Argument & Computation 11 (2020) 221–238
DOI 10.3233/AAC-190477
IOS Press

Before and after Dung: Argumentation in AI and Law

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

Abstract. Dung's abstract argumentation frameworks have had a very significant role in the rise in interest in argumentation throughout this century. In this paper we will explore the impact of this seminal idea on a specific application domain, AI and Law. Argumentation is central to legal reasoning and there had been a considerable amount of work on it in AI and Law before Dung's paper. It had, however, been rather fragmented. We argue that the abstract argumentation frameworks had a unifying effect by offering a means of relating previously diverse work. We also discuss how the particular demands of legal systems have led to developments building on the basic notions of abstract argumentation.

Keywords: Abstract argumentation, legal argumentation, argumentation frameworks

1. Introduction

From its very beginnings Artificial Intelligence and Law has been interested in argumentation. One of the first major AI and Law projects was McCarty's TAXMAN [74], which attempted to reproduce the argumentation in the majority and minority opinions in *Eisner v Macomber*, a well known tax law case. By 1987, when the first international conference on AI and Law was held, argumentation was clearly seen as an important part of AI and Law research, most notably through Rissland and Ashley's HYPO [102]. Mostly the argumentation concerned the use of precedent cases, but the need for argumentation was seen for rule based systems also [32]. There are several reasons why argumentation should be seen as of central importance in AI and Law:

- Law (at least in US and UK) is *adversarial*. The counsel for the parties to the dispute put forward *arguments* for their clients and the judges choose which to accept.
- Legal decisions must be *justified*. In law, the answer is not enough: the reasons for the answer must be given in order to guide future decisions, to ensure consistency of decisions, and to attempt to persuade the losing side of why they lost, perhaps leading to acceptance of the decision. If acceptance is not achieved, challenges to the argument of the decision will form the basis of an appeal to a higher court.
- Law is *defeasible*. Legal rules can be overturned by finding that an exception applies, or by finding a conflicting law, or by distinguishing the case so that the rule does not apply. Moreover, even once a decision has been made, it can be overturned on appeal, so that arguments which prevailed at one level of court may fail at a higher level. Sometimes also, as social values change, law has to change to adapt to the new values, and cases are decided differently.

- Law is *open textured*. According to Frege [47], logic requires that its concepts have sharp boundaries. Law, however, is not like that; many concepts are vague in that a new instance may or not fall under them, and require resolution in the light of arguments as to why and why not a given instance should be considered to do so [126]. Open texture was introduced by Waismann in [125] and adapted for use in law by Hart [66].

For these reasons AI and Law has always needed to concern itself with argumentation and there was a good deal of work before [42] appeared in 1995. AI and Law therefore offers an excellent domain in which to explore the impact of Dung's seminal paper [42] by examining the differences his ideas made.

2. Argumentation in AI and Law before Abstract Argumentation

In [13], the journal version of a 1995 Jurix keynote speech [12], we can find a survey of argumentation in AI and Law before Dung. It made a distinction between arguments based on a representation of cases and arguments based on rules. The former concentrated on the *generation* of arguments and counter arguments [7,75] and [103]. Argument based on rules was discussed under three headings: presentation and explanation [27], handling normative conflict and non-monotonicity [84] and modelling law as a process [52]. Dung's frameworks are principally concerned with *evaluation* of the status of sets of arguments, and so relate directly only to the second of the rule based topics. Argumentation frameworks were indeed mentioned in [13] under this topic, but Dung was only indirectly referenced: his IJCAI paper [41] on which [42] was based was cited in [85] which was in turn cited in [13].

2.1. Legal case based reasoning in the early 1990s

Thinking about reasoning with legal cases was dominated by HYPO [102] and [7], and subsequent work deriving from it, including CABARET [112], CATO [4] and BankXX [103]. For an overview of HYPO and its successors see [20]. In these systems cases could be seen as collection of factors.[1] The similarity between cases could then be assessed on the basis of these factors. Argumentation then was conducted using a three-ply argument structure.

- In the first ply the proponent cites the precedent with the desired outcome which most closely matches the current case, arguing that, on the basis of the similarities, the decision should be the same.
- In the second ply the opponent can respond in two ways:
  - ∗ By citing the precedent with a different outcome which most closely matches the current case as a *counter example*.
  - ∗ By citing features in the current case but not the precedent, or in the precedent but not the current case, to *distinguish* the current case from the precedent.
- In the third ply the proponent attempts to rebut the opposing arguments by distinguishing the counter examples and down playing the differences, arguing that they are not significant.

[1] A factor may be taken as a stereotypical pattern of facts with legal significance. As such they are Boolean: either present in, or absent from, a case. As discussed in [20], several of the systems use dimensions, which have magnitudes and range from an extreme pro-plaintiff point to an extreme pro-defendant point, rather than factors. For simplicity we will just use *factors* (which can be regarded as particular points or ranges on a dimension [101]).

This structure is a very natural way to organise legal arguments of the sort found in the high level US court decisions used by HYPO since it closely follows the structure of Oral Hearings in the US Supreme Court [1]. In [4] and [112] some additional structure was provided by the identification of the argument moves that the participants use to present their arguments within these three plies.

A key point about this approach is that there was no formal, theoretical, underpinning. Even though the model is computational (the algorithms involved in HYPO were given in Appendix A of [4]), the analysis relates to the informal argument tradition [46] which attempts to provide a description of legal argumentation from consideration of examples of actual practice. HYPO drew especially on the oral hearing stage of the US Supreme Court. Typical of this style of analysis we see the dialogue divided into a sequence of *stages* (the three plies of the HYPO model), and moves to present arguments characteristic of each of these stages.

CATO [4] has eight argument moves. Essentially in the first stage a case is *cited*, with the suggestion that its decision be followed in the current case. In the second stage this is challenged either by *providing* a counter example, at least as good as the cited case, or by *distinguishing* the cited case, and if possible, *emphasising* the distinction. In the third, rebuttal, stage, these challenges are met by *distinguishing* the counter examples, *emphasising* the distinction if appropriate; *downplaying* the distinctions made in the second stage and *citing* additional cases to emphasise strengths and downplay weaknesses.

A second feature of this kind of argumentation is that there was no attempt made to evaluate arguments: all the emphasis was on the generation of arguments, and the onus was on the user to decide which ones were strong and convincing. This stress on generation of arguments rather than their evaluation continued in CABARET [112] and BankXX [103] which generated its arguments through heuristic search of a collection of *argument pieces*. Only much later [40] were these techniques harnessed for the prediction of case outcomes.

2.2. Argument for presentation and explanation

Initially the output of rule based models of law relied on the standard *how* and *why* explanation, as pioneered in the MYCIN project [110]. However, dissatisfaction with the rather stilted explanations which were produced led to a search for some more natural form of explanation, and there were several suggestions that the form of explanations could be improved by presenting them as arguments. Typically this involved the use of Toulmin's well known argumentation scheme [114] to organise the explanation into a form more congenial to lawyers. This was independently proposed by Lutomski [71], Marshall [73] and Storrs [113]. Later a general compuational means of presenting the output of logic programs as Toulmin schemes was given in [27] and applied to law in [33]. The Toulmin scheme was also used to present the output from a neural network [134]. Like the case based systems described in the previous sections, these systems do not use argumentation to determine the outcomes of cases. The outcome is decided using other means, such as a logic program or a neural network, and then mapped into an argument scheme for the purposes of presentation.

2.3. Argument for modelling legal processes: Dialogue games

Dialogue Games were introduced in philosophy, mainly with the intention of modelling fallacies [65] and [72]. These were taken up in AI as a way of interacting with knowledge bases [25] and [80]. Dialogue games were popularised in AI and Law in by [64], [51] and [63]. Again, the point was not to predict the outcome of cases. In [63] the purpose was to model legal dialogues so that hard cases could be specified

in dialogical terms. In [51] arguments were evaluated to identify which issues were in dispute so as to model the process of pleading.[2] This idea rapidly became popular in AI and Law and resulted in the production of several dialogue games, based on various logics and modelling various legal processes (e.g. [70] and [26]). What is common to all these systems is a recognition of the importance of *how* an argument is conducted, and the desirability of making the rules governing such disputes (for example, *Robert's Rules of Order* [104] formalised in [87]) explicit so that participants could be aware of their rights and obligations in the dialogue, and to facilitate the interpretation of the utterances. Since the relevant dialogues were typically arguments between legal adversaries, they were effectively models of properly conducted legal argument. Some models were based on Toulmin's argumentation scheme (e.g. [24] and [14]).

