

Assessing Arguments in the Abstract: Some Problems with Argumentation Frameworks

James Palmer

**Balliol College, Oxford University
email: james.palmer@balliol.ox.ac.uk**

Abstract

Argumentation frameworks are logic systems that assess collections of conflicting arguments generated from a set of potentially inconsistent factual and rule-based premises. Argumentation frameworks were developed in response to criticisms that traditional formal logic did not do justice to the adversarial nature of legal reasoning, and that it ignored the problems caused by the ambiguity of language, incompleteness in the rule-set, and conflicts in the rule-set. In this paper I assess whether argumentation frameworks successfully overcome the shortcomings of traditional formal logic as a representation mechanism for Artificial Intelligence applications in the legal domain.

Although argumentation frameworks desirably extend traditional formal logic by allowing for inconsistent premises and non-monotonic reasoning the attempt to assess the conflicting arguments that are generated by such a system is unlikely to prove fruitful. There are two reasons for this. First, the formulation of realistic priority relationships seems an almost impossible task. Some of the most important factors in determining whether an argument is a ‘winner’ or a ‘loser’, such as the context of the case and the beliefs of the audience, resist formalisation. Also, the search for meta-priority relationships to resolve conflicts between priority relationships seems to lead to an endless recursion; it is always possible to find one more level of argument. Second, argumentation frameworks do not adequately capture the ‘room to argue’ created by, for example, flexibility in rule formulation, rule application and rule defeasibility.

In short, I will argue that while argumentation frameworks are well suited to modelling a piece of legal reasoning after the conflicting arguments have been presented and assessed by a human reasoner, they are not well suited as the basis for AI applications which attempt to generate and assess such arguments.

1. Introduction

Since the late 1980’s many Artificial Intelligence (‘AI’) and law researchers have been critical of the deficiencies associated with using traditional formal logic as a model of legal reasoning. Criticisms of this kind were first raised by Berman and Hafner (1987) and Rissland (1990) who voiced concerns that the rule-based formal logic approach did not do justice to the adversarial nature of legal reasoning and that it ignored the problems caused by the ambiguity of language, incompleteness in the rule-set, and conflicts in the rule-set. One response to these criticisms of formal logic has been the development of ‘argumentation frameworks’ in which conflicting

arguments are generated and assessed (Sartor 1994, at 200; and Prakken 1995, at 165).

Following Prakken (1995, at 165) an argumentation framework may be defined as a logic system that assesses a collection of conflicting arguments generated from a set of factual and rule-based premises. The current generation of argumentation frameworks is exemplified by two systems recently put forward by Hage (1996) and Prakken and Sartor (1996). These are systems based on formal logic which attempt to deal with problems such as contradictory rules, arguments about the validity and applicability of rules, rules with assumptions, and the need to weigh reasons and prioritise rules.

In Section 2, I describe the argumentation frameworks of Hage, and Prakken and Sartor. In Section 3, I consider the potential for using argumentation frameworks as the basis for building Legal Knowledge-Based Systems ('LKBSs'). My conclusion is that these argumentation frameworks are unlikely to produce LKBSs that are significantly more powerful than their ancestors based on traditional formal logic.

2. The Argumentation Frameworks of Hage, and Prakken and Sartor

The two primary distinguishing features that separate argumentation frameworks from systems based upon traditional formal logic are that (1) they allow for inconsistent premises and non-monotonic reasoning¹ and (2) they attempt to assess conflicting arguments by applying a set of ordering criteria. Two of the most important contemporary argumentation frameworks are those of Hage (1996) and Prakken and Sartor (1996).

2.1. Hage's Reason-Based Logic

Hage calls his argumentation framework 'Reason-Based Logic'. As the name suggests, it is organised around the idea that the essence of legal argument is the provision and assessment of reasons in favour of and against the proposition of interest. A 'reason' may be based on a rule, a principle or a goal. For a reason to come into play, a number of criteria must be satisfied. If the reason is based on a rule or a principle, it must be determined whether it is valid, applicable in terms of time and place, and the absence of excluding reasons, and whether it should be applied in terms of consistency between the effect of applying the reason and its underlying purpose. For a reason based on a goal, it must be determined whether the goal is valid, whether the proposition being argued for would contribute to the goal, and whether or not it is excluded in a particular case. When two or more conflicting reasons exist for a particular proposition, they must be weighed against each other before a conclusion can be reached.

