

Special Issue in Memory of Carole Hafner: Editor's Introduction

T.J.M. Bench-Capon

the date of receipt and acceptance should be inserted later

Abstract In this introduction I give an overview of Carole Hafner's work and discuss the papers in this volume. The final section offers some more personal reminiscences of Carole and her contribution to the AI and Law community, from myself and other colleagues.

1 Introduction

With Carole Hafner's death, on 5th April 2015, the AI and Law community lost one of its pivotal figures. As well as her pioneering research on conceptual retrieval, she contributed, with her colleague Don Berman, three papers on reasoning with legal cases which are still read and discussed today. Significant as this work was, however, her role as a community builder was perhaps even more important. She not only founded the series of International Conference on Artificial Intelligence and Law, the International Association for AI and Law and the Artificial Intelligence and Law journal, but her efforts as General and Programme Chair of the Conference (1987 and 2009 respectively), as Secretary and Treasurer of the Association (1992-2003 and 2006-7) and co-editor of the journal (1992-7) did much to ensure their success.

It is thus fitting that the AI and Law community pay tribute to Carole with this special issue of the journal she founded. The issue contains four original papers discussing her work. Importantly all four demonstrate its ongoing significance. This introduction will also provide a recapitulation of her research, and a number of personal memories contributed by some of her many friends and colleagues.

T.J.M. Bench-Capon
Department of Computer Science
University of Liverpool
Liverpool, UK
E-mail: tbc@csc.liv.ac.uk

2 Work

Carole's long career fell into a number of distinct phases, and addressed several different topics. In this section I give a largely chronological summary of the highlights.

2.1 Conceptual Retrieval

Carole's initial concern was with conceptual retrieval. Retrieving relevant cases was, and remains, a traditional part of the lawyer's task and it was a task that was identified very early as one with which computers could assist. But computers were very different in those days: main memories were tiny (measured in Kilobytes rather than Megabytes, let alone Gigabytes), storage was slow, predominately tape rather than disk, and only suitable for batch access: disks were small (in capacity), large (in physical size), fragile (requiring air-conditioning and liable to head crashes and other forms of corruption) and expensive; communications were limited (e-mail was on limited networks, and attachments were unknown - documents were transmitted by fax), and even VDUs were scarce¹. As for the internet, that was still more than twenty years away. A world where memory was no longer seen as a limitation, where one could store gigabytes of data on a memory stick, where all computers were linked for e-mail and the world wide web with digitised documents deliverable to one's own computer, and where screens were coloured and bit-mapped was beyond the imagination of most of us. In this context the best that could be done with text was to form an index by inverting the text files and search on keywords. A summary picture of the view from 1984 can be found in [21] which discusses Jon Bing's groundbreaking [17]. Problems with simple keyword search are fairly clear: precision is lowered by the fact that the meanings of words depend significantly on their context, and recall is affected by the use of synonyms and circumlocutions, both common in law. What was needed was (and still is) a way of retrieving cases according to the concepts they address (similar perhaps to the relevant dimensions (as in HYPO [3]) or factors (as in CATO [3]), but without the extensive analysis of cases required by these systems. The difficulties attendant on identifying the factors present in cases using a boolean expression composed from keywords can be seen by appreciating the difficulties in automatically assigning factors to cases even when sophisticated machine learning techniques are used [19].

After her doctoral thesis [24], Carole published a related paper, [25], and her doctoral project was one of the six AI and Law applications described in [20], which proved quite influential (see Marek Sergot's historical survey in [10]). This work also formed the basis of her paper at the first ICAIL [26], published after she had taken up a post at Northeastern University in Boston following a stint with General Motors Research Laboratories, where she seemed to have worked mainly on databases and natural language interfaces to them (e.g. [28]).

Carole's solution to the conceptual retrieval problem, which she called *Conceptual Organization of Cases* (COC) involved the construction of three structures:

¹ Even in 1979, when my wife started as a programmer in a national retail organisation, all programming was done on punched cards, and introduction of (monochrome, character-based) VDUs was considered a huge leap forward.

- a domain knowledge model that defines the basic concepts of a case law domain; including actors, events and relationships from the particular domain;
- a legal case frame that represents the properties, roles and relationships of elements of a case and which particular cases instantiate;
- Issue/Case Discrimination Trees that represent the significance of each case relative to a model of the normative relationships of the legal domain.

The first two would probably be termed ontologies nowadays, while the third reflects the argumentation which led to the decision. A discussion of [26] by Adam Wyner is included in [8].

At Northeastern Carole began to work with Don Berman and the next phase of her work was carried out with Don, and focused on the exploration of reasoning with legal cases. We will discuss this work in the next section.

2.2 Case Based Reasoning

Although their main contributions were to be in the field of legal case based reasoning, Carole's collaboration with Don Berman began with a consideration of rule-based approaches. Rule based systems, often based on logic programs representing legislation ([39] is an early inspiration although a less bullish account of such systems is given in [11]), were on the rise in Europe, and a very early Berman and Hafner publication [12] identified some limitations with such systems. Their main concerns were the monotonicity of logic contrasted with the ineluctable non-monotonicity of law and the relation with inconsistency. Logic was seen brittle in the face of inconsistency, whereas law needs to be robust, since inconsistencies in fact abound in law. This may well have overstated the role of the declarative as opposed to procedural semantics in the languages used to implement programs such as [39], but did express some widely held reservations². Another early paper by Berman and Hafner [13] gave a more pragmatic view of legal rule based systems, recognising their practical potential while not losing sight of their theoretical limitations.

It is, however, for their work on legal case based reasoning that Don and Carole's contribution will be remembered. This was presented through papers presented in three consecutive ICAIL conferences. The first, [14] was presented at Oxford in 1991, the second, [15] in Amsterdam in 1993 and the third [16] in Washington in 1995. The three papers were consolidated in the special issue of this journal [27] marking Don's death in 1997. All three papers reacted to the style of case based reasoning exemplified in HYPO and CATO: what has come to be known as factor based reasoning ([37], [9], [33]). In factor based reasoning cases are represented as bundles of factors (stereotypical patterns of fact, present or absent in a case and favouring one party or the other). While factor based reasoning does indeed address one very important aspect of reasoning with cases, this form of representation loses a lot of information about the cases which may well be necessary to account for the decision properly. Each of the papers looked

² Don and Carole presented [12] at the later of two workshops held in Houston in 1984 and 1985 under the auspices of Charles Walter. These workshops, together with the early Florence conferences (see Tom Gordon's memoir in section 4.3), brought AI researchers together and sowed the seeds which grew into the ICAIL series of conferences.

at the importance of a different sort of information discarded by the factor based approach.

