Artificial Intelligence and Law **5**: 249–261, 1997.

Argument in Artificial Intelligence and Law

TREVOR BENCH-CAPON
Department of Computer Science, University of Liverpool, P.O. Box 147, Liverpool L69 3BX, U.K.
E-mail: tbc@compsci.liverpool.ac.uk

Abstract. In this paper I shall discuss the notion of argument, and the importance of argument in AI and Law. I shall distinguish four areas where argument has been applied: in modelling legal reasoning based on cases; in the presentation and explanation of results from a rule based legal information system; in the resolution of normative conflict and problems of non-monotonicity; and as a basis for dialogue games to support the modelling of the process of argument. The study of argument is held to offer prospects of real progress in the field of AI and law, and the purpose of this paper is to provide an overview of work, and the connection between the various strands.

Key words: argument, AI and law, explanation, non-monotonic reasoning, dialogue games

1. Introduction

In this paper I will review the topic of argument, as addressed in Artificial Intelligence and Law. I believe argument to be a central topic in Artificial Intelligence in general: AI must concern itself with rationality and rational behaviour and belief, and an understanding of argument is central to an understanding of rationality. Argument can be studied in many places: we can turn to formal systems, such as mathematics or symbolic logic; or we might choose the unsystematic, undisciplined arguments that we find in everyday discourse. Law offers a middle way: it is not a formal system, but yet it does impose rules on the anarchy that is everyday argument. This enables questions as to how an argument should be conducted, how disputes can be resolved, what can be brought in aid of an argument and the like to have answers whereas no definitive response is possible in general argument. So Law is an excellent place in which to study argument.

Formal, symbolic, logic is an abstraction, designed to focus solely of the relationships between the truth values of propositions. Although "logic" can have a wider application, perhaps as "the science and art of reasoning correctly", in this paper I shall use the term with the narrower meaning of formal logic in mind. The propositions of logic are the bearers of truth values and nothing more. Frege argued that the denotation of a proposition was a truth value: this looks counter intuitive; both in that we might expect the denotation of "Trevor works in Liverpool" to be different from "Edwina works in Amherst", and in that we might think that the proposition is saying something about the world, that it denotes a state of affairs, in the manner of the picture theory of meaning of Wittgenstein's *Tractatus*. We

should, however, resist these intuitions: for the abstracted propositions of formal logic Frege is right – they are bearers of truth values and nothing more, whereas for the multifarious propositions of everyday use, general theories are inadequate, and cases must be looked at on their individual merits, as was realised by the Wittgenstein of the *Investigations*.

A consequence of this abstraction is that the propositions of logic lack many of the features that we might wish to ascribe to the propositions of informal argument and jurisprudence. The propositions of logic cannot be elegantly or inelegantly expressed; they cannot be well known or obscure; they cannot be interesting or useful. And just as the propositions are abstracted from these attributes, which may none the less be of great importance in a given context, so too are arguments abstracted into what are called "proofs". In a proof we move from axioms and premises given as true to a desired conclusion by a series of steps guaranteed to be truth preserving. The reader of a proof is convinced because it can be seen that the rules have been properly applied. But unlike an argument, everything must be explicit – there can be no appeal to a background of common, agreed, knowledge – and the organisation is in accordance with the rules of the calculus rather than in order to maximise the persuasive power. If we wish to study these extra-logical elements we move to the realm of rhetoric – which I shall use broadly to encompass all aspects of the organisation and presentation of an argument – rather than the realm of logic. Then it is possible to see that an argument may be convincing, although unsound in terms of logic, and unpersuasive, although sound. Logic may tell us how to reason *correctly*: rhetoric is required to tell us how to reason *well*.

As well as rhetorical features, the proofs of logic lack a notion of procedure. The axioms and premises are given at the outset, whereas in real arguments the initial premises must be established and new facts and rules can be discovered and introduced. Both how the initial premises are agreed and how new information is introduced into an argument is very important, and we lose the possibility of insight into the nature of dialectical reasoning if we stick with the abstraction of a proof. It is the procedural aspects the tell us how to reason *properly*. Also proofs lack a context. When I tell my students that the project counts for a quarter of the marks for the year, it is clear that I am talking about the University of Liverpool, and the Computer Science Department there. That what I have said is not in general true, does not matter: I need not qualify my statement as I would have to if I were talking to students from a range of Departments or a range of Universities. But for a proof to succeed these qualifications would be needed. So the fact that arguments are situated in a context, which determines what qualifications are necessary, and what can be left unsaid, is another important element that needs to recovered from the abstraction. This aspect tells us how to reason *appropriately*.

