Inconsistency Tolerance in Weighted Argument Systems

Paul E. Dunne
Dept. of Computer Science
University of Liverpool
Liverpool, UK
ped@csc.liv.ac.uk

Anthony Hunter
Dept. of Computer Science
University College London
London, UK
a.hunter@cs.ucl.ac.uk

Peter McBurney
Dept. of Computer Science
University of Liverpool
Liverpool, UK
mcburney@liverpool.ac.uk

Simon Parsons
Brooklyn College
CUNY, Brooklyn, NY, USA
parsons@sci.brooklyn.cuny.edu

Michael Wooldridge
Dept. of Computer Science
University of Liverpool
Liverpool, UK
mjw@csc.liv.ac.uk

ABSTRACT

We introduce and investigate a natural extension of Dung's well-known model of argument systems in which attacks are associated with a *weight*, indicating the relative strength of the attack. A key concept in our framework is the notion of an *inconsistency budget*, which characterises how much inconsistency we are prepared to tolerate: given an inconsistency budget β, we would be prepared to disregard attacks up to a total cost of β. The key advantage of this approach is that it permits a much finer grained level of analysis of argument systems than unweighted systems, and gives useful solutions when conventional (unweighted) argument systems have none. We begin by reviewing Dung's abstract argument systems, and present the model of weighted argument systems. We then investigate solutions to weighted argument systems and the associated complexity of computing these solutions, focussing in particular on weighted variations of grounded extensions.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; I.2.4 [**Knowledge representation formalisms and methods**]

General Terms

Theory

Keywords

Argumentation, handling inconsistency, complexity

1. INTRODUCTION

Inconsistency between the beliefs and/or preferences of agents is ubiquitous in everyday life, and yet coping with inconsistency remains an essentially unsolved problem in artificial intelligence [8]. One of the key aims of *argumentation* research is to provide principled techniques for handling inconsistency.

Although there are several different perspectives on argumentation (for a review see [9]), a common view is that argumentation starts with a collection of statements, called *arguments*, which are related through the notions of *support* and *attack*. Typically, argument α_1 supporting argument α_2 would be grounds for accepting α_2 if one accepted α_1, while argument α_1 attacking argument α_2 would be grounds for *not* accepting α_2 if one accepted α_1. Now, if we allow arguments to attack one-another, then such collections of arguments may be inconsistent; and the key question then becomes how to obtain a rationally justifiable position from such an inconsistent argument set. Various solutions have been proposed for this problem, such as *admissible sets*, *preferred extensions*, and *grounded extensions* [13]. However, none of these solutions is without drawbacks. A common situation is that, while a solution may guaranteed to give an answer, the answer may be the empty set. Conversely, several answers may be provided, with nothing to distinguish between them. These drawbacks limit the value of these solutions as argument analysis tools.

In part to overcome these difficulties, there is a trend in the literature on formalizations of argumentation towards considering the strength of arguments. In this work, which goes back at least as far as [16], it is recognized that not all arguments are equal in strength, and that this needs to be taken into account when finding extensions of a collection of arguments and counterarguments. We review this literature in Section 3, and we conclude that whilst it is clear that taking the strength of arguments into account is a valuable development, it is not just the strength of the arguments, per se, that is important. The strength of the attack that one argument (which may itself be very strong) makes on another, can be weak.

In this paper, we introduce, formalise, and investigate a natural extension of Dung's well-known model of argument systems [13], in which attacks between arguments are associated with a numeric *weight*, indicating the relative strength of the attack, or, equivalently, how reluctant we would be to disregard the attack. For example, consider the following arguments:

α_1: The house is in a good location, it is large enough for our family and it is affordable: we should buy it.

α_2: The house suffers from subsidence, which would be prohibitively expensive to fix: we should not buy it.

These arguments are mutually attacking: both arguments are credulously accepted, neither is sceptically accepted, and the grounded extension is empty. Thus the conventional analysis is not very useful for this scenario. However, the representation we are using surely misses a key point: the attacks are *not of equal weight*. We would surely regard the attack of α_2 on α_1 as being much stronger than the attack of α_1 on α_2, though both are very strong arguments

Cite as: Inconsistency Tolerance in Weighted Argument Systems, Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, Michael Wooldridge, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009),* Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 851–858

in their own right. Our framework allows us to take these differing weights of attack into consideration.