In [14] they explored the effect of procedural context. Their basic contention was that “procedural context is one of the factors that determines to what extent a prior case is relevant as a precedent in a current case” (page 17) and that whether two cases could be deemed similar was dependent on the procedural context as well as the factors present in the cases. As with all their papers, the claims were supported with well chosen examples. In [14] they identified thirty-two different procedural scenarios which they categorised into ten basic types of legal result. The paper was in particular concerned with whether, in the context, a particular question was a matter of fact or a matter of law.

In [15] they made the teleological considerations, the social purposes underlying the law, the focus. One abiding problem with factor based approaches is what to do when there is no precedent that matches. Although some (e.g. [31]) are willing to accept that, based on factors, only *a fortiori* reasoning is possible, most people have wanted to go beyond what is strictly implied by the set of available precedents. The answer to the problem given in [15] is that the answer is to be found in the purposes of the law. In the main case in the example, *Pierson v Post*, there is a conflict between the need to decide cases so that the law may be clear (often termed a *bright line*) and various activities that might be encouraged (hunting vermin in *Pierson*). The argument was that while hunting vermin was insufficiently worthy of encouragement, it was worth compromising the bright line in order to protect a person pursuing his livelihood (as in *Keeble v Hickeringill*). This paper has been enormously influential, perhaps because the question of how we can go beyond what is contained in the precedents has to be answered if we are to deploy factor based reasoning. Additionally the set of cases stemming from *Pierson* have become a widely used example in the AI and Law literature, e.g [4].

The third paper [16] concerned the temporal context and recognised that legal doctrine based on case law evolves over time. Whereas in factor based reasoning precedents are considered as a set, in practice they arrive in a sequence, and the order of their arrival may have an influence on the legal doctrine that emerges from the precedents (a topic not much discussed later but see [32]). As well as identifying the importance of taking this evolution into account when using precedents, they gave five warning signs (*red flags*) which could be used to alert people to the possibility of imminent change.

All of these papers followed a similar methodology. First a description of the legal problem, backed with real cases. Then an analysis of how the problem might be solved, and finally a discussion of how this solution could be made computational. This gives the papers a balance exceptional in AI and Law work: a real legal problem, described with sufficient detail, and a feasible computational solution. The papers thus show the strengths of close collaboration of a lawyer with a real passion for the role AI might play in law and a computer scientist with a commitment to providing practical assistance to lawyers. The result was a genuine synergy and a paradigm of what interdisciplinary research can be³. Unfortunately Don’s untimely death ended the collaboration and the 1995 paper was the last they were able to write together.

³ When presenting the papers at ICAIL both would participate, Don outlining the problem and Carole presenting the proposed computational solution.

2.3 Ontology

While Carole was working with Don on legal case based reasoning, one of the major developments in Knowledge Representation was the emergence of so-called *ontologies* [22]. Ontologies, in the sense of formal representations of domain conceptualizations, designed to enable reuse and sharing of knowledge bases, quickly established themselves as a tool for building knowledge based systems, and soon started to appear in AI and Law ([41], [47],[46]).

The ideas underlying ontologies were very close to some of the components of Carole's conceptual retrieval system. Indeed ontologies also came to be used for information retrieval [23]. Having worked on ontologies before they were even given a name, it is unsurprising that Carole took an interest in them. It is, however, perhaps surprising that she never applied them to AI and Law.

Carole did her work on ontologies while supervising a PhD student, Natalya Fridman Noy. Their publications include [30], [34] (Carole's most cited paper), [35], and [36]. Noy went on to join the prestigious ontology group at Stamford, where she had a very distinguished career before becoming a staff scientist at Google.

After this work with Noy, Carole became far less research active, although she did some work on legal drafting, and published in the Jurix conference of 2007 [29]. She did, however, continue to play a very active role in AI and Law community and was Programme Chair of ICAIL in 2009.

3 Papers in this Issue

This issue contains four papers discussing Carole's work: one on her conceptual retrieval work and three building on her work on reasoning with legal cases.

3.1 Taal et al

This first paper is written by a group of authors who are all based in New York and working for banks or a global law firm. The paper looks again at Carole's early work on conceptual retrieval. Their starting point is that the research question was very good and the techniques proposed had clear potential to provide a valuable solution, but the techniques were not taken up in practice, and were not applied by the legal case indexing services that came to dominate the electronic case law market. Since Carole did her research, however, the environment has changed greatly. Not only has the technology generally improved in the ways described in section 2.1, but three developments in particular suggest that the time may now be right for Carole's COC to move into real use. The three developments considered in particular are:

- The availability of the material in the form of digitised case law and supporting matter;
- The analytical protocols developed in computer science
- The strides made in Artificial Intelligence techniques, such as IBM's Watson and the very recent successes in the game of Go, as made by Google's Deep Mind initiative.

It is by no means rare for research prototype to have to wait for the technology to catch up with the ideas before it becomes commercially feasible. The Xerox Star point and click with a desktop metaphor was developed in 1981, but only began to take off with the launch of the Apple Mac. And even then the technology was being pushed beyond its limits: the need for an external drive militated against the general concept and the amount of physical disk swapping was, to say the least, irritating. But the technology did catch up, and now the basic interface idea is ubiquitous. What Taal *et al* show is that the problem tackled and the solution proposed remain commercially relevant.

This paper thus provides an interesting take of Carole's early work, seen from a contemporary, commercial, perspective.

3.2 Loui

The second paper in this volume centres on the second of Berman and Hafner's three legal case based reasoning papers [28], but much of what it says can also be applied to the other two papers. Loui sees the papers as providing an ontological challenge to the style of reasoning captured by HYPO, CATO and CABARET [40]⁴. The challenge should be considered *ontological* because each of the papers invites us to represent knowledge lost in a factor based representation of cases: knowledge which is, in each of the three types discussed, part of what makes legal reasoning distinctive. The knowledge identified in [28] is the purpose of the law, or of a decision in a case, the values that the decision promotes and the kind of society that the law is intended to produce. Loui offers a clear statement of the challenge posed in particular [28], and how AI and Law researchers have responded to that challenge and tried to incorporate this knowledge into their representations of cases.