For AI, which attempts to model general reasoning, not just formal reasoning, and for AI and Law which specifically attempts to capture the reasoning of lawyers, it is necessary to go back from proofs to argument and to reintroduce elements of rhetoric, procedure and context. To cite but two examples: the standard

explanations of expert systems such as MYCIN and its descendants are often held to be unsatisfactory. Since no more than a trace of the proof is offered, this should not be surprising. While logic can provide the raw materials for the explanation, we need all three of rhetoric, procedure and context as well if we are to organise the explanation in a convincing manner, and avoid cluttering the explanation with minor details so that the key elements are obscured. The second example is non-monotonic reasoning, which has come to dominate much of AI and knowledge representation: this depends on using assertions that are not really true ("All birds can fly" is the somewhat tired but classic example), and makes sense only if we allow for procedure and context. So important are these elements, that it appears to me that progress will be possible only if we turn away from the abstraction offered by logic which designedly excludes these features, and start to look seriously at argument. And I see no better context in which to look at argument than the field of AI and Law, in which we attempt to automate the disciplined arguments that are found in the legal context.

The notion of argument has, of course, received a great deal of attention in AI and Law. Here I want to distinguish two strands. First the older tradition, predominately American, which stems from a desire to model legal reasoning, and thus sees systems which construct, or support the construction of, arguments as being the natural style of system to build. This work, in which cases are fundamental, will be discussed in Section 2. The second strand is the more recent development of interest in arguments within the rule based tradition, predominant in Europe, in which argument is seen less of an end in itself and more of a means of analysing problematic notions such as normative conflict, non-monotonicity and the process of argument. Different facets of argument relating to rule based systems will be discussed in Section 3.

2. Modelling Legal Argument Based on Cases

Argument has long been at the centre of concern of leading AI and Law researchers in the U.S.A. Perhaps the best examples of this strand are provided by Thorne McCarty and Edwina Rissland. McCarty has focused his work on the TAXMAN project (see McCarty, 1995, for the most recent report on this work), which is a sustained attempt to produce a system which can reconstruct the arguments used in a celebrated Supreme Court decision. Rissland has also made argument her focus, in work with Kevin Ashley on the HYPO system (most extensively reported in Ashley, 1990), with David Skalak on CABARET (for which Skalak and Rissland, 1992, is the best reference), and with Skalak and Timur Friedman on the BankXX system (Rissland et al., 1996).

McCarty's approach is based on the notion of a prototype and deformations. The idea is that what is being argued is that a particular case is properly classified in one way rather than another. So the starting point is a set of facts describing a case and a set of prototypes describing paradigm instances of the various classifications.

An argument is produced by making a number of transformations of the prototype so that it fits the facts of the case. Thus in the case of *Eisner and Macomber*, which concerned whether a dividend distributed in the form of stock was taxable, he begins with a standard pattern of taxable receipts and a standard pattern of non-taxable receipts, and shows how deforming these patterns into the facts of *Eisner and Macomber* relate to the arguments actually advanced by the Supreme Court judges, Pitney and Brandeis. This work is of enormous interest, but since it has not been widely imitated I shall discuss it no further here, but refer the interested reader to (McCarty, 1995).

Rissland and her colleagues are primarily concerned with a notion of argument which requires the finding of a case which was decided in the past, establishing that the case under consideration is sufficiently like that previous case that the previous case provides a precedent which should be followed, and showing that there are no contrary cases which provide a better precedent.

In HYPO the form of argument is what Rissland and Ashley term a "three-ply argument", since it involves three turns, or plies, two by the proponent and one by the opponent:

- PROPONENT: Cite a case which matches the facts of the case in some respects, and which was decided in the way desired;
- OPPONENT: Distinguish this case, and cite other cases decided in the contrary way;
- PROPONENT: Distinguish these contrary cases.

