

Open Texture and Argumentation: what makes an argument persuasive?

Trevor Bench-Capon

Department of Computer Science, The University of Liverpool, Liverpool, UK

Abstract. Although Marek Sergot's contribution to Artificial Intelligence and Law is mainly associated with the formalisation of legislation as a logic program, he also wrote on an approach to the treatment of open textured concepts in law, using argumentation. That paper posed the question *what makes an argument persuasive?* This short paper considers the ideas of that paper and discusses developments in AI and Law over the subsequent 25 years, focusing on the progress made in answering this question in that domain.

1 Introduction

Marek Sergot is best known in AI and Law circles for his advocacy of the formalisation of legislation as a logic program as a method for building legal knowledge based systems. The position is stated at it sharpest with respect to the British Nationality Act (BNA) in [48]. Some issues with the approach are identified in respect of a large scale formalisation of UK Supplementary Benefit legislation in [15], and a further practical application is described in [47].

Such an approach undoubtedly has its strong points and works relatively well when the legislation is built around such readily understood and reasonably precise notions like *father*, *age* and *place of birth* as are found in the British Nationality Act. But very often legislators use rather vague terms, attempting to make their general intentions clear, but leaving the specific boundaries of the concept to be established by courts in the light of the facts of individual cases.

For example, it may be that a welfare benefit, *Heating Addition*, is introduced to supplement the benefits of pensioners who live in places which are considered difficult (and hence more costly) to heat. This could be done with a clause such as:

A Heating Addition shall be payable to a pensioner whose normal residence is hard to heat.

Now an expert system based on a formalisation of this legislation together with a Query the User [46] module will produce the following dialogue (user input in italics).

- *Is Peter entitled to Housing Addition?*
- Is Peter a Pensioner?

- *Yes.*
- What is Peter’s normal residence?
- *23, Acacia Avenue, East Cheam.*
- Is 23, Acacia Avenue, East Cheam hard to heat?
- *Yes.*
- Yes, Peter is entitled to Heating Addition.

This is fine as far as it goes, but the real difficulty in determining whether Peter is entitled is deciding whether 23, Acacia Avenue, East Cheam is hard to heat, and perhaps, whether 23, Acacia Avenue, East Cheam is Peter’s *normal* residence. To answer these questions (or at least to answer them with legal authority) it is essential to consider the case law on the topic, and to see what courts have considered important in deciding whether a house is hard to heat in the past. Subjective opinion is not enough, nor is common sense. (For example, it may turn out that any house containing a person over eighty years of age is considered by the courts to be *ipso facto* hard to heat.) Vague terms, the application of which is not fixed in advance of their use, are said to be *open textured* [51]: the boundaries of the concept are not sharply drawn, but are gradually specified (by persons empowered to specify them, such as benefits adjudicators) as and when particular cases arise. Sergot and Bench-Capon [8] addressed this difficult problem of the treatment of open textured concepts in logic programs¹, a solution to which is essential to supplement formalisations of legislation if useful systems are to be built.

That paper first considered and rejected some approaches which were current at the time. For example, *approximation*, using a sharp concept (e.g. heating bills) to replace the open textured one, was rejected as imposing an unjustified interpretation, which might well not conform to what the court decides in practice. Scorn was poured on the use of the numeric techniques of fuzzy logic and probabilities: the numbers to use could not be sensibly determined and both “the appearance of precision and the appearance of generality are spurious”. Instead, it was argued that we should embrace the adversarial nature of legal reasoning and recognise the application of open textured concepts is in practice resolved by an adjudicator who, having heard the *arguments* on both sides of the question, is empowered to decide which position should be accepted, and so how the term should be used subsequently. Thus we proposed that the logic program should represent the reasons for and against the application of the concept and we should regard the set of derivations, some saying yes and some saying no, as arguments for and against the application of the concept. The better arguments would then be chosen. At the time, in the then absence of anything better, this task could be left to the user. Of course, automated evaluation of arguments was the ultimate goal. This would require “a representation in computer intelligible terms of what it is that makes an argument persuasive”, reasons why an argument should be accepted or rejected by a given adjudicator.

¹ I am particularly fond of that paper since it was my first computer science publication.

The challenge to discover what makes an argument persuasive, both in general and in law in particular, thrown down in [8], has been an important and recurring theme running through AI and Law ever since. In this paper I will consider the question by looking at how it has been addressed from the perspective of AI and Law. Any such survey is necessarily somewhat selective, and there are, of course, other approaches in general AI which have not been much taken up in AI and Law and which will be therefore outside of the scope of this paper.