Following this Loui shows how the challenges like those posed by Berman and Hafner continue to stimulate researchers, and how they have spread beyond AI and Law to jurisprudence. In particular he discusses a very recent article [5] (still forthcoming as this goes to press), by two young jurisprudes. These authors, Loui argues, take inspiration for their account of how law should be interpreted from ideas that have developed, and become widespread, in AI and Law, and that addressing these problems could challenge AI and Law researchers to deeper ontological modelling, which will take us beyond the challenges posed by Berman and Hafner. For Loui the work of [5] raises (in the words of [28]) "many more questions than it answers questions that ought to be the subject of much more research by the AI and Law community".

3.3 Verheij

As the title of Verheij's paper suggests, he is interested in representing the values and purposes in [28], but, as he stresses, these ideas are also crucial for the account of evolving case law doctrine given in [16], since the evolution of doctrine is seen as bound up with the purposes and values that the law is supposed to serve, and

⁴ Loui has provided an excellent and succinct account of the reasoning exemplified by CABARET in section 3.4 of [8].

their evolution. Accordingly, the example explored in Verheij's paper is taken from [16]. Verheij applies techniques he has developed [42], [43], [44] connecting qualitative and quantitative primitives for evidential argumentation to value-guided argumentation.

Verheij gives a very detailed account of how his techniques apply to the series of cases in [16], which illuminates both the cases and the techniques themselves. This also casts light of the signs of precedent weakening ("red flags") identified by Berman and Hafner.

Verheij further underlines the importance of value-based reasoning by suggesting that it may enable the development of ethical system design, which will become increasingly important as autonomous agents become more pervasive and more powerful. The importance of values was recognised by its selection as a special theme of ECAI 2016⁵, and the workshop *AI for Justice*, held at ECAI 2016, which included several papers on this topic, including [45] and [6].

3.4 Al-Abdulkarim et al

The fourth paper also responds to [16]. The ANGELIC project developed a methodology [1] for the encapsulation of case law doctrine in the form of an Abstract Dialectical Framework [18]. Once a body of doctrine has been encapsulated, however, as is clear from [16], it must be possible to accommodate evolution and change, and not just to modify the representation *post hoc*, but also to represent possible impending changes. The paper by Al-Abdulkarim discusses these issues in terms of the ANGELIC methodology, by looking in detail at the variety of ways changes can modify the ADF.

ADFs provide a versatile knowledge representation format, which combines a tree-like structure, capable of capturing both the factor hierarchy of [2] and the decision tree structure used by Rissland and Friedman in their paper of changing case law [38], with acceptance conditions, which in the ANGELIC methodology take the form of a highly partitioned logic program. This means that the ADF can be modified by any of:

- Adding and Deleting nodes;
- Adding or Deleting acceptance conditions;
- Adding or Deleting elements of existing acceptance conditions;
- Reordering acceptance conditions.

These different responses are described and their appropriateness as a response to the various types of change noted in [16] and [38] is discussed. The change actions are also related to the red flags of [16], and to changing values and purposes, which is seen in [16] as a main driver of change in doctrine.

3.5 Discussion

What all four of the papers demonstrate is the continuing relevance of Carole's work for both practitioners and academics. Whether it seen as a source of techniques awaiting re-imagination in terms of technological developments, as a source

⁵ <http://www.ecai2016.org/calls/call-for-papers/>, accessed 21/10/16. Papers included [7].

of research challenges, or as raising issues that still need to be addressed in proposed developments in AI and Law methodologies and systems, Carole's body of work cannot (or at least, should not) be ignored by contemporary researchers.

4 Person

Carole's contribution to the field of AI and Law goes far beyond her publications. She was instrumental in founding several of its institutions, and it is to a considerable extent her doing that we have an AI and Law *community*. She played a key role in establishing the three important pillars of the community: the *International Conference on AI and Law*, the International Association for AI and Law and this journal, *AI and Law*. As Programme Chair of the fifth conference, President of the association from 2002-3 and a current editor-in-chief of the journal, I have particular reasons for be grateful to Carole for her part in establishing this organisational framework. Of these I personally feel the conference is the most significant: I have been to all fifteen so far held, and one of the great pleasures is meeting old friends and colleagues, which for the first fourteen included Carole. It is also a place to meet people who will become friends and colleagues: I first met many of my AI and Law friends at ICAIL. I feel that without the conference there would not be an international AI and Law community (although the Jurix conferences play an important role for Europeans, Americans often find it hard to attend. Carole did manage to get to Leiden in 2007, but that was the only time).

I first met Carole at the first ICAIL in Boston 1987. At the time I was only a post-doc, but because Marek Sergot, who was then a key figure in UK AI and Law, was unable to attend, I represented him at meetings to discuss the future of the conference. This gave me immediate access to the "inner circle" and I was surprised at how welcoming and receptive Carole and the others were to a newcomer to their field. This is not always the case at Computer Science conferences, and I hope this attitude continues to be a feature of ICAILs. ICAIL has always seemed to me a model of the way a small, focused conference should be organised. In section 4.2 Thorne McCarty identifies some of the guiding principles that Carole insisted on, and I think these are a significant reason for the success of the conference. When involved in setting up and organising the first COMMA (Computational Models of Argumentation) conference I was much influenced by these principles, especially that it should be every two years and comprise plenary rather than parallel sessions. COMMA was held for the sixth time in 2016, and I like to think that the adoption of Carole's principles has been one of the reasons for its continuing success.

Financial stability was an important concern of Carole, and when I was Programme Chair of ICAIL in 1995 we had some debate as to whether a poster session was desirable: Carole felt that this would increase attendance and offer more opportunity for people to present their work, but I was not in favour (I still don't like poster presentations, neither the giving nor the receiving). Our compromise was that I would have only one invited speaker, making more room for submitted papers, and reducing the expenses. But, with hindsight, Carole may well have been right: certainly poster presentations are an apparently popular part of many conferences today, and at least one invited speaker a day now seems to be the norm.

Many of us working in AI and Law have reason to be grateful to Carole, and many of us remember her with great affection. ICAIL in San Diego was not quite the same without her. In the remainder of this section other friends and colleagues contribute tributes and memories.

4.1 Tribute from IAAIL

The following tribute to Carole was written by the then President of the International Association for AI and Law, Jack Conrad⁶.