By using weights on attacks, we may be able to capture the relative strength of different attacks between arguments in a constellation. The use of strength of attack is wide-spread in informal argumentation, and real-world information is often available to judge the strength of the relations between arguments. To illustrate, in order to classify a compound according to potential toxicity, the U.S. Environmental Protection Agency needs to collect available scientific evidence on the compound and related compounds, and use this to construct arguments for and against a particular classification being applicable to the compound. Often, the evidence available is incomplete, and perhaps inconsistent, and to address this they systematically judge the result of attacks between arguments based on the nature of the evidence used. So for example, in their guidelines for the assessment of the health impacts of potential carcinogens, an argument for carcinogenicity that is based on human epidemiological evidence is considered to outweigh arguments against carcinogenicity that are based only on animal studies [32, 17]. This example indicates both the naturalness of considering strength of attack and of the availability of appropriate information for systematically evaluating the strength. Furthermore, in general, as we will discuss in Section 4, there are various semantics that we can apply to the weights assigned, and that these usefully reflect some of the usages of attack strength in real-world informal argumentation.

A key concept in our framework is the notion of an *inconsistency budget*, and this also distinguishes our approach from other methods of attaching weights to arguments. The inconsistency budget characterises how much inconsistency we are prepared to tolerate: given an inconsistency budget β, we would be prepared to disregard attacks up to a total cost of β. By increasing the inconsistency budget, we get progressively more solutions, and this in turn gives a preference ordering over solutions: we prefer solutions obtained with a smaller inconsistency budget. This approach permits a much finer-grained level of analysis of argument systems than is typically possible, and gives useful, non-trivial solutions when conventional (unweighted) argument systems have none. We begin by reviewing Dung's abstract argument systems, and present the framework of weighted argument systems. We then investigate solutions for weighted argument systems and the complexity of computing such solutions, focussing in particular on weighted variations of grounded extensions. Finally, we relate our work to the most relevant examples of systems that incorporate strengths.

2. ABSTRACT ARGUMENT SYSTEMS

Since weighted argument systems and their associated solutions generalise Dung's well-known abstract argument systems model, we begin by recalling some key concepts from this model. A Dung-style *abstract argument system* is a pair $D = \langle X, A\rangle$ where $X = \{\alpha_1, \ldots, \alpha_k\}$ is a finite set of *arguments*, and $A \subseteq X \times X$ is a binary *attack relation* on X [13]. Given a set of arguments X, let $\mathcal{D}(X)$ denote the set of all abstract argument systems over X, i.e., $\mathcal{D}(X) = \{\langle X, A\rangle : A \subseteq X \times X\}$. Note that Dung's model does not assume any internal structure for arguments, or give any concrete interpretation for them. The intended interpretation of the attack relation in Dung's model is also not completely defined, but intuitively, $(\alpha_1, \alpha_2) \in A$ means that if one accepts (in whatever solution one considers) α_1, then one should not accept α_2. In other words, *it would be inconsistent to accept α_2 if one accepted α_1*.

The next step is to define *solutions* for such argument systems. A solution for an argument system (over a set of arguments X) is a function $f : \mathcal{D}(X) \to \mathcal{P}(\mathcal{P}(X))$ i.e., a function that, given $\langle X, A\rangle$, will return a set of sets of arguments, such that each output represents a "position" that is in some sense rationally justifiable. Given $D = \langle X, A\rangle$ and $S \subseteq X$, we say that S is: *consistent* if $\nexists \alpha_1 \in S$ s.t. $\exists \alpha_2 \in X$ and $(\alpha_2, \alpha_1) \in A$; *internally consistent* (or *conflict free*) if $\nexists \alpha_1 \in S$ s.t. $\exists \alpha_2 \in S$ and $(\alpha_2, \alpha_1) \in A$; *defensive* if $\forall \alpha_1 \in X$ s.t. $\exists \alpha_2 \in S$ and $(\alpha_1, \alpha_2) \in A$, $\exists \alpha_3 \in S$ for which (α_3, α_1); *admissible* if it is both internally consistent and defensive; and a *preferred extension* if it is a maximal (wrt $\subseteq$) admissible set.

function $ge(X, A)$ returns a subset of X

$in \leftarrow out \leftarrow \emptyset$
while $in \neq X$ do
    $in \leftarrow \{\alpha \in X : \nexists \alpha' \in X$ s.t. $(\alpha', \alpha) \in A\}$
    $out \leftarrow \{\alpha \in X : \exists \alpha' \in in$ s.t. $(\alpha', \alpha) \in A\}$
    $X \leftarrow X \setminus out$
    $A \leftarrow A$ restricted to X
end-while
return X.

Figure 1: The function $ge(\cdots)$.

Consistency is the least problematic type of solution. However, while every argument system contains a consistent set of arguments, it may be that the only consistent set is the empty set. Such trivial solutions are typically unhelpful. If we do not have a non-empty consistent set of arguments, (which is the more general case), then we might look at the admissible sets, and the preferred extensions: a preferred extension is a maximal set of arguments that is both internally consistent and defends itself against all attacks. There will always be at least one preferred extension, although, again, this may be the empty set [13, p.327]. Note that non-empty preferred extensions may exist in argument systems for which the only consistent set of arguments is the empty set, and so we can usefully apply this solution in some situations where consistency is not a useful analytical concept. We say an argument is *credulously accepted* if it forms a member of at least one preferred extension, and *sceptically accepted* if it is a member of every preferred extension. Clearly, sceptical acceptance represents a stronger solution than credulous acceptance. Determining whether a given set of arguments is consistent or admissible can be solved in polynomial time; however, determining whether a set of arguments is a preferred extension is co-NP-complete, checking whether an argument is credulously accepted is NP-complete, while checking whether an argument is sceptically accepted is Π_2^p-complete [11, 15].

