

Identifying Audience Preferences in Legal and Social Domains

Paul E. Dunne and Trevor Bench-Capon
Dept. of Computer Science
University of Liverpool
Liverpool L69 7ZF, United Kingdom

1 Introduction

Although deductive techniques have proved successful in many areas, legal and social applications seem less amenable to this style of approach. In domains relating to practical reasoning, such as law, politics and ethics, it is not possible to demonstrate the acceptability of a position absolutely. Rationality in such fields depends on having an argument (not necessarily a proof) which provides a *prima facie* justification of the position, and then defending this argument against any counter arguments that may be offered. The process is thus open ended, and dependent on the context, which is supplied by arguments advanced against the position on the particular occasion.

While a variety of non-monotonic logics have been proposed which address problems arising from uncertain and incomplete information, these are not ideal for legal and social domains. There are two reasons for this. First, in legal and social domains counter arguments arise not only from additional and clarified information, but also from the perspective taken on existing information, the interpretation to be given to the available data. Second, arguments in legal and social domains typically exhibit a variety of argument schemas (for example, [19] identifies sixteen such schemas) and so are not always easily cajoled into some standard form.

For these reasons those modelling rationality in legal and social domains have increasingly been attracted to argumentation frameworks, as originally proposed in [10]. For example, in the legal domain [1], [16], and [14] all use these frameworks in particular ways. The usefulness of argumentation frameworks is that they have a very abstract notion of argument which allows them to accommodate the plurality of possible argumentation schemes, and that they consider conflict between a body of arguments, so as to identify arguments that are plausibly co-tenable, accounting for the notions of defeat and reinstatement that are typical of these areas of reasoning.

However, argumentation frameworks do not in general give a definitive answer as to which arguments should be accepted. This is as should be expected. As Searle remarks in [17, p. xv]

Assume universally valid and accepted standards of rationality. Assume perfectly rational agents operating with perfect information, and you will find that rational disagreement will still occur; because, for example, the rational agents are likely to have different and inconsistent values and interests, each of which may be rationally acceptable.

The point is also made by Perelman [15]:

If men oppose each other concerning a decision to be taken, it is not because they commit some error of logic or calculation. They discuss appro-

pos the applicable rule, the ends to be considered, the meaning to be given to values, the interpretation and characterisation of facts.

Since typically there are several sets of arguments which could be accepted, we need to try to motivate the choice between these positions in some well motivated way. We do this by recognising that legal and social arguments depend for their persuasiveness not only on their intrinsic merits, but also on the audience to which they are addressed. Which arguments are found persuasive depends, as Searle and Perelman suggest, on the opinions, values and, perhaps, even the prejudices of the audience to which they are addressed. Consider an example: it is an argument in favour of raising income tax that it promotes equality, and for decreasing income tax that it promotes enterprise. Most people would acknowledge that both arguments are valid, but which they will choose to follow depends on the importance they ascribe to equality as against enterprise in the given situation. Thus which of the arguments is found persuasive by a given audience will depend on the ordering of these two values by that audience. So if I claim that income tax should be raised to promote equality, I can defend my claim against the attack that so doing would discourage enterprise not only by attacking this counter argument but instead by declaring my preference for equality over enterprise. Whether this defence will be persuasive will depend on whether the audience shares my preference as to the values promoted by acceptance of the argument.

Having recognised that the acceptability arguments in these domains is dependent on values and needs to accommodate the notion of audiences with different values, Dung's original framework was extended in [2] and [3] to give what are termed there Value Based Argumentation Frameworks (VAFs). In those papers a number of properties of VAFs are demonstrated. For example, given an argument system, H , within which each distinct argument promotes one of a finite number of "values" and an ordering, α , of these values representing an individual's preferences, there is a unique maximal subset, $P(\alpha)$, of the arguments in H that is internally consistent and can defend itself against attacks from arguments of H that are not in $P(\alpha)$: in [2] an efficient algorithm of identifying the set $P(\alpha)$ given an value-based argument system H and value ordering α is also described.