Researcher, scholar, educator, facilitator, collaborator, community-minded human being. These are but a few of the expressions we would use to describe Carole Hafner, a doyenne of our field. Carole passed away on Sunday, April 5, 2015, following a prolonged illness. It is with deep sadness that we share the news of Carole's death as she was a long-standing proponent of and active participant in the Artificial Intelligence and Law community. Carole completed one of the first PhD theses in Artificial Intelligence and Law. Her dissertation, from the University of Michigan in 1978, was titled *An Information Retrieval System Based on a Computer Model of Legal Knowledge* [24]. Carole continued to pursue state-of-the-art approaches to conceptual information retrieval following her graduate studies. She joined the faculty at Northeastern University, where she formed a particularly productive relationship with her distinguished colleague, Professor Don Berman. In 1989, one of their joint articles was featured in a special edition of the Communications of the ACM that focused on the social aspects of computing: *The Potential of Artificial Intelligence to Help Solve the Crisis in our Legal System* [13]. The paper led to the founding of the journal Artificial Intelligence and Law, which began publication in 1992 with Carole and Don as its first co-editors. Carole also served as the chair of the very first ICAIL conference, held in May of 1987 in Boston at Northeastern. At early ICAILs, Carole and Don conducted tutorials on the field of AI and Law. Most importantly, they went on to co-author an acclaimed trilogy of ICAIL papers on case-based reasoning. These papers, published in 1991 [14], 1993 [15], and 1995 [16], articulated important aspects of real-world legal reasoning and problem solving that were largely absent from earlier CBR and argumentation models: the dependence of the precedential force of judicial decisions on (1) the procedural context in which the decision is made, (2) the rationale underlying the decision, and (3) the evolving network of other decisions that over time can extend or erode aspects of a decision's holding. At several later ICAILs, Carole again gave introductory tutorials. It was important to Carole to welcome new faces to the community and share her passion about enticing aspects of the field. She served as Program Chair at the 12th ICAIL, held in Barcelona in 2009. Carole was the "moving force" behind both the ICAIL conference and our organization, IAAIL, which was established in 1992 to provide the institutional structure to manage the conferences. She served as Secretary-Treasurer of IAAIL from its inception. At the last ICAIL that Carole attended, in Pittsburgh, the IAAIL organization gave her a lifetime service award, which noted:

⁶ with thanks to Kevin Ashley, Karl Branting, Anne Gardner and Thorne McCarty for their contributions.

“On the occasion of the Thirteenth International Conference on Artificial Intelligence and Law, ICAIL 2011, the International Association for Artificial Intelligence and Law recognizes Prof. Carole D. Hafner for her tireless efforts in organizing and sustaining the Conference, the Association and the journal of Artificial Intelligence and Law. June 8, 2011.”

We will be sharing a special tribute to Carole in the near future. We send our sincerest condolences to her family and friends. We will miss her dearly.

4.2 Thorne McCarty

A tribute to Carole was delivered by Thorne McCarty during the opening of the sixteenth International Conference on AI and Law, held in San Diego in 2015.

I worked with Carole Hafner on the first International Conference on Artificial Intelligence and Law, which was held at Northeastern University in Boston in 1987. She was the General Chair, and I was the Program Chair.

One question we faced was: How do you start a conference from scratch? Neither of us had any experience doing this, of course. But here’s what you do (or, at least, what we did):

- You send out an announcement, using whatever technology is available.
- You say: Here’s a date, here’s a place, here’s a list of topics, and here’s a Program Committee. “Send us your papers.”
- And then you wait . . .

And, amazingly, the papers arrived, and the people showed up! In fact, the attendance at that first ICAIL in Boston set a high-water mark that we did not surpass until Rome in 2013.

From the beginning, Carole insisted on several principles for the conference that were important (I think) for its success:

- Schedule it every two years, not every year. This was the IJCAI model, not the AAAI model.
- No parallel sessions! This was *not* the IJCAI or AAAI model, but it is the right model for a small conference. Carole said that she was inspired by the way Don Walker organized the conferences for ACL, the Association for Computational Linguistics.
- The banquet and the lunches should be included in the conference fee. We called this “the Italian model.”
- To make the conference truly international, alternate the location between continents.

We immediately violated this last principle in 1989, when we held the second conference in Vancouver. But at least that was a different country, if not a different continent.

After the first two ICAILs, we realized that we needed an institutional structure to organize future conferences, and this led to the formation of the International Association for Artificial Intelligence and Law (IAAIL). How do you do that? We collected all the by-laws we could find, from IJCAI, AAAI, ACL, etc., but nothing seemed quite right. Again, Carole insisted on several principles.

You will see some of these principles at work in the membership meeting for IAAIL on Wednesday. For example:

- The Vice President normally accedes to the Presidency.
- The ex-President and the ex-ex-President sit on the Nominating Committee.
- And several more ...

The purpose of these rules was to balance democracy (which is good) with stability and continuity (which is also good). I think these rules have worked well over the years.

But the real source of stability and continuity for ICAIL was Carole Hafner herself. She offered to run for the position of Secretary/Treasurer for many years, a position that she held until just a few years ago, when she was succeeded by Anne Gardner. And this was how the real work got done. Carole did it!

So we are very sad that Carole is no longer with us, but this conference is her legacy, and I hope you enjoy it!

4.3 Katie Atkinson

Katie Atkinson is current President of IAAIL and was programme chair of ICAIL 2015.

I write this tribute in two capacities: first as a researcher in the AI and Law community whose own work has been strongly influenced by Carole Hafner's; second as the current President of the International Association for AI and Law (IAAIL).

I entered the AI and Law community in 2003 when working on my PhD on the topic of computational argumentation for practical reasoning. Trevor Bench-Capon guided me to the topic of AI and law, and in particular Carole and Don Berman's work on teleological reasoning in AI and law. The idea of reasoning about values is a strong thread that runs through my own research and continues within it today. When talking to legal professionals about AI tools and techniques for representing and reasoning about the law, it is this aspect of subjective reasoning that often provokes keen interest and debate. The general AI community has in recent years begun to recognise more the importance of reasoning about values (see for example, the recent track on 'AI and Human Values' at the 2016 European Conference on Artificial Intelligence). The ideas presented in Don and Carole's early work on this topic remain relevant today for researchers who are seeking to build AI systems that address real world problems.

Not only did Carole contribute original research to the field of AI and law, but she was also one of the leaders in establishing the field as an academic community. The first International Conference on Artificial Intelligence and Law (ICAIL) was held in 1987 in Boston, Massachusetts and Carole served as Conference Chair. Following the success of this event and the international interest in the topic of AI and law it demonstrated, the International Association for AI and Law (IAAIL) was established, with Carole playing an instrumental role in this, as well as the establishment of the community's journal. The remit of IAAIL is to oversee the community's biennial conference and promote activities in AI and Law, and Carole served as IAAIL Secretary/Treasurer for many years. I served as Program Chair of the conference in 2015 and to recognise Carole's contributions to the field, in

that year an ICAIL best paper prize was introduced in her honour. The AI and Law community, which Carole did so much to build, looks forward to continuing to produce exciting research that will have an impact on both general AI and on Law.