Reasons may interact in a number of interesting ways. One reason may exclude the operation of another reason. For example, the operation of a criminal law rule may be excluded in a particular case because the accused was a foreign ambassador and there is a rule of diplomatic immunity which says foreign ambassadors are not

1) The traditional formal logic requirement of consistency mandates that the premise-set be free from contradictions. The need to preserve consistency constrains a LKBS's ability to faithfully represent many areas of the law. As Sartor (1992) argues, contradictions are not uncommon in sets of legal norms. They may arise for many reasons, including the existence of multiple norm-authorising sources, the ability to interpret norms in different ways, and the evolutionary nature of a legal system whereby new norms come into being and coexist (at least for a time) with older inconsistent norms. The requirement of consistency also prevents contradictions in the factual premises. This is problematic when there is evidence for and against a particular factual premise.

Traditional formal logic also imposes the requirement of monotonicity. Suppose that a certain conclusion can be derived from a particular set of premises. Monotonicity requires that the addition of new information cannot effect the validity of the original conclusion. This assumption is often at odds with ordinary ways of thinking and, in particular, legal reasoning.

subject to the host country's criminal law. Another type of complexity which can be formalised is the weighing of reasons. Suppose that Jane is a 12 year old who has just stolen a toy Ferrari from a store. We can imagine two conflicting reasons being applied to this case. The fact that Jane is a thief provides the basis for a reason in favour of Jane being punished; but the fact that she is a minor provides the basis for a reason against her being punished. Reason-Based Logic allows for relative weights to be assigned to reasons. So, for example, if the system has been told that being a minor 'outweighs' being a thief with respect to whether or not some one ought to be punished, then the system can conclude that Jane should not be punished.

2.2. Prakken and Sartor's Argumentation Framework

The framework presented by Prakken and Sartor has many characteristics in common with Hage's Reason-Based Logic. The most obvious difference, however, is that Prakken and Sartor focus on arguments rather than reasons. For Prakken and Sartor, an 'argument' consists of a series of chained rules, which are grounded in facts and which lead to a logical proof of some final proposition. Prakken and Sartor attempt to model the strategies of argument attack and defence that are typical of legal discourse. In particular, they look at contradictory rules, rules with assumptions, inapplicability statements and priority rules.

According to Prakken and Sartor's framework, one argument may 'attack' another by contradicting either the first argument's conclusion ('rebutting') or one of its assumptions ('undercutting'). When one argument is undercut by another, it is said to be 'defeated' outright. Where, however, two arguments rebut each other the relative strengths of the two arguments must be assessed. An argument is stronger than its rival if it is based on a higher priority rule.

To determine the relative priorities of two conflicting rules, Prakken and Sartor invoke meta-level rules such as *lex superior*, *lex specialis* and *lex posterior*. So, suppose we have an European Community law allowing any product lawfully put into commerce in one EC country to be freely sold in any other EC country, and an Italian law forbidding the sale of products called 'pasta' not being made from hard corn.² Under these rules, can 'pasta' that is sold in Germany, but not made from hard corn, be sold in Italy as 'pasta'? The conflict between the two rules can be resolved by adopting a priority rule that 'if x is an European Community law and y is an Italian law, then x is superior to y'. Conflicts between priority rules can be addressed through the use of still higher level priority rules.

The notion of an argument's 'status' is also developed. Since defeating arguments may be defeated by other arguments, comparing argument pairs is not sufficient to determine the final status of an argument. Prakken and Sartor present a tripartite division of arguments. 'Justified arguments' are not attacked at all, or are only attacked by an argument that is itself defeated by a justified argument, 'overruled arguments' are those that are defeated by a justified argument, and the remainder are 'defensible arguments'.

²) Case 407/85 3 *Glocken GmbH and Gertraud Kritzinger v USL Centro-Sud and Provincia autonoma di Bolzano* [1988] ECR 4233 (the 'pasta case').

