelles réelles ou supposées.

² 'Beschonken Meppeler rijdt slapend over de snelweg' (automatic translation: 'Drunken Meppeler sleeps on the highway'), RTV Drenthe, May 17, 2019.

³ 'Can AI be a fair judge in court? Estonia thinks so', Wired, March 25, 2019 (Eric Miller).



Fig. 2 A car breaching traffic law, automatically identified

[The identity data of magistrates and members of the registry cannot be reused with the purpose or effect of evaluating, analyzing, comparing or predicting their actual or alleged professional practices.]

The fears are real, as the fake news and privacy disasters that are happening show. Even the big tech companies are considering significant changes, such as a data diet.⁴ But no one knows whether that is because of a concern for the people’s privacy or out of fear for more regulation hurting their market dominance. Anyway, in China privacy is thought of very differently. Figure 2 shows an automatically identified car of which it is automatically decided that it is breaching traffic law—see the red box around it. And indeed with both a car and pedestrians on the zebra crossing something is going wrong. Just this weekend the newspaper reported about how the Chinese public thinks of their social scoring system.⁵ It seems that the Chinese emphasise the advantages of the scoring system, as a tool against crimes and misbehavior.

Against this background of the benefits and risks of contemporary AI, the AI community in the Netherlands has presented a manifesto⁶ emphasising what is needed: an AI that is aligned with human values and society. In Figure 3, key fields of research in AI are listed in rows, and in columns three key challenges are shown: first, AI should be social, and should allow for sensible interaction with humans; second, AI should be explainable, such that black box algorithms trained on data are made transparent by providing justifying explanations; and, third, AI should be responsible, in particular AI should be guided by the rules, norms, laws of society.

Also elsewhere there is more and more awareness of the need for a good, humane AI. For instance, the CLAIRE Confederation of Laboratories for AI Research in Europe⁷ uses the slogan

⁴ ‘Het nieuwe datadiet van Google en Facebook’ (automatic translation: ‘The new data diet from Google and Facebook’, nrc.nl, May 11, 2019.

⁵ ‘Zo stuurt en controleert China zijn burgers’ (automatic translation: ‘This is how China directs and controls its citizens’, nrc.nl, June 14, 2019.

⁶ bnvki.org/wp-content/uploads/2018/05/Dutch-AI-Manifesto.pdf

⁷ claire-ai.org

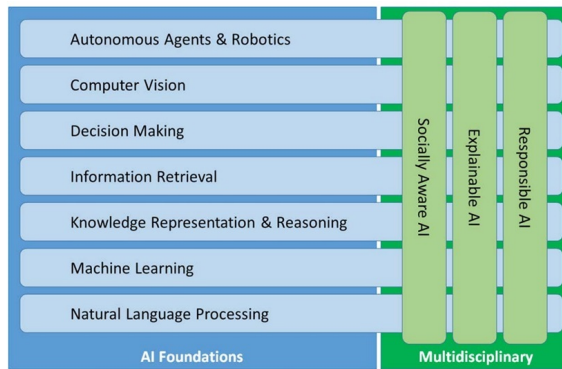


Fig. 3 Artificial Intelligence Grid: foundational areas and multidisciplinary challenges (source: Dutch AI Manifesto⁶)

Excellence across all of AI
 For all of Europe
 With a Human-Centered Focus.

In other words, this emerging network advertises a strong European AI with social, explainable, responsible AI at its core.

And now a key point for today: AI & Law has been doing this all along. At least since the start of its primary institutions—the biennial conference ICAIL (started in 1987 by IAAIL),⁸ the annual conference JURIX (started in 1988)⁹ and the journal *Artificial Intelligence & Law* (in 1992)—, we have been working on good AI. In other words, AI & Law has worked on the design of socially aware, explainable, responsible AI for decades already. One can say that what is needed in AI today is to do AI as we do law.

2 Legal technology today

But before explaining how that could go let us look a bit at the current state of legal technology, for things are very different when compared to the start of the field of AI & Law.

For one thing, all branches of government now use legal technology to make information accessible for the public and to provide services as directly and easily as possible. For instance, a Dutch government website¹⁰ provides access to laws, regulations and treaties valid in the Netherlands. The Dutch public prosecution provides an online knowledge-based system that gives access to fines and punishments in all kinds of offenses.¹¹ There you can for instance find out what happens when the police catch you with an amount of marihuana

⁸ iaail.org

⁹ jurix.nl

¹⁰ wetten.overheid.nl

¹¹ www.om.nl/onderwerpen/boetebase

between 5 and 30 grams. In the Netherlands, you have to pay 75 euros, and there is a note: also the drugs will be taken away from you. Indeed in the Netherlands all branches of government have online presence, as there is a website that gives access to information about the Dutch judicial system, including access to many decisions.¹²

An especially good example of successful legal technology is provided by the government's income tax services.¹³ In the Netherlands, filling out your annual tax form has become very simple. The software is good, it is easy to use, and best of all: in these days of big interconnected data much of what you need to fill in is already filled in for you. Your salary, bank accounts, savings, mortgage interest paid, the value of your house, it is all already there when you log in. In certain cases the tool even leaves room for some mild tax evasion—or tax optimisation if you like—since by playing with some settings a married couple can make sure that one partner has to pay just below the minimal amount that will in fact be collected, which can save about 40 euros.

One might think that such legal tech systems are now normal, but that is far from true. Many countries struggle with developing proper legal tech at the government level. One issue is that the design of complex systems is notoriously hard, and this is already true without very advanced AI.

Also the Netherlands has had its striking failures. A scary example is the Dutch project to streamline the IT support of population registers. One would say a doable project, just databases with names, birth dates, marriages, addresses and the like. The project was a complete failure.¹⁴ After burning 90 million euros, the responsible minister—by the way earlier in his career a well-recognized scientist—had to pull the plug. Today all local governments are still using their own systems.

Still legal tech is booming, and focuses on many different styles of work. The classification used by the tech index maintained by the CodeX center for legal informatics at Stanford university distinguishes nine categories (Marketplace, Document Automation, Practice Management, Legal Research, Legal Education, Online Dispute Resolution, E-Discovery, Analytics and Compliance).¹⁵ It currently lists more than a 1000 legal tech oriented companies.

And on the internet I found a promising graph about how the market for legal technology will develop. Now it is worth already a couple of 100s of millions of dollars, but in a few years time that will have risen to 1.2 billion dollars—according to that particular prediction. I leave it to you to assess what such a prediction really means, but we can be curious and hopeful while following how the market will actually develop.

So legal tech clearly exists, in fact is widespread. But is it AI, in the sense of AI as discussed at academic conferences? Most of it not really. Most of what we see that is successful in legal tech is not really AI. But there are examples.

¹² uitspraken.rechtspraak.nl

¹³ www.belastingdienst.nl

¹⁴ 'ICT-project basisregistratie totaal mislukt' (automatic translation: 'IT project basic registration totally failed'), nrc.nl, July 17, 2017.

¹⁵ techindex.law.stanford.edu



Fig. 4 Nederland ontwapent (The Netherlands disarm). Source: Nationaal Archief, 2.24.01.03, 918-0574 (Joost Evers, Anefo).

I don't know about you, but I consider the tax system just discussed to be a proper AI system. It has expert knowledge of tax law and it applies that legal expertise to your specific situation. True, this is largely good old-fashioned AI already scientifically understood in the 1970s, but by its access to relevant databases of the interconnected-big-data kind, it certainly has a modern twist. One could even say that the system is grounded in real world data, and is hence an example of situated AI, in the way that the term was used in the 1990s (and perhaps before). But also this is clearly not an adaptive machine learning AI system, as is today expected of AI.

