

A Coffeehouse Conversation on the Van den Herik Test

Bart Verheij*

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

Chris: Alex, I want to thank you for suggesting that I read Van den Herik's "Kunnen computers rechtspreken?" [4]. It's a wonderful piece and certainly made me think — and think about my thinking.

Alex: Glad to hear it. Are you still as much of a skeptic about applying artificial intelligence to the law as you used to be?

Chris: You've got me wrong. I'm not against artificial intelligence; I think it's wonderful stuff — perhaps a little crazy, but why not? I simply am convinced that you AI advocates have far underestimated the minds of lawyers, and that there are things a computer will never, ever be able to do. For instance, can you imagine a computer writing a volume in the Asser series? The richness in content, the complexity of the considerations —

Alex: Rome wasn't built in a day!

Floris: Hey, are you two going to clue me in as to what this text by Van den Herik is all about?² It's only available in Dutch, you know!

Alex: It is about the question whether computers can decide legal cases. Van den Herik discusses the nature of legal decision making and the prospects of automating it. The text is Van den Herik's 1991 inaugural address, delivered upon acceptance of his position in Leiden. It also contains what might be called the *Van den Herik test*.

Floris: Well, automating legal decision making should be easy! Just feed the computer all legal code books and you're done. And the computer will surely do better than a lawyer since computers will not forget a rule or make a mistake when applying the rules.

Alex : Wait a minute. It's not *that* easy and Van den Herik knows this. Although he is clearly an optimist, he is very well aware of obstacles on the road towards computer judges.

* Bart Verheij is assistant professor at the Artificial Intelligence department of the University of Groningen. Web page: <http://www.ai.rug.nl/~verheij/>.

¹ The characters are purely fictional.

² The foregoing part of this text is an almost literal citation of the beginning of Douglas R. Hofstadter's 1981 *A Coffeehouse Conversation on the Turing Test* [7].

Floris: I don't understand. The law is a formal system consisting of rules. Let a computer simply apply them.

Chris: If only it were *that* easy. Your view of the law is rather simplistic. I do not think of my work as simply applying the rules of the system.

Floris: Well ok, you also have to *find* them of course. And given all these code books — I've seen some law student friends struggling with them ... I feel lucky to have chosen another field! How can anyone remember all that?

Chris: No, its not just that — though you're right that sometimes finding the right legal sources is extremely tedious. We still use a lot of paper editions, you know! If you have an unlucky day, your opponent in court presents relevant material that you did not find, perhaps simply because you didn't have easy access to a particular verdict.

Floris: I assume that you use electronic databases which simplify your life?

Chris: Sure — but probably not to the extent one would expect today, with the ubiquity of the internet and all. Our profession likes its traditional routines. But anyway: things are changing in this respect. The availability of source material in electronic form is on the rise, and the accompanying possibility of full text search has significantly simplified finding what you need. The issue is more that it is not always clear what the rules really are.

Floris: Wait a minute! I may be naive, but isn't a fundamental principle of the law that the rules must be known in advance? We don't want that judges can make their own rules! They should be written down somewhere.

Chris: You are right. That's the heart of what is called the rule of law. The idea is that the state and its entities — which include judges — are bound by existing, written law. The state can do nothing except what is legal. So for instance a legislative body can only make laws within the boundaries of the law as it is. Moreover it must follow the appropriate legal procedures.

Floris: But this seems to support my point of view. If I understand you correctly, the rule of law implies that judges must decide by applying rules that already exist and by following existing procedures. So a judicial automat can be built by implementing these procedures and rules.

Alex: This makes me smile. Analogously, one could say that people are bound by the laws of physics — which I would say are fixed in some way or another — and that it is hence easy to construct a human automat. But the history of artificial intelligence has amply shown that that is far from easy. And I would say that most artificial intelligence researchers agree that our current knowledge of the laws of physics — though imperfect — is sufficient and at least is no real impediment to building artificially intelligent agents.