4.4 Tom Gordon

Tom Gordon has been a regular at ICAIL since 1987 and was programme chair of the seventh International Conference on AI and Law.

Carole and I first met on my honeymoon, at the second Logic, Informatics and Law conference organized by Antonio Martino, which took place in Florence, Italy, in September, 1985⁷. My wife, Ines, generous as she is, agreed to start our honeymoon by accompanying me to the conference. We were compensated with a wonderful Italian conference dinner on a terrace overlooking the city on a beautiful warm evening. And I recall us enjoying a nice organ concert in a cathedral with Kevin and Alida Ashley. In addition to Carole, several other founding members of the AI and Law community were at the conference, including Layman Allen, Kevin Ashley, Thorne McCarty, Edwina Rissland, and Charles Saxon, among others. I accompanied my mentor and director, Herbert Fiedler, who presented his vision of legal expert systems as a kind of interactive argumentation assistant, which set me on my path. (Sadly, he also passed away recently, just six months after Carole.)

During one of the breaks at the conference, Carole took me aside and asked me to join her for a walk around the fountain in front of the Palazzo dei Congressi, where the conference was being held. Although it was only the second Logic, Informatics and Law conference, Carole was concerned about the precedent being set, since both conferences were organized by the same people and held in the same location. She suggested organizing a more decentralized and international AI and Law conference, following the lead of other international AI conferences, such as IJCAI. Carole was lobbying people at the conference, surely the other (US) Americans, to support this new conference. Somehow it felt a bit like a conspiracy. Exciting. Carole was persistent and persuasive. Her efforts gave birth to the first ICAIL conference, held two years later in 1987 at Northeastern University in Boston.

Unfortunately, Carole and I never collaborated on AI and Law research. Her focus was case-based reasoning; mine was defeasible rule-based reasoning. It wasn't until much later that these strands began to come together as different forms of legal argumentation. But we did collaborate on organizational and educational matters in the AI and Law community. The International Association for Artificial Intelligence (IAAIL) was founded in 1992 to manage the ICAIL conference. Here too Carole played a leading role. She was the first Secretary-Treasurer of IAAIL and continued to serve in this capacity for about 15 years, to 2007, with a short break. I had the opportunity to work with her as a member of the Executive Committee of IAAIL for 12 of these years. We presented together the introductory tutorial on AI and Law at one of the ICAILs, I do not remember which one, but it must of have been after Donald Berman's death in 1997, since Carole and Don had

⁷ A list of presentations at this conference can be found at <http://www.ittig.cnr.it/Ricerca/Testi/atticonvegno1985.htm>.

given the tutorial together several times previously. And she was still Secretary-Treasurer during my two-year period as President of IAAIL from 2006 to 2007. The most intense period of collaboration was during the preparation of the 1999 ICAIL conference, which took place in Oslo, Norway. I had the honor of serving as Program Chair of the 1999 conference and cannot imagine having been able to perform this job without Carole's experience, support and guidance.

Carole and I both lived for many years in Michigan. I was born there, in infamous Flint, a distinction shared with Vern Walker, and lived there, mostly, until my sophomore year of high school in 1970. Back when one could drink the water⁸. I am not sure where Carole was born, but she went to college at the University of Michigan in Ann Arbor, where she completed her BA in Economics, as well as her Masters and PhD in Computer Science. Her first job after her PhD was a research position with General Motors. My grandfather was a tool and die man for Buick, a division of General Motors, in Flint.

Finally, while doing some research for this, I discovered that Carole wrote the MTS Lisp compiler while at the University of Michigan. As a long-time fan and user of Lisp dialects, especially Scheme, I am thoroughly impressed. Just another, if small, reason for my finding her thoroughly sympathetic. I miss you greatly, Carole. Rest in peace.

4.5 Ron Loui

Ron Loui was General Chair of the Eighth International Conference on AI and Law.

A University of Michigan PhD in 1972, CS professor for 30 years at Northeastern (recently emerita) specializing in natural language processing, ontologies, case based reasoning, legal reasoning, and legal informatics, Carole Hafner was for many years the operational backbone of AI and Law in the US. When we were planning the ICAIL conference in St. Louis in 2001, she said "Our budget is tight. Offer Obama⁹ \$300. He's a state senator. We'll pay him \$500 and he'll be happily surprised." I was so embarrassed. But apparently after keeping the budget in line, IAAIL went from financial fragility to strength. Sometimes you need someone who can say no to an expenditure. She did seem to like the dinner fare at Bar Italia Ristorante in the Central West End (of St Louis), which got us the green light to host ICAIL. She talked about our work in her AI and Law tutorial in St. Louis – the first and only time I saw that happen. We may have even talked about writing a grant proposal together at one point. Those who know Carole for her contributions to this community may not know the complete extent of her impact on ontologies, through her doctoral student, Natasha Noy. Even with Dr. Noy's scholar.google citation total passing 22,000, her 9th best cited paper is still her doctoral work with Carole Hafner. I remember how well Carole took care of things when her co-author, Don Berman, passed away. It seems clear that the editor of

⁸ This came as a shock to me, since previously I relied on the Clarence Williams and Jelly Roll Morton song which assured me that *Michigan water tastes like sherry wine*. - editor.

⁹ Ron was at school with Barack Obama from the 5th to 12th grades, and also spoke to him when he (Obama) was editing the *Harvard Law Review*. Ron hoped to get him as a speaker for the St Louis ICAIL, either at the conference dinner or the special event in the Old Courthouse. In the end Obama could not make it. His loss I'm afraid: it was the best of the fifteen conference dinners I have attended. - editor.

this special issue has taken up the same torch in this publication. AI and Law is as much a community as an academic area, and Carole was as much responsible for building this community as anyone.

4.6 Adam Wyner

Adam Wyner is a researcher active in AI and Law, originally from the North East of the USA. He has attended several ICAILs and is a member of the Editorial Board of this journal. He is currently a lecturer at the University of Aberdeen.