3 AI & Law is hard


The reason why much of the successful legal tech is not really AI is simple. AI & Law is hard, very hard. In part this explains why many of us are here in this room. We are brave, we like the hard problems. In AI & Law they cannot be evaded.

Let us look at an example of real law. We go back to the year when I was born when pacifism was still a relevant political attitude. In that year the Dutch Supreme court decided that the inscription 'The Netherlands disarm', mounted on a tower (Figure 4) was not an offense.¹⁶ The court admitted that indeed the sign could be considered a violation of Article 1 of the landscape management regulation of the province of North Holland, but the court decided that that regulation lacked binding power by a conflict with the freedom of speech, as codified in Article 7 of the Dutch constitution.

An example of a hard case. This outcome and its reasoning could not really be predicted, which is one reason why the case is still taught in law schools.

The example can be used to illustrate some of the tough hurdles for the development of AI & Law as they have been recognized from the start; here a list used by Edwina Rissland [1988] when reviewing Anne Gardner's pioneering

¹⁶ Supreme Court The Netherlands, January 24, 1967: Nederland ontwapent (The Netherlands disarm).



Montesquieu (1689-1755):
The judge as 'bouche de la loi'

Fig. 5 The subsumption model

book 'An AI approach to legal reasoning' [Gardner, 1987], a revision of her 1984 Stanford dissertation.¹⁷ I am happy that both are present in this room today.

1. *Legal reasoning is rule-guided, rather than rule-governed.* In the example, indeed both the provincial regulation and the constitution were only guiding, not governing. Their conflict had to be resolved. A wise judge was needed.
2. *Legal terms are open textured.* In the example it is quite a stretch to interpret a sign on a tower as an example of speech in the sense of freedom of speech, but that is what the court here did. It is the old puzzle of legally qualifying the facts, not at all an easy business, also not for humans. With my background in mathematics, I found legal qualification to be a surprisingly and unpleasantly underspecified problem when I took law school exams during my first years as assistant professor in legal informatics in Maastricht, back in the 1990s. Today computers also still would have a very hard time handling open texture.
3. *Legal questions can have more than one answer, but a reasonable and timely answer must be given.* I have not checked how quickly the supreme court made its decision, probably not very quickly, but the case was settled. The conflict was resolved. A solution that had not yet been there, had been created, constructed. The decision changed a small part of the world.
4. *The answers to legal questions can change over time.* In the example I am not sure about today's law in this respect, in fact it is my guess that freedom of speech is still interpreted as broadly as here, and I would not be surprised when it is now interpreted even more broadly. But society definitely has changed since the late 1960s, and what I would be surprised about is when I would today see such a sign in the public environment.

¹⁷ For more on the complexity of AI & Law, see for instance [Rissland, 1983, Sergot et al., 1986, Bench-Capon et al., 1987, Rissland and Ashley, 1987, Oskamp et al., 1989, Ashley, 1990, van den Herik, 1991, Berman and Hafner, 1995, Loui and Norman, 1995, Bench-Capon and Sartor, 2003, Sartor, 2005, Bench-Capon et al., 2012, Zurek and Araszkiwicz, 2013, Lauritsen, 2015, Ashley, 2017].




Fig. 6 The theory construction model [Verheij, 2003a, 2005]

One way of looking at the hurdles is by saying that the subsumption model is false. According to the subsumption model of law there is a set of laws, thought of as rules, there are some facts,—and you arrive at the legal answers, the legal consequences by applying the rules to the facts (Figure 5). The case facts are subsumed under the rules, providing the legal solution to the case. It is often associated with Montesquieu’s phrase of the judge as a ‘bouche de la loi’, the mouth of the law, according to which a judge is just the one who makes the law speak.

All hurdles just mentioned show that this perspective cannot be true. Rules are only guiding, terms are open-textured, there can be more answers, and things can change.

Hence an alternative perspective on what happens when a case is decided. Legal decision making is a process of constructing and testing a theory, a series of hypotheses that are gradually developed and tested in a critical discussion (Figure 6). The figure suggests an initial version of the facts, an initial version of the relevant rules, and an initial version of the legal conclusions. Gradually the initial hypothesis is adapted. Think of what happens in a court proceedings, and in what in the Netherlands is called the ‘raadkamer’, the internal discussion among judges, where after a careful constructive critical discussion—if the judges get the time for that of course—finally a tried and tested perspective on the case is arrived at, showing the final legal conclusions subsuming the final facts under the final rules. This is the picture I used in the 2003 AI & Law special issue of the AI journal, edited by Edwina Rissland, Kevin Ashley, and Ronald Loui, two of them here in this room. A later version with Floris Bex emphasises that also the perspective on the evidence and how it supports the facts is gradually constructed [Bex and Verheij, 2012]. In our field, the idea of theory construction in the law has for instance been emphasised by Thorne McCarty [1997], Carole Hafner and Don Berman [2002], Tom Gordon [1995], Trevor Bench-Capon and Giovanni Sartor [2003] and Jaap Hage, Ronald Leenes and Arno Lodder [1993].

4 AI as Law

Today's claim is that good AI requires a different way of doing AI, a way that we in the field of AI & Law have been doing all along, namely doing AI in a way that meets the requirements of the law, in fact in a way that models how things are done in the law. Let us discuss this perspective a bit further.

There can be many metaphors on what AI is and how it should be done, as follows.

1. AI as mathematics, where the focus is on formal systems;
2. AI as technology, where the focus is on the art of system design;
3. AI as psychology, where the focus is on intelligent minds;
4. AI as sociology, where the focus is on societies of agents.

And then AI as law, to which we return in a minute (Table 1).

In AI as mathematics, one can think of the logical and probabilistic foundations of AI, indeed since the start and still now of core importance. It is said that the namegiver of the field of AI—John McCarty—thought of the foundations of AI as an instance of logic, and logic alone. In contrast today some consider AI to be a kind of statistics 2.0 or 3.0.

In AI as technology, one can think of meticulously crafted rule-based expert systems or of machine learning algorithms evaluated on large carefully labeled data sets. In AI as technology, AI applications and AI research meet most directly.

In AI as psychology, one can think of the modeling of human brains as in cognitive modeling, or of the smart human-like algorithms that are sometimes referred to as cognitive computing.

In AI as sociology, one can think of multi-agent systems simulating a society and of autonomous robots that fly in flocks.

Perhaps you have recognized the list of metaphors as the ones used by Stephen Toulmin [1958] when he discussed what he thought of as a crisis in the formal analysis of human reasoning. He argued that the classical formal logic then fashionable was too irrelevant for what reasoning actually was, and

Logic	<i>AI as mathematics</i>	Probability theory
Expert systems	<i>AI as technology</i>	Machine learning
Cognitive modeling	<i>AI as psychology</i>	Cognitive computing
Multi-agent systems	<i>AI as sociology</i>	Autonomous robots
<i>AI as law</i>		
Hybrid critical discussion systems		

Table 1 AI metaphors

he arrived at a perspective of logic as law.¹⁸ What he meant was that counter-argument must be considered, that rules warranting argumentative steps are material—and not only formal—, that these rules are backed by factual circumstances, that conclusions are often qualified, uncertain, presumptive, and that reasoning and argument are to be thought of as the outcome of debates among individuals and in groups (see also Hitchcock and Verheij 2006, Verheij 2009). All of these ideas emphasised by Toulmin have now been studied extensively, with the field of AI & Law having played a significant role in the developments.¹⁹

The metaphors can also be applied to the law, exposing some key ideas familiar in law.

If we think of law as mathematics, the focus is on the formality of procedural rule following and of *stare decisis* where things are well-defined and there is little room for freedom.