Within VAFS arguments can be classified according as to whether they are accepted by at least one audience (so-called *subjective acceptance*); accepted by any audience, that is acceptable whatever the ordering of values, (so-called *objective acceptance*); or indefensible, that is unacceptable whatever the ordering on values. Ideally we would wish to have an efficient way of determining the status of a particular argument within a VAF: Unfortunately it is shown in [12] that this is unlikely to admit an efficient algorithmic solution: subjective acceptance being shown NP-complete, and objective acceptance CO-NP-complete.

These results raise the concern that the approach of using argumentation frameworks with values may, despite their attractions for modelling rationality in legal and social domains, present computational problems if applied to realistically large scale problems. Our aim in this paper is to present further indications that, notwithstanding the negative consequences of [12], the VAF formalism provides a practical model for the study and analysis of persuasive argument. This case is promoted by describing efficient algorithms on VAFs that could be exploited in arenas such as dialogue pro-

cesses on VAFs, in particular in uncovering points of disagreement regarding value precedence between the different audiences participating in the dialogue.

The remainder of this paper is structured as follows: section 2 reviews the definition of Argumentation Systems from [10] and their development into the Value Based Argumentation Frameworks of [2, 3], noting the computational properties established. In section 3 we discuss and motivate a number of algorithmic problems for VAFs and describe efficient solution methods for these: these problems relate to determining those values whose ranking is disputed. Finally, section 4 will offer some discussion and concluding remarks.

2 Basic Definitions: Argument Systems and VAFs

The basic definition below of an Argument System is derived from that given in [10].

Definition 1 *An argument system is a pair $\mathcal{H} = \langle \mathcal{X}, \mathcal{A} \rangle$, in which \mathcal{X} is a finite set of arguments and $\mathcal{A} \subset \mathcal{X} \times \mathcal{X}$ is the attack relationship for \mathcal{H} . A pair $\langle x, y \rangle \in \mathcal{A}$ is referred to as ‘ y is attacked by x ’ or ‘ x attacks y ’. For R, S subsets of arguments in the system $\mathcal{H}(\langle \mathcal{X}, \mathcal{A} \rangle)$, we say that*

- a. $s \in S$ is attacked by R if there is some $r \in R$ such that $\langle r, s \rangle \in \mathcal{A}$.
- b. $x \in \mathcal{X}$ is acceptable with respect to S if for every $y \in \mathcal{X}$ that attacks x there is some $z \in S$ that attacks y .
- c. S is conflict-free if no argument in S is attacked by any other argument in S .
- d. A conflict-free set S is admissible if every argument in S is acceptable with respect to S .
- e. S is a preferred extension if it is a maximal (with respect to \subseteq) admissible set.
- f. S is a stable extension if S is conflict free and every argument $y \notin S$ is attacked by S .
- g. \mathcal{H} is coherent if every preferred extension in \mathcal{H} is also a stable extension.

An argument x is credulously accepted if there is some preferred extension containing it; x is sceptically accepted if it is a member of every preferred extension.

Abstracting away concerns regarding the internal structure and representation of arguments affords a formalism which focuses on the relationship between individual arguments as a means of defining diverse ideas of acceptance. In particular the preferred extension is of interest as it represents a maximal coherent position that can be defended against all attackers.

While this approach offers a powerful tool for the abstract analysis of defeasible reasoning, there are, however, several potential problems. While every argument system has *some* preferred extension, this may simply be the empty set of arguments and although the use of *stable* extensions avoids such problems these in turn have the problem that there are systems which contain no stable extension. An additional problem is

	Problem	Decision Question	Complexity
1	$\text{ADM}(\mathcal{H}, S)$	Is S admissible?	P
2	$\text{STAB}(\mathcal{H}, S)$	Is S stable?	P
3	$\text{PREF-EXT}(\mathcal{H}, S)$	Is S preferred?	CO-NP-complete.
4	$\text{CA}(\mathcal{H}, x)$	Is x in a preferred S ?	NP-complete
5	$\text{STAB-EXIST}(\mathcal{H})$	Has \mathcal{H} a stable extension?	NP-complete
6	$\text{SA}(\mathcal{H}, x)$	Is x in every preferred S ?	$\Pi_2^{(p)}$ -complete
7	$\text{COHERENT}(\mathcal{H})$	Preferred \equiv stable?	$\Pi_2^{(p)}$ -complete