2 Arguing with Cases: Dimensions and Factors

After the presentation of [8], the primary focus of logic programming in AI and Law returned to logical models of legislation, and various issues of knowledge representation, software engineering and methodology. Approaches involving isomorphism [13] and the use of ontologies to underpin legal knowledge based systems (e.g. [17]) were to develop from this work.

Meanwhile argumentation to resolve open textured reasoning was mainly addressed in the context of reasoning with legal cases, most notably the work of Rissland and Ashley on HYPO (most fully reported in [4]) and of Ashley and Alevan on CATO (most fully reported in [1]). These systems described cases in terms of particular aspects representing patterns of fact appropriate to the particular area of law under consideration². In HYPO the facts gave rise to a position on a *dimension*. Dimensions favour the plaintiff at one extreme and defendant at the other, and so the party favoured in a particular case and the degree of support given depends on where on this range the particular case lies. In CATO the facts of cases are analysed so that *factors* can be ascribed to the case: if present a factor favours *either* the plaintiff or the defendant. Factors can be seen as particular points on dimensions. The idea was to match a past case to the current case on these dimensions (or factors) and use that match as a reason to decide the same way in the current case.

The notion of argumentation as developed in HYPO and CATO has an adversarial three-ply structure. First a precedent case is cited as a reason to decide the current case in a particular way. Then the other side attempts to distinguish the cases, by citing particular mismatches, or provide a counter example contradicting the precedent. In the third ply the original side has the opportunity to rebut the distinctions as unimportant, or distinguish the counter example.

This three-ply argumentation with dimensions, or more commonly factors, became established the dominant model of argumentation used to resolve problematic cases in AI and Law. The programs, however, do no more than present the arguments. Where there are arguments on both sides, the user must decide which will be accepted by the judge if the case comes to trial. The dominance of the case based approach led to a view that case based reasoning was needed for open texture and that the logic programming approach, or other rule based approaches, were suited only to Civil Law, or at least routine cases.

² Both HYPO and CATO took US Trade Secrets as their domain.

In the meantime logic programming in AI and Law tended to use argumentation for several purposes not specifically connected with resolution of open texture. One important application of rule based techniques was as the basis for dialogue systems, most notably the Pleadings Game [25], where the focus was on capturing the procedure used to make legal decisions. One view is that a legal decision is correct if it was made using the correct procedure, and one such procedure was modelled in [25] as the exchange of arguments.

Another use was for explanation, for example [18] in which the argumentation scheme of Stephen Toulmin [49] was used to structure the explanation from a logic program, annotated to indicate the various roles of the clauses in the bodies of the rules.

Finally argumentation was used by Prakken to resolve conflicts between norms (e.g. [39]). This work was, however, done within the European Civil Law tradition rather than the UK-US common law tradition, and so the focus was on conflicts derived from statutes, rather than arising from open texture and decided in the context of particular cases. In this context three principles could be used to determine which of two conflicting arguments should succeed.

- Prefer the more *specific* law: thus a norm referring to cats should (if the animal in the case is a cat) be preferred to a norm referring to animals.
- Prefer the more *recent* law: thus a norm taken from a later statute should be preferred to one taken from an earlier statute.
- Prefer the more *authoritative* law: thus a law passed by a national Government should be preferred to a local by-law.

These principles could readily be encoded, and so give rise to the possibility of automatic resolution of these conflicts. One problem, however, is that the principles may themselves conflict: a more recent law may come from a less authoritative source, and very often a local law will be more specific than a national one. Originally Prakken opted for a single principle, always preferring the most specific argument, but later he recognised that the principle to be given priority was itself something to argue about³.

