

Bart Verheij

Participants in the dialogue: Chris, a lawyer; Floris, a computer scientist; Alex.¹

Alex: Great to see you again! It has been a while ...

Chris: Indeed. Last time we discussed Van den Herik's 'Kunnen computers rechtspreken?' [1991], a wonderful piece that certainly made me think – and think about my thinking.²

Alex: Thanks for reminding me. Are you still as much of a sceptic about artificial intelligence as you then was?

Chris: On the contrary! I'm not at all a sceptic about artificial intelligence; I think it's wonderful stuff – actually rather scary. I simply am convinced that you AI advocates have far underestimated the risks of artificial intelligence, and that there are many things that computers can do that harm society. For instance, can you imagine a society with ever more computer autonomy, with ever less human control –

Alex: Robots haven't conquered the world – *yet!*

Floris: I am losing you two. Remind me a bit about last time?

Alex: In 2007, we discussed the question whether computers can decide legal cases. In Van den Herik's 1991 inaugural address, delivered upon acceptance of his position in Leiden, he discusses the nature of legal decision making and the prospects of automating it. He also presents what might be called the Van den Herik test.

1 The characters are purely fictional.

2 This text is a follow-up of an earlier coffeehouse conversation [Verheij, 2007] inspired by Douglas R. Hofstadter's *A Coffeehouse Conversation on the Turing Test* [1981].

Chris: And that test is this: if a computer gives flawless legal advice for three months, the machine is the judge, no matter what either thinks of that [van den Herik, 1991, p. 33].

Floris: Well, if you ask me, automating legal decision making should be easy! Just feed the computer all legal decisions and you're done. When machine learning algorithms have access to so much data, the computer will surely do better than a lawyer.

Alex: Wait a minute. It is not *that* easy and Van den Herik knows this. It was not easy in 1991, not in 2007, and is not now. Although he is clearly an optimist, he is very well aware of obstacles on the road towards computer judges. In [van den Herik, 1991], he mentions vague terms, open textured codes, interpretation, system maintenance, and later [van den Herik, 1999] he adds the creation of law and empathy. Such obstacles have been discussed since the early days of the field of AI & Law [Gardner, 1987].

Floris: I don't understand what is so hard. Judges have to apply the law. Let a computer simply apply legal codes and follow past decisions.

Chris: If only it were *that* easy. Your view of the law is rather simplistic. I do not think of my work as straightforwardly applying legal codes and following past decisions.

Floris: Well ok, it may be hard to find relevant decisions and codes of course. There are so many ... I have seen some law student friends struggle. I feel lucky to be in another field. How can anyone remember all that?

Chris: No, it is not really that – and these days finding the right sources has become much easier. The state of things is much different from when we met before. Legal codes are accessible,³ ever more legal decisions are available online,⁴ and the public prosecution publishes an online system, showing what fines can be expected.⁵ Some of that was also available in 2007, but importantly: now all lawyers are accessing such web sites using their phones, tablets and laptops as standard productivity tools. There is a web site curated by the CodeX center for legal informatics at Stanford University, where more than 1500 legal tech companies are listed.⁶

3 wetten.overheid.nl

4 www.rechtspraak.nl

5 boetibase.om.nl

6 techindex.law.stanford.edu

Floris: If I am correct, that fining database is a website that shows for all kinds of violations how much of a fine one can expect.

Alex: Indeed, each fine is just a few clicks away. After clicking ‘Traffic on the road’, ‘Bicycle’, ‘Vehicle requirements’, ‘Lighting’, ‘Inadequate lighting’, the system indicates that 60 euros fine can be expected when the lights on your bicycle don’t work. The system knows it all! And it is in that sense an AI system.

Floris: How is that an AI system? No data, no learning!

Alex: You make me smile. You know that AI is not only about learning. There is knowledge, reasoning, language, interaction. Intelligent behaviour exists in much variety, and all aspects are addressed in the field of AI. And the fining database uses expert knowledge about the applicable law in the Netherlands, handles that knowledge with a rudimentary form of reasoning, and communicates about its findings with elementary language skills.

Chris: Well, if THAT is AI! Nothing scary about that. It is just a useful web tool!