The final solution we consider is the *grounded extension* [13, p.328]. Roughly, the idea with grounded extensions is to iteratively compute the arguments whose status is beyond question, by first starting with arguments that have no attackers: we regard these as being unquestionably "in". Then, we eliminate arguments that these "in" arguments attack: since they are attacked by an argument whose status is unquestioned, we regard them as "out". We then eliminate the "out" arguments, and iterate, until we reach no change. The algorithm to compute the grounded extension of an argument system is given in Figure 1; basic properties of fixpoint algorithms tell us this algorithm is guaranteed to terminate in polynomial time. As a solution, grounded extensions are intuitively very appealing; an argument system will always have a unique grounded extension, although, again, this may be the empty set.

Notice that, while all of these solutions are guaranteed to give some "answer", it is possible that the only answer they give is the empty set. This is a key limitation of conventional systems.

3. TOWARDS ARGUMENT STRENGTH

There have been a number of proposals for extending Dung's framework in order to allow for more sophisticated modelling and analysis of conflicting information. A common theme among some of these proposals is the observation that not all arguments are equal, and that the relative strength of the arguments needs to be taken into account somehow.

The first such extension of Dung's work that we are aware of is [27], where priorities between rules are used to resolve conflicts ([16] was not based on Dung). These priorities seem best interpreted as relating to the strength of the arguments — indeed the strength of arguments are inferred from the strengths of the rules from which the arguments are constructed. A similar notion is at the heart of the argumentation systems in [1, 2], though here there is a preference order over all an agent's beliefs, and an argument has a preference level equal to the minimum level of the beliefs from which it is constructed.

Another early development of Dung's proposal with weights was Value-based Argumentation Frameworks (VAFs) [5]. In the VAF approach, the strength of an argument depends on the social values that it advances, and determining whether the attack of one argument on another succeeds depends on the comparative strength of the values advanced by the arguments concerned. Furthermore, some arguments can be shown to be acceptable whatever the relative strengths of the values involved are. This means that the agents involved in the argumentation can concur on the acceptance of arguments, even when they differ as to which social values are more important. One of the interesting questions that arise from this proposal is whether the notion of argument strength can be generalised from representing social values to representing other notions, and if so in what ways can the strength be harnessed for analysing argument graphs.

In a sense, a more general approach to developing Dung's proposal is that of bipolar argumentation frameworks (BAFs) which takes into account two kinds of interaction between arguments: a positive interaction (an argument can help, support another argument) and a negative interaction (an argument can attack another argument) [10]. The BAF approach incorporates a gradual interaction-based valuation process in which the value of each argument α only depends on the value of the arguments which are directly interacting with α in the argumentation system. Various functions for this process are considered but each value obtained is only a function of the original graph. As a result, no extra information is made available with which to ascertain the strength of an argument.

Recently, a game-theoretic approach, based on the minimax theorem, has been developed for determining the degree to which an argument is acceptable given the counterarguments to it, and by recursion the counterarguments to the counterarguments [19]. So given an abstract argument system, this game-theoretic approach calculates the strength of each argument in such a way that if an argument is attacked, then its strength falls, but if the attack is in turn attacked, then the strength in the original argument rises. Furthermore, the process for this conforms to an interpretation of game theory for argumentation. Whilst this gives an approach with interesting properties, and appealing behaviour, the strength that is calculated is a function of the original graph, and so like the BAF approach, no extra information is made available with which to determine the strength of each argument.

In another recent proposal for developing Dung's model, extra information representing the relative strength of attack is incorporated [18]. This is the only other approach that we are aware of which distinguishes the *strength of attack* from the strength of an argument. In this proposal, which we refer to as varied-strength attacks (or VSA) approach, each arc is assigned a type, and there is a partial ordering over the types. As a simple example, consider the following argument graph conforming to Dung's proposal, where α_1 is attacked by α_2 which in turn is attacked by α_3.

$$\alpha_3 \rightarrow \alpha_2 \rightarrow \alpha_1$$

Here, α_3 defends the attack on α_1, and as a result $\{\alpha_3, \alpha_1\}$ is the preferred, grounded and complete extension. Now, consider the following VSA version of the graph, where the attack by α_3 is of type i and the attack by α_2 is of type j.

$$\alpha_3 \rightarrow_i \alpha_2 \rightarrow_j \alpha_1$$

This gives us a finer grained range of defence depending on whether type j is higher, or lower, or equally, ranked than type i, or incomparable with it. Furthermore, this allows for a finer definition of acceptable extension that specifies the required level of the defence of any argument in the extension. For instance, it can be insisted in the VSA approach that every defence of an argument should be by an attack that is stronger, so in the above graph that would mean that the type of $\rightarrow_i$ needs to be stronger than the type of $\rightarrow_j$ in order for $\{\alpha_3, \alpha_1\}$ to be the preferred, grounded and complete extension.