2.4. *Non-monotonicity and normative conflict*

Legal reasoning is inherently defeasible. In the court of first instance the status of arguments may change in the light of new facts or new arguments. Even when the facts are agreed, a new interpretation of a rule may reverse a decision. To account for this legal systems typically offer a right of appeal, and often the right to appeal again against the appeal decision. Thus legal reasoning may be seen as non-monotonic. Moreover, we often find norms in conflict, often, but not always, because a legal norm is promulgated as a general rule and a series of exceptions. Further, since legal norms always require interpretation to apply them to specific fact situations, we may find conflicting interpretations, as in the case of open-textured concepts [66]. It is clear therefore that legal systems need to be able to handle non-monotonicity and conflicts between norms and between interpretations.

This could be accommodated by using a programming language that supported non-monotonicity, such as Prolog with negation as failure [68] or some more specialised language [49].[3] While these were able to cope with non-monotonicity arising from additional information, they were less able to handle conflicts which require a decision as to which rule to follow. Law does offer some principles to determine such choices: *lex specialis* (prefer the more specific law), *lex superior* (prefer the law issued by the higher authority) and *lex posterior* (prefer the more recent law). These principles formed part of Valente's ontology [115], and *lex specialis* in particular was advocated in [84].

The notion of argumentation seemed to some to provide excellent potential as a means of handling these problems. The idea is that the conflicting arguments can be generated and the user can be called upon to decide between them (effectively playing the role of judge). Presentation as arguments is preferable to presentation as rules, because it provides the surrounding context, support for the rule antecedents and the like. This was essentially the purpose of McCarthy's TAXMAN project [74], which was intended to generate the competing arguments in a particular case. Presentation of competing arguments was suggested as a general solution to problems of open texture in [32]. Although that paper only generated arguments, the desirability of computationally supporting the choice between them was acknowledged, suggesting that attention should be paid to producing "a representation in computer intelligible terms of what it is that makes an argument persuasive". The responsibility for deciding which argument to accept, however, remained for the present squarely on the user.

[2] "The process performed by the parties to a suit or action, in alternately presenting written statements of their contention, each responsive to what precedes, and each serving to narrow the field of controversy, until there evolves a single point, affirmed on one side and denied on the other, called the "issue", upon which they then go to trial" [36].

[3] Although Prolog was able to handle the exceptions in [109], there are aspects of non-monotonicity that it cannot deal with [50].

This role for argumentation can also be found in the hybrid rule and case system of CABARET [112]. In that paper, excellently summarised by Loui in Section 3.4 of [21], there is a top layer of rules, but the interpretive work required to determine whether or not the antecedents apply in a particular case is done by HYPO style case based reasoning. Thus both case-based and rule-based traditions saw argumentation as the key to resolving the problem of open texture.

2.5. Summary

Argumentation in AI and Law had a number of features before Dung.

- Its main inspiration came from informal logic traditions: notions of argumentation schemes, context, stages of dialogues and argument moves all come from that tradition. This imposes attention on the particular content and context of arguments. Arguments were rarely seen as convincing on the grounds of their form alone.
- The main role for technology was the generation of arguments, whether using case based reasoning, logic programs, non-monotonic logics, neural networks or through instantiating argumentation schemes. The evaluation of the status of arguments was normally left to the user.[4]
- In rule based approaches, argumentation was usually[5] assigned a specific role: explanation, presentation, resolving open texture or modelling a legal procedure.

The result was that argumentation in AI and Law taken as a whole was rather lacking in coherence. Arguments were generated in a variety of ways including: using instantiations of arguments schemes, ranging from generic schemes, such as that of Toulmin, to the highly specialised schemes such as those found in HYPO and its progeny; using the proof traces of logic programs; using other rule instantiations. This meant that arguments tended to take a rather diverse form, and that attacks on them could be highly application specific. The use of arguments, whether as the basis of dialogues or for explanation, tended to be very project specific, and approaches were often difficult to compare. What was lacking was a means of relating this diverse work within a common framework. This was achieved with the advent of Abstract Argumentation.

3. Argumentation in AI and Law AD (after Dung)

The idea that abstract argumentation could offer a way to relate the various strands of AI and Law on argumentation was first proposed in an important paper by Henry Prakken [85]. This paper draws its notion of abstract argumentation not from [42] but from the pre-journal version presented at IJCAI [41]. There had been argumentation frameworks before Dung: Prakken cites several examples in [85], including Simari and Loui [111] from general AI, and, from AI and Law, Sartor [106], Gordon [52] and his own [84]. But all of these are tied to a particular logic: in AI and Law, Gordon uses conditional entailment [48] while Sartor and Prakken rely on a version of default logic. Dung's insight was that arguments could be *abstract* and this freed them from any particular method of generation, whether using a particular logic, particular argument schemes, or case based methods. Once in the framework

[4] Casting technology in the role of supporting rather than making legal decisions was in part a consequence of resistance to the idea of a "computer judge". Some work on evaluation did go on, however. Arguments were evaluated in [52], although the purpose there was to identify disputed issues rather than make decisions on cases. There was also interest in using the principles described in Section 2.4. For example Prakken advocated preferring the more specific argument [83].

[5] Again there were exceptions, notably the work of Tom Gordon [52] and Henry Prakken [84].

all arguments were equal. This separated consideration of the status of arguments from the logic that produced them: much of Prakken's [85] is taken up with explaining why he was abandoning his previous support for non-monotonic logics as the basis of legal reasoning. Since the arguments were abstract, how they were generated became immaterial. Thus disputes as to the correct logic such as [50] and between rule and case based reasoning, such as the debate at the 1991 ICAIL panel [34], could be contained, and various approaches related at the level of *abstract* argumentation.

3.1. Levels in argumentation

The main point of [85]. is to distinguish three levels required for adversarial legal reasoning.

(1) The logic level, which generates arguments.
(2) The argument level, which organises these arguments and identifies attack relations between them. This determines the acceptability of arguments at a given point in the debate.
(3) The dialogical level, which determines how the arguments and attacks from the level below can be deployed in a dispute. This level moves the debate forward and refers to the level below to determine the acceptability of arguments at the current stage.

The use of *logic* at the first level is partly indicative of the author's formal bent, and partly because he is keen to stress that any logic can be used here: whether it is standard logic or some non-monotonic logic is of no importance: in fact in [85] he settles on a "monotonic underlying logic (possibly with a metalinguistic reason conditional)". The role of this stage is simply to identify arguments which may be considered sound in isolation. But any narrow interpretation of "logic" would be unnecessarily restrictive. As we have seen above arguments can also be generated through other means. Among these other means the instantiation of argumentation schemes[6] would become increasingly popular. There is no reason why such means of generating arguments should not also be used at this level. An example of aa approach which uses no logic at all, but generates all its arguments through the instantiation of an argument scheme and its critical questions is [23]. Moreover there is no need to choose one particular means of generating arguments: at this level we wish to find all the arguments that can be produced and any conflicts between the generated arguments can be sorted out at the second level.