Alex: You’d be surprised about how opinions about AI keep changing over time. Did you know that people considered it scary when computers started to play chess really well? Those days are behind us. Also machine chess is hardly recognized as AI these days. There is a tendency that when we have a good computational understanding of certain intelligent behavior, it is no longer considered AI.

Chris: There is much talk of automated legal decision making these days. It gives me the feeling I will be out of a job soon ... How can lawyers compete with omniscient machines that never get tired?

Floris: You are right that systems fed with much data can do remarkable things! I recently read that a computer model can reasonably predict decisions by the European Court of Human Rights.⁷

Alex: Well, I can do a reasonable prediction too: I predict that the next decision will be that there was a violation of the law. And I am not omniscient, and I do get tired.

⁷ NRC Handelsblad, December 30, 2020: ‘Robot knows which decision the court will make’, www.nrc.nl/nieuws/2020/12/30/robot-weet-welke-uitspraak-het-hof-zal-doen.

Floris: Without knowing anything about the case? How can that be a reasonable prediction?

Chris: I guess what Alex means is that most cases will be decided as violations anyway.

Alex: Indeed. I checked for February 2021 cases. The data show that always predicting a violation is correct in 78% of the cases (18 out of 23). Not so bad! Interestingly, the AI system you read about in the news got 61% right in that month (14 out of 23).

Floris: So the trivial baseline outperformed the AI system ... Let us apply the Van den Herik test: the baseline does not have flawless performance, as required by the test, but 78% is indeed not so bad.

Chris: I guess I am not yet out of a job! Also I would say that just predicting violation or not is not the same as giving flawless legal advice. Lawyers do more than just say ‘Violation!’ or ‘No violation!’.

Alex: There is still a long way to go. One important puzzle for today’s AI is that the law requires justifications for decisions, and successful machine learning algorithms aren’t well-equipped for providing a justification for their outcomes. A trained neural network is in a sense a black-box: its outcomes are often correct, but the reasons for the outcomes aren’t available. Consider for example the ‘predict-a-violation’ baseline just now: the data shows it is often correct but clearly not for the right reasons. Learning the right reasons from data turns out to be a complex puzzle anyway [Steging et al., 2021].

Chris: Are there no AI systems that decide for the right reasons?

Floris: Surely there are! Rule-based knowledge systems follow the rules that hold in a domain. So their decisions are based on the reasons that actually apply in the domain. That technology was already heavily investigated in the 1970s.

Alex: You are right. And rule-based knowledge systems perform well in practice, especially in stable domains with well-defined problems, such as in social security, finance, tax administration, with large caseloads.

Floris: But if such legal AI systems perform so well, why do we still need lawyers?

Chris: Ah! If I heard right, the adjectives ‘well-defined’ and ‘stable’ carry much meaning here. Domains change and problems often aren’t well-defined.

Alex: Yes. The law has its special properties, with the effect that current rule-based knowledge systems cannot address all legal problems. For instance, following [Rissland, 1988] (reviewing [Gardner, 1987]), legal reasoning is rule-guided, rather than rule-governed. Also legal terms are open textured. Moreover, legal questions can have more than one answer, but a reasonable and timely answer must always be given. And finally the answers to legal questions can change over time. Van den Herik [1991] was well aware of such properties.

Floris: But then what next? You say that machine learning methods aren't directly applicable in the law [Bench-Capon, 2020], for instance by the puzzles of explainable AI. And you also say that rule-based knowledge systems aren't enough. Would you say that AI methods are in principle insufficient for the law?

Alex: On the contrary. Indeed today's AI methods – though useful in specific settings – need foundational redirection for a next level of applicability in law. The stimulating point is this: the law can show how AI methods need to change. In a sense, we should do AI as we do law [Verheij, 2020].

Chris: Are you really telling me that my field can teach the engineers what to do?

Alex: Yes.

Floris: Computer scientists go to law school ... Who'd believe that!

Alex: You put it rather boldly, but indeed: computer scientists could learn a lot in law school.

Floris: For example?

Alex: Well, one thing is that many real life problems are underspecified but still need to be addressed. That is rather different from a mathematically specified puzzle with a well-determined outcome, such as the factorization of a given number.

Floris: I do not like open-ended problems! How can I design a correct and efficient algorithm for those?

Chris: Indeed lawyers are well aware that there is not always a guarantee that an answer to a legal problem is correct, and that sometimes inefficient detours have been made, and in a sense are unavoidable ... For instance, in appeal cases, it regularly happens that new information shows that a previous decision was wrong. Sometimes this is even acknowledged by both sides in the conflict.

