Artificial Intelligence and Law **9:** 59–71, 2001.

Book Review

George C. Christie, *The Notion of an Ideal Audience in Legal Argument*, Dordrecht: Kluwer Academic Publishers, 2000. 223 pages. ISBN 0-7923-6283-7.

1. Introduction

This book contains no mention of computers, let alone Artificial Intelligence. None the less it contains much of relevance to current concerns of Artificial Intelligence and Law, because it deals with the nature of legal argument. The realisation that an understanding of legal argument is a central challenge of AI and Law came very early; Thorne McCarty's pioneering work on TAXMAN aimed to reconstruct the arguments of the majority and dissenting judgements in a leading tax case. The work of Edwina Rissland, Kevin Ashley, Vincent Aleven and others has explored the ways in which arguments are constructed on the basis of case precedents. Jaap Hage, Henry Prakken, Giovanni Sartor, Bart Verheij, Ron Loui and others have used logical tools and the concept of argument to explain how rational decisions can be made when rules conflict. Tom Gordon, Arno Lodder, myself and others, have modelled the consideration of legal issues as a formal dialogue game in which disputing parties argue their case. These, and other uses of argument, such as for explanation of the reasoning of legal knowledge based systems, have confirmed this central role of argument in AI and Law. For an overview of argument in AI and Law up to the mid-1990s, see Bench-Capon (1997). Since then interest in argument has even increased.

All this excellent work has done much to suggest ways in which arguments can be generated, analysed and used in legal reasoning. What it has not addressed with much success is the question of what, given two conflicting arguments, makes one argument more convincing than another. In a 1985 paper, written with Marek Sergot, and published as Bench-Capon and Sergot (1988), the concluding sentence was:

> In the longer term, we hope to pursue what we have identified as a critical requirement: a representation in computer-intelligible terms of what it is that makes a legal argument persuasive.

In the decade and a half that has followed, I have been many times diverted from this pursuit. In recent years, however, I have returned to the chase, (e.g.,

Bench-Capon and Sartor 2000; Bench-Capon 2001), inspired by the work of Chaim Perelman (Perelman and Olbrechts-Tyteca 1969; Perelman 1980). I thus encountered this book at a most opportune time. Christie's book is dedicated to the memory of Chaim Perelman, and employs one of Perelman's key notions – *the universal audience* – to provide an exploration of what does, and what should, make a legal argument convincing, written from the perspective of a legal philosopher; in particular, a legal philosopher well versed in the law of several legal cultures.

Perhaps the central idea in Perelman's work is that an argument, indeed all speech, is directed towards an audience. But the audience towards which argument is directed is essentially a *construct* of the speaker:

> The essential consideration for the speaker who has set himself the task of persuading concrete individuals is that his construction of the audience should be adequate to the occasion. (Perelman and Olbrechts-Tyteca 1969, p. 19)

If we take this seriously we can see that whether an argument is accepted is not something that can be determined by considering the intrinsic qualities of the argument: rather it is a function both of the argument, and the audience to which it is addressed. But this starts to look dangerously *ad hominem*: while I may convince a particular group of people by correctly understanding and exploiting their prejudices and misconceptions, such a victory does not guarantee that my claim has any objective right to acceptance. I can convince some students that predicate calculus is worth learning by telling them that I intend to set an examination question on that subject; a very different argument is required to convince my colleagues that the topic *should* be in the syllabus. Perelman, of course, was well aware of this problem:

> Argumentation aimed exclusively at a particular audience has the drawback that the speaker, by the very fact of adapting to the views of his listeners, might rely on arguments that are foreign or even directly opposed to what is acceptable by persons other than those he is presently addressing. (Perelman and Olbrechts-Tyteca 1969, p. 31)

For this reason opinions that enjoy unanimous approval are considered stronger than those accepted by particular audiences only. But this casts no light on whether an argument *should* be accepted, only whether it is, as a matter of fact widely accepted. Many arguments aspire to more than this: to reflect this Perelman introduces the notion of the *universal* audience. This universal audience is not simply an aggregation of all particular audiences: acceptance by the universal audience refers