These three principles, however, represent a first systematic answer to the question: *what makes an argument persuasive?*

3 Rise of Abstract Argumentation

While these events were happening in AI and Law, elsewhere in AI the notion of abstract argumentation was being developed. This idea was presented first in [22] and in its fuller form in [23]. The idea here was to present arguments in a framework in which they would be entirely abstract, related only by a

³ In passing, one might remark how influential the 4th ICAIL, held in Amsterdam, was, especially the opening morning. The very first session contained both [39] and [25]. After coffee, [19], of which more later, and [28], another approach based on balancing reasons for and against a decision, were presented.

binary relation of attack. Often it is convenient to depict the framework as a directed graph. Intuitively an argument is acceptable if and only if it has no acceptable attacker⁴. Normally attackers will be rendered unacceptable by some other argument. Thus the acceptability of an argument is determined in the context provided by the other arguments in the framework, and is relative to a subset of these arguments. A set of arguments which contains no arguments attacking a fellow member of the set is known a *conflict free* set. A conflict free set in which every attacker of an argument in the set is attacked by some member of the set is known as an admissible set. We can now define acceptability in terms of a variety of semantics (see [6] for a selection), of which the most important are the *grounded* (a maximal admissible set where arguments are not permitted to defend themselves) and the *preferred* (a maximal admissible set where arguments are permitted to defend themselves). Whereas there is always a unique grounded extension (possibly the empty set), there may be several preferred extensions. An argument acceptable in all preferred extensions is said to be *sceptically* acceptable, while an argument acceptable in at least one, but not all, preferred extensions is *credulously* acceptable. Abstract argumentation was made known to AI and Law very early with [43] published in 1995 and [32] in 1996, but regrettably these papers did not make any significant impact at that time⁵. Abstract argumentation as proposed in [23] was, however, to prove enormously significant in the development of argumentation as a subfield of AI⁶, and was eventually to make an important impact of AI and Law, as well.

Application of these ideas to AI and Law is relatively straightforward. Suppose we have a knowledge base which describes the law (statute law and, where relevant, case law also) and the facts of a case. We may now compute all the arguments that can be derived from this knowledge base. So far this is exactly the process envisaged in [8]. Once we have the set of arguments we can determine the attack relations between them. It quickly became established that there were three ways to attack an argument: rebuttal where the conclusion of an argument \mathcal{A} is the negation of the conclusion of an argument \mathcal{B} ; premise defeat where the conclusion of an argument \mathcal{A} is the negation of a premise of an argument \mathcal{B} ; and - for defeasible reasoning - undercut, where the conclusion of an argument \mathcal{A} renders an argument \mathcal{B} inapplicable. The best current exposition of this approach is [42]. The approach was applied to a substantial body of case law in [9].

What does this approach add to the question *what makes an argument persuasive?* First it establishes that this is not a question that can be answered with respect to consideration of an argument in isolation: arguments are acceptable with respect to other arguments within a context given by the argumentation

⁴ For formalisations of this notion in the context of the semantics of logic programming see also [31] and [29], which could be regarded as the origin of abstract argumentation.

⁵ Of the 47 citations to [32] given on Scopus in September 2011, only three were before 2001.

⁶ Over 600 citations on Scopus as of September 2011.

framework. Secondly it indicates that acceptability needs to consider the attack and defence relations that arise in that context. Work such as [9] also suggested that, in this approach, there would typically be multiple preferred extensions, arising from a cycle in the graph representing the argumentation framework and representing opposing sides to the issue. Often such a cycle would be where a court made an important decision between two competing arguments, which subsequently influenced other decisions. The abstract argumentation approach was able to identify such choices, but not to record justifications for them⁷.

4 Purpose and Value

As indicated at the end of the last section, the acceptability of an argument in a body of case law could often be said to depend on a choice made between two arguments is some landmark case. But what motivated that choice? Moreover very often the conflict was resolved by a majority vote: while the majority was persuaded by the argument accepted, there was a minority who found its rival more persuasive. Thus both arguments were persuasive, but persuasive to different people. This fits very well with the ideas of Perelman’s *New Rhetoric* [38], which argued that the persuasiveness of an argument was relative to the *audience* to which it was presented. This gives rise to the question of how we are to characterise the audience.

The answer to this question that has received most attention in AI and Law⁸ has its origins in another paper from the 1993 ICAIL, [19]. In that paper the authors suggested that within the CATO approach the choice of one set of factors over the other was often motivated by the social purposes served by deciding one way rather than the other. As later developed in work such as [16], the social purposes were referred to as the *values* that will be promoted by deciding for the party favoured by a factor when that factor is present in a case. In this way, the choice can be seen as preferring one value to another, and therefore made according to a preference order on values. This enables us to characterise an audience as a preference ordering on values.