Meeting Carole in Boston Carole and I first met in August of 2003 in Boston, USA. I was a postgraduate student at King's College on a project to automate legal contracts, and I was looking for contacts in the AI and Law community who lived in the United States. Searching on the internet, I could see that at the time, there were few, and they were far between. Luckily for me, Carole lived in the Boston, Massachusetts area, where she was a professor at Northeastern University; this is "home territory" for me as I went to high school in the Boston area and knew Northeastern. As a senior researcher in AI and Law, she was someone to look up. So we met to chat about the field, what she was working on, and what I was working on. She filled me in on what AI and Law "action" was happening in the US. Her ongoing work was about legal drafting technology. I told her I was working on deontic concepts and language. As we had overlapping research interests, it was an easy and informative discussion. However, while it was sort of a friendly chat, she was very pessimistic about the feasibility of anything I proposed to work on - deontic logic had a long history of technical work but little consensus on solutions to long-standing problems; natural language processing of legal texts was well-nigh infeasible. She knew, since these are topic areas she had worked on or watched develop since the start of the AI and Law community. And she knew how to talk about the topics, practically squinting when hearing something she didn't believe and then waving off the point with a "Yeah, well..." It was a bracing conversation. She wished me well.

Ongoing discussions Subsequent to meeting that first time, I tried to arrange to meet with Carole every time I was in Boston, which was just about every year. Sometimes it worked out; sometimes it didn't. Or, we might meet at ICAIL or JURIX. It was always helpful to have someone at "home base", if only to gauge AI and Law research in the US. Over several years, we returned to talking about deontic concepts and language analysis; I grew to appreciate her skepticism better, yet sought ways to address specific problems. While she may never have "bought" my progress reports, I felt it was nonetheless helpful to have that internal voice to address, making counter-arguments and providing evidence. She was, however, supportive of auxiliary activities such as when I co-organised JURIX and ICAIL workshops in 2008/2009 on "Natural Language Engineering of Legal Argumentation: Language, Logic, and Computation". While the numbers of papers and participants were small, the interest in the topics was high.

We also had a bit of correspondence about developing legal corpora. I naively wrote a long complaint to her about problems finding legal corpora; I had spent 10 months researching corpora and finding it rather frustrating both to myself and

to the research field, which needed data to work with. On February 24th, 2009, she replied as follows:

Adam,

I am well aware of the political issues surrounding legal text in the U.S.

I have already been through this battle during the Clinton administration, including attending a meeting at the Department of Justice, another meeting at the Library of Congress, and writing a letter to the attorney general on behalf of the legal informatics research community. Other academics and also those trying to start new businesses to provide computer-based legal research tools or services were involved. Our pleas fell on deaf ears – at the time, the lobbyists from West included a former Minnesota congressman or senator (can't remember which), and a good friend of Al Gore.

There was a public interest group (one of Nader's) involved – their name for the effort was the “Crown Jewels project”, and the focus was trying to pressure the federal government to release federal legal documents for free access, either on line or on electronic media (such as CD's put out by the GPO).

Maybe it is worth another try under Obama, but I suspect his people have a lot of other things to worry about. I wish you good luck with this.

Regards,

Carole

I wrote a long reply, first apologising for not knowing this previous work and then acknowledging that it was a sensitive topic. Yet, I suggested that how perhaps the context, contacts, and approach were different. She didn't reply to that.

Shortly afterwards (March 2009), I published a short article “Text-Mining Case Law” in Legaltech news (now only available from Lexis-Nexis by subscription), which referred to Carole's problematic experience in accessing corpora. I was learning, but did have to lean on the experience of others. Moreover (and obviously), stating the problems did not solve them. While there were breaks in the access to corpora, e.g. the Legal Information Institutes and Law.gov, it was still very slow going. As well, it seemed to be to be difficult to “break out” from the academic arena - the organisations I approached were reluctant to take up text mining tools to extract detailed semantic information from legal texts. There was some academic research, yet it was relatively sparse.

I also spoke to Carole about the impact of those particular people with whom one works in a particularly productive way. For her, Donald Berman, a lawyer professor at Northeastern, was key; his untimely death forestalled her further progress on case-based reasoning.

Last contact In January of 2012, Carole declined an invitation to join (yet another) of my workshop program committees as she was on the cusp of retiring, had a heavy teaching load, and was taking care of an ageing parent. Our next contact was in November of 2012, when I was in Boston for a visit. She emailed that she had had surgery, but was mentally up for some discussion about a paper and legal drafting. We were to talk on the phone to make a plan for me to go to visit her at home. I called and emailed, but never heard back then or subsequently. I took it that she was ailing and unavailable.

Over the period of time from 2008-2013, I became more familiar with Carole's publications in large part from working with Trevor Bench-Capon at the University of Liverpool on case-based reasoning. In particular, Trevor and I discussed the ideas in Berman and Hafner: *Representing Teleological Structure in Case-based Legal Reasoning: The Missing Link* [15], which took a while for me to absorb. In 2012, I reviewed Hafner: *Conceptual Organization of Case Law Knowledge Bases* for the AI and Law journal issue giving "A History of AI and Law in 50 papers" [8]. I told Carole about the article as I was writing it and sent her a copy; she was glad to see it, but not particularly engaged.

A change in circumstance and drawing some morals I learned of Carole's death from Jack Conrad's email on the IAAIL list. I was very sad, but not surprised given that over time she had become quieter. I miss having her at "home base" as it was always a pleasure to have lunch in Boston with her and chat about work. Though she is gone, there are some professional "meta lessons" I learned from Carole.

First, pioneers in a field may not see the "promised land". Carole worked hard on case-based reasoning and information extraction from legal cases, making key, early contributions. Yet, her work was deeply frustrated by the lack of corpora to work with and the resistance to change amongst legal professionals and the legal informatics industry. How times have changed. Within the last two years, very, very large corpora of legal texts of all sorts have become freely available. Global-scale corporate enterprises, e.g. IBM, Thomson-Reuters, Lexis-Nexis, have taken a vigorous interest in mining the data. The US has start-ups in legal informatics and has active research centres. Industrial strength text analytic tools such as Watson are applied to legal texts at scale. The community of legal professionals have piled into the newly changed environment of legal practice. Workshops on e-discovery attract dozens of lawyers along with AI and Law academics who - surprise surprise - actually sit in the same lecture theatre. There is a formal role at law firms for 'legal knowledge engineers'. What would Carole have thought about and done with all these changes and new opportunities?

Second, considering Carole's working partnership with Donald, the people one closely collaborates with are essential. For researchers with a social disposition, creative work emerges from the collaboration and interaction, rather than isolation. The particular collaborator can fundamentally change one's research direction; by the same token, the loss of that person deprives the other partner of fruitful ongoing direction. So, if you work really well with someone, work more. The trick is to find such research partners.