Table 1: Decision Problems in Argument Systems

the computational complexity of a number of the associated decision problems that has been shown to range from NP-complete to Π_2^p -complete. A summary of these is given in Table 1 below. The classification of problems (3–5) follows from [9]; that of (6) and (7) has recently been demonstrated in [11]. Related problems arise with proof-theoretic mechanisms for establishing credulous acceptance, e.g. for the TPI-dispute mechanism proposed in [18], Dunne and Bench-Capon [13] show that this defines a weak propositional proof system under which proofs that arguments are not credulously accepted require exponentially many steps.

While the issues discussed above concern algorithmic and combinatorial properties of the standard argument system framework, there is also one interpretative issue of some importance. A typical argument system may contain many distinct preferred extensions and, in some cases, two different preferred extensions may define a partition of the argument set. Thus a single argument system can give rise to a number of disjoint internally consistent admissible argument sets. The abstract level at which Dung’s formalism operates avoids any mechanism for distinguishing notions of the relative merit of such mutually incompatible outcomes. Thus the situation arises in which we appear to have several coherent positions that could be adopted, and no well motivated way of choosing between them.

As mentioned above, [2] and [3] extend Dung’s framework to provide a semantics for distinguishing and choosing between consistent but incompatible belief sets through the use of *argument values*. Thus arguments are seen as grounded on one of a finite number of abstract values and the interpretation of which of a set of arguments to “accept” is treated in terms of preference orderings of the underlying value set according to the views held by a particular *audience*. Thus while in the standard Argumentation system the choice between preferred extensions is arbitrary, in a VAF we are able to motivate such choices by reference to the values of the audience. The formal definition of such *value-based argumentation frameworks* is given below.

Definition 2 A value-based argumentation framework (VAF), is defined by a triple $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$, where $\mathcal{H}(\mathcal{X}, \mathcal{A})$ is an argument system, $\mathcal{V} = \{v_1, v_2, \dots, v_k\}$ a set of k values, and $\eta : \mathcal{X} \rightarrow \mathcal{V}$ a mapping that associates a value $\eta(x) \in \mathcal{V}$ with each argument $x \in \mathcal{X}$. An audience, α , for a VAF $\langle \mathcal{H}, \mathcal{V}, \eta \rangle$, is a total ordering of the values \mathcal{V} . We say that v_i is preferred to v_j in the audience α , denoted $v_i \succ_\alpha v_j$, if v_i is ranked higher than v_j in the total ordering defined by α .

Using VAFs, ideas analogous to those of admissible argument in standard argument systems are defined in the following way. Note that all these notions are now relative to some audience.

Definition 3 Let $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$ be a VAF and α an audience.

- a. For arguments x, y in \mathcal{X} , x is a successful attack on y (or x defeats y) with respect to the audience α if: $\langle x, y \rangle \in \mathcal{A}$ and it is not the case that $\eta(y) \succ_{\alpha} \eta(x)$.
- b. An argument x is acceptable to the subset S with respect to an audience α if: for every $y \in \mathcal{X}$ that successfully attacks x with respect to α , there is some $z \in S$ that successfully attacks y with respect to α .
- c. A subset R of \mathcal{X} is conflict-free with respect to the audience α if: for each $\langle x, y \rangle \in R \times R$, either $\langle x, y \rangle \notin \mathcal{A}$ or $\eta(y) \succ_{\alpha} \eta(x)$.
- d. A subset R of \mathcal{X} is admissible with respect to the audience α if: R is conflict free with respect to α and every $x \in R$ is acceptable to R with respect to α .
- e. A subset R is a preferred extension for the audience α if it is a maximal admissible set with respect to α .
- f. A subset R is a stable extension for the audience α if R is admissible with respect to α and for all $y \notin R$ there is some $x \in R$ which successfully attacks y .