While this is a central style of argument in law, it is by no means the only one. Later work by Rissland and her colleagues on the CABARET project distinguished a number of other argument forms, and attempted to describe them in terms of a more restricted range of argument strategies, and the moves which can be used to effect them. Both of these approaches essentially generate arguments in a top down fashion. In BankXX a bottom up approach is used, where arguments are described in terms of a dozen building blocks or "argument pieces", and a strategy of heuristic search is used to retrieve such pieces from a representation of knowledge pertaining to the domain. All three of these systems are worthy of considerable study, and together they represent the richest computational characterisation of legal argument so far produced. Moreover, the systems are implemented and can be shown to be capable of generating legally plausible argument.

This work is clearly concerned with what I have termed rhetorical aspects – the way arguments are structured; with procedure, since it discusses argument strategies and tactics; and with context since many of the strategies are in fact concerned with extending or narrowing the context. The starting point is, however, the surface behaviour of how lawyers actually argue, and in consequence some of the fundamental connection with the underlying logic is lost: there is no real notion of the *soundness* of an argument, or of how the underlying logic provides the basis for the argument. This is, in contrast, at the heart of the approaches to argument

that come from the rule based tradition, and it is on these that I will concentrate in the remainder of this paper.

3. Approaches to Argument in Rule Based Systems

If we turn to rule based approaches we can distinguish three important strands. First we can see approaches which focus on argument as a means of presenting results; this strand is mainly concerned with the rhetorical aspects of argument. A second strand is concerned with problems that arise in rule based systems when there is an apparent conflict of rules, so that both a conclusion and its negation can be derived from the knowledge base. Linked with this strand is the desire to accommodate the apparent non monotonicity of legal reasoning, whereby a currently preferred argument can be rebutted by introducing new information, citing a different norm, or giving different weights to existing information. This strand is predominately concerned with context, and with the procedures which determine how the context can be modified. The third strand focuses on the interaction, claiming that a dialectical interaction is a useful way of using legal information systems. For this strand procedure is most important, although context is a concern, and rhetorical aspects may also be involved. I shall discuss each of these three strands in the following sections.

3.1. Presentation and Explanation

The use of argument for the presentation of reasoning and the explanation of the results of executing rule based systems started to receive widespread attention in the late 1980s when – independently – three different proposals were made to use the argument schema of Toulmin (1958) for this purpose. These proposals were made by Marshall (1989), Lutomski (1989) and Storrs (1991). Since then this very popular means of representing the structure of arguments has been used by, amongst others, Bench-Capon et al. (1991), Dick (1992) and Zeleznikow and Stranieri (1995). It is perhaps surprising that this diverse work has all found Toulmin's schema suitable for its purposes. It is perhaps, therefore, worth briefly stating what this schema is, and what makes it so attractive.

Toulmin's starting point was that arguments are not best illuminated by being cast into the mould of premises and conclusions. Instead he suggested a division into a more varied set of components that would highlight the role of the various kinds of assertions made in the course of an argument. The structure suggested by Toulmin may be represented diagrammatically, as in Figure 1.

The arguments are decomposed into

- claims whose truth we seek to establish by the argument,
- data that we appeal to as the grounds for the claim, and

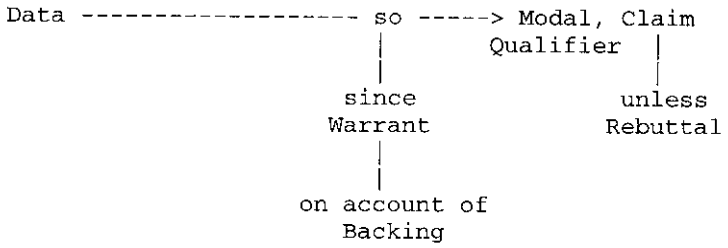

Figure 1. Toulmin's argument schema.

- warrants which provide the rules of inference that connect the data and the claim. Warrants can, in Toulmin's scheme, bestow varying degrees of support for the claim, and these degrees of support are indicated by the use of a
- modal qualifier, such as "necessarily" or "possibly". There may also be exceptional conditions which prevent the claim from being established: these are indicated by the
- rebuttal which contains circumstances acknowledged as requiring the authority of the warrant to be put aside. Warrants also require some justification: this is the role of the
- backing.