Finally, working for collegial acknowledgment rather than private satisfaction can be disappointing. When Carole and I talked about work, inevitably it would turn to who was not citing whom, but should, and how the line of precedent research ought to be cited. Perhaps this is a generational issue since the "first" generation of AI and Law researchers had a relatively small pool of work to cite, mostly from ICAIL, JURIX, and one journal. Now, keeping up with the flow of papers, never mind antecedents, is difficult. There seems so little time to read and think things through. Whatever the underlying cause, it is of little use to harp on about what one's colleagues are not attending to; rather one's energies are best put to use creating new and more relevant work, not resting on one's laurels.

Carole made substantive, early research contributions; she was a leading mover and shaker in the organisational underpinnings of AI and Law. At the moment,

recalling her, I appreciate her “meta lessons” to a junior colleague. I think she would think these were all incidental. If I told her, she likely would squint and wave me off, though I'd hope she would be slightly pleased.

4.7 John Zeleznikow

*John Zeliznikow was General Chair of both the sixth and the ninth International Conferences on AI and Law.*¹⁰

In the middle of 1990, I had my first encounter with the area of Artificial Intelligence and Law, when I read the seminal paper of Berman and Hafner 1989 [13]. At that time, I was a new computer science professor, having moved from the domain of algebra. To see this article in the *Communications of the Association for Computing Machinery*, which showed how computer science could help the legal community, greatly impressed me. With some trepidation, I contacted the author. Within a few months I was living in Don Berman's house and meeting with him and Carole. To me, Don and Carole were a team! He was the loquacious law professor, Carole was the more reserved, more demanding, computer scientist. From that time, no matter where I was traveling in the world, from Melbourne to New York, Tel Aviv or Paris, Boston Massachusetts was on the route. United Airlines and my Mileage Plus account (I am now a United Million Miler - with all its privileges) did very well in my determination to visit Carole and Don at every opportunity. Don taught me everything I know about Law. Because I had a fairly similar background to Carole (PhD in the same year and my mathematics qualification compared to her computer science PhD), my meetings with Carole were initially far less numerous than those with Don. This changed in 1995, when Dan Hunter and I offered to host the 1997 ICAIL Conference at the University of Melbourne. I had missed the 1995 ICAIL Conference in College Park Maryland, due to the birth of my fourth child Eva. So in December 1996, on my way to JURIX in Tilburg (in those days all JURIX conferences were held in the Netherlands), I spent an intense two days in Boston with Carole (by that time Don was very ill and unable to see me). As secretary of IAAIL (International Association for Artificial Intelligence and Law), Carole saw ICAIL as her child. The conference had to run well and not make a deficit. Whilst most attendees enjoyed ICAIL97 in Melbourne, it ran at a financial loss - fewer people came to the conference because travel to Melbourne was long and expensive (I am however used to travelling to Europe or USA approximately twice a year), it was in the middle of winter and we ran an excellent but expensive social program. Carole was furious. So when I offered to run ICAIL 2003 at the University of Edinburgh Law School, she placed major restrictions on me. The conference was very well attended and made a great profit - greatly boosting IAAIL finances. From then Carole was more relaxed! Over the quarter of a century since 1990, I have visited Boston annually: I had taught at Mt. Holyoke 1983-5 and was on sabbatical in Amherst in 2007. Plus I looked forward to my annual visits to Don (until 1995) and Carole. It was always great to visit Carole at Northeastern to discuss AI and Law, research, journal editorship and universities. On a few occasions my family and I would have a delightful brunch with Frank and Carole in Wayland. My last visit with Carole was in June 2013.

¹⁰ Also on half of Andrew Stranieri and George Vossos

We had dinner and visited the site of the Boston marathon bombing (I had run the marathon in both 2009 and 2010 and visited Carole beforehand). We were due to meet a few days later at ICAIL in Rome. But I had to rush back to Melbourne because my 89 year old father was on life support. By my 2014 Boston visit (to see my daughter receive her LLM at Harvard) I could no longer see Carole. By my visit in July 2015, she had sadly passed away. The doyens of AI and Law have departed this world and such research is no longer conducted at Northeastern University! How sad!

5 Concluding Remarks

In this introduction I have attempted to cover the three aspects of Carole’s legacy to AI and Law. First there is her body of work, and its continuing relevance. All four of the papers in this special issue describe contemporary research for which Carole’s work has lessons. Second there is Carole’s contribution to the AI and Law community: she founded and nourished all three of our major institutions, the conference, the association and this journal. Finally there is her contribution to individual members of the community, as described in the memories recounted in section 4. Her impact on our field, intellectual, organisational and personal, will long continue.

References

1. Latifa Al-Abdulkarim, Katie Atkinson, and Trevor Bench-Capon. A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. *Artificial Intelligence and Law*, 24(1):1–49, 2016.
2. V. Aleven. *Teaching case-based argumentation through a model and examples*. PhD thesis, University of Pittsburgh, 1997.
3. Kevin D Ashley. *Modeling legal arguments: Reasoning with cases and hypotheticals*. MIT press, Cambridge, Mass., 1991.
4. Katie Atkinson. Introduction to special issue on modelling popov v. hayashi. *Artificial Intelligence and Law*, 20(1):1–14, 2012.
5. William Baude and Stephen E. Sachs. The law of interpretation. *Artificial Intelligence and Law*, 130, Forthcoming, 2017.
6. T Bench-Capon and S Modgil. Rules are made to be broken. In *Proceedings of the Workshop on Artificial Intelligence for Justice (AI4J 2016) at ECAI 2016*, pages 18–21, 2016.
7. Trevor Bench-Capon. Value-based reasoning and norms. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1664–5. IOS Press, 2016.
8. Trevor Bench-Capon, Michał Araszkiwicz, Kevin Ashley, Katie Atkinson, Floris Bex, Filipe Borges, Daniele Bourcier, Paul Bourguine, Jack G Conrad, Enrico Francesconi, et al. 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.
9. Trevor Bench-Capon and Giovanni Sartor. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence*, 150(1):97–143, 2003.
10. Trevor JM Bench-Capon. *Knowledge-based systems and legal applications*. Academic Press, London, 2015.
11. Trevor JM Bench-Capon, Gwen O Robinson, Tom W Routen, and Marek 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*, pages 190–198. ACM, 1987.
12. D. Berman and C. Hafner. Obstacles to the development of logic-based models of legal reasoning. In *Walter, C. and Allen, L. E., editors, Computing Power and Legal Reasoning*, pages 45–64. West Pub Co. St Paul, 1985.