Floris: But if problems are so open-ended, does that preclude a computational approach to decision making in law?

Alex: I would not say that. But the computational paradigm for law should not be that a problem is given in a formally specified form and then computationally solved.

Chris: If only things were like that in law ... Sounds easy.

Floris: You may be surprised how hard formally specified computational problems are. In fact, a significant part of the history of theoretical computer science is about determining how hard formal problems are, computationally. Computational complexity theory includes a range of classes of complexity, so rich and diverse that people have referred to them as a zoo.

Chris: Interesting. From the start of my legal training, I have always been surprised how rich and diverse the problems addressed by the law are; also a bit like a varied zoo. It is not a coincidence that legal codes count many in number and in pages. And when judges apply the law to real cases, the richness and diversity does not diminish, instead grows and grows. In a sense, the academic work of lawyers in law schools to a large extent consists in maintaining at least some level of organization in that richness and diversity. Or one could say: they try to construct such a level of organization. Academic handbooks are very much needed to keep the law a workable discipline.

Floris: That sounds a lot like what happens in the sciences! I must admit that I do not normally think of the law as a science ...

Chris: Many of my colleagues would agree. The law is considered as a normative discipline, as opposed to the descriptive disciplines of the sciences.

Alex: I would say that there is an obviously scientific part to what happens in law schools: describing and organizing what the law is. Academic lawyers collect and analyze the empirical data in the form of the sources of law, including codes, decisions, parliamentary and other relevant discussions. They form theories about what the law is, and test them by considering the data, and by comparing them to other theories. In that sense, the law is definitely scientific, just like other descriptive disciplines.

Floris: But what about the normative side of the law? That sounds very distinctive.

Chris: Indeed my colleagues and I often discuss what is right, what is just, what the norms should be. Many of our publications are in a sense about whether a legal code or decision is as it should be.

Alex: Certainly. So in terms of subject matter, the law is very different from other scientific disciplines. Surely lawyers tend to have at most a general understanding of physics and astronomy. And the same holds in general for the physicists' and astronomers' understanding of the law.

Floris: But it is not just subject matter. Physicists and astronomers study the way the world is, lawyers how the world should be.

Alex: In the way that you phrase it, you actually confirm what I say: the difference is the subject matter that is being studied. Indeed, academic lawyers do not tend to study the stars, just as astronomers do not normally study the law. And studying the law includes studying how the law should be.

Chris: I am a bit surprised. I remember the nightly discussions in my sorority. I did not really have an answer to my physics friends who were boasting that they were applying the scientific method – and that surely lawyers weren't doing that.

Alex: Of course, academic lawyers are applying the scientific method, perhaps without being aware of it. You develop theories about what the law is and should be, and test and compare those theories.

Floris: Can you test what the law should be? For me, testing has to do with experiments and data collection. That does not sound like law school business.

Alex: Things often depend on how you phrase them, how they are framed. Here is a framing of the law, suggesting an analogy with the scientific method. One could say that a legal code is a theory how the law should be, in general, and every legal decision applying that theory is an experiment testing how the law should be, in specific cases. In this sense, a theory about how the law should be can be tested using experiments and data collection. Lawyers do that all the time.

Chris: Glad to hear that my discipline is also using the scientific method! And that on an everyday basis! I look forward to the next sorority lustrum dinner. I was not really thinking of my discipline as a science, and surely my physics friends weren't.

Floris: I like the analogy, but there are also clearly differences, not only subject matter. We in computer science like our formal methods, and in the empirical

sciences, the application of statistics provides quite a methodological headache to many students.

Chris: Indeed. Although we were taught some formal skills in law school – I remember a bit of logic and some discussion of statistical fallacies – , I do not see much of that used in law practice, and not even in the academic law research that I see.

Alex: Sure. I am all for more formal and methodological training in law schools, including computer programming. A good understanding of those are helpful for all, and especially in these days where data collection of our behavior is so easy, and algorithms are so widely affecting our lives.

Floris: I guess to refer to our online lives and what we do on our phones et cetera?

Alex: Yes. It is a good thing that Europe is trying hard to develop a normative framework for handling data, algorithms and artificial intelligence. Europe is working hard on a theory of what the law about those should be.