From these proposals for developing Dung's original approach, there is a common theme that arguments, or attacks by arguments, have variable strength. Some of these proposals are restricted to determining that strength is based on the other arguments available in the graph, together with their connectivity, and so the strength of an argument is a function solely of the graph. Others, in particular the VAF approach [5] and the VSA approach [18], use explicit ranking information over the arguments or the attacks by arguments. This ranking information requires extra information to be given along with the set of arguments and the attack relation. So, whilst there is gathering momentum for representing and reasoning with the strength of arguments or their attacks, there is not a consensus on the exact notion of argument strength or how it should be used. Furthermore, for the explicit representation of extra information pertaining to argument strength, we see that the use of explicit numerical weights is under-developed. So for these reasons, we would like to present weighted argument systems as a valuable new proposal that should further extend and clarify aspects of this trend towards considering strength, in particular the explicit consideration of strength of attack between arguments..

4. WEIGHTED ARGUMENT SYSTEMS

We now introduce our model of weighted argument systems, and the key solutions we use throughout the remainder of the paper. Weighted argument systems extend Dung-style abstract argument systems by adding numeric weights to every edge in the attack graph, intuitively corresponding to the strength of the attack, or equivalently, how reluctant we would be to disregard it. Formally, a *weighted argument system* is a triple $W = \langle X, A, w \rangle$ where $\langle X, A \rangle$ is a Dung-style abstract argument system, and $w : A \rightarrow \mathbb{R}_>$ is a function assigning real valued weights[1] to attacks. If X is a set of arguments, then we let $\mathcal{W}(X)$ denote the set of weighted argument systems over X. (In what follows, when we say simply "argument system", we mean "Dung-style (unweighted) abstract argument system".)

Notice that we require attacks to have a positive *non-zero* weight. There may be cases where it is interesting to allow zero-weight

[1] We let $\mathbb{R}_>$ denote the real numbers greater than 0, and $\mathbb{R}_\geq$ denote the real numbers greater than or equal to 0.

attacks, in which case some of the analysis of this paper does not go through. However, given our intuitive reading of weights (that they indicate the strength of an attack) allowing 0-weight attacks is perhaps counter-intuitive. For suppose by appealing to a particular 0-weight attack you were able to support some particular argument, then an opponent could discard the attack *at no cost*. So, we will assume attacks must have non-zero weight.

4.1 Where do Weights Come From?

We will not demand any specific interpretation of weights, and the technical treatment of weighted argument systems in the remainder of the paper does not require any such interpretation. However, from the point of view of motivation, it is important to consider this issue seriously (if only to convince the reader that weights are not a purely technical device). Note that these three examples do not exhaust the possibilities for the meaning of weights on attacks.

Weighted Majority Relations: In a multi-agent setting, one natural interpretation is that a weight represents *the number of votes in support of the attack*. This interpretation makes a link between argumentation and *social choice theory* – the theory of voting systems and collective decision making [3, 28].

Weights as Beliefs: Another interpretation would be to interpret weights as subjective beliefs. For example, a weight of $p \in (0, 1]$ on the attack of argument α_1 on argument α_2 might be understood as the belief that (a decision-maker considers) α_2 is false when α_1 is true. This belief could be modelled using probability, or any other model of belief [24].

Weights as Ranking: A simple and obvious interpretation is to use weights to rank the relative strength of attacks between arguments. In other words, a higher weight denotes a stronger attack, and so the absolute weight assigned to an attack is not important, just the relative weight compared to the weights assigned to other attacks. In this interpretation, we can consider subjective or objective criteria for ranking attacks. For instance, in the earlier example concerning arguments about the potential carcinogenicity of chemicals, arguments based on human epidemiological evidence are more compelling (at least to the USA EPA) than those based on animal studies, which are in turn more compelling than those based on bioassay evidence [32]. We might assign a weight of (say) 100 to an attack between two arguments which are both based on the same type of evidence, i.e., both human epidemiological studies, or both animal studies, or both bioassays. In the case where the attacking argument is based on human epidemiological studies and the attacked argument on animal studies, we may assign a weight of 125. In the case where the attacking argument is based on human epidemiological studies and the attacked argument on bioassay experiments, we may assign a weight of 150. For attacks between two such arguments in the reverse directions, we could assign weights of 75 and 50 (respectively). As mentioned, the absolute numbers here are not important; rather the weights are aiming to capture the relative degree of persuasive compulsion which a decision-maker believes when considering each type of attack. Clearly this interpretation has scope for a more finely-grained allocation of weights, for example to distinguish between attacks by arguments based on studies of different species of animals, or by arguments based on experimental studies with different levels of statistical power.