Its attractions, particularly for those working in the rule based tradition, can be summarised as follows:

- It is relatively simple, and yet able to accommodate a reasonable range of arguments. What it lacks in richness as compared with, say, Rissland's characterisation of argument, it gains in uniformity;
- It is essentially rule based: the warrant is just the sort of rule you find in a logic program or a production system, and thus it can be naturally adapted to the output of such systems;
- It provides a modal qualifier to support the use of heuristics and defeasible rules;
- It provides the rebuttal as a mechanism to deal with non-monotonic aspects;
- It nicely organises the conditions in the rule according to a view of their importance and their role in the argument.

Of course, those who have used the schema have found it necessary to make some adaptations. For example, in (Bench-Capon et al., 1991) it is extended to allow arguments to be chained together, so as to reflect the standard operation of a rule based system, and to make explicit preconditions which are often implicit in Toulmin's examples through the use of an appropriate sortal term in the warrant. Still the schema has proved extremely flexible, and given considerable leverage to those who wish to improve presentation by moving from the structure of a proof to the structure of an argument.

The essential goal of all the systems in this strand is to *organise* output so as to maximise understanding and to minimise obfuscating detail, and to demonstrate clearly the different roles of elements of the output. Toulmin's schema has been found an excellent and flexible starting point to progress towards these ends.

3.2. Normative Conflict and Non-Monotonicity

One persistent problem for rule based approaches to AI and Law is the existence of conflicting norms. That is we find that we need to represent two rules, each of equal *prima facie* validity, which lead to a contradiction if both are applied. It is hard to get away from these problems, especially in view of the fact that so much of law is structured as general rules and exceptions, and while exceptions are often explicitly signalled by some phrase such as "subject to" qualifying the general rule this is not always the case. A second kind of conflict arises when the conclusion of one rule negates a condition used by another rule, rendering it inapplicable, or where one rule has as its conclusion that another rule is inapplicable. This means that we must be able to argue about which rules are preferred as well as which conclusions are preferred.

The problem is by no means peculiar to AI and law. Since their inception production rule expert systems have felt the need to include general heuristics (such as "all birds can fly"), and particular exceptions ("penguins cannot fly"). The simplest way of doing this is to choose some appropriate conflict resolution principle, such as specificity, by which the more specific matching rule is fired in preference to the less specific rule. Specificity may be determined by purely syntactic means such as the number of conditions in the antecedent of the rule, or by the number of uninstantiated variables, or by appeal to some kind of type hierarchy which makes penguin more specific than bird. This is, however, a pragmatic and *ad hoc* solution, and those who want to really understand what is going on have been at some pains to formalise the notion: hence the vogue for non-monotonic logics. Two representative approaches are:

- Construe rules such as "All birds can fly" as *default* rules: that is the conclusion can only be drawn if it is consistent to do so. Since "penguins cannot fly" is not a default rule, its conclusion will be drawn in all cases, and so if we are considering a penguin, the default rule will not be applied, since to do so would result in an inconsistency. This kind of approach can be found in Reiter (Reiter, 1980).
- Read a rule such as "all birds can fly" as "if X is a bird and *X is not abnormal* then X can fly". Now we apply these rules so as to minimise the extension of the *abnormal* predicate: thus we assume it to be false unless we are forced to take it as true to avoid a contradiction. This style of approach, known as "circumscription" was introduced by McCarthy (1980).

In AI and Law we may take the work of Prakken and Sartor, both individually (Prakken, 1993; Sartor, 1992) and jointly (Prakken and Sartor, 1995), and the work

of Hage and his colleagues on reason based logic (Hage et al., 1993; Hage, 1995) as representative. Handling non-monotonicity is also a main feature of the work of Gordon (1993), although I want to discuss his work in the next subsection. Essentially this work deals with conflicts in the following way (note that Hage uses the term *reason* rather than the term *argument*; I shall use argument to cover both, and hope that I am not doing Hage an injustice thereby):

1. If two norms conflict they can produce an argument for both a conclusion and its negation.
2. We must then examine the arguments and see which we prefer.
3. The preferred argument is said to defeat the other argument, and its conclusion is drawn.