Chris: Many consider Europe as especially mastering their law making – and there is a long history: Roman law, the Napoleonic days of codification, the Habsburg bureaucracy. And then after the WW II disaster, a grand correction by a strong emphasis on human rights.

Alex: Surely the world is bigger than Europe in the mastery of law making. Chinese law also has a long history going back to ancient times, and the USA played a significant role in developing human rights. But indeed, today many think of the USA in terms of market and military dominance, China in terms of central government and production, and Europe in terms of law and social welfare.

Floris: I recognize those views of the USA, China and Europe. Heavily simplifying, for sure, but helpful for understanding our world. Now that Europe is working so hard on regulating technology and AI, I am very worried that Europe cannot compete against the USA and China in its tech development.

Alex: Things might turn out different than you think.

Chris: Indeed! Perhaps you think – like many – that the law is hindering efficient markets and governance. But in fact the opposite is true. A significant success factor of contemporary states is how well their legal system works. Good law supports good markets and good governance. Pushing the earlier

metaphor a bit further: a European approach supports American and Chinese approaches.

Alex: I agree, but in the case of AI there is perhaps even more to it. I believe that regulating AI will lead to a new surge of innovation.

Floris: Instead of hindering it? Are you serious? If I look at what some say tech companies will need to do ... Their algorithms should be fair and unbiased, their AI applications should be transparent, their systems should be compliant to all kinds of law ...

Alex: You say it! And that will require significant innovation. Fairness, bias, transparency, compliance: these are just a few of the big technological puzzles that need to be solved. AI will have to be changed to be responsible, explainable and social; or in other words: to be good for us and our societies.

Floris: I agree that, although there are many helpful AI applications, we aren't as yet so very good at such a human-aligned AI. That requires an entirely new level of what AI can do.

Alex: Which brings us back to the field of AI & Law: it is exactly a responsible, explainable and social AI that the field of AI & Law has been working on since its early days [Verheij, 2020]. Hard problems indeed, but relevant steps have been made.

Chris: You already suggested earlier in this conversation that engineers have some things to learn in law school.

Alex: Indeed, and the other way around too. But I am thinking of something more specific: I claim that the law can lead the way in showing how two rather distinct technologies can go together: on the one hand knowledge technology (such as the rule-based expert systems mentioned earlier), and on the other hand data technology (such as the neural networks of machine learning).

Floris: Many of us are working on trying to bridge that gap. Interestingly, the sides have rather different communities. In part this may be caused by the mathematical foundations where knowledge technology traditionally has close ties with logic, and data technology with probability theory. Connecting qualitatively and quantitatively oriented methods has proven to be rather hard.

Alex: I believe that the development of argumentation technology is the way to bridge the gap [Verheij, 2018].

Chris: Argumentation? I see the law connection with its emphasis on debate. But the argumentation skills that I am aware of do not have much to do with logic, statistics or computation.

Alex: Well-noted, and it is actually the case that key developments in the 20th century history of argumentation theory are caused by the struggles to understand argumentation from a logical, a statistical and a computational perspective [van Eemeren et al., 2014]. For instance, Toulmin noted in the 1950s that the logical and probabilistic methods of that period could not properly handle real-life argumentation [Toulmin, 1958]. A growth of alternative approaches followed, with less emphasis on formal methods, and instead emphasising an informal or semi-formal study of argumentation.

Chris: If I understand you correctly, these new insights were not so much caused by computational considerations, but by the analysis of everyday reasoning and argumentation?

Alex: True. The computational perspective became relevant around the 1980s. Then it was noted that the handling of inconsistent, incomplete and uncertain information was at the same time necessary for intelligent agents (a term used for AI programs with some level of autonomous agency), and hard to handle computationally. The topic inspired the study of nonmonotonic logic, which turned out to be connected to what philosophers referred to as defeasible reasoning [Reiter, 1980, Pollock, 1987]. In a nonmonotonic logic, it can happen that a conclusion that initially seems well-supported by given information, may need to be withdrawn when further information is discovered.

Chris: In the law that is rather normal. There can always be exceptions or specific circumstances that give rise to counterarguments outweighing the initial support for a claim.