In law as technology, one can think of the art of doing law in a jurisdiction with either a focus on rules, as in civil law systems, or with a focus on cases, as in common law systems.

In law as psychology, one can think of the judicial reasoning by an individual judge, and of the judicial discretion that is to some extent allowed, even wanted.

In law as sociology, the role of critical discussion springs to mind, and of regulating a society in order to give order and prevent chaos.

And finally the somewhat pleonastic metaphor of law as law, but now as law in contrast with the other metaphors. I think of two specific and essential ideas in the law, namely that government is to be bound by the rule of law, and that the goal of law is to arrive at justice, thereby supporting a good society and a good life for its citizens.

Note how this discussion shows the typically legal, hybrid balancing of different sides: rules and cases, regulations and decisions, rationality and interpretation, individual and society, boundedness and justice. And as we know this balancing best takes place in a constructive critical discussion.

Which brings us to bottom line of the list of AI metaphors (Table 1).

5. AI as law, where the focus is on hybrid critical discussion.

In AI as law, AI systems are to be thought of as hybrid critical discussion systems, where different hypothetical perspectives are constructed and evaluated until a good answer is found.

In this connection, I recently explained what I think is needed in AI (Figure 7), namely the much needed step we have to make towards hybrid

¹⁸ Toulmin [1958] speaks of logic as mathematics, as technology, as psychology, as sociology and as law (jurisprudence).

¹⁹ See for instance the research by Henry Prakken [1997], Giovanni Sartor [2005], Tom Gordon [1995], Trevor Bench-Capon [2003] and Katie Atkinson [2006]. Argumentation research in AI & Law is connected to the wider study of formal and computational argumentation, see for instance [Simari and Loui, 1992, Pollock, 1995, Vreeswijk, 1997, Chesñevar et al., 2000]. See also the handbooks [Baroni et al., 2018, van Eemeren et al., 2014]

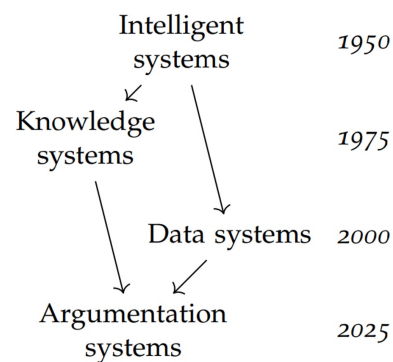


Fig. 7 Bridging the gap between knowledge and data systems in AI [Verheij, 2018]

systems that connect knowledge representation and reasoning techniques with the powers of machine learning. In this diagram I used the term argumentation systems. But since argumentation has a very specific sound in this community, and perhaps to some feels as a too specific, too limiting perspective, I today speak of AI as Law in the sense of the development of hybrid critical discussion systems.

5 Topics in AI

Let me continue with a discussion of core topics in AI with the AI as Law perspective in mind. My focus is on reasoning, knowledge, learning and language.

5.1 Reasoning

First reasoning. I then indeed think of argumentation where arguments and counterarguments meet [van Eemeren et al., 2014, Atkinson et al., 2017, Baroni et al., 2018]. This is connected to the idea of defeasibility, where arguments become defeated when attacked by a stronger counterargument. Argumentation has been used to address the deep and old puzzles of inconsistency, incomplete information and uncertainty.

Here is an example argument about the Dutch bike owner Mary whose bike is stolen (Figure 8). The bike is bought by John, hence both have a claim to ownership—Mary as the original owner, John as the buyer. But in this case the conflict can be resolved as John bought the bike for the low price of 20 euros, indicating that he was not a bona fide buyer. At such a price, he could have known that the bike was stolen, hence he has no claim to ownership as the buyer, and Mary is the owner.

It is one achievement of the field of AI & Law that the logic of argumentation is by now well understood, so well that it can be implemented in

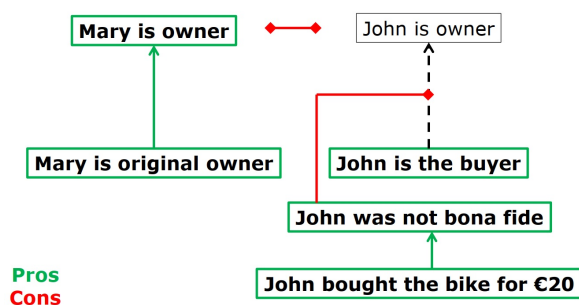


Fig. 8 Argumentation

argumentation diagramming software that applies the logic of argumentation, for instance the ArguMed software that I implemented long ago during my postdoc period in the Maastricht law school [Verheij, 2003a, 2005].²⁰ It implements argumentation semantics of the stable kind in the sense of Dung’s abstract argumentation that was proposed some 25 years ago [Dung, 1995], a turning point and a cornerstone in today’s understanding of argumentation, with many successes. Abstract argumentation also gave new puzzles such as the lack of standardization leading to all kinds of detailed comparative formal studies, and more fundamentally the multiple formal semantics puzzle. The stable, preferred, grounded and complete semantics were the four proposed by Dung [1995], quickly thereafter extended to 6 when the labeling-based stage and semi-stable semantics were proposed [Verheij, 1996]. But that was only the start because the field of computational argumentation was then still only emerging.

For me, it was obvious that a different approach was needed when I discovered that after combining attack and support 11 different semantics were formally possible [Verheij, 2003b], but practically almost all hardly relevant. No lawyer has to think about whether the applicable argumentation semantics is the semi-stable or the stage semantics.

One puzzle in the field is the following, here included after a discussion on the plane from Amsterdam to Montreal with Trevor Bench-Capon and Henry Prakken. A key idea underlying the original abstract argumentation paper is that derivation-like arguments can be abstracted from, allowing to focus only on attack. I know that for many this idea has helped them in their work and understanding of argumentation. For me, this was—from rather early on—more a distraction than an advantage as it introduced a separate, seemingly spurious layer. In the way that my PhD supervisor Jaap Hage put it: ‘those cloudy formal structures of yours’—and Jaap referred to abstract graphs in the sense of Dung—have no grounding in how lawyers think. There is no separate category of supporting arguments to be abstracted from before considering

²⁰ For some other examples, see [Gordon et al., 2007, Loui et al., 1997, Kirschner et al., 2003, Reed and Rowe, 2004, Scheuer et al., 2010, Lodder and Zelznikow, 2005].

attack; instead, in the law there are only reasons for and against conclusions that must be balanced. Those were the days when Jaap Hage was working on Reason-Based Logic [1997] and I was helping him [Verheij et al., 1998]. In a sense, the ArguMed software based on the DefLog formalism was my answer to removing that redundant intermediate layer (still present in its precursor the Argue! system), while sticking to the important mathematical analysis of reinstatement uncovered by Dung (see [Verheij, 2003a, 2005]). For background on the puzzle of combining support and attack, see [van Eemeren et al., 2014, Section 11.5.5].

But as I said from around the turn of the millenium I thought a new mathematical foundation was called for, and it took me years to arrive at something that really increased my understanding of argumentation: the case model formalism [Verheij, 2017a,b], but that is not for now.

5.2 Knowledge

The second topic of AI to be discussed is knowledge, so prominent in AI and in law. I then think of material, semi-formal argumentation schemes such as the witness testimony scheme, or the scheme for practical reasoning, as for instance collected in the nice volume by Doug Walton, Chris Reed and Fabrizio Macagno [2008].