A standard consistency requirement which we assume of the VAFs considered is that every directed cycle of arguments in these contains *at least two* differently valued arguments. We do not believe that this condition is overly restricting, since the existence of such cycles in VAFs can be seen as indicating a flaw in the formulation of the framework. While in standard argumentation frameworks cycles arise naturally, especially if we are dealing with uncertain or incomplete information, in VAFs odd length cycles in a single value represent paradoxes and even length cycles in a single value can be reduced to a self-defeating argument. Given the absence of cycles in a single value the following important property of VAFs and audiences was demonstrated in [2].

Fact 1 For every audience, α , $\langle \mathcal{H}(\langle \mathcal{X}, \mathcal{A} \rangle), \mathcal{V}, \eta \rangle$ has a unique non-empty preferred extension, $P(\mathcal{H}, \eta, \alpha)$ which can be constructed by an algorithm that takes $O(|\mathcal{X}| + |\mathcal{A}|)$ steps. Furthermore $P(\mathcal{H}, \eta, \alpha)$ is a stable extension with respect to α .

From Fact 1 it follows that, when attention is focused on one specific audience, the decision questions analogous to those described in Table 1 become much easier. There are, however, a number of new issues that arise in the value-based framework from the fact that that the relative ordering of different values promoted by distinct audiences results in arguments falling into one of three categories.

- C1. Arguments, x , that are in the preferred extension $P(\mathcal{H}, \eta, \alpha)$ for some audiences but not all. Such arguments being called *subjectively acceptable*.
- C2. Arguments, x , that are in the preferred extension $P(\mathcal{H}, \eta, \alpha)$ for every audience. Such arguments being called *objectively acceptable*.

- C3. Arguments, x , that do not belong to the preferred extension $P(\mathcal{H}, \eta, \alpha)$ for any choice of audience. Such arguments being called *indefensible*.

To show the advantages of taking values into account, consider the following ethical debate, discussed in, e.g. [5]. Hal, a diabetic, loses his insulin and can save his life only by breaking into the house of another diabetic, Carla, and using her insulin. We may consider the following arguments:

- A. Hal should not take Carla's insulin as he may be endangering her life.
- B. Hal can take the insulin as otherwise he will die, whereas there is only a potential threat to Carla.
- C. Hal cannot take Carla's insulin because it is Carla's property.
- D. Hal can replace Carla's insulin once the emergency is over.

Now B attacks A , C attacks B and D attacks C . Moreover Christie argues in [5] that A attacks D , since if even if Hal were unable to replace the insulin he would still be correct to act so as to save his life. The argumentation system can be depicted as a directed graph as shown in Figure 1.

Figure 1: VAF Example Argument System

Considered as a standard Argumentation System there is no non-empty preferred extension, and it seems we have no coherent position, which is why it is seen and discussed as an ethical dilemma. If, however, we consider it as a VAF, we can see that arguments A and B rely on the importance of preserving life, whereas C and D depend on respect for property. We will now have two preferred extensions, depending on whether life or property is preferred. If we prefer life, we will accept $\{B, C\}$: whilst we respect Carla's property rights, we regard Hal's need as paramount. In contrast if we prefer property to life, the preferred extension is $\{B, D\}$: the property claim can be discharged if restitution is made. Thus B is objectively acceptable, C and D are subjectively acceptable and A is indefensible. This small example shows how we can use explicit value preferences to cut through what would otherwise be an irresolvable dilemma.

In [12] the following questions specific to the value-based setting are considered.

Definition 4 *The decision problem Subjective Acceptance (SBA) takes as an instance a VAF $\langle \mathcal{H}, \mathcal{V}, \eta \rangle$ and an argument x in this. The instance is accepted if there is at least one audience, α , for which $x \in P(\mathcal{H}, \eta, \alpha)$. The decision problem Objective Acceptance (OBA) takes as an instance a VAF $\langle \mathcal{H}, \mathcal{V}, \eta \rangle$ and an argument x in this. The instance is accepted if $x \in P(\mathcal{H}, \eta, \alpha)$ for every audience α .*