These ideas were integrated with abstract argumentation in [10], producing *Value Based Argumentation Frameworks* (VAFs). Here every argument was associated with a particular value (through a function *val*) and the attack relation of [23] supplemented by a defeat-for-audience relation. An argument \mathcal{A} defeats \mathcal{B} for an audience \mathcal{P} if and only if \mathcal{A} attacks \mathcal{B} and *val*(\mathcal{B}) is not preferred to *val*(\mathcal{A}) by the audience \mathcal{P} . This enables the definition of sets which are *conflict*

⁷ As I will discuss in the following sections, justifying the choice represented by an even cycle in terms of case law and the preferences expressed in precedent decisions has been an important concern of AI and Law: there may be potential for this work to feed back into abstract argumentation generally.

⁸ There are other approaches using conditional priorities and higher order priorities to resolve conflicts abouts priority arguments in the AI and argumentation literature, but in this paper we are considering the issues from an AI and Law perspective, and so we will not consider them here.

free for audience \mathcal{P} , admissible for \mathcal{P} , acceptable to \mathcal{P} , etc. We can also distinguish between arguments acceptable to all audiences (*objectively* acceptable) and arguments acceptable to at least one audience (*subjectively* acceptable). Note that an argumentation framework specific to audience \mathcal{P} can be produced from a VAF by removing all the attacks which fail for \mathcal{P} : i.e do not correspond to a defeat for \mathcal{P} . In order to associate arguments with values, deciding for a particular party was construed as an example of practical reasoning using the argumentation scheme of [5]. The whole approach to using VAFs in legal reasoning is described in [12].

This work then defines another component to add to our answer to *what makes an argument persuasive?* We now recognise that the persuasiveness of the argument depends on a context comprising both the other relevant arguments and the *audience* evaluating these arguments. Different audiences will, because of different aspirations and preferences, find different arguments persuasive: thus the persuasiveness of an argument will always be relative to the value ordering of the audience to which it is addressed.

5 Extending Argumentation

The use of purposes and values gave a means of justifying the preference of one argument over another. Value-Based Argumentation supports, however, consideration of only one property of an argument, its value, or the purpose it promotes. Greater flexibility can be given by generalising this idea. This was done in Extended Argumentation Frameworks (EAFs) [35]. In that important paper a new kind of attack was introduced, so that arguments could now attack and defeat *attacks* as well as arguments. This generalises VAFs, for example, in that we can now say that an argument which concludes that $val(\mathcal{B})$ is preferred to $val(\mathcal{A})$ defeats the attack of \mathcal{A} on \mathcal{B} . Arguments attacking attacks need not be restricted to simple value preferences, however. The use of EAFs allows for other properties of arguments to be used to resist attacks, and, even more importantly, value preferences to be argued for, rather than simply stated. EAFs also generalise not only VAFs, but also other systems of preference between arguments as is shown in [37].

Using EAFs allows the choice of audience to itself be justified. For example, in law, we may decide that while finding for the defendant would provide a *bright line* and so promote clarity of the law, deciding for the plaintiff would give encouragement to a socially useful activity, and so have benefits of its own. Deciding the case now requires us to choose whether we prefer the value of clarity to social usefulness. The case in mind is the famous (and much discussed in AI and Law) property law case of *Pierson v Post*. Post was chasing a fox with horse and hounds but before he could catch it, Pierson clubbed the fox to death with a fence post. Among the judges deciding the case were Tompkins and Livingston. The majority (in an opinion written by Tompkins) held that to claim possession of a wild animal required the bodily seizure of that animal, and any less stringent requirement would make the law too unclear to apply in practice, but a minority

(in an opinion written by Livingston) claimed that Post should win so that the socially useful (in their view) activity of fox hunting should be encouraged. We thus have initial arguments \mathcal{P} (*find for Post to promote SocialUtility*) and \mathcal{D} (*find for Pierson to promote Clarity*). \mathcal{P} and \mathcal{D} attack one another. The arguments of Tompkins for the majority then attack a claim that $SocialUtility \succ Clarity$, which attacks the attack of \mathcal{P} on \mathcal{D} . Livingston’s arguments for the minority attack the claim that $Clarity \succ SocialUtility$ which attacks the attack of \mathcal{D} on \mathcal{P} . The arguments of Tompkins and Livingston also attack one another, but we can use the decision in *Pierson v Post* to attack the attack of the preference for the value of the minority argument on the preference for the value of the majority argument. This EAF, shown in Figure 1, can then be used to justify finding for the defendant, on the basis of clarity, in future similar cases. Thus EAFs give us a way of relating cases, arguments presented in cases, with precedents and the arguments presented in precedents, and the principles used to adjudicate between them. This approach is fully expounded in [14].