I also think of norms, in our community often studied with a Hohfeldian or deontic logic perspective on rights and obligations as a background.²¹ And then there are the ontologies that can capture large amounts of knowledge in a systematic way.²²

One lesson that I have taken home from working in the domain of law—and again don't forget that I started in the field of mathematics where things are thought of as neat and clean—one lesson is that in the world of law things are always more complex than you think. One could say that it is the business of law to find the exactly right level of complexity, and that is often just a bit more complex than one's initial idea. And if things are not yet complex now, they can become tomorrow. Remember the dynamics of theory construction that we saw earlier (Figure 6).

Figure 9 (left) shows how in the law different categories of juristic facts are distinguished. Here juristic facts are the kind of facts that are legally relevant, that have legal consequences. They come in two kinds: acts with legal consequences, and bare juristic facts, where the latter are intentionless events such as being born, which still have legal consequences. And acts with legal consequences are divided in on the one hand juristic acts aimed at a legal consequence (such as contracting), and on the other factual acts, where although there is no legal intention, still there are legal consequences. Here the

²¹ See for instance [Sartor, 2005, Gabbay et al., 2013, Governatori and Rotolo, 2010].

²² See for instance [McCarty, 1989, Valente, 1995, van Kralingen, 1995, Visser, 1995, Visser and Bench-Capon, 1998, Hage and Verheij, 1999, Boer et al., 2002, 2003, Breuker et al., 2004, Hoekstra et al., 2007, Wyner, 2008, Casanovas et al., 2016].

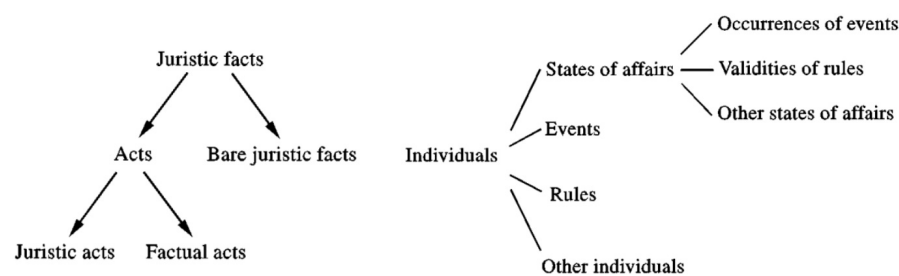


Fig. 9 Types of juristic facts (left); tree of individuals (right) [Hage and Verheij, 1999]

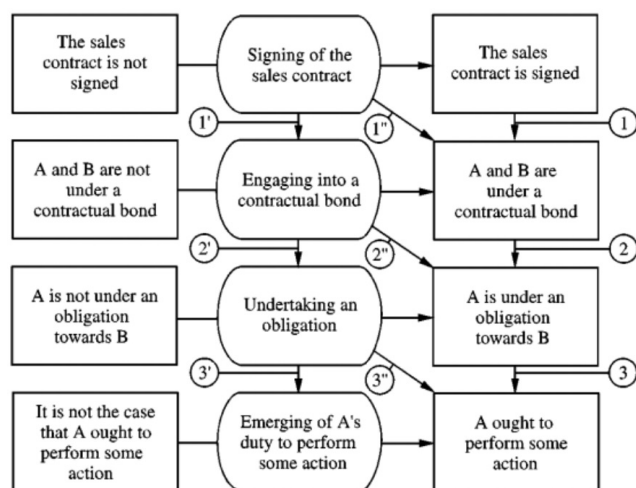


Fig. 10 Signing a sales contract [Hage and Verheij, 1999]

primary example is that of unlawful acts as discussed in tort law. I am still happy that I learnt this categorization of juristic facts in the Maastricht law school, as it has relevantly expanded my understanding of how things work in the world. And of how things should be done in AI. Definitely not purely logically or purely statistically, definitely with much attention for the specifics of a situation.

Figure 9 (right) shows another categorization, prepared with Jaap Hage, that shows how we then approached the core categories of things, or 'individuals' that should be distinguished when analyzing the law: states of affairs, events rules, other individuals, and then the subcategories of event occurrences, rule validities and other states of affairs. And although such a categorization does have a hint of the baroque-ness of Jorge Luis Borges' animal taxonomy (that included those animals that belong to the emperor, mermaids and innumerable animals), the abstract core ontology helped us to analyze the relations between events, rules and states of affairs that play a role when signing a con-

tract (Figure 10). Indeed at first sight a complex picture. For now it suffices that at the top row there is the physical act of signing—say when the pen is going over the paper to sign—and this physical act counts as engaging in a contractual bond (shown in the second row), which implies the undertaking of an obligation (third row), which in turn leads to a duty to perform an action (at the bottom row). Not a simple picture, but as said, in the law things are often more complex than expected, and typically for good, pragmatic reasons.

The core puzzle for our field and for AI generally that I would like to mention is that of commonsense knowledge. This remains an essential puzzle, also in these days of big data; also in these days of cognitive computing. Machines simply don't have commonsense knowledge that is nearly good enough. A knowledgeable report in the *Communications of the ACM* explains that progress has been slow [Davis and Marcus, 2015]. It goes back to 2015, but please do not believe it when it is suggested that things are very different today. The commonsense knowledge problem remains a relevant and important research challenge indeed and I hope to see more of the big knowledge needed for serious AI & Law in the future. Only brave people have the chance to make real progress here, like the people in this room.

One example of what I think is an as yet underestimated cornerstone of commonsense knowledge is the role of globally coherent knowledge structures—such as the scenarios and cases we encounter in the law. Our current program chair Floris Bex took relevant steps to investigate scenario schemes and how they are hierarchically related, in the context of murder stories and crime investigation [Bex, 2011].²³ Our field would benefit from more work like this, that goes back to the frames and scripts studied by people such as Roger Schank and Marvin Minsky.

My current favorite kind of knowledge representation uses the case models mentioned before. It has for instance been used to represent how an appellate court gradually constructs its hypotheses about a murder case on the basis of the evidence, gradually testing and selecting which scenario of what has happened to believe or not [Verheij, 2019], and also to the temporal development of the relevance of past decisions in terms of the values they promote and demote [Verheij, 2016].

5.3 Learning

Then we come to the topic of learning. It is the domain of statistical analysis that shows that certain judges are more prone to supporting democrat positions than others, and that as we saw no longer is allowed in France. It is the domain of open data, that allows public access to legal sources and in which our community has been very active [Biagioli et al., 2005, Francesconi and Passerini, 2007, Francesconi et al., 2010b,a, Sartor et al., 2011, Athan

²³ For more work on evidence in AI & Law, see for instance [Keppens and Schafer, 2006, Bex et al., 2010, Keppens, 2012, Fenton et al., 2013, Vlek et al., 2014, Di Bello and Verheij, 2018].



Fig. 11 Convictions in criminal cases in the Netherlands; source: Central Bureau of Statistics (www.cbs.nl), data collection of September 11, 2017

et al., 2013]. And it is the realm of neural networks, back in the days called perceptrons, now referred to as deep learning.

The core theme to be discussed here is the issue of how learning and the justification of outcomes go together, using a contemporary term: how to arrive at an explainable AI, an explainable machine learning. We have heard it discussed at all career levels, by young PhD students and by a Turing award winner.

The issue can be illustrated by a mock prediction machine for Dutch criminal courts. Imagine a button that you can push, that once you push it always gives the outcome that the suspect is guilty as charged. And thinking of the need to evaluate systems [Conrad and Zeleznikow, 2015], this system has indeed been validated by the Dutch Central Bureau of Statistics, that has the data that shows that this prediction machine is correct in 91 out of a 100 cases (Figure 11). The validating data shows that the imaginary prediction machine has become a bit less accurate in recent years, presumably by changes in society, perhaps in part caused by the attention in the Netherlands for so-called dubious cases, or miscarriages of justice, which may have made judges a little more reluctant to decide for guilt. But still: 91% for this very simple machine is quite good. And as you know, all this says very little about how to decide for guilt or not.