Let $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$ be a VAF, and $\langle v, v' \rangle$ be an ordered pair of distinct values from \mathcal{V} . The pair $\langle v, v' \rangle$ is critical with respect to an argument $x \in \mathcal{X}$ if there is an audience α for which $v \succ_{\alpha} v'$ and $x \in P(\mathcal{H}, \eta, \alpha)$, but for every audience β for which $v' \succ_{\beta} v$ it holds that $x \notin P(\mathcal{H}, \eta, \beta)$. The decision problem Critical Pair (CP) takes

as an instance a triple $\langle \langle \mathcal{H}, \mathcal{V}, \eta \rangle, \langle v, v' \rangle, x \rangle$ comprising a VAF, ordered pair of values $\langle v, v' \rangle$ within this, and argument x . An instance is accepted if $\langle v, v' \rangle$ is critical with respect to x .

Regarding these [12] prove,

Theorem 1

- a. SBA is NP-complete.
- b. OBA is CO-NP-complete.
- c. CP is D^P -complete.

We have now arrived at the position where we can detect efficiently the arguments acceptable to any particular given audience, but cannot guarantee that we will be able to determine the status of an argument with respect to the totality of audiences. In the next section we consider another problem relating to VAFs which admits of an efficient solution, namely finding an audience for whom a subset of arguments represents a preferred extension, if one exists.

3 Efficient Algorithms for VAFs

We begin by giving a formal statement of our problem:

- a. Given a VAF $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$ and a subset S of \mathcal{X} find an audience α for which $S = P(\mathcal{H}, \eta, \alpha)$ or report that no such audience exists.

In this section we address this problem raised and some related applications. On first inspection, it might appear that, given Theorem 1 (a), this too would be an intractable problem. We will show, however, that this pessimistic view is ill-founded: the critical difference between the two problems is that subjective acceptance concerns the existence of an audience with respect to which a *single* given argument is accepted; whereas the current problem asks for an audience with respect to which a given *set* of arguments defines the totality of what that audience accepts. In the context of standard argument systems, subjective acceptance is analogous to deciding *credulous acceptance* whereas (a) to deciding if a given set defines a *preferred extension*. While there is no reduction in difficulty in moving from the standard to value-based frameworks, it turns out that the nature of acceptance following the introduction of value associations allows (a) to be solved by an efficient method. To see this we consider the following algorithm:

FIND AUDIENCE

Instance: VAF $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$; $S \subseteq \mathcal{X}$.

Returns: Audience α for which $S = P(\mathcal{H}, \eta, \alpha)$ or FAIL if no such audience exists. The audience is returned as a set of pairs of $\langle v_i, v_j \rangle$ for which $v_i \succ_\alpha v_j$ must hold.

- 1. PARTIAL- $\alpha := \emptyset$;

2. **for each** $\langle x, y \rangle \in S \times S$:
 - 2.1. **if** $\langle x, y \rangle \in \mathcal{A}$ **then**
 - a. **if** $\eta(x) = \eta(y)$ **then** report FAIL **else**

$$\text{PARTIAL-}\alpha := \text{PARTIAL-}\alpha \cup \{\langle \eta(y), \eta(x) \rangle\}$$
3. Replace PARTIAL- α with the transitive closure of PARTIAL- α .
4. **if** PARTIAL- α is “inconsistent” **then** report FAIL **else**
 5. **for each** $z \notin S$
 - a. **if** $\eta(z) = \eta(x)$ for some $x \in S$ **then**

Find some $y \in S$ for which $\langle y, z \rangle \in \mathcal{A}$ and for which $\langle \eta(z), \eta(y) \rangle \notin \text{PARTIAL-}\alpha$.

report FAIL if no suitable $y \in S$ is found.
 - b. **else** – $\eta(z)$ does not occur as the value of any $x \in S$

Choose any $y \in S$ for which $\langle y, z \rangle \in \mathcal{A}$ and add $\langle \eta(y), \eta(z) \rangle$ to PARTIAL- α .

report FAIL if no $y \in S$ attacks z .
6. **return** any audience that is consistent with the (partial) ordering given by PARTIAL- α .