In deciding which argument to prefer Sartor and Prakken make considerable use of three legal principles: the specific law derogates the general law; the higher law derogates the lower law; and the recent law derogates the older law. These principles may be applied individually or in combination, and orderings are suggested, e.g. in (Prakken, 1993). Of these principles the first has a strong resemblance to the conflict resolution principle of specificity which is the traditional conflict resolution principle of production systems: one might perhaps include also analogues of the other two principles in production systems with good effect. Hage makes less use of these principles, but still requires that a relation *outweighs* exists between reasons, allowing us to prefer one to the other: there is no reason not to, although it not compulsory to, use the general principles to infer this relation.

Two things are particularly interesting. First that there is recognition of the need for a plurality of, possibly competing, principles to resolve conflicts. This is also recognised in the wider AI non-monotonic logic community, which spends considerable time on considering examples which tend to show one better than another, and in efforts to combine them into a single universal default theory. But as Doyle and Wellman persuasively argue (Doyle and Wellman, 1991)

> is there a universal theory that automatically combines correctness criteria in a rational manner. The answer is no: *any universal theory of default inference based on combining correctness criteria must sometimes produce irrational conclusions* (italics theirs)

Their conclusion from this is that we

> must continue to investigate special theories of reasoning and the conditions under which each of these is to be preferred or avoided.

Perhaps legal reasoning can profitably be seen as one such special theory of reasoning. The second point – particularly emphasised by Hage and Gordon – is that the relative priorities may themselves be the subject of argument, and it is possible to reason explicitly about the relative merits of conflict resolution principles in particular cases. This seems to me essential, and something which should be taken on board by the non-monotonic reasoning community. Another issue, explored mainly by Hage, is whether several weak arguments should be allowed to succeed

against an argument which is stronger than any of them individually. Again this may well be something to which no general solution exists.

For me the problem with non-monotonic logic is that it attempts to use a narrowly construed, formal, logic, to deal with a problem that requires "logic" to be construed broadly, as meaning the science of rational argument. Essentially a formalism is being asked to do work which requires more than a formalism. A formal logic is supposed to abstract from content – to make an argument sound *in virtue of its form alone*. Yet it is precisely content that leads us to prefer one argument to another in many cases. The form of "all birds fly" is the same as that of "all mammals are viviparous". Yet an argument based on the former is certainly weaker than one based on the latter. Whilst only a couple of obscure mammals found only in Australia lay eggs, many well-known birds are innately flightless, and personal accidents may render others flightless also. Only by allowing us to take specific content into account – as Hage's mechanism does – can we come to rational conclusions: but this takes us beyond a logic, and back into the world of reasoning at a less abstracted level. An alternative approach for indicating content attaches numeric weights or certainty factors to default rules. This I believe creates many more problems than it solves: I have argued against it and in favour of evaluating arguments elsewhere (Bench-Capon and Sergot, 1989), and I shall not repeat that discussion here.

In his most recent work (Prakken, 1995), Prakken attempts to provide a framework in which formal logic can be allocated its proper place. In this important paper he proposes a three layer model in which an *argument framework* is built around an underlying logic – which may be either classical or non-monotonic – and the argument framework is used in a dialectical protocol. This usefully separates the logic from the context in which deductions are made and conclusions are used, and from the considerations of procedure for establishing and changing the context. It is thus possible to obtain non-monotonic behaviour from the system as a whole even when the underlying logic is classical. This can be compared with the role of deductive systems that I argued for in (Bench-Capon, 1994), where the role of the system is only to make deductions and it is the user's understanding of the system, and is the premises that have been supplied by the user and the use the user makes of these deductions which gives the point to these deductions. Here the user is supplying the argument framework and the protocols; the system as a whole has Prakken's three layers, but only the lowest of them is automated. Of course, automating the argument framework and making it part of the system will lead to a richer system, but one which will go beyond the simple "first generation" systems described in that paper, and lead on to systems which put computational dialectic in the centre of the stage. My belief is that it is by separating the logic from the context and construction of context that we will make progress here: otherwise we will turn in vain to non-monotonic logics to do for us what no logic can be expected to do.

The three layer model is not universally accepted. As Prakken notes (Brewka and Gordon, 1994) proposes a different three layer model with the logic and argument framework combined into a single logic layer, and the higher levels being a speech act level, determining which speech acts are possible and a protocol level, which determines how they should be used. I prefer Prakken's model because it makes what I consider to be the right distinction in the logic level, allowing "logic" both its narrow and its broad use: If we want the speech act protocol distinction as well – and for some purposes it may be important – we could go to a four layer model.