How hard judicial prediction really is, also when using serious machine learning techniques, is shown by some recent examples. Katz, Bommarito and Blackman showed [2017] that their US Supreme Court prediction machine could achieve a 70% accuracy. A mild improvement over the baseline of the historical majority outcome (to always affirm a previous decision) which is 60%, and even milder over the 10 year majority outcome which is 67%. The system based its predictions on features such as judge identity, month, court of origin and issue, so modest results are not surprising.




Fig. 12 Neural networks and open texture [Bench-Capon, 1993]

In another study Aletras and colleagues [2016] studied European Court of Human Rights cases. They used n-grams and topics as the starting point of their training, and used a prepared dataset to make a cleaner baseline of 50% accuracy by random guessing. They reached 79% accuracy using the whole text, and noted that by only using the part where the factual circumstances are described already an accuracy of 73% is reached.

Naively taking the ratios of 70 over 60 and of 79 over 50, one sees that factors of 1.2 and of 1.6 improvement are relevant research outcomes, but practically modest. And more importantly these systems only focus on outcome, without saying anything about how to arrive at an outcome, or about for which reasons an outcome is warranted or not.

And indeed and as said before learning is hard, especially in the domain of law.²⁴ I am still a fan of an old paper by Trevor Bench-Capon on neural networks and open texture [Bench-Capon, 1993]. In an artificially constructed example about welfare benefits, he included different kinds of constraints: boolean, categorical, numeric. For instance, women were allowed the benefit after 60, and men after 65. Trevor found that after training, the neural network could achieve a high overall performance, but with somewhat surprising underlying rationales. In Figure 12, on the left, one can see that the condition starts to be relevant long before the ages of 60 and 65 and that the difference in gender is something like 15 years instead of 5. On the right, with a more focused training set using cases with only single failing conditions, the relevance started a bit later, but still too early, while the gender difference now indeed was 5 years.

What I have placed my bets on is the kind of hybrid cases and rules systems that for us in AI & Law are normal.²⁵ I now represent Dutch tort law in terms of case models validating rule-based arguments [Verheij, 2017b] (cf. Figure 13 below).

²⁴ See also recently [Medvedeva et al., 2019].

²⁵ See for instance work by [Branting, 1991, Skalak and Rissland, 1992, Branting, 1993, Prakken and Sartor, 1996, 1998, Stranieri et al., 1999, Roth, 2003, Brünighaus and Ashley, 2003, Atkinson and Bench-Capon, 2006, Čyras et al., 2016].

5.4 Language

Then language, the fourth and final topic of AI that I would like to discuss with you. Today the topic of language is closely connected to machine learning. I think of the labeling of natural language data to allow for training; I think of prediction such as by a search engine or chat application on a smartphone, and I think of argument mining, a relevant topic with strong roots in the field of AI & Law.

The study of natural language in AI, and in fact of AI itself, got a significant boost by IBM's Watson system that won the Jeopardy! quiz show. For instance, Watson correctly recognized the description of 'A 2-word phrase [that] means the power to take private property for public use'. That description refers to the typically legal concept of eminent domain, the situation in which a government disowns property for public reasons, such as the construction of a highway or windmill park. Watson's output showed that the legal concept scored 98%, but also 'electric company' and 'capitalist economy' were considered with 9% and 5% scores, respectively. Apparently Watson sees some kind of overlap between the legal concept of eminent domain, electric companies and capitalist economy, since $98+9+5$ is more than a 100 percent.

And IBM continued, as Watson was used as the basis for its debating technologies. In a 2014 demonstration,²⁶ the system is considering the sale of violent video games to minors. The video shows that the system finds reasons for and against banning the sale of such games to minors, for instance that most children who play violent games do not have problems, but that violent video games can increase children's aggression. The video remains impressive, and for the field of computational argumentation that I am a member of it was somewhat disconcerting that the researchers behind this system were then outsiders to the field.

The success of these natural language systems leads one to think about why they can do what they do. Do they really have an understanding of a complex sentence describing the legal concept of eminent domain; can they really digest newspaper articles and other online resources on violent video games?

These questions are especially relevant since in our field of AI & Law we have had the opportunity to follow research on argument mining from the start. Early and relevant research is by Raquel Mochales Palau and Sien Moens, who studied argument mining in a paper at the 2009 ICAIL conference [2009, 2011]. And as already shown in that paper, it should not be considered an easy task to perform argument mining. Indeed the field has been making relevant and interesting progress, as also shown in research presented at this conference, but no one would claim the kind of natural language understanding needed for interpreting legal concepts or online debates.²⁷

²⁶ Milken Institute Global Conference 2014, session 'Why Tomorrow Won't Look Like Today: Things that Will Blow Your Mind', youtu.be/6fJ0tAzICzw?t=2725.

²⁷ See for instance [Schweighofer et al., 2001, Wyner et al., 2009, 2010, Grabmair and Ashley, 2011, Ashley and Walker, 2013, Grabmair et al., 2015, Tran et al., 2020].

So what then is the basis of apparent success? Is it simply because a big tech company can do a research investment that in academia one can only dream of? Certainly that is a part of what has been going on. But there is more to it than that as can be appreciated by a small experiment I did, this time actually an implemented online system. It is what I ironically called Poor Man's Watson,²⁸ which has been programmed without much deep natural language technology, just some simple regular expression scripts using online access to the Google search engine and Wikipedia. And indeed it turns out that the simple script can also recognize the concept of eminent domain: when one types 'the power to take private property for public use' the answer is 'eminent domain'. The explanation for this remarkable result is that for some descriptions the correct Wikipedia page ends up high in the list of pages returned by Google, and that happens because we—the people—have been typing in good descriptions of those concepts in Wikipedia, and indeed Google can find these pages. Sometimes the results are spectacular, but also they are brittle since seemingly small, irrelevant changes can quickly break this simple system.

And for the debating technology something similar holds since there are web sites collecting pros and cons of societal debates. For instance, the web site procon.org has a page on the pros and cons of violent video games.²⁹ Arguments it has collected include 'Pro 1: Playing violent video games causes more aggression, bullying, and fighting' and 'Con 1: Sales of violent video games have significantly increased while violent juvenile crime rates have significantly decreased'. The web site Kialo has similar collaboratively created lists.³⁰ Concerning the issue 'Violent video games should be banned to curb school shootings', it lists for instance the pro 'Video games normalize violence, especially in the eyes of kids, and affect how they see and interact with the world' and the con 'School shootings are, primarily, the result of other factors that should be dealt with instead'.

Surely the existence of such lists typed in, in a structured way, by humans is a central basis for what debating technology can and cannot do. It is not a coincidence that—listening carefully to the reports—the examples used in marketing concern curated lists of topics. At the same time this does not take away the bravery of IBM and how strongly it has been stimulating the field of AI by its successful demos. And that also for IBM things are sometimes hard is shown by the report from February 2019 when IBM's technology entered into a debate with a human debater, and this time lost.³¹ But who knows what the future brings.

What I believe is needed is the development of an ever closer connection between complex knowledge representations and natural language explanations, as for instance in work by Charlotte Vlek on explaining Bayesian Net-

²⁸ Poor Man's Watson, www.ai.rug.nl/~verheij/pmw.

²⁹ videogames.procon.org

³⁰ kialo.com

³¹ 'IBM's AI loses debate to a human, but it's got worlds to conquer', cnet.com, February 11, 2019.

works [Vlek et al., 2016], which had nice connections to the work discussed by Jeroen Keppens yesterday [2019].

6 Conclusion

As I said I think the way to go for the field is to develop an AI that is much like the law, an AI where systems are hybrid critical discussion systems.