Floris: There you have a point. Let me try to reshape my doubts. Isn't legal thinking simpler than human thinking in general? I would say that we have made the law ourselves, in contrast with the laws of physics.

Chris: Now you've got me. The law is at least to a large extent a human construction. We have a term for that: positive law. Positive law is law that can be traced back to a written source. I do believe there is a place for what is called — often rather loathingly — natural law, but these days the major role of positive law is obvious. I don't have a quick answer to this one.

Alex: It is this intuition — namely that the law consists of rules made by humans — that formed the basis of early work on building legal expert systems. In fact, expert systems are a well-understood technology and are one of the success stories of artificial intelligence in general, not only in the law. And they already were successful at the time of Van den Herik's inaugural address. Expert systems, in the law and elsewhere, are a useful tool in areas that consist of explicit, specifiable rules. Unfortunately, this phrasing also expresses the limitations of the technology. Classical expert systems are successful to the extent that the knowledge and procedures implemented in them are specifiable by rules. It is not a coincidence that legal expert systems are most successful in areas where the rules are precisely and explicitly available, for instance in certain parts of administrative law, such as social security and welfare.

Floris: I see what you mean, Alex. And you seem to suggest that the law is not in general specifiable in terms of rules?

Alex: That question is not really answerable. Perhaps the law is specifiable in terms of rules — in principle. But what counts here and now is that as yet the law has not been fully specified in terms of formal rules, and to the extent that it has been, the process wasn't easy.

Chris: This reminds of my legal theory courses. There was talk about vague terms and things like that. For instance, in the Netherlands, the legal consequences of contracts are constrained by what is reasonable and fair. Can you specify what is reasonable, let alone what is fair? I can't. Perhaps the most interesting lesson that I learnt from legal theory is that the law uses vague terms, not because of laziness or sloppiness, but by necessity. One could say that it is an asset of the law that it is not fully specified.

Floris: In computer science there's a stock phrase for this: that's not a bug in the program, it's a feature. So you say that the underspecification of the law is a feature and not a bug? Doesn't that contradict the idea of the rule of law?

Chris: Yes, it's a feature, and, yes, there is a tension with the rule of law. And this is not only because the law uses evaluative terms such as reasonableness and justness, that are by their very nature open. Also apparently straightforward terms such as "vehicle" are open-textured. The rule "No vehicles are permitted in the park" is a well-known example that occurred in a discussion between

the legal theorists H.L.A. Hart and Lon L. Fuller. Cars are easy. But bicycles? Roller skates? Airplanes? Fully functional World War II trucks that are used as a memorial exhibit? And to use a modern example: what about the omnipresent mobility scooters?

Floris: Good point. Have you heard about the Segway? That's a nice example that has been in the news recently. At first the Segway was tolerated on Dutch public roads, but since the beginning of this year it was banned. Only recently the Minister of Transport, Public Works and Water Management made a promise to the House of Representatives that people with lung and heart diseases are allowed to use the Segway on the pavement, awaiting an investigation concerning traffic safety.

Chris: Your example is a nice one. But we are moving away from our subject: the possibility of computer judges.

Alex: You are right, Chris. By the way, we have stumbled upon the first two obstacles to building computer judges that Van den Herik mentions in his inaugural address: vague terms and open textured codes.

Chris: Is there a way to circumvent these obstacles?

Alex: There is. These days researchers focus more on decision-making aids instead of on systems that can decide cases autonomously.

Floris: And do they succeed in these scaled-down ambitions?

Alex: This has certainly led to new and interesting directions of research. For instance, several researchers are developing tools that can visualize reasoning and argumentation. It is hoped that argument diagramming can support lawyers to develop and organize their lines of argument [15]. These researchers can build on a long tradition of argument diagramming, of which perhaps Stephen Edelston Toulmin's argument model is most widely known [6]. In the field of evidence law, John Henry Wigmore's diagramming approach, developed in the early twentieth century, has found followers in the so-called New Evidence Scholarship.