Theorem 2 *Given an instance $\langle \mathcal{H}(\mathcal{X}, \mathcal{A}), \mathcal{V}, \eta \rangle$ and $S \subseteq \mathcal{X}$ the algorithm FIND AUDIENCE returns an audience α for which $S = P(\mathcal{H}, \eta, \alpha)$ or reports FAIL if no such audience exists. Furthermore the time taken is $O(|\mathcal{X}|^2)$.*

Steps (2) and (3) of FIND AUDIENCE construct a partial ordering of the values in S that satisfies the requirement that S must be conflict-free with respect to the audience: thus each $\langle x, y \rangle \in \mathcal{A}$ for which both x and y are in S forces an ordering of the values $\eta(x), \eta(y)$ according to the constraints specified in Definition 3(c). All constraints arising thus are added by the loop comprising (2), resulting in the set of constraints PARTIAL- α upon completion. At step (3), this set is extended to include all of the additional pair-wise orderings arising through the property that if $\langle \eta(x), \eta(y) \rangle \in \text{PARTIAL-}\alpha$ and $\langle \eta(y), \eta(z) \rangle \in \text{PARTIAL-}\alpha$ then any audience α consistent with PARTIAL- α must have $\eta(x) \succ_\alpha \eta(z)$: constructing all of the pair-wise ordering that should be included simply involves computing the transitive closure of the relations identified after (2) has completed. Step (4) deals with the requirement that any audience α for which $S = P(\mathcal{H}, \eta, \alpha)$ defines a *total* ordering of the value set \mathcal{V} and thus the set of pairs PARTIAL- α cannot contain both $\langle v_i, v_j \rangle$ and $\langle v_j, v_i \rangle$: this would happen if, for example, there were $\{x, y, z\} \in S$ with $\langle x, y \rangle \in \mathcal{A}$, $\langle y, z \rangle \in \mathcal{A}$ and $\eta(x) = \eta(z)$. Since (3) has formed the transitive closure of the constraint relationship identified in (2), the “consistency” test in (4) involves checking that for each $x \in S$ the pair $\langle \eta(x), \eta(x) \rangle$ has not been added to PARTIAL- α . Step (5) is concerned with checking that S is *maximal* with respect to the partial audience that has been constructed in the earlier stages. Again, from Definition 3, this simply involves testing for each argument $z \notin S$, that z cannot be added to S without creating a conflict. There

are two possibilities. The value $\eta(z)$ is among those considered in S : in which case it suffices to ensure that z is successfully attacked by some $y \in S$. The value $\eta(z)$ is distinct from any value used in S : in which it suffices to find any $y \in S$ that attacks z .

4 Discussion

Considered abstractly it is tempting to see the problem of resolving a dispute as a matter of computing a "grand" VAF - the union of the VAFs recognised by every participant to the dispute - with a fixed ordering of values for each of the audiences. This picture is, however, inappropriate for many practical settings in two ways. First, as Searle [?] says, participants do not always come to a dispute with a fixed value order; rather they work out what values they prize the most as the dispute progresses. Second the argumentation framework relevant to the dispute is typically constructed during the dispute as the participants introduce arguments.

In this setting we can see the importance of the algorithm established in section 3. When contributing an argument to the growing framework representing the dispute, the contributor knows which of the arguments already put forward he wishes to accept, but is perhaps unsure as to his exact values preferences. Given a state of the dispute, before contributing an additional argument, a participant should first consider if he is content with the current state. If the arguments he favours form a preferred extension for some audiences, and he is willing to accept a value ordering that this requires, he should contribute no further. Should a contribution be necessary, he must be sure that the argument added to the dispute will produce a preferred extension containing the arguments he desires: otherwise the opponents could choose to terminate the dispute by remaining silent leaving him with a less than optimal preferred extension.

Given the ability efficiently to compute the audience for which a given subset of the current arguments is a preferred extension, we can suggest the following strategy.

Suppose AF to be the current argumentation framework. Let S be the set of arguments desired to be the preferred extension. Is S a preferred extension of AF? If yes, remain silent. If no, consider arguments, A , which extend the current framework. If $S + A$ yields a preferred extension of $AF + A$, advance A . Otherwise consider arguments, B , such that S is a preferred extension in $AF + B$. Advance any such B .