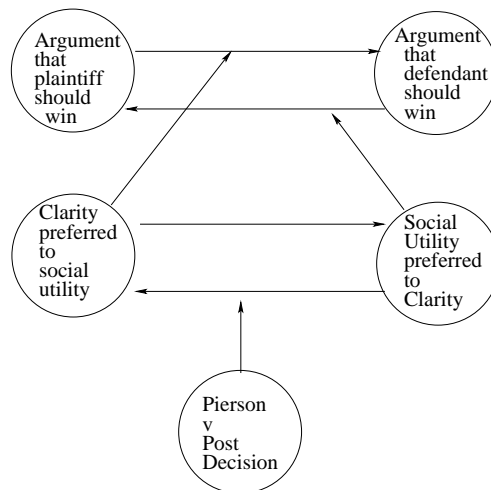


Figure 1: EAF applying Pierson v Post

EAFs provide a very useful piece of machinery which has been used in AI and Law⁹ to provide answers to our question *what makes an argument persuasive?*. They enable us to specify whole theories recording expertise about that question, and we can make these theories context specific if we wish. Such theories can be based on general principles (*prefer the more specific law or provide a bright line distinction*), or very specific reasons (*Bench-Capon often confuses US film directors of the 50s*), or precedents (*Pierson v Post showed that Clarity*

⁹ There are other approaches to reasoning about priorities to support general decision making in the argumentation literature, such as [30], which explores how different personalities based on Maslow’s hierarchy of needs [34] will prioritise different tasks differently. Despite the interest and importance of this work, since it has not been explored in AI and Law, it is outside the scope of this paper.

was preferred to Social Utility) or whatever else we wish to use as arguments to arbitrate between the persuasiveness of competing arguments. EAFs do not answer the question of what makes an argument persuasive themselves, but they do enable us to deploy computationally any answers we may find.

6 Current Issues

The previous sections have shown that our understanding of computational argument, and how to address the question of *what makes an argument persuasive?*, has come a very long way since [8]. It is, however, far from complete, and is still an active topic of research in AI in general and in AI and Law in particular. I will look at two issues here.

6.1 Accrual

One thorny issue is that abstract argumentation uses a relation of attack, but has no relation of *support*. Intuitively, however, in a collection of arguments some will oppose a given argument and some will support it. There have been attempts to incorporate an explicit support relation (e.g. [20]), but these have not been entirely successful, and represent a departure from the spirit of abstract frameworks as proposed in [23]. Moreover, the notion of support can be captured by considering the minimal admissible set containing the desired argument. The arguments in that set support the argument in that they are necessary for its acceptance.

Even if, however, one can reject the necessity for an explicit relation for logical support, there is, once one admits values, a problem of arguments which promote several values, and of multiple arguments for the same conclusion. This is not a concern in logical argumentation or proof, where arguments are either conclusive or useless. If an argument strictly entails its conclusion one argument is enough, and in this style of reasoning a flawed argument is useless, and a hundred flawed arguments are equally useless. But not so once we think in terms of values. Suppose I can choose to go to Paris or Benidorm on holiday. Benidorm has better weather, but Paris has better attractions and better cuisine. I may regard weather as the most important value, and so, considering the values and arguments individually, Benidorm is preferred. But if I take into account not only the attractions of Paris but also the quality of cuisine, that is I allow the reasons to accumulate or *accrue*, I should prefer Paris. As this example shows, it seems that the persuasiveness should be considered as something that can be summed across arguments when deciding which conclusion to choose¹⁰.

Value Based arguments as introduced in [10] only allowed an argument to be associated with a single value. Sets of values were used in [40]. There the idea was to compare a set according to the most preferred value of the sets, ignoring any values in common between them. In practice this gives the same result as

¹⁰ This issue of accrual is quite long standing, going back at least to [50].

VAFs as used in [12], where arguments with the same conclusion but different values mutually attack. A more formal approach is taken in [41], where accrual gives rise to a new super-argument, with the combined force of the accruing arguments. An alternative approach using EAFs was given in [36], where the *preferences* accrue so that the combination of the weaker values is preferred to the individually stronger value, and this preference defeats the attacks of the argument based on the stronger value on the arguments based on the weaker value. This allows the original arguments themselves to succeed and removes the necessity to combine them into new arguments. Looking beyond AI and Law, another approach, based on Defeasible Logic programming, can be found in [33] and application of argumentation to multi-criteria decision making has also been explored in [2] and [3]. How to treat accrual remains an important open issue in computational argumentation, and one which is considerable significance for AI and Law.