Chris: That indeed sounds interesting. If I think of the piles of paper that some of my dossiers consist of! Organizational aids sure would come in handy.

Alex: In fact, there is also ongoing work that integrates argument diagramming with case management tools. In this connection, there is not only attention for the application of legal rules, but also for investigating and establishing the facts of a case [1].

Floris: I am sorry to repeat myself, but my question whether legal thinking is easier than thinking in general has still not been answered.

Alex: My answer to that question is: no. I would say that the problem of building judicial automats is “AI complete”. If you can do that, you can build any kind of artificial intelligence. The key reason why I think this is that practically any kind of human endeavor can become the subject of a legal debate. Hence understanding, dealing with and evaluating all kinds of human behavior is part of legal thinking. Learning, creativity and use of language are also necessary skills.

Chris: From what I know of computers: we have a long way to go before something like a computer judge will ever be built. For instance, have you ever seen computers heatedly debating an issue, as we do now? But that’s the law’s core business one could say! The law is an adversarial domain, in which often two parties defend their own position and attack each other’s.

Floris: Good point. This makes me wonder whether there is a useful way in which a computer can change its mind, or in which one machine persuades another. People do that all the time!

Alex: One way of looking at this would be that change of mind in a computer is just change of the data stored in its memory. Similarly one could say that a computer can learn new behavior when new software is loaded. Persuasion of one machine by another would involve that machines can communicate with each other and that as a result of such communication a machine changes its mind. Today’s agent-oriented software technology is in fact based on taking this idea of machines talking to each other seriously.

Floris: Sounds far too trivial for me. Is that all there is to change of mind and persuasion?

Alex: Of course not, but I would say that something like this is the core of change of mind and persuasion. What is missing is filling this core with content, and that has for long been one of the most stubborn issues in the field of artificial intelligence. We call that the knowledge acquisition bottleneck.

Floris: I must think of my logic courses. It was presented as formalized reasoning, but the connection with real reasoning wasn’t all that clear. I recall that it was often not easy at all to express a simple piece of reasoning in logic.

Alex: Don’t forget that you have only scratched the surface of logic. Logic is a vast field these days. And things that are hard to express in one system or are not all expressible, can be easy in another.

Chris: Would you say that there is a logical system capturing legal reasoning?

Alex: There certainly are systems that are closer to legal reasoning than for instance standard first order predicate logic [2,3,10]. And the argumentative, adversarial nature of the law is also better understood nowadays [12,13]. In a sense, we have reached the stage where speaking of a legal logic makes sense [11,14]. I must admit that much of this work remains rather theoretical in nature. There is not much that has been fleshed out in implemented software.

Chris: This all sounds very interesting, but also hard to believe, and also a bit beside the point. For until now we have presumed that the law exists, and can be applied once we have found it. But what we did not discuss is that the law changes all the time. For instance, let's for the sake of argument assume that we have a computer judge that can apply existing law. But then we do not have all of legal decision making! For instance, the Supreme Court does not simply apply existing law. It makes a habit of *changing* the law every now and then.

Alex: Sure. But in a limited sense. You know very well that the Supreme Court is not a formally legislative body. Hence its law creation capacity is strongly limited to say the least.

Chris: True. But still. Every now and then, these limitations leave enough room to lead to a landmark case with an unexpected result. Of course the implied change of such a case is normally limited by for instance known legal principles.

Alex: Did you know that there exist logical analyses of the role of principles in legal reasoning [16]? Even case analogies can be treated [8,9]!

Chris: I did not. But you are not going to tell me that these analyses of principles and analogy can *predict* landmark cases by the Supreme Court.

Alex: No, but we can reconstruct what has been going on. It is possible to *redo* the essential steps in the reasoning that underlie a landmark case. Surely this is always looking back: the reconstruction can only be done since the core ingredients of the reasoning — such as the rules, principles, value judgments and reason weighings — have been made explicit.