> not to an experimentally proven fact, but to a universality *imagined* by the speaker, to the agreement of an audience which *should* be universal, since for *legitimate* reasons, we need not take into consideration those which are not part of it. (Perelman and Olbrechts-Tyteca 1969, p. 31, *italics mine*)

Those who address the universal audience

> think that all who understand their reasons will have to accept their conclusions. *The agreement of a universal audience is thus a matter, not of fact, but of right.* (Perelman and Olbrechts-Tyteca 1969, p. 31, italics theirs)

At this point we might be tempted to believe that the universal audience will accept only claims that can be logically proven. If we are doing mathematics this may even be correct. But this is not Perelman's intention: that would limit the realm of what can be argued before the universal audience too drastically.

> Everyone constitutes the universal audience from what he knows of his fellow men, in such a way as to transcend the few oppositions he is aware of. Each individual, each culture, has thus its own conception of the universal audience. **The study of these variations would be very instructive**, as we would learn from it what men, at different times in history, have regarded as *real*, *true*, and *objectively valid.* (Perelman and Olbrechts-Tyteca 1969, p. 31, *italics theirs*, **bold font mine**)

Despite the sentence in bold in the above quotation, Perelman does not perform such a study. Christie, however, in the book under review, does just this, for arguments presented in legal contexts. It does indeed prove "very instructive".

At some point Christie shifts from Perelman's notion of a *universal* audience to one of an *ideal* audience. He intends this to be no more than terminology – he says that he uses the terms interchangeably (p. 41). Christie is never entirely comfortable with this shift – the phrase "ideal or universal audience" occurs frequently, but I think it is a good move: it emphasises that the ideal audience is not "everyone, as a matter of fact", but that it is a construct capturing "an audience which *should* be universal, since for *legitimate* reasons, we need not take into consideration those which are not part of it. Christie writes:

> Some aspects of the ideal audience that manifest themselves in legal argumentation seem to have truly universal acceptance; others, although thought of as universal, seem more closely tied to particular cultures. (p. 41)

To my mind, this notion of an audience *seen as universal from within a culture* is better expressed by "ideal" than by "universal". The book is essentially an attempt to identify the various aspects of an ideal audience for legal argumentation, both those which are universally acceptable and those which are culturally dependant. The key point to remember is that we cannot determine the ideal or universal audience by counting heads. The majority may reject an argument, but for the wrong reasons: such an argument would remain acceptable to the universal audience. Many social reforms, such as the abolition of the death penalty in the UK, have been carried through against the wishes the majority, but validated by the ideal audience which embodies the aspirations of society, rather than what is currently the case.

2. Contents of the Book

Following an introduction in chapter 1 and some motivating analysis using the notion of a universal/ideal audience in chapter 2, chapter 3 attempts to identify constraints on an ideal audience, without limiting the context to legal argument. These turn out to be very few: the universal audience expects that the speaker should be truthful (although they can use humour and sarcasm), and expects the speaker to be open to opposing arguments. Additionally all participants must share the rejection of physical violence as a way of winning the argument. There is little to dispute here: these things remind one of Alexy's rules of legal discourse (Alexy 1989), but do not go even as far. And Alexy's rules give little guide to what to do in practice: Gordon's criticisms that they are insufficiently precise, too general, and fail to restrict judicial discretion in a clear way (Gordon 1993), apply equally to the characterisation of the universal audience of chapter 3. Christie is well aware of this: his aim is not to stop here, but to focus first on what universal features can be added if we focus on a legal context, and then to start to explore cultural variations. The bulk of the book is in fact concerned with delineating the sources of differences in conceptions of the ideal legal audience.

The universal features in the legal context seem to focus mainly on an idealised view of the legal process. Judges should be impartial, be patient enough to hear everybody, base decisions on the evidence presented, and be independent of executive and legislature. Counsel should respect client confidentiality and avoid conflicts of interest. The defendant should not be forced to testify against himself, and should not be subject to torture. The legislature should strive for generality, promulgate their laws, and not pass retroactive legislation. Laws should not change too frequently, and there should be a "congruence between official action and the law" (quotation from Lon Fuller). This forms the discussion in chapter 4, and this still seems too general to be readily applicable in practice. The chief interest of the book lies in the following chapters which discuss different conceptions of the ideal audience.