The second level is the argumentation framework itself, and it here that Dung's framework makes its impact. Whereas in previous work the argument framework had been dependent of the way in which the arguments were generated, Dung's proposal was to abstract away all these generation-specific aspects. Thus, using Dung's framework at this level, enabled arguments to be abstracted from the means used to produce them and from their particular forms. Dung's framework does, however require the identification of the attack relations between the arguments (in [42] an argument is "abstract entity whose role is solely determined by its relations to other arguments" and the attack relation is the only such relation in [42]). This means that nature of attack is abstract, and no distinction is made between different kinds of attack. In [85] only two types of attack, rebuttal and undercut, were recognised. This is in keeping with the focus on logic, and excluded underminers (attacks on arguments represented by attacks on their premises) because they would be attacks on sub-arguments. But there is no need to restrict the nature

[6] In the early 90s, Toulmin's scheme was almost the only explicit scheme used. The explicit use of a variety of argumentation schemes to represent legal rules was proposed in [117], drawing on [127]. After that the explicit use of schemes became widespread: e.g. [29,53,56,95,99,120,131]. But the use of argument schemes was also implicit in earlier work: [89] argues "that much AI and Law research in fact employs the argument-scheme approach, although it usually is not presented as such" and provides an excellent retrospective discussion of this earlier work.

of attacks. If we are using the instantiation of argumentation schemes to generate arguments then it is appropriate to include attacks deriving from so called *critical questions* [129] in the attack relation for such arguments. This could lead to a large number of different types of attack: for example [61] identifies fifteen different kinds of attack, some with several variants. This is quite normal in approaches such as [23] which generate arguments from argumentation schemes and their critical questions and organise them into an argumentation framework for evaluation. The variety of attacks from critical questions characteristic of argumentation schemes can, however, be related to the standard three types of attack of rebuttal, undercut and undermining since critical questions can be seen as having one of three different roles: some give rise to rebuttals, others give rise to premise attacks and yet others to undercutters [120].

The third level, the level of dialogue, is able to draw on the abstract argumentation framework for its moves. Although in many cases the dialogue operates on a fixed argumentation framework, it is possible to return to the lower levels to seek additional arguments if they are needed. Because the arguments are abstract, different classes of dialogue such as persuasion, deliberation and inquiry [128] can all be implemented on the same structure. Using argumentation frameworks gave rise to a new type of dialogue game (e.g. [124]). When the dialogues are based on abstract frameworks, the content of the move is changed from that found in most earlier dialogue games[7] and the focus is on the relationship between arguments. In most earlier games (e.g. [24,70]) the content of the various moves had been statements, but now arguments could be presented in a single move. So whereas previously the arguments had emerged over a series of moves (e.g. claim, challenge, reason, as in games like those described in [90]), now they could themselves be the subject of the moves. This, its proponents argued, improved the naturalness of the interaction, since providing the argument as a whole meant that it did not need to be teased out over a series of moves, reducing the number of turns in the dialogue.[8] In games in which the argumentation framework is complete (e.g. [28] and [132]), the effect of various moves could be better anticipated, offering more scope for a strategic choice of attacks on them. In other uses, the framework is dynamically extensible, with it being possible to seek further arguments as the need arose, with the search focussed on arguments that would help in the dialogue. Examples of dialogue games based on argumentation frameworks and with arguments as locutions can be found in [44,124] and, applied to law, [28]. Importantly also the argumentation framework layer separated the dialogue from the underlying logic used to generate the arguments. Previously systems had been forced to commit to a particular logic (e.g conditional entailment [48] in the Pleadings Game [52], Prolog in PLAID [33], and DEFLOG [119] in ARGUMED [118]). As mentioned earlier, it is this abstraction provided by Dung's framework that enabled a better understanding of the relationships between these various systems. The separation by a mediating abstract framework also meant that the same game could be played with the arguments generated using different logics, or combinations of them. This separation can, however, pose problems for dynamic systems which may require awareness of the argument structure to request additional arguments and to explore other dialogical aspects.

Thus while some dialogue games such as [124] and [28] operated at the argument level, and were based on complete argumentation frameworks derived from the underlying knowledge base, issues relating to dialogical aspects, such as the structure of arguments and the way in which information is introduced into the discussion, required that arguments retain their structure. Therefore several researchers continued to use the more traditional style of games to explore these important issues. These include, amongst others, Prakken [91,98], Lodder [69] and Verheij [118].

[7] There were some exceptions, including [52].

[8] A similar effect is achieved, by different means, in [52].

3.2. Conflicting arguments

A detailed examination of the three level approach is given in [96], and the approach developed in that paper was used to model reasoning with legal precedents in [98]. The second of these papers is especially significant since it enables reasoning with precedents to be rendered in terms of arguments. The key idea, reflecting the adversarial nature of legal dispute, and the fundamental principle of precedential reasoning, *stare decisis*, is that the effect of a precedent can be rendered in terms of three arguments:

- The strongest argument in favour of the plaintiff.
- The strongest argument in favour of the defendant.
- An argument that the precedent shows that one of the preceding arguments is preferred, depending on the outcome of the precedent case.

If we represent cases as sets comprising the factors in a case in the manner of CATO [4], the three arguments come from three rules (if Plaintiff factors then find for Plaintiff, if Defendant factors then find for Defendant and a priority rule between these two rules). In this way a body of case law can be represented as a knowledge base of rules or a logical theory. This representation has proved fruitful in developing a logic of precedential reasoning in [67] and [100].

While the structure of arguments was important in [98], a body of case law was represented as a framework of purely abstract arguments in [16].[9] That paper considered seven landmark cases concerning property law relating to wild animals. These cases often form part of the introduction to property law for US law students, and, since their introduction to AI and Law in [35], have proved a popular test bed for ideas about reasoning with cases in AI and Law (e.g. [31] and [8]). The paper identified 26 arguments used in one or more of the cases and organised them into an abstract argumentation framework, as shown in Fig. 1.

Depending on which arguments can be made in a given case, argument A (affirming that the pursuer had a right to the animal) will be in or out of the chosen extension, and the case is decided accordingly. Note, however, the framework contains cycles and so there may be multiple preferred extensions, some with A and some without. The two-cycle M-O concerns whether or not Justinian provides an authority that should be followed, and is capable of different resolutions in different jurisdictions or at different times. In the actual series of cases, the very first, *Pierson v Post*, decided that Justinian should be followed, although the contrary was argued in the minority opinion. The other important cycle, T-S-E-B concerns an allegation of unfair competition (argument T) which arose in *Young v Hitchens*. This presents a dilemma: if we accept that the competition was unfair (T) we will also accept E (that Young had done enough to establish possession) and Young will win. Alternatively we can accept the other two arguments in the cycle and Hitchens will win. Unlike the two-cycle, this required more than a simple preference between arguments in the cycle. In the framework there is an argument U which attacks one of the arguments in the cycle (T) and so, if accepted, will break it. In *Young v Hitchens* the court broke the cycle by accepting U and holding that it was not within their remit to rule on whether competition was unfair and so Hitchens won. In another case, *Ghen v Rich*, concerning whaling, where there were well established conventions governing what was unfair competition and what constituted possession of the whale, an argument that the conventions should be respected V defeats U, which restores the cycle and prevents the defendant from winning by using Hitchen's arguments. This case was resolved because

[9] This is, of course, not an example of the three level model of [85] (unless one counts the analyst as constituting the logic level).

V is also able to break the 4-cycle since it attacks an argument in the cycle, namely B, and so the plaintiff won. This representation of a reasonably substantial body of case law used in teaching property law showed the usefulness of abstract argumentation frameworks for analysing a body of case law so as to identify the issues that needed to be resolved, which manifest themselves in the form of cycles. This in turn made it important to provide the means of providing a resolution of these issues.

4. Using preferences

In Dung's abstract argumentation framework of [42] all attacks always succeed. This gives rise to several possible semantics to determine acceptance. Grounded, preferred and stable were mentioned in [42], and since then there have been many other proposals [11]. Among the properties of the original semantics are:

- The grounded extension is unique, but may be the empty set.
- If there are no odd length cycles, there will always be at least one non-empty preferred extension, but there may be many preferred extensions. Multiple preferred extensions arise out of even length cycles.
- There may be no stable extension.