6.2 Degrees of Promotion

Apart from values having a cumulative effect, they may also be promoted to differing *degrees*. Thus over eating damages your health, over drinking significantly damages your health and smoking severely damages your health. So if you can only give up one of these activities, you should make it smoking. Where we are comparing arguments on the same value this is a help, not a problem. But if we have different values there is a problem: does significant promotion of a less favoured value outweigh slight promotion of a more favoured value? Of course, these problems can add complication to the accrual problems discussed above. A fully arithmetic treatment was given in [21] in which values were given integer weights and different ways of promoting these values weights were given between 0 and 1 so as to modify the effect of the values. While this gave pragmatically good results, it contributed little to an understanding of what is involved in making an argument persuasive. It is difficult, however, to see a real alternative at present.

This issue is a current topic in AI and Law. In particular it is noted that often a sequence of cases, such as those dealing with the so-called automobile exception to the Fourth Amendment¹¹, considers a clash between two key values, such as *LawEnforcement* and *Privacy*, but sometimes one seems preferred and sometimes the other, so that there is no discernible ordering on them. There have been several attempts to resolve and explain this phenomenon.

In [7] instead of comparing values, each of the values is compared against a threshold, and the threshold requirements must be satisfied in order for the decision to be made. Applying this approach to *Pierson v Post*, we would conclude either that Tompkins did not consider that the social utility was sufficient, or that too much unclarity in ascribing possession would result from a decision for Post, or both. More recently, Sartor [45] has argued that rather than a series of

¹¹ This topic was introduced to AI and Law in [44], but the current interest was sparked by [27].

thresholds, the resolution should be seen in terms of a trade-off, so as to strike an appropriate and proportional balance between the values. Thus in *Pierson v Post* the question would be whether the social utility was sufficient to support the degree of unclarity that would be required to decide for Post. A similar effect could be produced in [7] if the threshold for one value were stated as a function of the threshold for the other. Should a balance be required, the question arises as to whether the balance should reflect (for US Constitution issues) that struck by the founding fathers, or whether it can respond to societal changes. Most recently, rather than looking for a balance between two values, Grabmair and Ashley [26] have argued that what is needed is what they term a *value judgement*, in which all the relevant values are considered together. The judgement is whether overall the values are promoted to a greater extent that they are demoted. In [26], promotion is estimated qualitatively (*somewhat, greatly* and *overwhelmingly*) rather than quantitatively as in e.g. [21].

From this body of work we can draw out the following questions:

1. Is promotion and demotion of values to be regarded as boolean (either promoted or not) as in [16], ordinal (instances of promotion can be ordered), qualitative (instances of promotion can be placed into bands, as in [26]), or quantitative (instances of promotion can be assigned specific numeric values) as in [21]?
2. Should we see promotion and demotion as relative to thresholds as in [7], or to be considered as trading-off against one another as in [45]?
3. Should values be considered separately as in [7], pairwise, as in [16], or collected together as in [26] and [21]?

These questions are taken from [11], which discusses them in the light of a sequence of Fourth Amendment automobile cases. That paper can find no evidence in the decisions that anything more sophisticated than the threshold approach is required, but whether this view can be justified in terms of legal theory is a matter on which the jury is still out, and one looks forward to much subsequent work on this topic.

7 Conclusion

This paper has undertaken a very rapid survey of work done on answering the question posed in [8] as to *what makes an argument persuasive?* The perspective taken here is that of work done in AI and Law, the original context in which [8] posed the question. The question itself is of great significance, going right to the heart of legal reasoning, and indeed much other reasoning. The question remains unanswered, but in the 25 years since [8] we have developed valuable insights into what is relevant, especially the other arguments in the context and the audience to which it is addressed, and useful computational methods to support these insights, notably Abstract Argumentation, combining structured and abstract argumentation using argumentation schemes, and Extended Argumentation Frameworks. Outside of AI and Law, there are still other important

approaches to the question, which I have not been able to discuss here. In return AI and Law may be able to contribute to argumentation generally, perhaps through its work on values, purposes and precedents. We have come a long way, but there is still a long way to go. Argumentation is like an ocean: the further one wades into it, the deeper it gets.