Chris: You seem to suggest that the Supreme Court specifies the essential elements of its reasoning in its decisions. If only that were true... The hermeneutics of Supreme Court decisions is an industry that earns legal professors a good living.

Floris: Wait a minute! This suggests that we demand more of computers than what we expect from us humans. A landmark case is not predictable by humans, so it is strange to ask that of machines. What seems to count more is whether a machine can *decide* a landmark case.

Alex: I agree. Perfect prediction is not a necessity for acceptable legal decision making. The test is more whether the outcome is good.

Chris: You forget that it is not just the outcome that counts. In the law it is also important that the procedures are followed correctly. Or perhaps it is better to say: that they are followed acceptably. We already discussed the open-texture of legal norms. Procedural norms can be similarly underspecified.

Floris: But doesn't that preclude the building of computer judges? After all, the computation in a machine is deterministic when we know its initial state. This suggests that machines can only follow perfectly specified procedures.

Alex: True, but don't forget that this only implies that what a computer *actually* does is deterministic in this sense. This does not imply that the *legal procedure* of which it is an implementation is fully deterministic. As long as what a computer does remains within the boundaries of acceptable legal procedures everything is ok.

Floris: But this makes things even harder! We already discussed that the rules about legal content are underspecified. Now we also must work with underspecified procedures. Doesn't this mean that the whole project of the building of computer judges falls apart?

Alex: It doesn't. It only means that the computer metaphor is not as simply relevant as one might naively hope. And that is no surprise. Moreover, as I already said when we discussed AI completeness: building a computer judge is equally hard as building machines that perform general tasks that require intelligence. For instance, both require basic linguistic competence. Automating that is still out of our reach.

Chris: I would say that it is then better to start with ordinary intelligence. When *that* problem has been solved you AI folks can turn to legal intelligence.

Alex: I understand your point, but disagree. It is my conviction that research in AI & law can *help* solving the problems of AI in general.

Chris: I like that!

Floris: Sure, you take pride in being a lawyer. And this suggests that legal thinking somehow includes or supersedes thinking in general.

Alex: In a way I would say that it does, but that should not increase your professional pride, Chris. I mean something else. I believe that the field of AI can learn from the field of law since in a way *the law has struggled with AI problems throughout its whole existence*.

Floris: Excuse me?

Chris: This sounds too good to be true... I didn't know that I was an AI whiz!

Alex: Well, I guess that that is not true of course, but I don't know what you read in your spare time. What I want to say is that the fields of AI and law struggle with an analogous difficulty: both try to get a grip on the real world using formal tools. Clearly the two fields do that in different ways, with different biases and they have different kinds of means at their disposal. Whereas AI uses computers and algorithms, the law uses rules and procedures. Nevertheless it can be said that both AI and law look for ways to deal with the limitations of formalism. The law has been doing that for a couple of millennia, though, and AI only for some decades. Let me mention one other trump card of the field of law: the legal profession does not have much time for armchair philosophizing. There are actual problems that have to be addressed. Cases must be decided and clients don't want to pay too much.

Floris: In this connection there is a nice radio add of a law firm these days. The slogan they use is: do you want to pay for the development of case law or do you want a solution of your problem?

Alex: That's exactly what I mean. As a result, lawyers tend to solve hard problems concerning the tension between the formal and the material world in an extremely pragmatic way. They are satisficers by nature.

Chris: Satisficers?

Floris: Herbert Simon coined that term in the early days of AI. It means something like choosing what is good enough, not what is best. The term is used to highlight the importance of resource constraints for decision making.

Alex: I am reminded of the slogan: the better is the enemy of the good. Can be a hard one to apply!

Chris: I have to go in a minute or so. I like the discussion, but there is a client waiting for me.

Floris: One final question: did we address all obstacles to building computer judges mentioned by Van den Herik?

Chris: Let's see. He discusses the vagueness of norms, the open-texture character of legal codes, the role of interpretation, and the requirement of system maintenance. I would say that we have at least touched upon each of these!