For after phases of AI as mathematics, as technology, as psychology, and as sociology—all still important and relevant—, an AI as Law perspective provides fresh ideas for designing an AI that is good (Table 1). And in order to build the hybrid critical discussion systems that I think are needed, lots of work is waiting in reasoning, in knowledge, in learning and in language, as follows.

For reasoning (Section 5.1), the study of formal and computational argumentation remains relevant and promising, while work is needed to arrive at a formal semantics that is not only accessible for a small group of experts.

For knowledge (Section 5.2), we need to continue working on knowledge bases large and small, and on systems with embedded norms. But I hope that some of us are also brave enough to be looking for new ways to arrive at good commonsense knowledge for machines. In the law we cannot do without wise commonsense.

For learning (Section 5.3), the integration of knowledge and data can be addressed by how in the law rules and cases are connected and influence one another. Only then the requirements of explainability and responsibility can be properly addressed.

For language (Section 5.4), work is needed in interpretation of what is said in a text. This requires an understanding in terms of complex, detailed models of a situation, like what happens in any court of law where every word can make a relevant difference.

Lots of work to do. Lots of high mountains to conquer.

The perspective of AI as Law discussed here today can be regarded as an attempt to broaden what I said in the lecture on ‘Arguments for good AI’ where the focus is mostly on computational argumentation [Verheij, 2018]. There I explain that we need a good AI that can give good answers to our questions, give good reasons for them, and make good choices. I projected that in 2025 we will have arrived at a new kind of AI systems bridging knowledge and data, namely argumentation systems (Figure 7). Clearly and as I tried to explain today, there is still plenty of work to be done. I expect that a key role will be played by work in our field on connections between rules, cases and arguments, as in the set of cases formalizing tort law (Figure 13, left) that formally validate the legally relevant rule-based arguments (Figure 13, right).

By following the path of developing AI as Law we can guard against technology that is bad for us, and that—unlike the guillotine I started with—is a really humane technology that directly benefits society and its citizens.

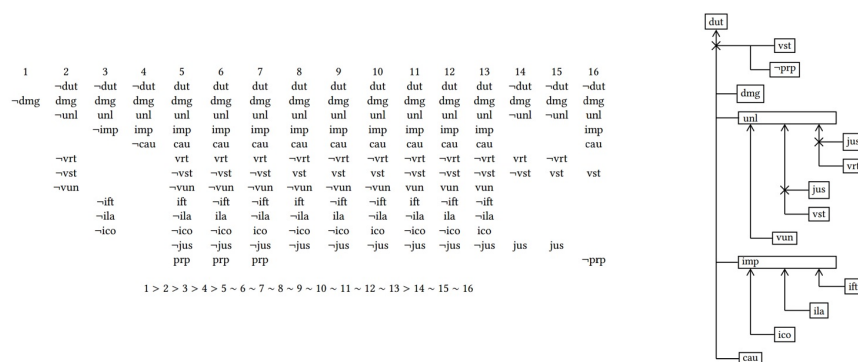


Fig. 13 Arguments, rules and cases for Dutch tort law [Verheij, 2017b]

In conclusion, in these days of dreams and fears of AI and algorithms, our beloved field of AI & Law is more relevant than ever. We can be proud that AI & Law has worked on the design of socially aware, explainable, responsible AI for decades already.

And since we in AI & Law are used to address the hardest problems across the breadth of AI (reasoning, knowledge, learning, language)—since in fact we cannot avoid them—, our field can inspire new solutions. In particular, I discussed computational argumentation, schemes for arguments and scenarios, encoded norms, hybrid rule-case systems and computational interpretation.

We only need to look at what happens in the law. In the law, we see an artificial system that adds much value to our life. Let us take inspiration from the law, and let us work on building Artificial Intelligence that is not scary, but that genuinely contributes to a good quality of life in a just society. I am happy and proud to be a member of this brave and smart community and I thank you for your attention.

References

N. Aletras, D. Tsarapatsanis, D. Preotiuc-Pietro, and V. Lampos. Predicting judicial decisions of the European Court of Human Rights: a natural language processing perspective. *Peer J Computer Science*, 2:1–19, 2016. URL <https://doi.org/10.7717/peerj-cs.93>.

K. D. Ashley. *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals*. The MIT Press, Cambridge (Massachusetts), 1990.

K. D. Ashley. *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge University Press, Cambridge, 2017.

K. D. Ashley and V. R. Walker. Toward constructing evidence-based legal arguments using legal decision documents and machine learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 176–180. ACM, New York (New York), 2013.

- T. Athan, H. Boley, G. Governatori, M. Palmirani, A. Paschke, and A. Wyner. OASIS LegalRuleML. In *Proceedings of the 14th International Conference on Artificial Intelligence and Law (ICAIL 2013)*, pages 3–12. ACM Press, New York (New York), 2013.
- K. Atkinson and T. J. M. Bench-Capon. Legal case-based reasoning as practical reasoning. *Artificial Intelligence and Law*, 13:93–131, 2006.
- K. Atkinson, P. Baroni, M. Giacomini, A. Hunter, H. Prakken, C. Reed, G. Simari, M. Thimm, and S. Villata. Toward artificial argumentation. *AI Magazine*, 38(3):25–36, 2017.
- P. Baroni, D. Gabbay, M. Giacomini, and L. van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, London, 2018.
- T. J. M. Bench-Capon. Neural networks and open texture. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law*, pages 292–297. ACM Press, New York (New York), 1993.
- T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.
- T. J. M. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1):97–143, 2003.
- T. J. M. Bench-Capon, G. O. Robinson, T. W. Routen, and M. J. Sergot. Logic programming for large scale applications in law: A formalisation of supplementary benefit legislation. In *Proceedings of the 1st International Conference on Artificial Intelligence and Law (ICAIL 1987)*, pages 190–198. ACM, New York (New York), 1987.
- T. J. M. Bench-Capon, M. Araszkiwicz, K. D. Ashley, K. Atkinson, F. J. Bex, F. Borges, D. Bourcier, D. Bourguine, J. G. Conrad, E. Francesconi, T. F. Gordon, G. Governatori, J. L. Leidner, D. D. Lewis, R. P. Loui, L. T. McCarty, H. Prakken, F. Schilder, E. Schweighofer, P. Thompson, A. Tyrrell, B. Verheij, D. N. Walton, and A. Z. Wyner. A history of AI and Law in 50 papers: 25 years of the International Conference on AI and Law. *Artificial Intelligence and Law*, 20(3):215–319, 2012.
- D. H. Berman and C. L. Hafner. Understanding precedents in a temporal context of evolving legal doctrine. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 42–51. ACM Press, New York (New York), 1995.
- F. J. Bex. *Arguments, Stories and Criminal Evidence: A Formal Hybrid Theory*. Springer, Berlin, 2011.
- F. J. Bex and B. Verheij. Solving a murder case by asking critical questions: An approach to fact-finding in terms of argumentation and story schemes. *Argumentation*, 26:325–353, 2012.
- F. J. Bex, P. J. van Koppen, H. Prakken, and B. Verheij. A hybrid formal theory of arguments, stories and criminal evidence. *Artificial Intelligence and Law*, 18:1–30, 2010.
- C. Biagioli, E. Francesconi, A. Passerini, S. Montemagni, and C. Soria. Automatic semantics extraction in law documents. In *Proceedings of the 10th International Conference on Artificial Intelligence and Law (ICAIL 2005)*,