Floris: I remember the heyday of nonmonotonic logic, rather confusing. Many formalisms were proposed, often presented and investigated with a lot of mathematical intricacy, while it was not always clear what the practical import was. There was a lot of talk about Tweety, a bird first assumed to fly, but not flying after all since it was a penguin.

Alex: Good points. In the 1990s, things changed: a famous paper by Dung [1995] showed that several approaches to nonmonotonic reasoning (including Reiter's and Pollock's) could be studied by focusing only on the attack relations between the arguments used. The perspective in that paper also contained a subtle and relevant formal theory of argument attack that gave a boost to the computational study of argumentation [Baroni et al., 2020, 2018, Atkinson

et al., 2017, Rahwan and Simari, 2009, Bench-Capon and Dunne, 2007, Chesñevar et al., 2000].

Chris: But if that was all needed for modeling a bird that cannot fly ... Not so inspiring.

Alex: Things also improved for practical relevance (what in particular also Toulmin was looking for), and complex argumentation settings were addressed. For instance, it was shown that an argumentative discussion format made sense for negotiation, persuasion, deliberation and a variety of other settings [Walton and Krabbe, 1995], and that also real-life argumentation followed schematic patterns [Walton et al., 2008]. Especially also the law proved to provide much inspiration for the computational study of rich examples of the complexity of real-life argumentation and discussion [Rissland et al., 2003, Prakken and Sartor, 2015, Hage et al., 1993]. And the connections between argumentation, logic and probabilities are also ever better understood, interestingly in the context of the handling of evidence in law [Verheij, 2017].

Floris: Ok, I believe you when you say that now legally relevant examples can be computationally supported, and that that can help lawyers.

Alex: Yes. There have been proposals for argumentation support software [van Gelder, 2003, Verheij, 2005, Gordon et al., 2007]. Some nowadays speak of ‘hybrid intelligence’ when AI is used to support human intelligence, thereby augmenting what humans can do by themselves [Akata et al., 2020].

Floris: Such hybrid intelligence indeed requires a change of AI to a set of tools supporting what humans need, or what we earlier referred to as a human-aligned AI.

Chris: And, Alex, earlier you suggested in this connection that AI will have to change by looking at what lawyers do.

Alex: Yes, that is what I suggested. Very briefly, the point is that we need to develop argumentation systems that can construct and test possible answers to their problems in a critical discussion [Verheij, 2018]. In that sense, AI can and should take inspiration from what happens in everyday life, research, and law. What makes law special is that it is strongly connected to what is good for us and our society. Possible answers to legal problems have actual impact for people and for how we live together. So the stakes are significantly higher in law than in a family discussion or in a science debate.

Floris: Not *only* in law, the stakes are so high, I would say.

Alex: I agree. Also for instance in medicine the stakes are very high, and in fact there often a matter of life and death. But something that is special about the law is that more or less every aspect of human life and society is addressed in legal discussion. And if there is a situation that has not yet been discussed in law, that could always happen tomorrow. As a consequence, a deep understanding of human life and society is a core part of what makes a good lawyer. And that is still far beyond today's computers. They simply do not understand us and our world.

Floris: Commonsense knowledge remains a fundamental bottleneck for AI [Davis and Marcus, 2015].

Alex: Indeed. One more reason why I believe we should do AI as we do law. We have done AI as mathematics, as technology, as psychology, and as sociology; and now it is time to do AI as law.⁸

Floris: I remember that in 2007 we discussed that AI & Law is 'AI-complete': having solved the problems of AI & Law implies having solved all AI problems. You now suggest that in fact we *have* to find inspiration in law.

Alex: Yes. There are significant puzzles concerning knowledge, reasoning, learning and language; all key topics of AI. And how these are handled in the law can and should inspire new AI solutions. Concrete innovations can be expected in computational argumentation, schemes for arguments and scenarios, encoded norms, hybrid rule-case systems and computational interpretation [Verheij, 2020]. Like in the law, AI will require permanent checks and balances for optimally supporting us and our society. And before AI systems will pass the Van den Herik test, they will become much more like us, much more human-aligned.

Chris: Very nice. I look forward to those developments. I started today's conversation with an expression of fear, but I come out of it intrigued and inspired – and no longer scared.

Alex: Van den Herik will be happy to hear this – after his long history and unfading stamina of spreading the word on AI & Law – and pushing its boundaries. By AI & Law's longstanding efforts for developing a responsible AI, the field is today more relevant than ever.