Alex: I agree. Let me add two more obstacles: the creation of law and empathy. Van den Herik discusses these in later texts [5].

Floris: We did not address empathy yet! And I think of something else: we started talking about some kind of test. We didn't talk about that either.

Alex: Correct. The Van den Herik 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 [4, p. 33].

Floris: That is indeed a nicely phrased analogy with the Turing test.

Chris: And it gives me an opportunity to end the discussion with the obstacle of empathy: I am not sure whether I like it to be judged by a machine, even one that passes the test...

Floris: Me neither.

Alex: Don't forget that it is not so nice to be judged by any judge, whether human or machine!

References

1. Bex, F. J., Van den Braak, S. W., Van Oostendorp, H., Prakken, H., Verheij, B., and Vreeswijk, G. (to appear 2007). Sense-making software for crime investigation: how to combine stories and arguments? *Law, Probability & Risk*.
2. Hage, J. C. (1997). Reasoning with Rules. An Essay on Legal Reasoning and Its Underlying Logic. Dordrecht: Kluwer Academic Publishers.
3. Hage, J. C., and Verheij, B. (1994). Reason-based Logic: a logic for reasoning with rules and reasons. *Law, Computers & Artificial Intelligence* 3(2/3), pp. 171-209.
4. van den Herik, H. J. (1991). *Kunnen computers rechtspreken?* Arnhem: Gouda Quint.
5. van den Herik, H. J. (1999). Computers kunnen *rechtspreken*. In H. Franken (ed.), *Privaatrecht in de 21e eeuw - Rechtsinformatica* (pp. 11-26). Den Haag: Kluwer.
6. Hitchcock, D. L., and Verheij, B. (eds.). (2006). *Arguing on the Toulmin Model. New Essays in Argument Analysis and Evaluation* (Argumentation Library: Vol. 10). Dordrecht: Springer-Verlag.
7. Hofstadter, D. R. (1981). A Coffeehouse Conversation on the Turing Test, *Metamagical themas: questing for the essence of mind and pattern* (pp. 492-525). New York (New York): Basic Books.
8. Roth, B. (2003). Case-based reasoning in the law. A formal theory of reasoning by case comparison. Dissertation Universiteit Maastricht.
9. Roth, B. and Verheij, B. (2004). Dialectical arguments and case comparison In T.F. Gordon (ed.), *Legal Knowledge and Information Systems. JURIX 2004: The Seventeenth Annual Conference* (pp. 99-108). Amsterdam: IOS Press.
10. Verheij, B. (1996). Rules, Reasons, Arguments. Formal studies of argumentation and defeat. Dissertation Universiteit Maastricht.
11. Verheij, B. (1999). Logic, context and valid inference. Or: Can there be a logic of law? In H. J. van den Herik, M.-F. Moens, J. Bing, B. van Buggenhout, J. Zeleznikow and C. A. F. M. Grütters (eds.), *Legal Knowledge Based Systems. JURIX 1999: The Twelfth Conference* (pp. 109-121). Nijmegen: Gerard Noodt Instituut.
12. Verheij, B. (2003). Artificial argument assistants for defeasible argumentation. *Artificial Intelligence* 150(1-2), pp. 291-324.
13. Verheij, B. (2003). DefLog: on the logical interpretation of prima facie justified assumptions. *Journal of Logic and Computation* 13(3), pp. 319-346.
14. Verheij, B. (2003). Dialectical argumentation with argumentation schemes: An approach to legal logic. *Artificial Intelligence and Law* 11(1/2), pp. 167-195.
15. Verheij, B. (2005). Virtual arguments. On the design of argument assistants for lawyers and other arguers. The Hague: TMC Asser Press.
16. Verheij, B., Hage, J. C., and van den Herik, H. J. (1998). An integrated view on rules and principles. *Artificial Intelligence and Law* 6(1), pp. 3-26.