Christie begins chapter 5 by emphasising that there are differences in how different cultures envision the ideal form of argument.

> That different cultures, even within the Western legal tradition, should have different conceptions of the ideal or universal audiences to which legal arguments are addressed is of course not surprising. It merely reflects that our notion of an ideal or universal audience is basically a construct, even if it is a construct that we constantly presuppose and could not do without. (p. 49)

His first examples concern the interpretation of statutes and treaties. He contrasts the narrow, cramped and particularistic view of interpretation of some courts (US in the examples), with the more expansive interpretation of others (Israeli and German). He concludes that

> The United States Supreme Court is obviously torn between two competing visions of what the ideal or universal audience expects of it. One vision, the

> traditional and still predominant one, stresses fidelity to the traditional role of common-law judges, and focuses on discovering the precisely intended meaning of statutes, constitutions and other legal instruments. This vision eschews completely the spirit behind Article 1 of the Swiss Civil Code that when neither statutory provision nor customary law cover the case, the judge should decide 'according to the rules he would lay down if he himself had to act as legislator'. (pp. 58–9)

Chapter 6 attempts to account for these differences in terms of competing visions of the way the state should be organised, and competing visions of the purpose of the state. This contrasts states which exercise authority through a hierarchically organised set of officials, with those which exercise authority through a large number of co-ordinate officials. Also there is the contrast between the "reactive" state which exists to resolve conflicts, and the "activist" state, which has its own notion of the good which it attempts to promote. Moreover, even accepting that the law exists to promote *the common good*, should we see the common good as composed of individual goods, or as an independent entity? This last is a particularly interesting question, since the tension becomes evident in particular cases, and Christie cites some telling examples, in which the choice needs to be made. These are discussed at some length in chapter 7, with reference to "private necessity" whereby the property or even well-being of an innocent person is sacrificed to preserve the property and/or well being of one or more other persons. This chapter I found a fascinating discussion of some hard choices, and how different intuitions can be explained by appeals to the ideal audience presupposed by those faced with them.

Chapters 8 and 9 focus on another tension:

> that between the intellectual urge to greater generality, and the cautionary counsel of experience to focus as much as possible on the particular. The ideal or universal audiences to which legal argument is directed are constantly torn by these conflicting impulses. (p. 108)

Chapter 8 offers some legal background, drawn mainly from the law of tort, while chapter 9 offers some theoretical perspectives, basing the case for general principles mainly on the work of Ronald Dworkin, before deploying some counter arguments. His own sympathies, as he admits, tend towards the particular, but his purpose is to

> set forth the historical and theoretical reasons why the ideal audience which we construct as the auditor of legal and even moral argumentation might vacillate between the two views and why, at any one time, it might choose one rather than the other. (p. 151)

Chapter 10 looks at another source of difference between universal audiences: the amount of freedom that they conceive a decision maker such as a judge should have. Here again we have a tension: there are good reasons both to seek to broaden discretion, and to narrow it. In this chapter Christie uses examples from jury selec-

tion and tort law to show how there is movement to and fro between the extreme positions. He concludes:

> While no legal system exclusively favours one style of legal reasoning over another at any given period of its historical development, for reasons I have indicated, a legal system will display a systematic preference for one style of decision making over another. In the situations [tort] we have been discussing there has been a preference for a more open-ended, a more discretionary, a more legislative style, if you will, of legal decision making. (p. 178)

In chapter 11 Christie turns to consider consistency - the degree to which diverse outcomes are tolerated. In this chapter the examples are mainly drawn from UK cases subsequently successfully appealed to the European Court of Human Rights. Such cases are of particular interest because they involve cases originally decided in a common-law culture, appealed to a court grounded in a civil-law tradition. His argument is that a civil-law tradition is less likely to accept that there are many reasonable answers to a problem, and to be less tolerant of inconsistency.