8 Acknowledgments

I would, of course, like to thank Marek Sergot for introducing me to the Logic Programming approach to AI and Law in the first place. Although Marek himself has moved on from Law to other, and possibly better, things, the topic has kept me gainfully employed for over a quarter of a century. Thanks also to the others that I have discussed open texture and argumentation in AI and Law with, especially (in no particular order) Frans Coenen, Henry Prakken, Giovanni Sartor, Katie Atkinson, Alison Chorley, Sanjay Modgil, Edwina Rissland, Kevin Ashley and Tom Gordon. I should also mention David Moore, Floriana Grasso and Paul Dunne, insights from whom have also made key contributions to the work reported above.

I would also like to thank the anonymous reviewers for their comments and for drawing my attention to some interesting additional references.

References

1. V. Aleven. *Teaching Case Based Argumentation Through an Example and Models*. Phd thesis, University of Pittsburgh, Pittsburgh, PA, USA, 1997.
2. L. Amgoud, J.-F. Bonnefon, and H. Prade. An argumentation-based approach to multiple criteria decision. In Godo [24], pages 269–280.
3. L. Amgoud and H. Prade. Using arguments for making and explaining decisions. *Artif. Intell.*, 173(3-4):413–436, 2009.
4. K. D. Ashley. *Modeling Legal Argument*. MIT Press, Cambridge, MA, USA, 1990.
5. K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Synthese*, 152(2):157–206, 2006.
6. P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artif. Intell.*, 171(10-15):675–700, 2007.
7. T. Bench-Capon and H. Prakken. Using argument schemes for hypothetical reasoning in law. *Artif. Intell. Law*, 18(2):153–174, 2010.
8. T. Bench-Capon and M. Sergot. Towards a rule-based representation of open texture in law. In C. Walter, editor, *Computer Power and Legal Language: The Use of Computational Linguistics, Artificial Intelligence, and Expert Systems in the Law*, chapter 6, pages 39–60. Quorum, New York, 1988. From the Second Annual Conference on Law and Technology, University of Houston, 24–28 June 1985.
9. T. J. M. Bench-Capon. Representation of case law as an argumentation framework. In *Proceedings of Jurix 2001*, pages 103–112, 2001.
10. T. J. M. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. Log. Comput.*, 13(3):429–448, 2003.
11. T. J. M. Bench-Capon. Relating values in a series of supreme court decisions. In *Proceedings of Jurix*, page 13, 2011.

12. T. J. M. Bench-Capon, K. Atkinson, and A. Chorley. Persuasion and value in legal argument. *J. Log. Comput.*, 15(6):1075–1097, 2005.
13. T. J. M. Bench-Capon and F. P. Coenen. Isomorphism and legal knowledge based systems. *Artificial Intelligence and Law*, 1:65–86, 1992.
14. T. J. M. Bench-Capon and S. Modgil. Case law in extended argumentation frameworks. In *The 12th International Conference on Artificial Intelligence and Law*, pages 118–127, 2009.
15. T. J. M. Bench-Capon, G. O. Robinson, T. Routen, and M. J. Sergot. Logic programming for large scale applications in law: A formalisation of supplementary benefit legislation. In *Proceedings of the 1st International Conference on AI and Law*, pages 190–198, 1987.
16. T. J. M. Bench-Capon and G. Sartor. A model of legal reasoning with cases incorporating theories and values. *Artif. Intell.*, 150(1-2):97–143, 2003.
17. T. J. M. Bench-Capon and P. R. S. Visser. Ontologies in legal information systems: The need for explicit specifications of domain conceptualisations. In *Proceedings of the 6th International Conference on AI and Law*, pages 132–141, 1997.
18. T. BenchCapon, F. Coenen, and P. Orton. Argument based explanation of the british nationality act as a logic program. *Information and Communications Technology Law*, 2(1):53–66, 1993.
19. D. H. Berman and C. D. Hafner. Representing teleological structure in case-based legal reasoning: The missing link. In *Proceedings of the 4th International Conference on AI and Law*, pages 50–59, 1993.
20. C. Cayrol and M.-C. Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In Godo [24], pages 378–389.
21. A. Chorley and T. J. M. Bench-Capon. An empirical investigation of reasoning with legal cases through theory construction and application. *Artif. Intell. Law*, 13(3-4):323–371, 2005.
22. P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning and logic programming. In *IJCAI*, 1993.
23. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.*, 77(2):321–358, 1995.
24. L. Godo, editor. *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6-8, 2005, Proceedings*, volume 3571 of *Lecture Notes in Computer Science*. Springer, 2005.
25. T. F. Gordon. The pleadings game: Formalizing procedural justice. In *Proceedings of the 4th International Conference on AI and Law*, pages 10–19, 1993.
26. M. Grabmair and K. Ashley. Facilitating case comparison using value judgments and intermediate legal concepts. In *Proc. of the 13th ICAIL*, pages 161–170, 2011.
27. M. Grabmair and K. D. Ashley. Argumentation with value judgments - an example of hypothetical reasoning. In R. Winkels, editor, *JURIX*, volume 223 of *Frontiers in Artificial Intelligence and Applications*, pages 67–76. IOS Press, 2010.
28. J. Hage. Monological reason-based logic: A low level integration of rule-based reasoning and case-based reasoning. In *Proceedings of the 4th International Conference on AI and Law*, pages 30–39, 1993.
29. A. C. Kakas, P. Mancarella, and P. M. Dung. The acceptability semantics for logic programs. In *Proceedings of the Eleventh International Conference on Logic Programming*, pages 504–519, 1994.
30. A. C. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In *AAMAS*, pages 883–890. ACM, 2003.