Before leaving this subsection I just wish to draw attention to the work of Farley and Freeman (1995) which highlights a rather different aspect of the argument context that needs to be used in evaluating arguments – the standard of proof required. Thus what may be acceptable "on the balance of probabilities" may not be acceptable "beyond all reasonable doubt", and the winning argument may even differ in the two cases. Once again we see that priorities between rules and arguments are not absolute, but must be evaluated in the light of the context in which they are presented.

3.3. Argument as a Process

The previous section discussed the importance of context in assessing arguments, and hence what we should believe when there are conflicting arguments. But of equal importance is how that context becomes established and modified, and in this section I shall mention some of the work which has explored this issue, focusing on the procedural aspects of argument.

This work has typically been couched in the form of the implementation of a *dialogue game*, emphasising both the adversarial, dialectic, nature of the process, and the rule governed nature of the activity. As three examples we can cite Bench-Capon et al. (1992), Gordon's *Pleadings Game* (1993) and the work of Arno Lodder and his colleagues (Hage et al., 1993; Lodder and Herczog, 1995). These examples provide a range of systems: (Bench-Capon et al., 1992) is general, not targeted specifically at law, and employs a standard logic programming framework and Toulmin's argument structure; (Gordon, 1993) is targeted at a specific legal practice, and is based on a well-known non-monotonic logic; and (Lodder and Herczog, 1995) is aimed at law in general and based on reason based logic. These differences are not very important for our present purposes, but rather show that a wide range of design choices can be accommodated in this framework. The commonalities are much more important: the recognition of the importance of how an argument is conducted, and the choice to represent this in a formal way by codifying it as a set of rules that makes all the participants aware of their rights and obligations, and enables them to interpret the moves made by the other participants. This work is directed towards the third layer identified by Prakken (1995), and I for one think it to be perhaps the most critical layer of all.

This work is still in its infancy – only Gordon as yet offers a fully fledged implementation – and much needs to be investigated. In particular, in addition to matters of detail as to the rules of the various games, the following questions need to be answered:

- How natural are the resulting dialogues? Do they have empirical validity in that they produce arguments that resemble those produced in practice?
- Do they foster fair and efficient decisions?
- Can the games accommodate and replicate the kinds of argument move identified by Rissland and her colleagues?
- How does this work relate to the choice of underlying logics, and the notion of an argument framework?

Much remains to be done, but it is much that is very well worth doing.

4. Conclusion

In this paper I have attempted to give an overview of the different ways in which argument is currently being used in AI and Law. That it is currently a "hot topic" is evinced by the recent nature of so many of the references, and the number of papers that addressed the topic in one way or another in the latest International Conference on AI and Law (1995). But it is not just fashionable: it is genuinely important, and offers, I believe, a prospect of making some real progress on some of the thorniest problems of AI and Law, and indeed automated reasoning in general. In particular I have tried to bring out the connections between the various strands of work. It is the dialectical process that I see as central, although some have come at it through a desire to improve presentation, and others through worries about normative conflicts and non-monotonicity. It is important not to forget the work on modelling legal reasoning done by McCarty and Rissland and her colleagues also: ultimately we will have succeeded in capturing the nature of the process if we can account for the phenomena they describe.

There is not as yet a comprehensive computational account of argument. But work towards it is important for AI and Law and, indeed AI in general. The foundations have been laid, and it is time to build upon them.

Acknowledgement

This paper is a revised version of (Bench-Capon, 1995), which was presented at JURIX95. I would like to thank Tom Gordon, Jaap Hage and Pepijn Visser for their comments on that earlier version.

References

Ashley, Kevin D. (1990). *Modeling Legal Argument*. MIT Press: Cambridge, MA.

Bench-Capon, T.J.M. (1994). Legal Theory and Legal KBS: A Computer Scientist's Perspective. In Prakken, H., Muntjewerff, A.J. & Soetman, A. (eds.) *Legal Knowledge Based Systems: The Relation with Legal Theory (JURIX 1994)*, 33–42. Koninklijke Vermade: Lelystad.

Bench-Capon, T.J.M., (1995). Argument in AI and Law. In Hage, J.C. et al. (eds.) *Legal Knowledge Based Systems: Telecommunication and AI and Law*, 5–14. Koninklijke Vermade: Lelystad.