If we consider a legal dispute, we will have arguments presented by two opposing parties. These opposing arguments may manifest themselves as cycles, giving rise to competing positions. In Fig. 1, for example we see two even length cycles: M and O represents a difference of opinion as to whether Justian is too old be an acceptable authority, and B-T-S-E represents a dilemma turning on whether deciding what constitutes unfair competition is within the remit of the Court. In courts where the facts

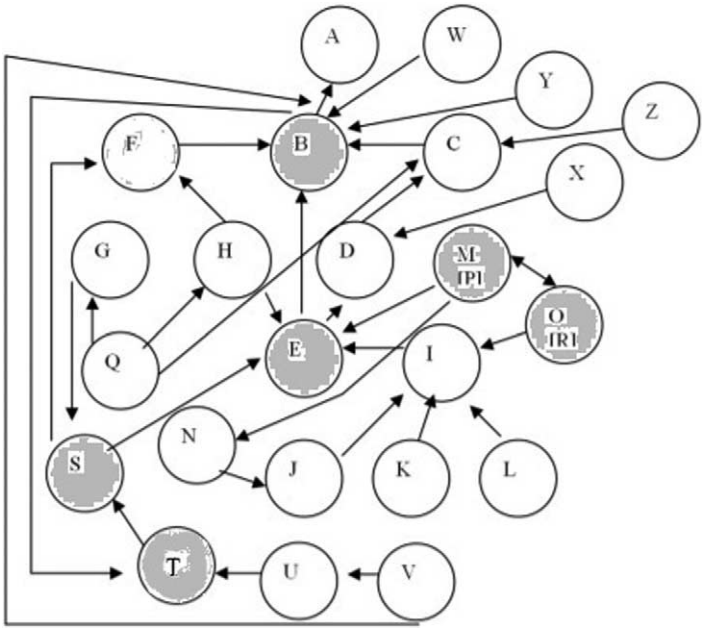

Fig. 1. Abstract Argumentation for the wild animals domain as represented in [16]. Shaded nodes indicate cycles. P and R are arguments to accept M and O respectively.

are at issue, 2-cycles naturally arise from conflicting testimony and other evidence. Even length cycles are almost inescapable in an AF representation of a legal dispute. Odd length cycles almost never occur since odd length cycles represent paradoxes and so suggest a problem with the representation [19].

If the dispute is live, either side *may* win. Moreover one of the sides *must* win: the judge must choose one of the sides. For this reason, unless we wish to extend the framework with an explicit argument breaking each cycle, the notion that attacks always succeed must be modified. In the simplest case we have two mutually attacking arguments. Once the judge has ruled on the case one of the attacks will have succeeded and the other have failed. If using preferred semantics this enables us to select a preferred extension, while in grounded semantics it allows us to decide that the prosecution case is sufficient to discharge its burden of proof and overcome the presumption of innocence.

In this sense legal reasoning has more in common with practical reasoning than theoretical reasoning ([61] and [9]). Consider Searle's notion of *direction of fit* [108]. Whereas in theoretical reasoning people are obliged to fit their beliefs to the world, in practical reasoning they attempt to fit future states of the world to their desires. In legal reasoning, at least in cases where the answer is not already constrained by precedent, there may be no pre-existing right answer: the judge chooses what *should* be the answer, for the current case, and for future cases. Where the decision is constrained by precedents, this will reflect the past choices of judges. This applies to conflicts on points of law: conflicts of fact must also be resolved, although this is often left to a lay jury.

For this reason it is important that there is a means of presenting arguments for and against accepting one argument against another: essentially these are the arguments that will be presented in the judge's decision. This means that some attacks will fail. The need to enable a distinction between successful and unsuccessful attacks was recognised by several researchers, both before and after Dung: for example by Pollock [82], Vreeswijk, [123] and [122], Simari and Loui [111] and Prakken and Sartor [97]. A way of distinguishing between successful and unsuccessful attacks as a direct extension of Dung frameworks was given in [5]. This approach augmented abstract argumentation frameworks by the addition of a preference relation between arguments. Now an attack fails if the attacked argument is preferred to its attacker. The effect is essentially the same as [98] which derived its preference relation from precedents. There are some problems with this approach, however, in that the preferences are simply given, not explained. Although preferences derive from many sources, such as social and moral values, and these motivated the precedent decisions, in [98] there is no scope to give reasons for the preferences other than the precedents themselves. Moreover the preferences are between arguments at a very coarse granularity: each case gives rise to only two arguments each amalgamating *all* the reasons to decide for a particular party. More importantly, for [98], preferences only come from precedents and so only *a fortiori* reasoning is possible. If only a subset of the factors from the precedent are present in a case, no preference concerning the corresponding argument is available. There are, therefore, limitations in simply considering a preference relation between arguments which is given and unchallengeable.

To address this, an alternative extension to Dung's framework was provided by value based reasoning [17], with the intention of avoiding the reduction of reasoning with precedents to *a fortiori* reasoning. This was inspired by work in AI and Law, namely [35]. That paper had confronted the issue of how to choose between factor based arguments, and argued that the choice should be governed by the social purposes of the law concerned, and which choice would better serve these purposes. This work was revived by a special issue of *Artificial Intelligence and Law* on the work of Don Berman, in which three papers, [15], [107] and [88], discussed this issue, and explicitly made reference to social values. It could be then related to the work of Perelman [81] by equating the notion of an audience with an ordering on these values. A detailed discussion of value based reasoning in AI and Law was given in [62]. These ideas

were provided with a formal treatment building on abstract argumentation frameworks in [17], which spoke only of *values* rather than *purposes*. In a value based framework, arguments have an additional property. As well as the arguments they attack, they now also have a set of values they promote. Now the preference between arguments can be expressed as a preference between the values promoted by accepting them, with the value preferences motivated by precedents which had exhibited them [31], or other considerations [28]. The advantages of this were that it gave reasons for the preferences in terms of the values, that it enabled arguments to be considered at the granularity of a single factor, and that it could go beyond what was already present in the preferences, by enabling preferences between as yet undecided combinations of factors to be used, provided the combinations of values were found in a precedent. The approach of using value based precedential reasoning is fully described in [31] and [23]. Another advantage of this approach is that the preferred arguments are relative to an ordering of values, and so this can be used to account for the existence of majority and minority opinions in the Supreme Court and explain differences across jurisdictions, or across times within a jurisdiction, in terms of different, or changing, value orders.

There are, however, as shown in [78], some strong limitations on the approaches of [5] and [17], namely that they only work correctly if all attacks are preference-dependent (which is why most work on value based argumentation follows [10] and resolves factual disputes before forming the value based argumentation framework) and if they are all direct (so that can be no subargument attacks). Moreover, while value preferences gave reasons for preferring arguments, no reasons for preferring values (other than the audience) could be given in the framework of [17], and so the problem had been merely moved up a level. Moreover as well as precedents and values, there may be other reasons for preferring a precedent: thus a general way of arguing for a preference is needed.

Arguments about preferences had been recognised as necessary in AI and Law, and addressed in [97], [86] and [54]. The problem of relating arguments about preferences to Preference and Value Based Argumentation Frameworks was solved by the notion of Extended Argumentation Frameworks [76], which allowed arguments to attack attacks as well as other arguments, and so enabled reasoning about preferences. These ideas were applied to reasoning with legal cases in [28], which applied the approach to the abstract framework of [16]. In fact, as was shown in [77], Extended Argumentation Frameworks can be rewritten as standard abstract argumentation frameworks, using arguments with statements about arguments such as "A is accepted", "A defeats B" and the like as their conclusions, enabling the work associated with Dung's framework to be applied to extended argumentation frameworks also. In EAFs, preferences are explicitly expressed as the attackers of attacks: for an example see [28].

The above discussion refers to attempts to handle preferences as an extension to Dung's framework. There are also, approaches in which, like the pre-Dung work cited above, attacks are defined in terms of a more basic notion of conflict plus the use of preferences. The influential ASPIC+ [92] and [79] represents a contemporary example of that approach.

5. Argument structure

Another line of investigation in AI and Law is the structure of arguments. Although abstraction for use in a Dung style or Value Based Argumentation Framework is very useful, especially for evaluating competing positions, law does need to deal with particular concrete arguments. This can be very important: knowing the type of attack can be crucial since while preferences can be used to defend against rebuttals and underminers, they provide no defence against undercutters [92].