⁸ Toulmin [1958] used this series of metaphors for the study of reasoning. The adaptation to the study of AI is by Verheij [2020].

REFERENCES

- Z. Akata, D. Balliet, M. de Rijke, F. Dignum, V. Dignum, G. Eiben, A. Fokkens, D. Grossi, K. Hindriks, H. Hoos, H. Hung, C. Jonker, C. Monz, M. Neerincx, F. Oliehoek, H. Prakken, S. Schlobach, L. van der Gaag, F. van Harmelen, H. van Hoof, B. van Riemsdijk, A. van Wynsberghe, R. Verbrugge, B. Verheij, P. Vossen, and M. Welling. A research agenda for hybrid intelligence: Augmenting human intellect with collaborative, adaptive, responsible, and explainable artificial intelligence. *Computer*, 53(08):18–28, August 2020.
- K. Atkinson, P. Baroni, M. Giacomin, 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. Giacomin, and L. van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, London, 2018.
- P. Baroni, F. Toni, and B. Verheij. Introduction to the special issue ‘On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games: 25 years later’. *Argument & Computation*, 11(1–2): 1–14, 2020.
- T.J.M. Bench-Capon. The need for good old fashioned AI and Law. *Jusletter IT*, 21:23–35, 2020.
- T.J.M. Bench-Capon and P.E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.
- C.I. Chesñevar, A.G. Maguitman, and R.P. Loui. Logical models of argument. *ACM Computing Surveys*, 32(4):337–383, 2000.
- E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in Artificial Intelligence. *Communications of the ACM*, 58(9):92–103, 2015.
- 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.
- A. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. The MIT Press, Cambridge (Massachusetts), 1987.
- 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.
- J.C. Hage, R. Leenes, and A.R. Lodder. Hard cases: a procedural approach. *Artificial intelligence and law*, 2(2):113–167, 1993.
- D.R. Hofstadter. A coffeehouse conversation on the Turing test. In *Metamagical Themas: Questing for the Essence of Mind and Pattern*, pages 492–525. Basic Books, New York (New York), 1981.
- J.L. Pollock. Defeasible reasoning. *Cognitive Science*, 11(4):481–518, 1987.
- H. Prakken and G. Sartor. Law and logic: A review from an argumentation perspective. *Artificial Intelligence*, 227:214–245, 2015.
- I. Rahwan and G. R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer, Dordrecht, 2009.
- R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- 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, K.D. Ashley, and R.P. Loui. AI and Law: A fruitful synergy. *Artificial Intelligence*, 150(1–2):1–15, 2003.

- C. Steging, S. Renooij, and B. Verheij. Discovering the rationale of decisions: Towards a method for aligning learning and reasoning. In *Proceedings of the 18th International Conference on Artificial Intelligence and Law (ICAIL 2021)*, pages 235–239. ACM Press, New York (New York), 2021.
- S.E. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, 1958.
- H.J. van den Herik. *Kunnen Computers Rechtspreken?* Gouda Quint, Arnhem, 1991.
- H.J. van den Herik. Computers kunnen *Recht* spreken. In H. Franken, editor, *Privaatrecht in de 21e eeuw – Rechtsinformatica*, pages 11–26. Kluwer, Den Haag, 1999.
- 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.
- T. van Gelder. Enhancing deliberation through computer supported argument visualization. In P.A. Kirschner, S.J.B. Shum, and C.S. Carr, editors, *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*, pages 97–115. Springer, 2003.
- 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. A coffeehouse conversation on the Van den Herik test. In *Liber Amicorum ter gelegenheid van de 60e verjaardag van Prof.dr. H. Jaap van den Herik*, pages 155–163. Maastricht ICT Competence Center, Maastricht, 2007.
- 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, 2017.
- B. Verheij. *Arguments for Good Artificial Intelligence*. University of Groningen, Groningen, 2018. URL http://www.ai.rug.nl/_verheij/oratie/.
- B. Verheij. Artificial intelligence as law. Presidential address to the Seventeenth International Conference on Artificial Intelligence and Law. *Artificial Intelligence and Law*, 28(2):181–206, 2020.
- D.N. Walton and E.C.W. Krabbe. *Commitment in Dialogue. Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany (New York), 1995.
- D.N. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, Cambridge, 2008.