3. A Pragmatic Evaluation of Argumentation Frameworks in AI and Law

3.1. Introduction

In this section I consider whether the current generation of argumentation frameworks is likely to prove of practical significance in the development of AI applications in the legal domain. This is not a claim that the designers of these systems would necessarily make. Nevertheless, the potential for these systems to produce practical LKBSs should be of general interest within the AI and law community.

An initial observation is that the argumentation frameworks put forward by Hage, and Prakken and Sartor, are powerful systems for representing arguments in a retrospective sense. That is, an argumentation framework can successfully be used to formally describe the reasoning procedure and conflict resolution strategies of an existing piece of legal reasoning. For instance, an argumentation framework is well suited to reconstructing and describing the decision of the European Court of Justice that 'pasta' sold in Germany, but not made from hard corn, can be sold in Italy as 'pasta' because the permissive European Community law is superior to the prohibitive Italian law. In this sense, argumentation frameworks 'have shown their usefulness' (Verheij et al, 1997, at 248). However, it is questionable how useful this really is in the AI and law context. First, all the real analytical work is done by the person analysing the argument, not by the argumentation framework. Second, it is not at all obvious that the argumentation framework representation of the argument will be more edifying than a narrative explanation of the argument structure.³ And finally, the most apparent benefits of formally representing an argument – in terms of finding flaws and characteristic structures – are obtained from the pre-computational analysis when the knowledge engineer identifies the rules, facts and priority relationships; so what does the computational implementation achieve?

Using an argumentation framework to represent an argument in an AI application really only makes sense if the application is going to reason with the representation in some way, and apply it to factual settings other than the one in which the argument was first made. This is consistent with Sartor (1993, at 193) who says that inference procedures taking into account an ordering relation are needed in order 'to derive reasonable consequences out of a legal system' with inconsistent premises. This places us in the realm of prospective applications that are capable of applying a knowledge base of rules and reasons to new factual situations. Here the acid test for assessing the utility of argumentation frameworks in AI and law is: 'Are they more useful as the basis of forward-looking AI applications than the traditional logic systems they might replace?'

3) For example, the following statements would be necessary in Prakken and Sartor's argumentation framework to represent the narrative explanation that "'pasta' sold in Germany, but not made from hard corn, can be sold in Italy as 'pasta' because the permissive European Community law is superior to the Italian law which would otherwise prohibit the sale":

$r_1(x, y, z)$: x can be sold in $y \wedge y$ is an EC country $\wedge z$ is an EC country $\Rightarrow x$ can be sold in z

$r_2(x, \text{Italy})$: x is called 'pasta' $\wedge \sim(x$ is made from hard corn) $\Rightarrow (x$ can be sold in Italy)

$r_3(r_1(x, y, z))$: $r_1(x, y, z)$ is an EC law

$r_4(r_2(x, \text{Italy}))$: $r_2(x, \text{Italy})$ is an Italian law

$r_5(x, y)$: x is Italian law $\wedge y$ is an EC law $\Rightarrow x < y$; where ' $<$ ' means has priority over

$f_6(\text{Germany})$: \rightarrow Germany is an EC country

$f_7(\text{Italy})$: \rightarrow Italy is an EC country

$f_8(P_0, \text{Germany})$: $\rightarrow P_0$ can be sold in Germany

$f_9(P_0)$: $\rightarrow P_0$ is called 'pasta'

$f_{10}(P_0)$: $\rightarrow \sim(P_0$ is made from hard corn)

3.2. Advantages associated with argumentation frameworks

Four main advantages might be claimed by an advocate of argumentation framework use in LKBS design ahead of traditional formal logic. First, argumentation frameworks allow for the presence of inconsistent premises in the knowledge base and permit non-monotonic reasoning. Second, they go beyond traditional formal logic's exclusive focus on formal validity and define different relations of conflict which may exist between arguments. Third, argumentation frameworks model conflicts between competing conflict resolution principles in the same way as the conflicts between the arguments themselves (Bench-Capon, 1997; and Verheij et al, 1997, at 244). Fourth, argumentation frameworks assess conflicting arguments in an attempt to determine the overall outcome from a set of conflicting premises.