While there have been proposals to include a support relation as well as an attack relation [6] in abstract argumentation frameworks, the arguments in such a framework do not capture sufficient structure for successful modelling of legal argument. Another suggestion to capture the structure of arguments in an abstract argumentation framework was made in [133], but this had limitations as shown in [94]. Instead the approach taken has been to model arguments with structure (at least premises and conclusions). Although this represents a step back from abstraction, one popular approach, ASPIC+ [79], is intermediate in its level of abstraction between concrete logics and the fully abstract level and accommodates a broad range of instantiating logics. ASPIC+ is intended to generate abstract argumentation frameworks to enable its arguments to be evaluated. Another approach to capturing structure, Carneades [55], has its own method for evaluating arguments and does not typically use abstract argumentation frameworks. The relation of Carneades to abstract frameworks (and to ASPIC+) is shown in [116].

These frameworks have been used for some quite concrete modelling. Different argumentation based approaches to representing a particular case was the topic of a special issue of *AI and Law* journal [8,18,57] and [93]. Elsewhere, a number of argumentation schemes have been proposed for modelling reasoning with precedents, both with [30] and without [99] values. For a concrete use of Carneades to model statutory interpretation see [130]. For another distinctive approach, particularly directed towards evidential reasoning, see [121].

6. Concluding remarks: Dung and AI and Law

The impact of Dung on AI in general is unarguable. Before [42] argumentation was very much a minority taste. It rarely appeared as a keyword in AI conferences and there were no specialist conferences or workshops. Now, however, it is an established subfield of AI: argumentation sessions are a feature of AI conferences, there is a specialised conference (COMMA) and several regular workshops, such as Computational Models of Natural Argument (CMNA) and Theory and Applications of Formal Argument (TAFA), and a specialist journal in *Argument and Computation*. Without Dung's abstract argumentation it is hard to believe that this would have happened: [42] has itself 2454 citations on Scopus[10] and an h-index of 77, which is a clear indication of its importance. While, however, it is clear that it has had extensive influence, the nature of that influence on *applications* of argumentation is harder to determine. In this paper we have looked at a particular field, AI and Law, which had worked on argumentation for a long period before Dung's paper, to see the differences that abstract argumentation made.

Before abstract argumentation there had been a lot of work directed towards argumentation in AI and Law, as described in Section 2. However, this work was lacking in coherence and unity. Projects were pursued independently and were not easy to relate to one another. Even projects with a common origin such as Aleven and Ashley's CATO and Skalak and Rissland's CABARET went in separate, quite different, directions, as described in [20]. This diversity manifested itself in a variety of approaches, each of which were individual to a particular project, in source material (cases or statutes or commentaries), and in the particular part of the process addressed (generation, conflict identification, explanation, etc). What abstract argumentation offered was a level of abstraction at which it was possible to relate these strands. In addition to the elegance and accessibility of Dung's framework, with an abstract argumentation framework as the target differences between approaches such as the source of the arguments, the means used to generate them, the particular aspect of the task addressed, and the use to which they

[10] On 30th September 2019. The second most cited computer science paper with argumentation in its title, abstract or keywords [17] has only 505 citations.

would be put could be ignored. This greatly reduced the idea that different approaches were in conflict. Different approaches could even contribute to the same framework, and different approaches could be compared at the framework level. The result was that debates as to the "correct" logic to use largely disappeared and the exploration of legal reasoning with rules and cases could be seen as synergistic rather than competitive alternatives.

This influence lay more in the perspective it offered than the explicit use of Dung's frameworks. Of course AI and Law research on argumentation continues to exhibit a great deal of diversity, and by no means everyone uses argumentation frameworks. Defeasible logic has continued to be the basis of the work of Governatori and his colleagues (e.g. [58] and [105]) and Carneades has its own particular approach. Verheij has developed two novel approaches using DEFLOG [119] and case models [121]. Abstract Dialectical Frameworks (ADFS: see [39] and [38]) form the basis of the methodology advocated in [2]. None the less it is possible to view much of this work from the perspective of abstract argumentation: defeasible logic can provide arguments in an abstract argumentation framework, and the relation between Carneades and abstract argumentation frameworks was shown in [116]. Verheij's DEFLOG focusses on statements rather than arguments, but is said in [119] to have "close formal relations" to Dung, although providing a richer language. ADFs can be seen as a generalisation of AFs [39].[11] Also, work on case based reasoning has continued. For example the work of Grabmair ([60] and [59]), provides a sophisticated approach to case based reasoning. This work uses argumentation methods making extensive use of argument schemes and argument graphs, which it says are similar to those of Carneades, to model the value judgements involved in deciding cases. The formal accounts of precedential constraint in [67] and [100] depend crucially on the relation of case and rule based reasoning advanced in [98].

This diversity is to be expected, and welcomed. The unifying perspective of abstract argumentation applies at the middle level of a three level model. Different approaches are still to be expected at the argumentation generation and argument deployment levels, and some may choose to use a less abstract representation at the second level also. Nevertheless the abstract level remains a way of seeing connections between different approaches.

AI and Law has also been able to contribute to abstract argumentation in general. Much of the work on abstract argumentation is at a theoretical level: comparison of different semantics [11] and their complexity [43] for example. Any examples used in such work tend to very simple, even simplified, sometimes even simplistic. Law, in contrast, demands the ability to represent substantial frameworks, and requires them to be reasonably faithful to the cases they represent which militates against simplifications. Thus many of the more substantial representations of domains using argumentation frameworks have been in the domain of law. For this reason, legal instantiations have proved useful to those wishing to provide computational implementations of abstract argumentation, e.g. [45]. Three other developments were driven by the needs of legal applications. First, the need to provide reasons for choosing between positions and determining which arguments should be preferred led to value-based [17], extended [76] and metalevel [77] argumentation frameworks. Second, the need to examine the structure of arguments and to determine the nature of the attacks between them led to structured argumentation frameworks such

[11] Abstract Dialectical Frameworks may in future, like Abstract Argumentation Frameworks, provide a way of bringing together different strands of AI and Law research. The methodology of [2] attempts to encapsulate case based knowledge and is used to represent the knowledge of CATO [4] (and other domains) in an ADF. Subsequently the ADF was extended to represent the US Trade Secrets domain with factors with magnitudes [22]. Recently it has been successfully applied to the domain of compensation for Noise Induced Hearing Loss [3]. ADFs have also been used to reconstruct Carneades [37]. They may provide a fruitful avenue for future exploration in AI and Law.

as ASPIC+ [79] and Carneades [55]. Third, the widespread use of informal logic techniques already a feature of AI and Law, especially argumentation schemes and dialogue games, could now be used in the context of abstract argumentation, with mutual benefit.

Thus Dung's landmark paper [41] played an important role in the development of the study of argumentation in AI and Law, by offering a level of abstraction at which the previously disparate lines of work could be related, for comparison and common enlightenment, even where subsequent work did not explicitly use Dung's frameworks. In return AI and Law offered the opportunity to explore and develop abstract argumentation is a domain where faithful modelling of argumentation and debate is of central importance.

References

[1] L. Al-Abdulkarim, Representation of case law for argumentative reasoning, PhD thesis, University of Liverpool, 2017.

[2] L. Al-Abdulkarim, K. Atkinson and T. Bench-Capon, A methodology for designing systems to reason with legal cases using Abstract Dialectical Frameworks, *Artificial Intelligence and Law* **24**(1) (2016), 1–49. doi:10.1007/s10506-016-9178-1.

[3] L. Al-Abdulkarim, K. Atkinson, T. Bench-Capon, S. Whittle, R. Williams and C. Wolfenden, Noise induced hearing loss: Building an application using the ANGELIC methodology, *Argument and Computation* **10**(1) (2019), 5–22. doi:10.3233/AAC-181005.

[4] V. Aleven, Teaching case-based argumentation through a model and examples, PhD thesis, University of Pittsburgh, 1997.