- pages 133–140. ACM Press, New York (New York), 2005.
- A. Boer, R. Hoekstra, and R. Winkels. METAlex: Legislation in XML. In T. J. M. Bench-Capon, A. Daskalopulu, and R. Winkels, editors, *Legal Knowledge and Information Systems. JURIX 2002: The Fifteenth Annual Conference*, pages 1–10. IOS Press, Amsterdam, 2002.
- A. Boer, T. van Engers, and R. Winkels. Using ontologies for comparing and harmonizing legislation. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 60–69. ACM, New York (New York), 2003.
- L. K. Branting. Building explanations from rules and structured cases. *International Journal of Man-Machine Studies*, 34(6):797–837, 1991.
- L. K. Branting. A computational model of ratio decidendi. *Artificial Intelligence and Law*, 2(1):1–31, 1993.
- J. Breuker, A. Valente, and R. Winkels. Legal ontologies in knowledge engineering and information management. *Artificial Intelligence and Law*, 12(4):241–277, 2004.
- S. Brüninghaus and K. D. Ashley. Predicting outcomes of case based legal arguments. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law (ICAIL 2003)*, pages 233–242. ACM, New York (New York), 2003.
- P. Casanovas, M. Palmirani, S. Peroni, T. van Engers, and F. Vitali. Semantic web for the legal domain: the next step. *Semantic web*, 7(3):213–227, 2016.
- C. I. Chesñevar, A. G. Maguitman, and R. P. Loui. Logical models of argument. *ACM Computing Surveys*, 32(4):337–383, 2000.
- J. G. Conrad and J. Zeleznikow. The role of evaluation in ai and law: an examination of its different forms in the ai and law journal. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law (ICAIL 2015)*, pages 181–186. ACM, New York (New York), 2015.
- E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in Artificial Intelligence. *Communications of the ACM*, 58(9):92–103, 2015.
- M. Di Bello and B. Verheij. Evidential reasoning. In G. Bongiovanni, G. Postema, A. Rotolo, G. Sartor, C. Valentini, and D. N. Walton, editors, *Handbook of Legal Reasoning and Argumentation*, pages 447–493. Springer, Dordrecht, 2018.
- P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- N. E. Fenton, M. D. Neil, and D. A. Lagnado. A general structure for legal arguments about evidence using Bayesian Networks. *Cognitive Science*, 37: 61–102, 2013.
- E. Francesconi and A. Passerini. Automatic classification of provisions in legislative texts. *Artificial Intelligence and Law*, 15(1):1–17, 2007.
- E. Francesconi, S. Montemagni, W. Peters, and D. Tiscornia. Integrating a bottom-up and top-down methodology for building semantic resources for the multilingual legal domain. In *Semantic Processing of Legal Texts*, pages

- 95–121. Springer, Berlin, 2010a.
- E. Francesconi, S. Montemagni, W. Peters, and D. Tiscornia. *Semantic processing of legal texts: Where the language of law meets the law of language*. Springer, Berlin, 2010b.
- D. Gabbay, J. Horty, X. Parent, R. Van der Meyden, and L. van der Torre. *Handbook of Deontic Logic and Normative Systems*. College Publication, London, 2013.
- A. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. The MIT Press, Cambridge (Massachusetts), 1987.
- T. F. Gordon. *The Pleadings Game: An Artificial Intelligence Model of Procedural Justice*. Kluwer, Dordrecht, 1995.
- T. F. Gordon, H. Prakken, and D. N. Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10–15):875–896, 2007.
- G. Governatori and A. Rotolo. Changing legal systems: Legal abrogations and annulments in defeasible logic. *Logic Journal of the IGPL*, 18(1):157–194, 2010.
- M. Grabmair and K. D. Ashley. Facilitating case comparison using value judgments and intermediate legal concepts. In *Proceedings of the 13th international conference on Artificial intelligence and law*, pages 161–170. ACM, New York (New York), 2011.
- M. Grabmair, K. D. Ashley, R. Chen, P. Sureshkumar, C. Wang, E. Nyberg, and V. R. Walker. Introducing LUIMA: an experiment in legal conceptual retrieval of vaccine injury decisions using a uima type system and tools. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, pages 69–78. ACM, New York (New York), 2015.
- C. L. Hafner and D. H. Berman. The role of context in case-based legal reasoning: Teleological, temporal, and procedural. *Artificial Intelligence and Law*, 10(1–3):19–64, 2002.
- J. C. Hage. *Reasoning with Rules. An Essay on Legal Reasoning and Its Underlying Logic*. Kluwer Academic Publishers, Dordrecht, 1997.
- J. C. Hage and B. Verheij. The law as a dynamic interconnected system of states of affairs: a legal top ontology. *International Journal of Human-Computer Studies*, 51(6):1043–1077, 1999.
- J. C. Hage, R. Leenes, and A. R. Lodder. Hard cases: a procedural approach. *Artificial intelligence and law*, 2(2):113–167, 1993.
- D.L. Hitchcock and B. Verheij, editors. *Arguing on the Toulmin Model. New Essays in Argument Analysis and Evaluation (Argumentation Library, Volume 10)*. Springer, Dordrecht, 2006.
- R. Hoekstra, J. Breuker, M. Di Bello, and A. Boer. The lkif core ontology of basic legal concepts. In P. Casanovas, M. A. Biasiotti, E. Francesconi, and M. T. Sagri, editors, *Proceedings of LOAIT 2007. Second Workshop on Legal Ontologies and Artificial Intelligence Techniques*, pages 43–63. CEUR-WS, 2007.
- D.M. Katz, M. J. Bommarito II, and J. Blackman. A general approach for predicting the behavior of the Supreme Court of the United States. *PLoS*

- ONE*, 12(4):1–18, 2017. URL <https://doi.org/10.1371/journal.pone.0174698>.
- J. Keppens. Argument diagram extraction from evidential Bayesian networks. *Artificial Intelligence and Law*, 20:109–143, 2012.
- J. Keppens. Explainable bayesian network query results via natural language generation systems. In *Proceedings of the 17th International Conference on Artificial Intelligence and Law (ICAIL 2019)*, pages 42–51. ACM, New York (New York), 2019.
- J. Keppens and B. Schafer. Knowledge based crime scenario modelling. *Expert Systems with Applications*, 30(2):203–222, 2006.
- P. A. Kirschner, S. J. B. Shum, and C. S. Carr. *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Springer, Berlin, 2003.
- M. Lauritsen. On balance. *Artificial Intelligence and Law*, 23(1):23–42, 2015.
- A. R. Lodder and J. Zelznikow. Developing an online dispute resolution environment: Dialogue tools and negotiation support systems in a three-step model. *Harvard Negotiation Law Review*, 10:287–337, 2005.
- R. P. Loui and J. Norman. Rationales and argument moves. *Artificial Intelligence and Law*, 3:159–189, 1995.
- R. P. Loui, J. Norman, J. Altepeter, D. Pinkard, D. Craven, J. Linsday, and M. Foltz. Progress on room 5: a testbed for public interactive semi-formal legal argumentation. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pages 207–214. ACM Press, 1997.
- L. T. McCarty. A language for legal discourse. i. basic features. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law (ICAIL 1989)*, pages 180–189. ACM, New York (New York), 1989.
- L. T. McCarty. Some arguments about legal arguments. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law (ICAIL 1997)*, pages 215–224. ACM Press, New York (New York), 1997.
- M. Medvedeva, M. Vols, and M. Wieling. Using machine learning to predict decisions of the european court of human rights. *Artificial Intelligence and Law*, pages 1–30, 2019.
- R. Mochales Palau and M. F. Moens. Argumentation mining: the detection, classification and structure of arguments in text. In *Proceedings of the 12th International Conference on Artificial Intelligence and Law (ICAIL 2009)*, pages 98–107. ACM Press, New York (New York), 2009.
- R. Mochales Palau and M. F. Moens. Argumentation mining. *Artificial Intelligence and Law*, 19(1):1–22, 2011.
- A. Oskamp, R. F. Walker, J. A. Schrickx, and P. H. van den Berg. PROLEXS divide and rule: a legal application. In *Proceedings of the Second International Conference on Artificial intelligence and law*, pages 54–62. ACM, New York (New York), 1989.
- J. L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person*. The MIT Press, Cambridge (Massachusetts), 1995.
- H. Prakken. *Logical Tools for Modelling Legal Argument. A Study of Defeasible Reasoning in Law*. Kluwer Academic Publishers, Dordrecht, 1997.