> It seems, for example, to be a characteristic of civil-law jurisdictions that, as cases wind their way up the hierarchy of appellate courts, they come to lose their character as a discreet dispute between the parties and become the vehicle for sweeping and ahistorical assertions about the "law" to a greater extent than is characteristic of cases winding their way up the appellate ladder in a common-law jurisdiction. (p. 187)

And this is important:

> the accepted style of legal reasoning has a profound effect on the substance of the decisions the courts will make. Style is not a matter of mere form. It sometimes has a decisive influence on the decision of cases. That is one more reason the way the members of a particular legal community view the form, structure and methods of the ideal legal argument – the argument that would be approved by the ideal or universal audiences presupposed by those legal communities – is an important subject for careful study. (p. 192)

The final chapter is a conclusion which contains some reflections on globalisation and the increasing trend for international conventions which cross legal cultures.

The summary of the book I have given shows how Christie looks at a number of specifically legal factors which characterise the ideal audience for legal argument. These are a series of tensions: between narrow and broad construction; between a reactive and an activist state; between the common good as an aggregation of individual goods and as an independent notion; between the particular and the general; between wide discretion and limited discretion; and between consistency and diversity. The position between these extremes ascribed to the ideal audience will vary across cultures and within cultures over time. The lawyer constructing an argument gauges the position of the ideal audience to be addressed by being part of the relevant culture. This summary follows the organisation of the book. There is,

however, another aspect to the ideal audience which emerges from the examples – that the ideal audience embodies a set of values, and of preferences between those values when they conflict.

In *Gregg v. Georgia* (428 US 153 1976), in which the Supreme Court upheld a sentence of death, there were two dissenting arguments, from Justices Brennan and Marshall. Marshall's dissent is quoted:

> In *Furman* [an earlier case] I observed that the American people are largely unaware of the information critical to a judgement on the morality of the death penalty, and concluded that if they were better informed they would consider it shocking, unjust and unacceptable ... the death penalty is unconstitutional because it is excessive. An excessive penalty is invalid under the Cruel and Unusual Punishments Clause even though popular sentiment may favour it. (quoted on pp. 22–3)

This dissent is clearly addressed to an ideal audience, one which is "better informed" than the audience that will receive it. It is addressed in a sense to a future audience, which will hold more enlightened values. It accepts that it will fail to convince its immediate auditors, but seeks to promote a change in their values, so that it will be accepted at a later date.

Again, discussing the *Restatement of Torts* and the *Restatement (Second) of Torts*, in the context of whether property can be destroyed to save human life, he says

> For those who accept the Restatements' positions, compensation must be paid whenever property is destroyed to save human life. Such an assertion presupposes that the universal audience addressed in the course of legal argument values life more than property, but not so much as to absolve innocent persons from the obligation to pay compensation for property unavoidably destroyed in the process of saving lives. (p. 84)

This is a very clear statement that the ideal audience subscribes to values, and is capable of comparing and ranking them.

A third example concerns pre-Civil war cases in the US in which judges were called upon to enforce fugitive slave laws. Some refused to enforce them, offering legally suspect reasons, but appealing to the rejection of slavery they ascribed to the ideal audience. Others did enforce them, not because they condoned slavery, but because they "felt that preserving the integrity of the federal union, which was threatened with secession by the slave states, *was the greater good*". (p. 45, *italics mine*)

Thus in addition to embodying the various positions that can be taken as to the organisation and purpose of the state, and various preferences for styles of legal reasoning, the ideal audience must be seen also as embodying a set of values and a system of preferences amongst those values.

3. Implications for AI and Law

What are the implications of a study of the ideal audience for AI and Law? How is the ideal audience represented in a computer system? There are several perspectives we can take, but first we should notice that cognisance of the ideal audience is not usually explicit in AI and Law systems. Rather the system builder's construct of the ideal audience tends to be built into the system, incorporated in the very principles according to which the system is constructed. The system does not construct the ideal audience, rather the system is constructed according to the system builder's presupposition of an ideal audience. Nevertheless, it is instructive to use the notion of the ideal audience to analyse some issues in AI and Law.

First we can use the notion of ideal audience to account for differences between systems and approaches in AI and Law. Let us remember that the nature of the ideal audience is *presupposed* within particular legal cultures. Thus people building AI and Law systems will have a tendency to picture the ideal audience in terms of the culture of which they are a part.