[5] L. Amgoud and C. Cayrol, On the acceptability of arguments in preference-based argumentation, in: *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., 1998, pp. 1–7.

[6] L. Amgoud, C. Cayrol, M.-C. Lagasquie-Schiex and P. Livet, On bipolarity in argumentation frameworks, *International Journal of Intelligent Systems* **23**(10) (2008), 1062–1093. doi:10.1002/int.20307.

[7] K.D. Ashley, *Modeling Legal Arguments: Reasoning with Cases and Hypotheticals*, MIT Press, 1990.

[8] K. Atkinson, Introduction to special issue on modelling Popov v. Hayashi, *Artificial Intelligence and Law* **20**(1) (2012), 1–14. doi:10.1007/s10506-012-9122-y.

[9] K. Atkinson and T. Bench-Capon, Legal case-based reasoning as practical reasoning, *Artificial Intelligence and Law* **13**(1) (2005), 93–131. doi:10.1007/s10506-006-9003-3.

[10] K. Atkinson and T. Bench-Capon, Practical reasoning as presumptive argumentation using Action based Alternating Transition Systems, *Artificial Intelligence* **171**(10–15) (2007), 855–874. doi:10.1016/j.artint.2007.04.009.

[11] P. Baroni and M. Giacomin, Semantics of abstract argument systems, in: *Argumentation in Artificial Intelligence*, G. Simari and I. Rahwan, eds, Springer, 2009, pp. 25–44. doi:10.1007/978-0-387-98197-0_2.

[12] T. Bench-Capon, Argument in Artificial Intelligence and Law, in: *Proceedings of JURIX 1995*, 1995, pp. 5–14.

[13] T. Bench-Capon, Argument in Artificial Intelligence and Law, *Artificial Intelligence and Law* **5**(4) (1997), 249–261. doi:10.1023/A:1008242417011.

[14] T. Bench-Capon, Specification and implementation of Toulmin Dialogue Game, in: *Proceedings of JURIX 1998*, 1998, pp. 5–20.

[15] T. Bench-Capon, The missing link revisited: The role of teleology in representing legal argument, *Artificial Intelligence and Law* **10**(1–3) (2002), 79–94. doi:10.1023/A:1019501830692.

[16] T. Bench-Capon, Representation of case law as an argumentation framework, in: *Proceedings of JURIX 2002*, 2002, pp. 103–112.

[17] T. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, *Journal of Logic and Computation* **13**(3) (2003), 429–448. doi:10.1093/logcom/13.3.429.

[18] T. Bench-Capon, Representing Popov v Hayashi with dimensions and factors, *Artificial Intelligence and Law* **20**(1) (2012), 15–35. doi:10.1007/s10506-012-9118-7.

[19] T. Bench-Capon, Dilemmas and paradoxes: Cycles in argumentation frameworks, *Journal of Logic and Computation* **26**(4) (2014), 1055–1064. doi:10.1093/logcom/exu011.

[20] T. Bench-Capon, HYPO's legacy: Introduction to the virtual special issue, *Artificial Intelligence and Law* **25**(2) (2017), 1–46. doi:10.1007/s10506-017-9201-1.

[21] T. Bench-Capon, M. Araszkiewicz, K. Ashley, K. Atkinson, F. Bex, F. Borges, D. Bourcier, P. Bourgine, J. Conrad, E. Francesconi et al., A history of AI and Law in 50 papers: 25 years of the International Conference on AI and Law, *Artificial Intelligence and Law* **20**(3) (2012), 215–319. doi:10.1007/s10506-012-9131-x.

[22] T. Bench-Capon and K. Atkinson, Lessons from implementing factors with magnitude, in: *Proceedings of JURIX 2018*, 2018, pp. 11–20.
[23] T. Bench-Capon, K. Atkinson and A. Chorley, Persuasion and value in legal argument, *Journal of Logic and Computation* **15**(6) (2005), 1075–1097. doi:10.1093/logcom/exi058.
[24] T. Bench-Capon, F. Coenen and P. Orton, Argument-based explanation of the British Nationality Act as a logic program, *Information and Communications Technology Law* **2**(1) (1993), 53–66. doi:10.1080/13600834.1993.9965668.
[25] T. Bench-Capon, P. Dunne and P. Leng, Interacting with knowledge-based systems through dialogue games, in: *Proceedings of the Eleventh International Conference on Expert Systems and Applications*, 1991, pp. 123–140.
[26] T. Bench-Capon, T. Geldard and P. Leng, A method for the computational modelling of dialectical argument with dialogue games, *Artificial Intelligence and Law* **8**(2–3) (2000), 233–254. doi:10.1023/A:1008308627745.
[27] T. Bench-Capon, D. Lowes and A. McEnery, Argument-based explanation of logic programs, *Knowledge-Based Systems* **4**(3) (1991), 177–183. doi:10.1016/0950-7051(91)90007-O.
[28] T. Bench-Capon and S. Modgil, Case law in Extended Argumentation Frameworks, in: *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, ACM, 2009, pp. 118–127.
[29] T. Bench-Capon and H. Prakken, Using argument schemes for hypothetical reasoning in law, *Artificial Intelligence and Law* **18**(2) (2010), 153–174. doi:10.1007/s10506-010-9094-8.
[30] T. Bench-Capon, H. Prakken, A. Wyner and K. Atkinson, Argument schemes for reasoning with legal cases using values, in: *Proceedings of the 14th International Conference on Artificial Intelligence and Law*, ACM, 2013, pp. 13–22.
[31] T. Bench-Capon and G. Sartor, A model of legal reasoning with cases incorporating theories and values, *Artificial Intelligence* **150**(1–2) (2003), 97–143. doi:10.1016/S0004-3702(03)00108-5.
[32] T. Bench-Capon and M. Sergot, Towards a rule-based representation of open texture in law, in: *Computer Power and Legal Language*, C. Walter, ed., Quorum Books, New York, 1988, pp. 38–61.
[33] T. Bench-Capon and G. Staniford, PLAID: Proactive legal assistance, in: *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, ACM, 1995, pp. 81–88.
[34] D.H. Berman, Developer's choice in the legal domain: The Sisyphean journey with CBR or down hill with rules, in: *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, ACM, 1991, pp. 307–309.
[35] 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 Artificial Intelligence and Law*, ACM, 1993, pp. 50–59.
[36] H.C. Black, B. Garner, B. McDaniel, D. Schultz and W.P. Company, *Black's Law Dictionary*, West Group, St. Paul, MN, 1999.
[37] G. Brewka and T.F. Gordon, Carneades and Abstract Dialectical Frameworks: A reconstruction, in: *Proceedings of COMMA 2010*, IOS Press, 2010, pp. 3–12.
[38] G. Brewka, H. Strass, S. Ellmauthaler, J.P. Wallner and S. Woltran, Abstract Dialectical Frameworks revisited, in: *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
[39] G. Brewka and S. Woltran, Abstract Dialectical Frameworks, in: *Twelfth International Conference on the Principles of Knowledge Representation and Reasoning*, 2010.
[40] S. Brüninghaus and K.D. Ashley, Predicting outcomes of case based legal arguments, in: *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, ACM, 2003, pp. 233–242.
[41] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning and logic programming, in: *Proceedings of the International Joint Conference on Artificial Intelligence*, 1993, pp. 852–857.
[42] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and *n*-person games, *Artificial Intelligence* **77**(2) (1995), 321–357. doi:10.1016/0004-3702(94)00041-X.
[43] P. Dunne, The computational complexity of ideal semantics, *Artificial Intelligence* **173**(18) (2009), 1559–1591. doi:10.1016/j.artint.2009.09.001.
[44] P. Dunne and T. Bench-Capon, Two party immediate response disputes: Properties and efficiency, *Artificial Intelligence* **149**(2) (2003), 221–250. doi:10.1016/S0004-3702(03)00076-6.
[45] U. Egly, S.A. Gaggl and S. Woltran, ASPARTIX: Implementing argumentation frameworks using answer-set programming, in: *International Conference on Logic Programming*, Springer, 2008, pp. 734–738. doi:10.1007/978-3-540-89982-2_67.
[46] E. Feteris, *Fundamentals of Legal Argumentation*, Springer, 2017. doi:10.1007/978-94-024-1129-4.
[47] G. Frege, *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, Northwestern University Press, 1968, first published 1879.
[48] H. Geffner and J. Pearl, Conditional entailment: Bridging two approaches to default reasoning, *Artificial Intelligence* **53**(2–3) (1992), 209–244. doi:10.1016/0004-3702(92)90071-5.
[49] T.F. Gordon, Oblog-2: A hybrid knowledge representation system for defeasible reasoning, in: *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, ACM, 1987, pp. 231–239.
[50] T.F. Gordon, Some problems with Prolog as a knowledge representation language for legal expert systems, *International Review of Law, Computers & Technology* **3**(1) (1987), 52–67. doi:10.1080/13600869.1987.9966253.