The first three of these claims are, I believe, uncontroversial and genuine benefits associated with the development of argumentation frameworks. I would only note as an aside that the often remarked upon problem of *ex falso sequitur quidlibet* (Berman and Hafner, 1987, at 31; Sartor, 1992, at 211; and Prakken, 1993a, at 94), that is, the ability to derive any proposition whatsoever from a system containing inconsistent axioms, has rarely caused problems in practice for the developers of logic-based LKBSs. The force of this surprising sounding result can be avoided if only conclusions which are grounded, that is, which can be logically derived from a consistent subset of the system, are considered (Rescher, 1964; Brewka, 1991; Sartor, 1992; and Prakken, 1993a).⁴ When this is done, the more intuitive outcome results and only proposition which are supported by deductively valid reasoning chains are derived. Therefore the development of inconsistency-tolerant logic systems is not as problematic as has been suggested in the past. This is illustrated by the development of such systems by Indurkha (1996), Freeman and Farley (1996) and St.-Vincent et al. (1995) amongst others.

The fourth claim – that the assessment of conflicting arguments is a feasible and desirable operation – is, however, more controversial. I consider this issue next, and then discuss a set of problems associated with traditional formal logic which are left unaddressed by argumentation frameworks.

3.3. The unfruitful path of argument assessment

Argumentation frameworks are intended to model the adversarial thrusts and parries of legal reasoning and to determine the end status of the combatant arguments. Imagine that the entire collection of arguments supporting or rebutting a particular proposition which can be derived from a set of (potentially inconsistent) rules and associated factual predicates are represented as the branches of a tree. What an argumentation framework attempts to do is prune this tree of its dead wood – those branches that are defeated by other arguments. In particular a branch may be pruned either because it is excluded by some other rule (or reason), or because it is based on a rule (or reason) which is outweighed by another rule (or reason).

The first point to note is that this pruning process cannot be implemented using only very general priority rules. General propositions such as the principle of *lex posterior*, which asserts that a recent law has priority over an older law, are disregarded as often as they are applied. Another problem is that competing priority rules are almost always available. This is particularly obvious with respect to the

4) What is almost never explained is that this problem is really the result of a technicality. Traditional formal logic uses 'resolution proofs by refutation' (Nikolopoulos, 1997). In essence to prove a theorem, its negation is assumed to be true and added to the set of axioms in the system. The next step is to attempt to derive a contradiction between those axioms. If a contradiction exists, then it is assumed that the added premise must be false, and therefore its negation, the theorem being tested, must be true. Now, if the axiom set already contains a contradiction before the procedure begins, then any and every theorem fed into it will give rise to a contradiction!

principles of case analysis and statutory interpretation which, as Llewellyn (1951 and 1960) has shown, exist in equally valid opposed pairs. For example, pairs of opposing canons of statutory interpretations, such as: 'a statute cannot go beyond its text' and 'to effect its purpose a statute may be implemented beyond its text'; and '[the] expression of one thing at the exclusion of another' and 'the language may fairly comprehend many different cases where some only are expressly mentioned by way of example' happily coexist in legal discourse.

The abstract priority relationships available to a judicial decision maker necessitate discretion and choice. So although it may be easy to model the use of a particular principle after the fact, abstract priority principles do not provide a reliable guide to conflict resolution in a forward-looking sense. The so-called dead wood on an argument tree cannot be separated from the vibrant flourishing branches merely by specifying priorities between rules and reasons at such a high level of abstraction.

The designers of argumentation frameworks acknowledge that these kinds of general priority rules are not valid in a universal sense. Therefore, an attempt is made to take into account those contextual factors that determine whether or not a priority principle will be applied in a particular case by allowing for argument to take place over which priority rule to apply. However, this does not provide a viable solution. To begin with, it is hard to even imagine how one would construct realistic meta-level priority rules that incorporated the relevant contextual factors, such as coherence with the surrounding body of law, the fine grained details of a case that determine its substantive merits (for example, in our case of young Jane the thief, the decision maker would want to know: Had she stolen before? What punishments are available if she is convicted? Did she steal cynically knowing that minors are not normally punished?), the current tradition of the court, the judge's basic sense of justice, and her perception of how the decision will be received by the legal community. The meta-priority rules suggested by Prakken and Sartor (1996, at 355), such as *lex superior* has priority over *lex posterior*, and those suggested by Hage (1996, at 260-261), such as a decision of a superior court outweighs a more recent decision of a lower court unless there has been a relevant social change, do not even come close to fitting the bill.