There used to be a lively debate in AI and Law between proponents of a case based approach and proponents of a rule based approach. Panels addressed this topic at both the First and Third International Conferences on AI and Law, held in 1987 and 1991. Without too much of a caricature, we can relate these approaches to some of the positions identified as characteristics of the ideal audience above. Preferring the particular over the general gives a tendency towards a case based approach. Case based approaches can be seen as reactive and rule based approaches as activist. Rule based systems tend to promote consistency, and minimise discretion. In sum, case base approaches might be expected to emerge from a common law culture, and rule based systems from a civil-law culture. This is borne out by the fact that case based approaches were more popular in the US, and rule based approaches were more popular in Europe. To meet objections that rule based approaches were also popular in the UK, one may point to the fact that much of the UK work concerned administrative welfare law, where the activist state, lay adjudication, and the need for consistency, push it towards the civil-law position. In the late 1980s and early 1990s, when AI and Law was only beginning to develop an international community, people would be prepared to argue that one position or the other was *right.* Such argument was based on the feeling that the ideal audience was the one which adopted the position of the ideal audience from the culture from which its proponent came. The debate has now died down, and there is far more of a tendency to see the approaches as complementary, or suitable for use in different circumstances, than as either right or wrong absolutely. In part this is a matter of recognising that there is no single, universal ideal audience, and in part because the AI and Law community has developed its own construct of the ideal audience, which is capable of recognising merit in both approaches. There have been gains and losses in this development. Gains stem from the ability to learn from other approaches (Prakken and Sartor (1998), in which they model a

classic case based system, HYPO, in their preferred rule based terms, is an excellent example), and a more receptive climate for discussion within the AI and Law community. Losses have come when we attempt to return results from AI and Law to the legal community. Because the systems are no longer so tied to the ideal audiences presupposed by particular communities, they speak less clearly to them, and assumptions embodied in the system which are not those of the legal community mean that they can be seen as actually *wrong* by those communities. Thus the relaxation of the links between AI systems and particular constructs of the ideal audience have facilitated progress in AI and Law, but militated against acceptance of them by legal communities.

From this we should expect the systems which gain most acceptance outside of AI and Law to be those which both have a clear conception of the ideal audience of the legal community at which they are targeted, and which most closely reflect this audience.

We can illustrate this by reference to several systems. First consider the fielded systems in the Netherlands, which tend to be rather straightforward rule based expert systems dealing with administrative law, especially welfare benefits, and based on shells such as TESSEC (de Bakker and Wassink 1991) and JURICAS (van Noortwijk and Stubbe 1986). These were deployed in a country which has a civil-law tradition, and in an application domain where emphasis is on consistency and accuracy. Moreover the applications were developed in close consultation with the users. Thus both the predisposition of the techniques used in systems, and the implemented details were in harmony with their customer's vision of the ideal audience.

A very different example is supplied by CATO (Aleven 1997). Here the idea was to produce a system to support the teaching of case based argument. The system was evaluated according to improvements in student performance, and its own answers to coursework questions. The ideal audience here is very specific: the person grading the papers. Most academics have a clear idea of the structure they expect, the cases they want cited, and the points they want made about them. Thus building a system directed towards this audience is well placed to produce what is wanted, as CATO did. This is in no way intended to be a criticism of CATO, which is an excellent system which fulfils its aims. On the other hand it is possible that the reason why this system has not had wider acceptance among the US legal community may be that its target ideal audience was too specific. There may well be a difference between the argument of a bright student which will attract a high grade, and the argument that an experienced attorney will produce.

A third example is given by the Split-Up system (Stranieri and Zeleznikow 1999). This system relies heavily on the use of an artificial neural net, which one might doubt would gain ready acceptance in the legal community. But the application is a rather special one: apportioning resources following a family breakdown. In this domain where the ideal audience seems to require *factor balancing* as the

central reasoning technique, ANNs seem a natural technique to use. Again it was developed in very close co-operation with its potential users.

These three examples use very different techniques, but in each case the technique seems to have considerable affinity to the ideal audience that will receive its output. Moreover, and perhaps more importantly, the systems were developed from within the legal culture they were targeted at: there was therefore no mis-match between the construct of the ideal audience of the system builder and the user. I think there are some clear lessons here for those who wish to develop systems for practical use.