[51] T.F. Gordon, The pleadings game: Formalizing procedural justice, in: *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, ACM, 1993, pp. 10–19.
[52] T.F. Gordon, The pleadings game, *Artificial Intelligence and Law* **2**(4) (1994), 239–292. doi:10.1007/BF00871972.
[53] T.F. Gordon, Constructing arguments with a computational model of an argumentation scheme for legal rules: Interpreting legal rules as reasoning policies, in: *Proceedings of the 11th International Conference on Artificial Intelligence and Law*, ACM, 2007, pp. 117–121.
[54] T.F. Gordon and N.I. Karacapilidis, The Zeno argumentation framework, in: *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, ACM, 1999, pp. 10–18.
[55] T.F. Gordon, H. Prakken and D. Walton, The Carneades model of argument and burden of proof, *Artificial Intelligence* **171**(10–15) (2007), 875–896. doi:10.1016/j.artint.2007.04.010.
[56] T.F. Gordon and D. Walton, Legal reasoning with argumentation schemes, in: *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, ACM, 2009, pp. 137–146.
[57] T.F. Gordon and D. Walton, A Carneades reconstruction of Popov v Hayashi, *Artificial Intelligence and Law* **20**(1) (2012), 37–56. doi:10.1007/s10506-012-9120-0.
[58] G. Governatori and A. Rotolo, A computational framework for institutional agency, *Artificial Intelligence and Law* **16**(1) (2008), 25–52. doi:10.1007/s10506-007-9056-y.
[59] M. Grabmair, Modeling purposive legal argumentation and case outcome prediction using argument schemes in the value judgment formalism, PhD thesis, University of Pittsburgh, 2016.
[60] M. Grabmair and K.D. Ashley, Facilitating case comparison using value judgments and intermediate legal concepts, in: *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, ACM, 2011, pp. 161–170.
[61] K. Greenwood, T. Bench-Capon and P. McBurney, Towards a computational account of persuasion in law, in: *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, ACM, 2003, pp. 22–31.
[62] J.C. Hage, Comparing alternatives in the law, *Artificial Intelligence and Law* **12**(3) (2004), 181–225. doi:10.1007/s10506-005-6926-z.
[63] J.C. Hage, R. Leenes and A.R. Lodder, Hard cases: A procedural approach, *Artificial Intelligence and Law* **2**(2) (1993), 113–167. doi:10.1007/BF00871759.
[64] J.C. Hage, G.P.J. Span and A.R. Lodder, A dialogical model of legal reasoning, in: *Proceedings of JURIX 1992*, 1992, pp. 135–145.
[65] C.L. Hamblin, *Fallacies*, Methuen, London, 1970.
[66] H.L.A. Hart, *The Concept of Law*, Oxford University Press, 1961.
[67] J. Horty and T. Bench-Capon, A factor-based definition of precedential constraint, *Artificial Intelligence and Law* **20**(2) (2012), 181–214. doi:10.1007/s10506-012-9125-8.
[68] R. Kowalski, The treatment of negation in logic programs for representing legislation, in: *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, ACM, 1989, pp. 11–15.
[69] A.R. Lodder, *DiaLaw: On Legal Justification and Dialogical Models of Argumentation*, Law and Philosophy Library, Vol. 42, Springer, 2001.
[70] A.R. Lodder and A. Herczog, DiaLaw: A dialogical framework for modeling legal reasoning, in: *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, ACM, 1995, pp. 146–155.
[71] L.S. Lutomski, The design of an attorney's statistical consultant, in: *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, 1989, pp. 224–233.
[72] J.D. Mackenzie, Question-begging in non-cumulative systems, *Journal of Philosophical Logic* **8**(1) (1979), 117–133.
[73] C.C. Marshall, Representing the structure of a legal argument, in: *Proceedings of the 2nd International Conference on Artificial Intelligence and Law*, ACM, 1989, pp. 121–127.
[74] L.T. McCarty, Reflections on TAXMAN: An experiment in Artificial Intelligence and legal reasoning, *Harvard Law Review* **90** (1976), 837–893. doi:10.2307/1340132.
[75] L.T. McCarty, An implementation of Eisner v. Macomber, in: *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, ACM, 1995, pp. 276–286.
[76] S. Modgil, Reasoning about preferences in argumentation frameworks, *Artificial Intelligence* **173**(9–10) (2009), 901–934. doi:10.1016/j.artint.2009.02.001.
[77] S. Modgil and T. Bench-Capon, Metalevel argumentation, *Journal of Logic and Computation* **21**(6) (2011), 959–1003. doi:10.1093/logcom/exq054.
[78] S. Modgil and H. Prakken, A general account of argumentation with preferences, *Artificial Intelligence* **195** (2013), 361–397. doi:10.1016/j.artint.2012.10.008.
[79] S. Modgil and H. Prakken, The ASPIC+ framework for structured argumentation: A tutorial, *Argument and Computation* **5**(1) (2014), 31–62. doi:10.1080/19462166.2013.869766.
[80] D. Moore, Dialogue game theory for intelligent tutoring systems, PhD thesis, Leeds Metropolitan University, 1993.
[81] C. Perelman, *Justice, Law, and Argument: Essays on Moral and Legal Reasoning*, Reidel, Dordrecht, 1980. doi:10.1007/978-94-009-9010-4.