Even if more realistic meta-level priority rules were devised one could never have confidence that they would be determinate. Arguments can always be made against any priority rule. The whole enterprise of searching for firmly grounded, transcendental, meta-priority rules is misguided. Meta-level priority rules are no better grounded than the priority rules that they adjudicate over.⁵ The goal of finding definitive priority relationships that can be applied in a prospective sense is simply not tractable. The novelty of each moment and every situation means that the application of a priority principle must always be haunted 'by the doubt that some unique features of the present situation may prove crucial in the end.' (Billig, 1996, at 92.) Since no priority principle can be expected to discriminate accurately

5) The view that priority relationships can be grounded on some stable foundation if one looks deep enough brings to mind the following story from Cramton (1986, at 1-2):

A prominent scientist had just given a brilliant lecture on the foundations of the universe. During the question period an elderly woman suggested that there was a problem with the professor's analysis. "What is that?" asked the professor cautiously. "It's all wrong", the woman replied, "because the universe actually rests on the back of a giant turtle." The professor, taken aback, forced a smile and then countered: "If that's the case there is still the question, what is that turtle standing on?" The audience tittered, but the woman, undaunted, replied: 'Another, much larger turtle.' "But" objected the professor. "I'm sorry, Professor, it's turtles all the way down."

One may say that when it comes to using an argumentation framework for determining the priority of conflicting arguments, and the priority of conflicting priority relationships, and the priority of conflicting priority relationships of conflicting priority relationships, ... "it's turtles all the way down."

between winning and losing arguments, the risk is that a lot of winners will be mistakenly pruned from the argument tree and not made available to the user.

3.4. Aspects of legal argumentation that are not captured by argumentation frameworks

The derivation of robust and reliable priority relationships which could be applied in a LKBS does not seem feasible. However, even if it was, there are other aspects of legal argumentation not captured by the current generation of argumentation frameworks. This is not the place for a full discussion, but I will mention three important sources of the 'room to argue' in legal discourse that seem to resist treatment within any system based on formal logic. These are all problems which have been recognised with respect to traditional formal logic and remain unanswered in the argumentation framework context.

First, argumentation frameworks tacitly assume that legal rules (as well as principles and goals) have precise verbal formulations that can be found in a straightforward and uncontroversial fashion. In reality, however; lawyers often (perhaps typically) are able to argue for different formulations of the rule under consideration that would produce different results. If these competing formulations could be listed in advance then they could be modelled in an argumentation framework. However, it is almost impossible to imagine the creation of such a knowledge base for even a single legal rule.

Second, even if the formulation of a particular rule is agreed upon, it must be remembered that rules are not self-applying. Rather, the application of a rule to a set of facts requires a human 'judgement', 'decision', or 'act of will' (Wittgenstein, 1974). This judgement is required because arguments may be made for and against the use of a word or phrase as an appropriate characterisation of a given situation. For example, in the pasta case, the Court had to determine whether the challenged Italian law could be reprieved as being necessary to ensure 'fair trading' in Italy. The argumentation frameworks of Hage and Prakken and Sartor provide no guide as to how the arguments relevant to this rule application issue could be generated and assessed.

Third, even if a verbal formulation of a rule is agreed upon and it is accepted that the facts come under the explicit requirements of a rule, there may be some other facts which, if present, will invalidate the operation of the rule. The presence of the defeating event makes the apparent solidity and applicability of the rule disappear. This possibility is described as 'rule defeasibility'. Again this is illustrated by the pasta case where the Court was faced with the claim that even if the EC law applied, it was defeated because the Italian law was necessary to protect Italian producers of hard corn. If a defeasing condition is known then it can be modelled in an argumentation framework, but the real problem is that all the potential exceptions to a rule can never be fully enumerated in advance (MacCormick, 1995).

While an argumentation framework may be capable of illustrating the resolution of these kinds of arguments, argumentation frameworks are not well suited to automating the generation and resolution of such disputes.