So far I have used the notion of the ideal audience to explain differences in approach amongst workers in AI and Law, and why some systems seem to have proved more acceptable to the target users than others. I now want to consider how the notion of an ideal audience can throw light on some problems that arise for researchers in an inter-disciplinary field like AI and Law.

The major output of academic research is research publications, papers in journals and conferences and books. These are written for an audience comprising largely academic research peers. Such an audience is very different from that of the target users of systems. In an inter-disciplinary field such as AI and Law, moreover, we must be aware that that the ideal audience for a computer science paper is very different from the ideal audience for a law paper, (different footnoting conventions are a superficial but real sign of this), and that the interdisciplinary audience will contain features of both. The danger here is that work may be directed too much at one sector of the audience. For example, papers intended for an AI and Law audience may address computer science concerns, with insufficient acknowledgement of the legal audience. The famous British Nationality Act paper (Sergot et al. 1996) was addressed first and foremost at a logic programming audience. The hostile reception of some lawyers, who pointed to the simplicity of its "black letter" legal reasoning and found it incapable of addressing the subtleties of legal reasoning should have been expected, and perhaps addressed at the outset. The problem is that attention can be diverted from the promise of the technique because it is presented in a way which lays it open to attacks based on a misconstruction of its arguments. Influential as this paper has been, an impression that computer scientists working in this field try to impose their techniques without sufficient regard to legal concerns remains in some quarters (e.g., Moles 1992). Similarly papers on legal theory are often ignored by computer scientists when such papers make no effort to include them in the audience addressed: for example, any such proposals need to be aware of issues such as computational feasibility. Such cases illustrate failure to take into account the range of the potential audience. Other examples can be seen where techniques interesting to computer science, such as objects, neural networks, agents, and the like, have been applied to AI and Law in ways interesting for the sake of the technique rather than for any application leverage. These techniques have, perhaps in consequence, attracted less attention than might have been expected in AI and Law.

There are, however, times when computer science concerns rightly take centre stage. If one intends to build a system to model legal reasoning, one must be aware not only of the ideal legal audience at which the output will be targeted, but also that computer science concerns are very relevant because one is *building a computer system.* Thus the kind of software engineering concerns that underlie much AI and Law work on the representation of source documents (e.g., Johnson and Mead 1991; Bench-Capon and Coenen 1992) are legitimate for software developers, even though they may not seem interesting to lawyers. Fidelity to legislative text is primarily to enable better systems to be built, not to allow better reasoning to be performed. Both writers and readers of work in AI and Law need to recognise this spread of concerns and the different constructs of the ideal audience that may be present in an inter-disciplinary field, and to which papers may sensibly be directed. Writers need to be aware of the different perspectives from which their work will be evaluated, and readers must be prepared to allow for the fact that work may address concerns that they do not share. Only in this way can the interchange of ideas be maximised.

Finally I want to make some brief remarks about how we might make the notion of an (ideal) audience explicit in AI and Law systems. Most of the dichotomies identified by Christie have been and will continue to be built into systems implicitly. There have, however, been some efforts, by myself and others, to make explicit the construction of the values and system of preferences embodied by the ideal audience as a way of deciding between competing arguments. This work has been explicitly motivated by the views of Perelman, and has drawn inspiration from a seminal paper of Berman and Hafner (1993), which drew attention to the need to look beyond the facts of precedent cases to the value preferences they revealed in order to resolve conflicting arguments. One such attempt is Bench-Capon and Sartor (2000). The idea here is that competing arguments are preferred according to the values of the audience to which they are addressed. These values, and their ordering, are revealed in previous decisions accepted by that audience. The idea is that reasoning about a case involves first constructing, from past decisions, an ordering on values, and a set of rules related to those values, which together form a theory, attributed to the ideal audience, which explains past cases and which should be applied in the current case. In the current context we can see this approach as explicitly recognising that the ideal audience is a construct which is built from the experience of the legal culture through the past decisions made within that culture. I believe this to be a very stimulating idea which opens new vistas for AI and Law.