4.2 Inconsistency Budgets and β-Solutions

A key idea in what follows is that of an *inconsistency budget*, $\beta \in \mathbb{R}_{\geq}$, which we use to characterise *how much inconsistency we are prepared to tolerate*. The intended interpretation is that, given an inconsistency budget β, we would be prepared to *disregard attacks up to a total weight of* β. Conventional abstract argument systems implicitly assume an inconsistency budget of 0. However, by relaxing this constraint, allowing larger inconsistency budgets, we can obtain progressively more solutions from an argument system.

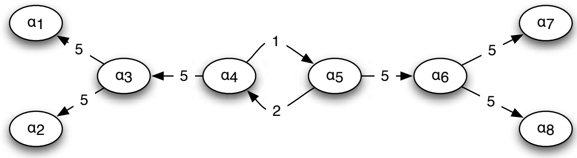

Figure 2: Weighted argument system W_1 from Example 1.

To make this idea formal, we first define a function $sub(\cdots)$, which takes an attack relation A, weight function $w : A \to \mathbb{R}_{>}$, and inconsistency budget $\beta \in \mathbb{R}_{\geq}$, and returns the set of sub-graphs R of A such that the edges in R sum to no more than β:

$$sub(A, w, \beta) = \{R : R \subseteq A \;\&\; \sum_{e \in R} w(e) \leq \beta\}.$$

We now use inconsistency budgets to introduce weighted variants of the solutions introduced for abstract argument systems, above. Given a weighted argument system $\langle X, A, w\rangle$, a solution $f : \mathcal{D}(X) \to \mathcal{P}(\mathcal{P}(X))$, and a set of arguments $S \subseteq X$, we say that S is β-f if $\exists R \in sub(A, w, \beta)$ such that $S \in f(\langle X, A \setminus R\rangle)$. So, for example, S is *β-admissible* if $\exists R \in sub(A, w, \beta)$ such that S is admissible in the argument system $\langle X, A \setminus R\rangle$.

EXAMPLE 1. *Consider the weighted argument system W_1, illustrated in Figure 2. The only consistent set of arguments in W_1 is the empty set; however, $\{\alpha_5\}$ is 1-consistent, since we can delete the edge (α_4, α_5) with $\beta = 1$. If $\beta = 2$, we have two consistent sets: $\{\alpha_4\}$ and $\{\alpha_5\}$. Table 1 shows consistent sets (and other β-solutions) for some increasing values of β.*

Now, weighted argument systems straightforwardly generalise unweighted argument systems: each unweighted solution f is directly realised by the weighted solution 0-f. However, weighted solutions have a number of advantages over unweighted solutions. Consider for example the notion of consistency. We know that in unweighted systems, there is always a consistent set, but this could be empty. As we noted above, this may be undesirable – if an argument system only has a trivial solution, then we obtain no information from it. In contrast, weighted argument systems have the following, (readily proved), property:

PROPOSITION 1. *Let $W = \langle X, A, w\rangle$ be a weighted abstract argument system. For every set of arguments $S \subseteq X$, $\exists\beta$ such that S is contained in a β-consistent set in W.*

Thus, intuitively, every set of arguments is consistent at some cost, and the cost required to make a set of arguments consistent immediately gives us a preference ordering over sets of arguments: we prefer sets of arguments that require a smaller inconsistency budget. Notice that a similar observation holds true for admissibility, preferred extensions, credulous acceptance, and sceptical acceptance.

Now, consider how grounded extensions are generalised within weighted systems. The first observation to make is that while in unweighted argument systems the grounded extension is unique, this will not necessarily be the case in weighted argument systems: *in*

$\beta =?$	β-consistent sets	β-preferred extensions	β-grounded extensions
0	$\{\emptyset\}$	$\{\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\},\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\}\}$	$\{\emptyset\}$
1	$\{\emptyset,\{\alpha_5\}\}$	$\{\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\},\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\}\}$	$\{\emptyset,\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\}\}$
2	$\{\emptyset,\{\alpha_4\},\{\alpha_5\}\}$	$\{\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\},\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\}\}$	$\{\emptyset,\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\},\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\}\}$
3	$\{\emptyset,\{\alpha_4\},\{\alpha_5\},\{\alpha_4,\alpha_5\}\}$	$\{\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\},\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\},$ $\{\alpha_1,\alpha_2,\alpha_4,\alpha_5,\alpha_7,\alpha_8\}\}$	$\{\emptyset,\{\alpha_3,\alpha_5,\alpha_7,\alpha_8\},\{\alpha_1,\alpha_2,\alpha_4,\alpha_6\},$ $\{\alpha_1,\alpha_2,\alpha_4,\alpha_5,\alpha_7,\alpha_8\}\}$