Bench-Capon, T.J.M. & Sergot, M.J. (1989). Towards a Rule Based Representation of Open Texture in Law. In Walter, Charles (ed.) *Computing Power and Legal Reasoning*, 39–60. Greenwood Press: New York.

Bench-Capon, T.J.M., Lowes, D. & McEnery, A.M. (1991). Using Toulmin's Argument Schema to Explain Logic Programs. *Knowledge Based Systems* **4**(3): 177–183.

Bench-Capon, T.J.M., Dunne, P.E. & Leng, P.H. (1992). A Dialogue Game for Dialectical Interaction with Expert Systems. In *12th Annual Conference on Expert Systems and Their Applications*, 105–116. Avignon, EC2: Paris.

Brewka, G. & Gordon T.F. (1994). How to Buy a Porsche: An Approach to Defeasible Decision Making. In Proceedings of the *AAAI-94 Workshop on Computational Dialectics*, 28–38. Seattle.

Dick, Judith P. (1992). *A Conceptual Case Relation Representation of Text for Intelligent Retrieval.* Technical Report CSR1-265, University of Toronto.

Doyle, Jon & Wellman, Michael P. (1991). Impediments to Universal Preference Based Default Theories. *Artificial Intelligence* **49**: 97–128.

Farley, Arthur, M. & Freeman, Kathleen (1995). In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 156–164. ACM Press: New York.

Gordon, Thomas F. (1993), The Pleadings Game – Formalising Procedural Justice. In Proceedings of *The Fourth International Conference on AI and Law*, Amsterdam, 10–19. ACM Press: New York.

Hage. Jaap C. (1995). Teleological Reasoning in Reason Based Logic/ In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 11–20. ACM Press: New York.

Hage, Jaap C., Leenes, Ronald & Lodder, Arno R. (1993). Hard Cases: A Procedural Approach. *Artificial Intelligence and Law* **2**(2): 113–167.

Lodder, Arno R. & Herczog, Aimee (1995). DIALAW: A Dialogical Framework for Modelling Legal Reasoning. In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 11–20. ACM Press: New York.

Lutomski, Leonard S. (1989). The Design of an Attorney's Statistical Consultant. In Proceedings of *The Second International Conference on AI and Law*, Vancouver, 224–233. ACM Press: New York.

McCarthy, J. (1980). Circumscription – A Form of Non-Monotonic Reasoning. *Artificial Intelligence* **13**: 27–39.

McCarty, L. Thorne (1995). An Implementation of Eisner v Macomber. In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 276–286. ACM Press: New York.

Marshall, Catherine C. (1989). Representing the Structure of a Legal Argument. In Proceedings of *The Second International Conference on AI and Law*, Vancouver, 121–127. ACM Press: New York.

Prakken, Henry (1993). *Logical Tools for Modelling Legal Argument.* Ph.D. Thesis, Free University of Amsterdam.

Prakken, Henry (1995). From Logic to Dialectics in Legal Argument. In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 165–174. ACM Press: New York.

Prakken, Henry & Sartor, Giovanni (1995). On the Relation between legal Language and Legal Argument. In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 1–10. ACM Press: New York.

Reiter, R. (1980). A Logic for Default Reasoning. *Artificial Intelligence* **13**: 81–132.

Rissland, Edwina L., Skalak, David B. & Friedman, M. Timur (1996). BankXX: Supporting Legal Argument Through Heuristic Retrieval. *Artificial Intelligence and Law* **4**(1): 1–71.

Sartor, Giovanni (1992). Normative Conflicts in Legal Reasoning. *Artificial Intelligence and Law* **1**(2–3): 209–235.

Skalak, David B. & Rissland, Edwina L. (1992). Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law* **1**(1): 3–42.

Storrs, Graham (1991). The Policy System. In Bench-Capon, T.J.M. (ed.) *Knowledge Based Systems and Legal Applications*, 165–182. Academic Press: London.

Toulmin, S. (1958). *The Uses of Argument*. Cambridge University Press: Cambridge.

Zeleznikow, John & Stranieri, Andrew (1995). The Split-up System: Integrating Neural Nets and Rule Based Reasoning in the Legal Domain. In Proceedings of *The Fifth International Conference on AI and Law*, University of Maryland, 185–194. ACM Press: New York.