13. Donald H Berman and Carole D Hafner. The potential of artificial intelligence to help solve the crisis in our legal system. *Communications of the ACM*, 32(8):928–938, 1989.
14. Donald H Berman and Carole D Hafner. Incorporating procedural context into a model of case-based legal reasoning. In *Proceedings of the 3rd international conference on Artificial intelligence and law*, pages 12–20. ACM, 1991.
15. Donald H. Berman and Carole D. Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the Fourth International Conference on Artificial intelligence and Law*, pages 50–59, 1993.
16. Donald H Berman and Carole D Hafner. Understanding precedents in a temporal context of evolving legal doctrine. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 42–51. ACM, 1995.
17. Jon Bing. *Handbook of legal information retrieval*. Elsevier Science Inc., Amsterdam, 1984.
18. Gerhard Brewka, Stefan Ellmauthaler, Hannes Strass, Johannes Peter Wallner, and Stefan Woltran. Abstract dialectical frameworks revisited. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 803–809. AAAI Press, 2013.
19. Stefanie Brüninghaus and Kevin D Ashley. Generating legal arguments and predictions from case texts. In *Proceedings of the 10th international conference on Artificial intelligence and law*, pages 65–74. ACM, 2005.
20. Sandra Cook, Carole D Hafner, L Thorne McCarty, Jeffrey A Meldman, Mark Peterson, James A Sprowl, NS Sridharan, and Donald A Waterman. The applications of artificial intelligence to law: A survey of six current projects. In *Proceedings of the May 4-7, 1981, national computer conference*, pages 689–696. ACM, 1981.
21. Graham Greenleaf. Jon Bing and the history of computerised legal research—some missing links. *Et tilbakeblikk på fremtiden (Looking back at the future)*, pages 61–75, 2004.
22. Thomas R. Gruber. The role of common ontology in achieving sharable, reusable knowledge bases. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). Cambridge, MA, USA, April 22-25, 1991.*, pages 601–602, 1991.
23. Nicola Guarino. Formal ontology and information systems. In *Proceedings of Formal Ontology in Information Systems*, volume 98, pages 81–97, 1998.
24. Carole D Hafner. An information retrieval system based on a computer model of legal knowledge, University of Michigan. *PhD Thesis, Ann Arbor, MI*, 1978.
25. Carole D Hafner. Representation of knowledge in a legal information retrieval system. In *Proceedings of the 3rd annual ACM conference on Research and development in information retrieval*, pages 139–153. Butterworth & Co., 1980.
26. Carole D Hafner. Conceptual organization of case law knowledge bases. In *Proceedings of the 1st international conference on Artificial intelligence and law*, pages 35–42. ACM, 1987.
27. Carole D Hafner and Donald 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.
28. Carole D Hafner and John D Joyce. Using natural language descriptions to improve the usability of databases. In *Proceedings of the first conference on Applied natural language processing*, pages 132–135. Association for Computational Linguistics, 1983.
29. Carole D Hafner and Marc Lauritsen. Extending the power of automated legal drafting technology. In *JURIX*, volume 165, pages 59–68, 2007.
30. Carole D. Hafner and Natalya Fridman Noy. Ontological foundations for biology knowledge models. In *Proceedings of the Fourth International Conference on Intelligent Systems for Molecular Biology, St. Louis, MO, USA, June 12-15 1996*, pages 78–87, 1996.
31. Jaap Hage. Monological reason-based logic: a low level integration of rule-based reasoning and case-based reasoning. In *Proceedings of the 4th international conference on Artificial intelligence and law*, pages 30–39. ACM, 1993.
32. John Henderson and Trevor Bench-Capon. Dynamic arguments in a case law domain. In *Proceedings of the 8th international conference on Artificial intelligence and law*, pages 60–69. ACM, 2001.
33. John F Horty and Trevor J M Bench-Capon. A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20(2):181–214, 2012.
34. Natalya Fridman Noy and Carole D Hafner. The state of the art in ontology design: A survey and comparative review. *AI magazine*, 18(3):53, 1997.

35. Natalya Fridman Noy and Carole D. Hafner. Representing scientific experiments: Implications for ontology design and knowledge sharing. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence and Tenth Innovative Applications of Artificial Intelligence Conference, AAAI 98, IAAI 98, July 26-30, 1998, Madison, Wisconsin, USA.*, pages 615–622, 1998.
36. Natalya Fridman Noy and Carole D Hafner. Ontological foundations for experimental science knowledge bases. *Applied Artificial Intelligence*, 14(6):565–618, 2000.
37. Henry Prakken and Giovanni Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artif. Intell. Law*, 6(2-4):231–287, 1998.
38. Edwina L Rissland and M Timur Friedman. Detecting change in legal concepts. In *Proceedings of the 5th international conference on Artificial intelligence and law*, pages 127–136. ACM, 1995.
39. Marek J. Sergot, Fariba Sadri, Robert A. Kowalski, Frank Kriwaczek, Peter Hammond, and H Terese Cory. The British Nationality Act as a logic program. *Communications of the ACM*, 29(5):370–386, 1986.
40. David B Skalak and Edwina L Rissland. Arguments and cases: An inevitable intertwining. *Artificial intelligence and Law*, 1(1):3–44, 1992.
41. André Valente. *Legal Knowledge Engineering: A Modelling Approach*. PhD thesis, University of Amsterdam, 1995.
42. B. Verheij. Jumping to conclusions. a logico-probabilistic foundation for defeasible rule-based arguments. In *Proceedings of the 13th European Conference on Logics in Artificial Intelligence, JELIA*. Springer, 2012.
43. B. Verheij. Arguments about values. In K Atkinson, H Prakken, and A Wyner, editors, *Argumentation in AI, Law and Policy Making. A Festschrift in Honour of Trevor Bench-Capon on the Occasion of his 60th Birthday*, pages 243–57. College Publications, London, 2013.
44. B. Verheij. To catch a thief with and without numbers: Arguments, scenarios and probabilities in evidential reasoning. *Law, Probability and Risk*, 13:307–325, 2014.
45. B. Verheij. Formalizing correct evidential reasoning with arguments, scenarios and probabilities. In *Proceedings of the Workshop on Artificial Intelligence for Justice (AI4J 2016) at ECAI 2016*, pages 87–95, 2016.
46. Pepijn RS Visser and Trevor JM Bench-Capon. A comparison of four ontologies for the design of legal knowledge systems. *Artificial Intelligence and Law*, 6(1):27–57, 1998.
47. Pepijn RS Visser, Robert W Van Kralingen, and Trevor JM Bench-Capon. A method for the development of legal knowledge systems. In *Proceedings of the 6th international conference on Artificial intelligence and law*, pages 151–160. ACM, 1997.