Table 1: Solutions for W_1, for some increasing values of β.

weighted systems there may be many β-grounded extensions. Formally, let $wge(X, A, w, \beta)$ denote the set of β-grounded extensions of the weighted argument system $\langle X, A, w\rangle$ (recall that the function $ge(\cdots)$, which computes the unweighted grounded extension, is defined in Figure 1):

$$wge(X, A, w, \beta) = \{ge(X, A \setminus R) : R \in sub(A, w, \beta)\}.$$

Table 1 shows β-grounded extensions for some increasing values of β for system W_1 of Figure 2.

We conclude this section with another possible interpretation for weights, and an associated example.

EXAMPLE 2. *Suppose we interpret the weight on an edge (α_i, α_j) as a* costed risk. *By this, we mean that the weight of (α_i, α_j) is the cost/penalty that is incurred if α_i is true, normalized by the probability that α_i actually is true. To illustrate, consider the following arguments where α_2 attacks α_1, α_3 attacks α_2, and α_4 attacks α_2.*

(α_1) *The patient needs bypass surgery now*
(α_2) *The patient will die in theatre*
(α_3) *The patient will die within a week without surgery*
(α_4) *The patient will have impaired heart functionality*

Assume a probability function p over arguments, so $p(\alpha)$ is the probability that α is true. Now, suppose p is such that $p(\alpha_2) = 0.5$, $p(\alpha_3) = 0.9$, and $p(\alpha_4) = 1$. Let the penalty of α_2 (respectively α_3 and α_4) being true be 100 (resp. 99.9 and 5). Then $w(\alpha_2, \alpha_1) = 50$, $w(\alpha_3, \alpha_2) = 89.9$, and $w(\alpha_4, \alpha_2) = 5$. For all $\beta < 94.9$, α_1 is in every β-grounded extension. This seems reasonable, since α_3 has a sufficiently high penalty and probability of occurrence to defeat α_2 hence allow α_1 to be undefeated.

Now, let us change α_2 to α'_2 and α_3 to α'_3, with p giving $p(\alpha'_2) = 0.9$ and $p(\alpha'_3) = 0.1$, and let the penalty of α'_2 be the same as α_2 and the penalty of α'_3 be the same as α_3. Then $w(\alpha'_2, \alpha_1) = 90$, and $w(\alpha'_3, \alpha_2) = 10$, and hence, for any $\beta \geq 15$, α_1 there is some β-grounded extension not containing α_1. This also is reasonable, since if we are prepared to overlook some costed risk, then we are safe against the much greater costed risk that comes from α_2. In a sense, via inconsistency tolerance, we are trading one costed risk against another.

From this example, we can see how the uncertainty and potential negative ramifications of counterarguments can be intuitively captured using weighted argument systems.

5. COMPLEXITY OF SOLUTIONS

An obvious question now arises. *Prima facie*, it appears that weighted argument systems offer some additional expressive power over unweighted argument systems. So, does this apparently additional power come with some additional computational cost? The β versions of the decision problems for consistency, admissibility, checking preferred extensions, sceptical, and credulous acceptance are in fact no harder (although of course no easier) than the corresponding unweighted decision problems – these results are easy to establish. However, the story for β-grounded extensions is more complicated, since there may be *multiple* β-grounded extensions. Since there are multiple β-grounded extensions, we can consider credulous and sceptical variations of the problem, as with preferred extensions. Consider the credulous case first:

PROPOSITION 2. *Given weighted argument system $\langle X, A, w\rangle$, inconsistency budget β, and argument $\alpha \in X$, the problem of checking whether $\exists S \in wge(X, A, w, \beta)$ such that $\alpha \in S$ is* NP*-complete. The problem remains* NP*-complete even if the attack relation is planar and/or tripartite and/or has no argument which is attacked by more than two others.*