3.5. Recommendations

I conclude this section with three recommendations for the improvement of argumentation frameworks. First, argumentation frameworks should not be used to assess conflicts between arguments as this cannot be done reliably. The assessment of conflicting arguments is more appropriately left to the user. Second, a richer standard form for representing arguments which distinguishes between the different roles that different types of premises may play in an argument would

seem appropriate. One such model, which is very familiar in AI and law research, is that of Toulmin (1958). Third, as I have argued elsewhere (Palmer, 1997), in addition to the 'rebutting' and 'undercutting' relationships found in contemporary argumentation frameworks, more complicated relationships of structural support and opposition between different argument building-blocks should be explored as a means of enriching the framework.

4. Conclusions

Argumentation frameworks are a step in the right direction. They recognise the importance of, and provide computational implementations of, premise inconsistency and non-monotonic reasoning. The incorporation of these features and the definition of relations of conflict which may exist between arguments make argumentation frameworks much more powerful for representing legal reasoning than traditional formal logic. Argumentation frameworks provide a powerful representational language when used to model past reasoning and should be acknowledged as such. However, when it comes to AI and law applications which operate in the prospective sense they are unable to cope with the creativity and choice that characterise argument generation and assessment in legal discourse.

While legal reasoning often has the appearance of inevitability after the fact – after the judgment is given, or after the textbook is written – the forward looking process of legal reasoning is one of creative invention. Priority relations are not passively applied by judges, rather they are created by judicial decisions. The corpus juris is an organised chaos of cases, rules and principles with which multiple conflicting arguments can be constructed. Some arguments will be clearly available, others may not be. Only more research and thought will tell. Other avenues of argument may seem open but close off as more possibilities are considered. Which ever way the process goes, the key is to broaden the range of arguments under consideration, not restrict it by imposing abstract priority relationships that may or may not, hold in any given context. There are no transcendental principles for assessing arguments, assessment is always dependent upon context and audience. Therefore, the basic approach of argumentation frameworks towards argument assessment and pruning branches from the argument tree is unsound.

The inability of argumentation frameworks to deal with the issues and complexities of real world arguments is reminiscent of the failure of attempts to extend early AI 'micro-world' programs into more realistic domains. Micro-worlds are small, simple, artificial domains which were the concern of a branch of early AI programs. Winograd's (1972) SHRDLU was the classic micro-world program. Its world consisted of a tabletop, a box and a number of coloured blocks. SHRDLU was able to follow natural language instructions which required it to move its blocks around, and it was able to answer questions about its world. At the time it was thought that these same abilities could be implemented in real world domains. This never happened. As Crevier (1993, at 102) has explained '[i]t soon became clear that [SHRDLU] could not be extended beyond the Blocks Micro World. The simplicity, logic, and isolation of this domain allowed the appearance of intelligent dialogue by simply dodging difficult language issues.'

Similarly, the issues of context, rule formulation, rule application and rule defeasibility, are all dodged when argumentation frameworks are used in a retrospective sense. The rules, facts and priority relationships to be modelled are all 'cleaned up' before being brought anywhere near the argumentation framework. They are presented to the argumentation framework as clear, discrete, well-labelled interlocking entities, ready to be manipulated; much like the coloured play-blocks of a micro-world. So, when argumentation frameworks are demonstrated by modelling

an existing piece of legal reasoning they seem very powerful. However, when an argumentation framework is applied in a prospective sense its fragility becomes apparent. Priority rules that were clear *ex post* are unreliable *ex ante*, rules that were clearly expressed become subject to multiple formulations, and the application of rules which seemed clear are now subject to the fluidity of language. To deal with these issues requires a degree of common-sense that was beyond the state of the art when micro-world programs were being extended into the real world, that is beyond the current state of the art, and that may always exceed it.