Now the participant can be sure that he is only contributing when it is necessary to so, and that any contributions he makes will successfully defend his preferred set of arguments. The algorithm in section 3 ensures that the required information can be calculated efficiently.

5 Conclusion

In this paper we have argued that VAFs are suitable to capture the features characteristic of argumentation in legal and social domains. Additionally, we have shown that there is an efficient algorithm to supply the information required to drive an heuristic for deciding whether to continue to contribute to a dispute, and if so what would make a sensible continuation.

References

- [1] T. J. M. Bench-Capon 'Representation of Case Law as an Argumentation Framework' in *Legal Knowledge and Information Systems*, eds., T. Bench-Capon, A. Daskalopoulou and R. Winkels, IOS Press: Amsterdam. 103-112 (2002)
- [2] T. J. M. Bench-Capon, 'Agreeing to differ: modelling persuasive dialogue between parties with different values', *Informal Logic*, **22** 3 (2003).
- [3] T. J. M. Bench-Capon, 'Persuasion in Practical Argument Using Value Based Argumentation Frameworks', *Journal of Logic and Computation*, **13** 3 429-48 (2003).
- [4] A. Bondarenko, P. M. Dung, R. A. Kowalski, and F. Toni, 'An abstract, argumentation-theoretic approach to default reasoning', *Artificial Intelligence*, **93**(1-2), 63-101, (1997).
- [5] G. Christie, *The Notion of an Ideal Audience in Legal Argument*, Kluwer Academic, Dordrecht, (2000).
- [6] Y. Dimopoulos, B. Nebel, and F. Toni, 'Preferred arguments are harder to compute than stable extensions', in *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol1)*, ed., T. Dean, pp. 36-43, San Francisco, (1999). Morgan Kaufmann Publishers.
- [7] Y. Dimopoulos, B. Nebel, and F. Toni, 'Finding admissible and preferred arguments can be very hard', in *KR2000: Principles of Knowledge Representation and Reasoning*, eds., A. G. Cohn, F. Giunchiglia, and B. Selman, pp. 53-61, San Francisco, (2000). Morgan Kaufmann.
- [8] Y. Dimopoulos, B. Nebel, and F. Toni, 'On the computational complexity of assumption-based argumentation by default reasoning', *Artificial Intelligence*, **141**, 57-78, (2002).
- [9] Y. Dimopoulos and A. Torres, 'Graph theoretical structures in logic programs and default theories', *Theoretical Computer Science*, **170**, 209-244, (1996).
- [10] P. M. Dung, 'On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming, and N -person games', *Artificial Intelligence*, **77**, 321-357, (1995).
- [11] P.E. Dunne and T.J.M. Bench-Capon, 'Coherence in finite argument systems', *Artificial Intelligence*, **141**, 187-203, (2002).
- [12] P.E. Dunne and T.J.M. Bench-Capon, 'Complexity in value-based argument systems'. Technical Report, Dept. of Comp. Sci., Univ. of Liverpool, February 2004.
- [13] P.E. Dunne and T.J.M. Bench-Capon, 'Two party immediate response disputes: properties and efficiency', *Artificial Intelligence*, **149**, 221-250, (2003).

- [14] H. Jakobovits and D. Vermeir, 'Dialectic semantics for argumentation frameworks', in *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, ACM SIGART, pp. 53–62, N.Y., (June 1999). ACM Press.
- [15] C. Perelman, *Justice, Law and Argument*, Reidel: Dordrecht, 1980.
- [16] H. Prakken, *Logical Tools for Modelling Legal Argument*, Kluwer Academic Publishers, 1997.
- [17] J.R. Searle, *Rationality in Action*. MIT Press, Cambridge Mass., 2001
- [18] G. Vreeswijk and H. Prakken, 'Credulous and sceptical argument games for preferred semantics.', in *Proceedings of JELIA'2000, The 7th European Workshop on Logic for Artificial Intelligence.*, pp. 224–238, Berlin, (2000). Springer LNAI 1919, Springer Verlag.
- [19] D.N. Walton, *Argumentation Schemes for Presumptive Reasoning*, Erbaum: Mahwah, NJ.