PROOF. For membership, a conventional "guess and check" approach suffices. For NP-hardness, we reduce from 3-SAT. Given an instance $\varphi(Z_n)$ of 3-SAT with m clauses C_j over propositional variables $Z_n = \{z_1, \ldots, z_n\}$, form the weighted argument system $\langle X_\varphi, A_\varphi, w_\varphi\rangle$, illustrated in Figure 4. Specifically, X_φ has $3n + m + 1$ arguments: an argument C_j for each clause of $\varphi(Z_n)$; arguments $\{z_i, \neg z_i, u_i\}$ for each variable of Z_n, and an argument φ. The relationship, A_φ, contains attacks (C_j, φ) for each clause of φ, $(z_i, \neg z_i)$, $(\neg z_i, z_i)$, (z_i, u_i), $(\neg z_i, u_i)$, and (u_i, φ) for each $1 \leq i \leq n$. Finally, A_φ contains an attack (z_i, C_j) if z_i is a literal in C_j, and $(\neg z_i, C_j)$ if $\neg z_i$ occurs in C_j. The weighting function w_φ assigns cost 1 to each of the attacks $\{(z_i, \neg z_i), (\neg z_i, z_i)\}$ and cost $n + 1$ to all remaining attacks. To complete the instance the available budget is set to n and the argument of interest to φ. We claim that $\varphi \in S$ for some $S \in wge(X_\varphi, A_\varphi, w_\varphi, n)$ if and only if $\varphi(Z_n)$ is satisfiable. We first note that φ is credulously accepted in the (unweighted) system $\langle X_\varphi, A_\varphi\rangle$ if and only if $\varphi(Z_n)$ is satisfiable.[2] We deduce that if $\varphi(Z_n)$ is satisfied by an instantiation $\langle a_1, a_2, \ldots, a_n\rangle$ of Z_n then φ is a member of the *grounded* extension of the *acyclic* system $\langle X_\varphi, A_\varphi \setminus B\rangle$ in which B contains $(\neg z_i, z_i)$ (if $a_i = \top$) and $(z_i, \neg z_i)$ (if $a_i = \bot$). Noting that B has total weight n, and that the subset $\{y_1, y_2, \ldots, y_n\}$ in which $y_i = z_i$ (if $a_i = \top$) and $\neg z_i$ (if $a_i = \bot$) is unattacked, it follows that from $\varphi(Z_n)$ satisfiable we may identify a suitable cost n set of attacks, B, to yield $\varphi \in ge(X_\varphi, A_\varphi \setminus B)$

On the other hand, suppose that $\varphi \in S$ for some S belonging to $wge(X_\varphi, A_\varphi, w_\varphi, n)$. Consider the set of attacks, B, eliminated from A_φ in order to form the system $\langle X_\varphi, A_\varphi \setminus B_\varphi\rangle$ with grounded extension S. Since $\varphi \in S$, exactly one of $(z_i, \neg z_i)$ and $(\neg z_i, z_i)$ must be in B for *every* i. Otherwise, if for some i, neither attack is in B then $\{z_i, \neg z_i\} \cap S = \emptyset$, and thus φ has no defence to the attack by u_i, contradicting $\varphi \in S$; similarly if *both* attacks are in B then, from the fact B has total cost at most n, for some other variable, z_k, both $(z_k, \neg z_k)$ and $(\neg z_k, z_k)$ would be in $A_\varphi \setminus B$. In total, from S in $wge(X_\varphi, A_\varphi, w_\varphi, n)$ and $\varphi \in S$ for each $1 \leq i \leq n$ we identify exactly one unattacked argument, y_i from $\{z_i, \neg z_i\}$, so that $S = \{\varphi, y_1, \ldots, y_n\}$. That the instantiation $z_i = \top$ (if

[2] This follows from [11] which uses a similar construction for which the u_i arguments and associated attacks do not occur.

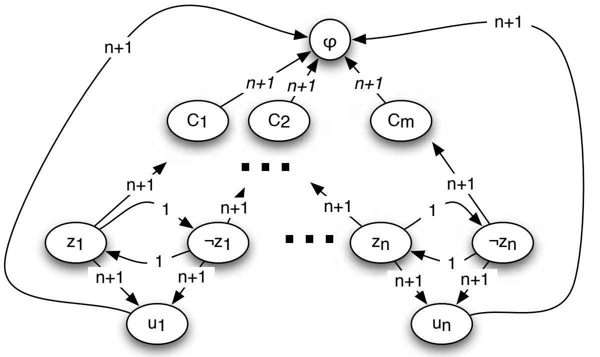

Figure 3: The reduction used in Proposition 2.

$y_i = z_i$) and $z_i = \bot$ (if $y = \neg z_i$) satisfies $\varphi(Z_n)$ is immediate from [11].

The remaining cases (for planar, tripartite graphs, etc.) can be derived from the reductions from 3-SAT given in [14]. □

Now consider the "sceptical" version of this problem.