4. Summary

I found this book instructive and valuable. I was already persuaded of the importance of Perelman's insights, but Christie takes them further with a thorough and well informed analysis of the variations that can be found in the ideal audiences for legal argument that can be found in different legal cultures. The discussion is

always interesting, and well illustrated with apt examples drawn from a variety of legal cultures. The value for AI and Law perhaps lies less in specific, readily applicable techniques, than in the way it frees the mind by breaking down the notion that legal culture and legal reasoning is a single entity, rather than a construct determined by particularities of tradition, time and place. With this in mind we can make useful distinctions between disagreements within a culture and disagreements across cultures.

The concept of an ideal audience is also, as I hope I illustrated in the last section, a useful analytical tool for thinking about the practice of AI and Law. I finished the book even more convinced of the need to recognise that arguments are directed towards an audience, and of the need to characterise an ideal audience if we are to escape the traps of regarding arguments as subjective, or deriving their authority only from the fact of widespread acceptance. Anyone who is concerned with legal argument, either as a practitioner of law, or because they wish to model it with computers, should find this a fascinating and thought provoking book.

References

Aleven, V. (1997). Teaching Case-Based Argumentation Through a Model and Examples. Ph.D. Thesis, The University of Pittsburgh.

Alexy, R. (1989). A Theory of Legal Argumentation. Clarendon Press: Oxford.

Bench-Capon, T. J. M. and Sergot, M. J. (1988). Towards a Rule Based Representation of Open Texture in Law. In Walter, C. (ed.), Computer Power and Legal Language, Quorum Books: Westport, Connecticut, pp. 39–60.

Bench-Capon, T. J. M. (1997). Argument in artificial intelligence and law, Artificial Intelligence and Law 5(4), 249–261.

Bench-Capon, T. J. M. and Coenen, F. P. (1992). Isomorphism and legal knowledge based systems, Artificial Intelligence and Law 1(1), 65–86.

Bench-Capon, T. J. M. and Sartor, G. (2000). Using values and theories to resolve disagreement in law. In Breuker, J., Leenes, R. and Winkels, R. (eds), Proceedings of JURIX 2000, IOS Press: Amsterdam, pp. 73–84.

Bench-Capon, T. J. M. (2001). Truth and consequence: complementing logic with values in legal reasoning. Information and Communications Technology Law 10(1), 2001.

de Bakker, K. F. C. and Wassink, J. G. J. (1991). Development, implementation and impact of the TESSEC expert system. In Proceedings of the European Group of Public Administration 1991.

Berman, D. H. and Hafner, C. D. (1993). Representing teleological structure in case based reasoning. In Proceedings of the Fourth International Conference on AI and Law, ACM Press: New York, pp. 50–59.

Gordon, T. F. (1993). The Pleadings Game. Ph.D. Thesis, Technische Hochschule, Darmstadt.

Johnson, P. and Mead, D. (1991). Legislative systems for public administration – some practical issues. In Proceedings of the Third International Conference on AI and Law. ACM Press: New York, pp. 108–117.

Moles, R. N. (1992). Expert systems - the need for theory. In Proceedings of JURIX 1992. Koninklijke Vermade: Lelystad, pp. 113–122.

Perelman, Ch. and Olbrechts-Tyteca, L. (1969). The New Rhetoric. University of Notre Dame Press: Notre Dame.

Perelman, Ch. (1980). Justice, Law, and Argument. Reidel: Dordrecht.

Prakken, H. and Sartor, G. (1998). Modelling reasoning with precedents in a formal dialogue game. Artificial Intelligence and Law 6(2–4), 231–287.

Sergot, M. J., Sadri, F., Kowalski, R. A., Kriwaczek, F., Hammond, P, and Cory, T. (1986). The British Nationality Act as a Logic Program. Communications of the ACM 29, 5, pp. 370–386.

Stranieri, A. and Zeleznikow, J. (1999). The evaluation of legal knowledge based systems. In Proceedings of the Seventh International Conference on AI and Law. ACM Press: New York, pp. 18–24.

van Noortwijk, C. and Stubbe, A. L. W. M. (1986). JURICAS: Erfrecht, een juridisch computeradviessystem. Lelystad: Vermade.

Department of Computer Science
University of Liverpool
United Kingdom

Trevor J.M. Bench-Capon