31. A. C. Kakas and F. Toni. Computing argumentation in logic programming. *J. Log. Comput.*, 9(4):515–562, 1999.
32. R. A. Kowalski and F. Toni. Abstract argumentation. *Artif. Intell. Law*, 4(3-4):275–296, 1996.
33. M. J. G. Lucero, C. I. Chesñevar, and G. R. Simari. On the accrual of arguments in defeasible logic programming. In C. Boutilier, editor, *IJCAI*, pages 804–809, 2009.
34. A. Maslow. *Motivation and Personality*. Harper and Row, New York, 1954.
35. S. Modgil. Reasoning about preferences in argumentation frameworks. *Artif. Intell.*, 173(9-10):901–934, 2009.
36. S. Modgil and T. J. M. Bench-Capon. Integrating dialectical and accrual modes of argumentation. In P. Baroni, F. Cerutti, M. Giacomin, and G. R. Simari, editors, *COMMA*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 335–346. IOS Press, 2010.
37. S. Modgil and T. J. M. Bench-Capon. Metalevel argumentation. *J. Log. Comput.*, 21(6):959–1003, 2011.
38. C. Perelman and L. Olbrechts-Tyteca. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press, June 1969.
39. H. Prakken. A logical framework for modelling legal argument. In *Proceedings of the 4th International Conference on AI and Law*, pages 1–9, 1993.
40. H. Prakken. An exercise in formalising teleological case-based reasoning. *Artif. Intell. Law*, 10(1-3):113–133, 2002.
41. H. Prakken. A study of accrual of arguments, with applications to evidential reasoning. In *Proceedings of the Tenth International Conference on AI and Law*, pages 85–94, 2005.
42. H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2011.
43. H. Prakken and G. Sartor. On the relation between legal language and legal argument: Assumptions, applicability and dynamic priorities. In *Proceedings of the 5th International Conference on AI and Law*, pages 1–10, 1995.
44. E. L. Rissland. Dimension-based analysis of hypotheticals from supreme court oral argument. In *Proceedings of the Second International Conference on AI and Law*, pages 111–120, 1989.
45. G. Sartor. Doing justice to rights and values: teleological reasoning and proportionality. *Artif. Intell. Law*, 18(2):175–215, 2010.
46. M. Sergot. A query-the-user facility for logic programming. In *New horizons in educational computing*, pages 145–163, New York, NY, USA, 1984. Halsted Press.
47. M. J. Sergot, A. S. Kamble, and K. K. Bajaj. Indian central civil service pension rules: A case study in logic programming applied to regulations. In *Proceedings of the 3rd International Conference on AI and Law*, pages 118–127, 1991.
48. M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The british nationality act as a logic program. *Commun. ACM*, 29(5):370–386, 1986.
49. S. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, England, 1958.
50. B. Verheij. Accrual of arguments in defeasible argumentation. In *Proceedings of the Second Dutch/German Workshop on Nonmonotonic Reasoning*, pages 217–24, 1995.
51. F. Waismann. *The Principles of Linguistic Philosophy*. St Martins Press, NY, 1965.