- H. Prakken and G. Sartor. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law*, 4:331–368, 1996.
- H. Prakken and G. Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.
- C. Reed and G. Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal of AI Tools*, 14(3–4):961–980, 2004.
- E. L. Rissland. Examples in legal reasoning: Legal hypotheticals. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence (IJCAI 1983)*, pages 90–93. 1983.
- E. L. Rissland. Book review. An Artificial Intelligence Approach to Legal Reasoning. *Harvard Journal of Law and Technology*, 1(Spring):223–231, 1988.
- E. L. Rissland and K. D. Ashley. A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, pages 60–66. ACM Press, New York (New York), 1987.
- B. Roth. *Case-Based Reasoning in the Law. A Formal Theory of Reasoning by Case Comparison*. Dissertation Universiteit Maastricht, Maastricht, 2003.
- G. Sartor. *Legal reasoning: a cognitive approach to the law. Vol 5 of Treatise on legal philosophy and general jurisprudence*. Springer, Berlin, 2005.
- G. Sartor, M. Palmirani, E. Francesconi, and M. A. Biasiotti. *Legislative XML for the semantic web: principles, models, standards for document management*. Springer, Berlin, 2011.
- O. Scheuer, F. Loll, N. Pinkwart, and B. M. McLaren. Computer-supported argumentation: A review of the state of the art. *International Journal of Computer-Supported Collaborative Learning*, 5(1):43–102, 2010.
- E. Schweighofer, A. Rauber, and M. Dittenbach. Automatic text representation, classification and labeling in european law. In *Proceedings of the 8th International Conference on Artificial intelligence and law*, pages 78–87. ACM, New York (New York), 2001.
- M.J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The british nationality act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.
- G. R. Simari and R. P. Loui. A mathematical treatment of defeasible reasoning and its applications. *Artificial Intelligence*, 53:125–157, 1992.
- D. B. Skalak and E. L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial Intelligence and Law*, 1(1):3–44, 1992.
- A. Stranieri, J. Zeleznikow, M. Gawler, and B. Lewis. A hybrid rule–neural approach for the automation of legal reasoning in the discretionary domain of family law in australia. *Artificial Intelligence and Law*, 7(2-3):153–183, 1999.
- S. E. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, 1958.
- V. Tran, M. Le Nguyen, S. Tojo, and K. Satoh. Encoded summarization: Summarizing documents into continuous vector space for legal case retrieval. *Artificial Intelligence and Law*, pages 1–27, 2020.

- A. Valente. *Legal Knowledge Engineering. A Modelling Approach*. IOS Press, Amsterdam, 1995.
- H. J. van den Herik. *Kunnen Computers Rechtspreken?* Gouda Quint, Arnhem, 1991.
- F. H. van Eemeren, B. Garssen, E. C. W. Krabbe, A. F. Snoeck Henkemans, B. Verheij, and J. H. M. Wagemans. *Handbook of Argumentation Theory*. Springer, Berlin, 2014.
- R. W. van Kralingen. *Frame-based Conceptual Models of Statute Law*. Kluwer Law International, The Hague, 1995.
- K. Ćyras, K. Satoh, and F. Toni. Abstract argumentation for case-based reasoning. In *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, pages 549–552. AAAI Press, 2016.
- B. Verheij. Two approaches to dialectical argumentation: Admissible sets and argumentation stages. In J. J. Meyer and L. C. van der Gaag, editors, *Proceedings of NAIC'96*, pages 357–368. Universiteit Utrecht, Utrecht, 1996.
- B. Verheij. Artificial argument assistants for defeasible argumentation. *Artificial Intelligence*, 150(1–2):291–324, 2003a.
- B. Verheij. DefLog: on the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation*, 13(3):319–346, 2003b.
- B. Verheij. *Virtual Arguments. On the Design of Argument Assistants for Lawyers and Other Arguers*. T.M.C. Asser Press, The Hague, 2005.
- B. Verheij. The Toulmin argument model in artificial intelligence. or: How semi-formal, defeasible argumentation schemes creep into logic. In I. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 219–238. Springer, Berlin, 2009.
- B. Verheij. Formalizing value-guided argumentation for ethical systems design. *Artificial Intelligence and Law*, 24(4):387–407, 2016.
- B. Verheij. Proof with and without probabilities. Correct evidential reasoning with presumptive arguments, coherent hypotheses and degrees of uncertainty. *Artificial Intelligence and Law*, 25(1):127–154, 2017a.
- B. Verheij. Formalizing arguments, rules and cases. In *Proceedings of the 16th International Conference on Artificial Intelligence and Law (ICAIL 2017)*, pages 199–208. ACM Press, New York (New York), 2017b.
- B. Verheij. *Arguments for Good Artificial Intelligence*. University of Groningen, Groningen, 2018. URL <http://www.ai.rug.nl/~verheij/oratie/>.
- B. Verheij. Analyzing the Simonshaven case with and without probabilities. *Topics in Cognitive Science*, pages 1–25, 2019. URL <https://doi.org/10.1111/tops.12436>.
- B. Verheij, J. C. Hage, and H. J. van den Herik. An integrated view on rules and principles. *Artificial Intelligence and Law*, 6(1):3–26, 1998.
- P. R. S. Visser. *Knowledge Specification for Multiple Legal Tasks; A Case Study of the Interaction Problem in the Legal Domain*. Kluwer Law International, The Hague, 1995.
- P. R. S. Visser and T. J. M. Bench-Capon. A comparison of four ontologies for the design of legal knowledge systems. *Artificial Intelligence and Law*, 6

- (1):27–57, 1998.
- C. S. Vlek, H. Prakken, S. Renooij, and B. Verheij. Building Bayesian Networks for legal evidence with narratives: a case study evaluation. *Artificial Intelligence and Law*, 22(4):375–421, 2014.
- C. S. Vlek, H. Prakken, S. Renooij, and B. Verheij. A method for explaining Bayesian Networks for legal evidence with scenarios. *Artificial Intelligence and Law*, 24(3):285–324, 2016.
- G. A. W. Vreeswijk. Abstract argumentation systems. *Artificial Intelligence*, 90:225–279, 1997.
- D. N. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, Cambridge, 2008.
- A. Wyner. An ontology in OWL for legal case-based reasoning. *Artificial Intelligence and Law*, 16(4):361, 2008.
- A. Wyner, K. Angelov, G. Barzdins, D. Damljanovic, B. Davis, N. Fuchs, S. Hoefler, K. Jones, K. Kaljurand, T. Kuhn, et al. On controlled natural languages: Properties and prospects. In *International Workshop on Controlled Natural Language*, pages 281–289. Berlin, 2009.
- A. Wyner, R. Mochales-Palau, M. F. Moens, and D. Milward. Approaches to text mining arguments from legal cases. In *Semantic processing of legal texts*, pages 60–79. Springer, Berlin, 2010.
- T. Zurek and M. Araszkievicz. Modeling teleological interpretation. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*, pages 160–168. ACM, New York (New York), 2013.