References

- Bench-Capon, T.J.M. (1997) *Argument in Artificial Intelligence and Law*, forthcoming in *Artificial Intelligence and Law*
- Billig M. (1996) *Arguing and Thinking*. Cambridge University Press: Cambridge.
- Berman, D.H. and Hafner, C.D. (1987) Indeterminacy: A Challenge to Logic-based Models of Legal Reasoning, 3 *Yearbook of Law, Computers and Technology* 1-35.
- Brewka, G. (1991) *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press: Cambridge.
- Cramton, (1986) Demystifying Legal Scholarship, 75 *Georgetown Law Journal* 1.
- Crevier, D. (1993) *AI: The Tumultuous History of the Search for Artificial Intelligence*. Basic Books: New York.
- Freeman K. and Farley A.M. (1996) A Model of Argumentation and its Application to Legal Reasoning, 4 *Artificial Intelligence and Law* 163-197.
- Hage, J. (1996) A Theory of Legal Reasoning and a Logic to Match, 4 *Artificial Intelligence and Law* 199-273.
- Hage, J. (1997) *Reasoning with Rules*. Kluwer Academic: Dordrecht.
- Indurkha, B. (1996) Using Set-of-Support Control Strategy to Deal with Indeterminacy in Legal Reasoning, in *Proceedings of the Ninth International Conference on Legal Knowledge Based Systems, JURIX-96* 123-134.
- Llewellyn, K.N. (1951) Remarks on the Theory of Appellate Decision and the Rules or Canons About How Statutes Are to Be Construed, 3 *Vanderbilt Law Review* 395.
- Llewellyn, K.N. (1960) *The Common Law Tradition: Deciding Appeals*. Little, Brown.
- MacCormick, N. (1995) Defeasibility in Law and Logic, in Zenon Bankowski, Ian White and Ulrike Hahn, *Informatics and the Foundations of Legal Reasoning* 99-117.
- McCarty, L.T. (1997) Some Arguments about Legal Arguments, in *Proceedings of the Sixth International Conference on Artificial Intelligence and Law* 215-224.
- Palmer, J. (1997) Legal Merit Arguments, Legal Semiotics and the Design of Legal Knowledge-Based Systems, in *Proceedings of the Sixth International Conference on Artificial Intelligence and Law* 198-206.
- Peczenik, A. (1996) Jumps and Logic in the Law, 4 *Artificial Intelligence and Law* 297-329.
- Poole, D.L. (1988) A Logical Framework for Default Reasoning, 36 *Artificial Intelligence* 27-47.
- Prakken, H. (1993a) An Argumentation Framework in Default Logic, 9 *Annals of Mathematics and Artificial Intelligence* 93-132.
- Prakken, H. (1993b) A Logical Framework for Modelling Legal Argument, in *Proceedings of the Fourth International Conference on Artificial Intelligence and Law* 1-9.
- Prakken, H. (1995) From Logic to Dialectics in Legal Argument, in *Proceedings of the Fifth International Conference on Artificial Intelligence and Law* 165-174.

- Prakken, H. and Sartor, G. (1996) A Dialectical Model of Assessing Conflicting Arguments in Legal Reasoning 4 *Artificial Intelligence and Law* 331-368.
- Prakken, H. and Sartor, G. (1997) Reasoning with Precedents in a Dialogue Game, in *Proceedings of the Sixth International Conference on Artificial Intelligence and Law* 1-9.
- Reiter, (1980) A Logic for Default Reasoning 13 *Artificial Intelligence* 81-132.
- Rescher, N. (1964) *Hypothetical Reasoning*. North Holland: Amsterdam.
- Rissland, E.L. (1990) *Artificial Intelligence and Law: Stepping Stones to a Model of Legal Reasoning* 99 *Yale Law Journal* 1957-1981.
- St.-Vincent P., Poulin D. and Bratley P. (1995) A Computational Framework for Dialectical Reasoning, in *Proceedings of the Fifth International Conference on Artificial Intelligence and Law* 137-145.
- Sartor, G. (1992) Normative Conflicts in Legal Reasoning 1 *Artificial Intelligence and Law* 209-235.
- Sartor, G. (1993) A Simple Computational Model for Nonmonotonic and Adversarial Legal Reasoning in *Proceedings of the Fourth International Conference on Artificial Intelligence and Law* 192-201.
- Toulmin, S. (1958) *Uses Of Argument*. Cambridge University Press: Cambridge.
- Verheij, B., Hage, J. and Lodder, A. (1997) Logical Tools for Legal Argument: A Practical Assessment in the Domain of Tort, in *Proceedings of the Sixth International Conference on Artificial Intelligence and Law* 243-249.
- Winograd, T. (1972) *Understanding Natural Language*. Academic Press: New York.
- Wittgenstein, L. (1974) *Philosophical Investigations* (G.E.M. Anscombe trans.)