[82] J.L. Pollock, *Cognitive Carpentry: A Blueprint for How to Build a Person*, MIT Press, 1995.
[83] H. Prakken, A tool in modelling disagreement in law: Preferring the most specific argument, in: *Proceedings of the 3rd International Conference on Artificial Intelligence and Law*, ACM, 1991, pp. 165–174.
[84] H. Prakken, Logical tools for modelling legal reasoning, PhD thesis, Free University of Amsterdam, 1993.
[85] H. Prakken, From logic to dialectics in legal argument, in: *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, ACM, 1995, pp. 165–174.
[86] H. Prakken, Dialectical proof theory for defeasible argumentation with defeasible priorities (preliminary report), in: *ModelAge Workshop on Formal Models of Agents*, Springer, 1997, pp. 202–215.
[87] H. Prakken, *Formalizing Robert's Rules of Order: An Experiment in Automating Mediation of Group Decision Making*, GMD – Forschungszentrum Informationstechnik, 1998.
[88] H. Prakken, An exercise in formalising teleological case-based reasoning, *Artificial Intelligence and Law* **10**(1–3) (2002), 113–133. doi:10.1023/A:1019536206548.
[89] H. Prakken, AI & Law, logic and argument schemes, *Argumentation* **19**(3) (2005), 303–320. doi:10.1007/s10503-005-4418-7.
[90] H. Prakken, Formal systems for persuasion dialogue, *The Knowledge Engineering Review* **21**(2) (2006), 163–188. doi:10.1017/S0269888906000865.
[91] H. Prakken, A formal model of adjudication dialogues, *Artificial Intelligence and Law* **16**(3) (2008), 305–328. doi:10.1007/s10506-008-9066-4.
[92] H. Prakken, An abstract framework for argumentation with structured arguments, *Argument and Computation* **1**(2) (2010), 93–124. doi:10.1080/19462160903564592.
[93] H. Prakken, Reconstructing Popov v. Hayashi in a framework for argumentation with structured arguments and Dungean semantics, *Artificial Intelligence and Law* **20**(1) (2012), 57–82. doi:10.1007/s10506-012-9117-8.
[94] H. Prakken, Relating ways to instantiate abstract argumentation frameworks, in: *From Knowledge Representation to Argumentation in AI, Law and Policy Making. A Festschrift in Honour of Trevor Bench-Capon on the Occasion of His 60th Birthday*, K. Atkinson, H. Prakken and A. Wyner, eds, College Publications, 2013, pp. 167–189.
[95] H. Prakken, C. Reed and D. Walton, Argumentation schemes and generalisations in reasoning about evidence, in: *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, ACM, 2003, pp. 32–41.
[96] H. Prakken and G. Sartor, A dialectical model of assessing conflicting arguments in legal reasoning, *Artificial Intelligence and Law* **4**(3–4) (1996), 331–336. doi:10.1007/BF00118496.
[97] H. Prakken and G. Sartor, Argument-based extended logic programming with defeasible priorities, *Journal of Applied Non-Classical Logics* **7**(1–2) (1997), 25–75. doi:10.1080/11663081.1997.10510900.
[98] H. Prakken and G. Sartor, Modelling reasoning with precedents in a formal dialogue game, *Artificial Intelligence and Law* **6**(2–4) (1998), 231–287. doi:10.1023/A:1008278309945.
[99] H. Prakken, A. Wyner, T. Bench-Capon and K. Atkinson, A formalization of argumentation schemes for legal case-based reasoning in ASPIC+, *Journal of Logic and Computation* **25**(5) (2015), 1141–1166. doi:10.1093/logcom/ext010.
[100] A. Rigoni, An improved factor based approach to precedential constraint, *Artificial Intelligence and Law* **23**(2) (2015), 133–160. doi:10.1007/s10506-015-9166-x.
[101] A. Rigoni, Representing dimensions within the reason model of precedent, *Artificial Intelligence and Law* **26**(1) (2018), 1–22. doi:10.1007/s10506-017-9216-7.
[102] E.L. Rissland and K.D. Ashley, A case-based system for Trade Secrets law, in: *Proceedings of the 1st International Conference on Artificial Intelligence and Law*, ACM, 1987, pp. 60–66.
[103] E.L. Rissland, D.B. Skalak and M.T. Friedman, BankXX: Supporting legal arguments through heuristic retrieval, *Artificial Intelligence and Law* **4**(1) (1996), 1–71. doi:10.1007/BF00123994.
[104] H.M. Robert, *Robert's Rules of Order Revised for Deliberative Assemblies*, Scott, Foresman, 1915.
[105] A. Rotolo, G. Governatori and G. Sartor, Deontic defeasible reasoning in legal interpretation: Two options for modelling interpretive arguments, in: *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, ACM, 2015, pp. 99–108.
[106] G. Sartor, A simple computational model for nonmonotonic and adversarial legal reasoning, in: *Proceedings of the 4th International Conference on Artificial Intelligence and Law*, ACM, 1993, pp. 192–201.
[107] G. Sartor, Teleological arguments and theory-based dialectics, *Artificial Intelligence and Law* **10**(1–3) (2002), 95–112. doi:10.1023/A:1019589831118.
[108] J.R. Searle, *Rationality in Action*, MIT Press, 2003.
[109] M.J. Sergot, F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond and H.T. Cory, The British Nationality Act as a logic program, *Communications of the ACM* **29**(5) (1986), 370–386. doi:10.1145/5689.5920.
[110] E.H. Shortliffe and B.G. Buchanan, *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Addison-Wesley Publishing Company, 1985.
[111] G.R. Simari and R.P. Loui, A mathematical treatment of defeasible reasoning and its implementation, *Artificial Intelligence* **53**(2–3) (1992), 125–157. doi:10.1016/0004-3702(92)90069-A.

[112] D.B. Skalak and E.L. Rissland, Arguments and cases: An inevitable intertwining, *Artificial Intelligence and Law* **1**(1) (1992), 3–44. doi:10.1007/BF00118477.
[113] G. Storrs, Group decision making, in: *Knowledge-Based Systems and Legal Applications*, T. Bench-Capon, ed., Academic Press, London, 1991, pp. 295–308. doi:10.1016/B978-0-12-086441-6.50022-2.
[114] S.E. Toulmin, *The Uses of Argument*, Cambridge University Press, 1958.
[115] A. Valente, *Legal Knowledge Engineering: A Modelling Approach*, Penn State Press, 1995.
[116] B. van Gijzel and H. Prakken, Relating Carneades with abstract argumentation via the ASPIC+ framework for structured argumentation, *Argument and Computation* **3**(1) (2012), 21–47. doi:10.1080/19462166.2012.661766.
[117] B. Verheij, Legal decision making as dialectical theory construction with argumentation schemes, in: *Proceedings of the 8th International Conference of Artificial Intelligence and Law*, ACM, 2001, pp. 225–226.
[118] B. Verheij, Artificial argument assistants for defeasible argumentation, *Artificial Intelligence* **150**(1–2) (2003), 291–324. doi:10.1016/S0004-3702(03)00107-3.
[119] B. Verheij, DefLog: On the logical interpretation of prima facie justified assumptions, *Journal of Logic and Computation* **13**(3) (2003), 319–346. doi:10.1093/logcom/13.3.319.
[120] B. Verheij, Dialectical argumentation with argumentation schemes: An approach to legal logic, *Artificial Intelligence and Law* **11**(2–3) (2003), 167–195. doi:10.1023/B:ARTI.0000046008.49443.36.
[121] B. Verheij, Proof with and without probabilities, *Artificial Intelligence and Law* **25**(1) (2017), 127–154. doi:10.1007/s10506-017-9199-4.
[122] G. Vreeswijk, Abstract argumentation systems, *Artificial Intelligence* **90**(1–2) (1997), 225–279. doi:10.1016/S0004-3702(96)00041-0.
[123] G. Vreeswik, Studies in defeasible argumentation, PhD thesis, Free University of Amsterdam, 1993.
[124] G. Vreeswik and H. Prakken, Credulous and sceptical argument games for preferred semantics, in: *European Workshop on Logics in Artificial Intelligence*, Springer, 2000, pp. 239–253. doi:10.1007/3-540-40006-0_17.
[125] F. Waismann, Verifiability, *Proceedings of the Aristotelian Society* **19** (1945), 118–150.
[126] F. Waismann, *The Principles of Linguistic Philosophy*, Macmillan, London, 1965.
[127] D. Walton, *Argumentation Schemes for Presumptive Reasoning*, Lawrence Erlbaum Associates, Inc., 1996.
[128] D. Walton and E. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, SUNY Press, 1995.
[129] D. Walton, C. Reed and F. Macagno, *Argumentation Schemes*, Cambridge University Press, 2008. doi:10.1017/CBO9780511802034.
[130] D. Walton, G. Sartor and F. Macagno, An argumentation framework for contested cases of statutory interpretation, *Artificial Intelligence and Law* **24**(1) (2016), 51–91. doi:10.1007/s10506-016-9179-0.
[131] A. Wyner and T. Bench-Capon, Argument schemes for legal case-based reasoning, in: *Proceedings of JURIX 2007*, 2007, pp. 139–149.
[132] A. Wyner, T. Bench-Capon and K. Atkinson, Arguments, values and baseballs: Representation of Popov v. Hayashi, in: *Proceedings of JURIX 2007*, 2007, pp. 151–160.
[133] A. Wyner, T. Bench-Capon and P. Dunne, On the instantiation of knowledge bases in abstract argumentation frameworks, in: *International Workshop on Computational Logic in Multi-Agent Systems*, Springer, 2013, pp. 34–50.
[134] J. Zeleznikow and A. Stranieri, The SPLIT-UP system: Integrating neural networks and rule-based reasoning in the legal domain, in: *Proceedings of the 5th International Conference on Artificial Intelligence and Law*, ACM, 1995, pp. 185–194.