PROPOSITION 3. *Given weighted argument system $\langle X, A, w\rangle$, inconsistency budget β, and argument $\alpha \in X$, the problem of checking whether, $\forall Y \in wge(X, A, w, \beta)$, we have $\alpha \in Y$ is co-NP-complete.*

PROOF. Membership of co-NP is immediate from the algorithm which checks for every $B \subseteq A$ that if $\sum_{e \in B} w(e) \leq \beta$ then $x \in ge(X, A \setminus B)$. For co-NP-hardness, we use a reduction from UNSAT, assuming w.l.o.g. that the problem instance is presented in CNF. Given an m-clause instance $\varphi(Z_n)$ of UNSAT, we construct a weighted argument system $\langle X_\varphi, A_\varphi, w_\varphi\rangle$ as follows. The set X_φ contains $4n+m+3$ arguments: $\{\varphi, \psi, \chi\}$; $\{z_i,\ \neg z_i, u_i, v_i\ :\ 1 \leq i \leq n\}$; and $\{C_j\ :\ 1 \leq j \leq m\}$. The attack set A_φ comprises: $\{(\varphi, \psi),\ (\chi, \varphi)\}$; $\{(v_i, z_i),\ (v_i, \neg z_i),\ (z_i, u_i),\ (\neg z_i, u_i),\ (u_i, \varphi)\}$ for each $1 \leq i \leq n$; $\{(C_j, \varphi)\ :\ 1 \leq j \leq m\}$; $\{(z_i, C_j)\ :\ z_i \in C_j\}$ and $\{(\neg z_i, C_j)\ :\ \neg z_i \in C_j\}$ The attacks are weighted so that $w_\varphi((\chi, \varphi)) = 1$; $w_\varphi((v_i, z_i)) = w_\varphi((v_i, \neg z_i)) = 1$; all remaining attacks have weight $n+2$. The instance is completed using ψ as the relevant argument and an inconsistency tolerance of $n+1$. (See Figure 4 for an illustration of the construction.)

Now, suppose that $\varphi(Z_n)$ is satisfied by an instantiation $\alpha = \langle a_1, \ldots, a_n\rangle$ of Z_n. Consider the subset B_α of A_φ given by $\{(\chi, \varphi)\}$ together with $\cup\{(v_i, z_i)\ :\ a_i = \top\} \cup \{(v_i, \neg z_i)\ :\ a_i = \bot\}$. The weight of B_α is $n+1$ and (since α satisfies $\varphi(Z_n)$ it follows that $ge(X_\varphi, A_\varphi \setminus B_\alpha)$ contains exactly the arguments $\{\chi, \varphi\} \cup \{v_1, \ldots, v_n\} \cup \{z_i\ :\ a_i = \top\} \cup \{\neg z_i\ :\ a_i = \bot\}$. Hence $\psi \notin ge(X_\varphi, A_\varphi \setminus B_\alpha)$ as required.

Conversely, suppose $\langle\langle X_\varphi, A_\varphi, w\rangle, \psi, n+1\rangle$ is *not* accepted. We show that we may construct a satisfying instantiation of $\varphi(Z_n)$ in such cases. Consider $B \subseteq A_\varphi$ of cost at most $n+1$ for which $\psi \notin ge(X_\varphi, A_\varphi \setminus B)$. It must be the case that $(\chi, \varphi) \in B$ for otherwise the attack by φ on ψ is defended so that ψ would belong to the grounded extension. The remaining elements of B must form a subset of the attacks $\{(v_i, z_i),\ (v_i, \neg z_i)\}$ (since all remaining attacks are too costly). Furthermore, exactly one of $\{(v_i, z_i),\ (v_i, \neg z_i)\}$ must belong to B for each $1 \leq i \leq n$: otherwise, some u_i will be in $ge(X_\varphi, A_\varphi \setminus B)$, thus providing a defence to the attack on ψ by φ and contradicting the assumption

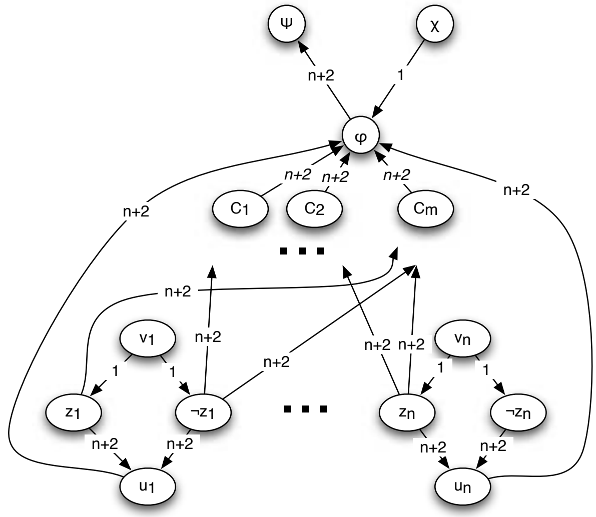

Figure 4: The reduction used in Proposition 3.

$\psi \notin ge(X_\varphi, A_\varphi \setminus B)$. Now consider the instantiation, α, with $a_i = \top$ if $(v_i, z_i) \in B$, $a_i = \bot$ if $(v_i, \neg z_i) \in B$. We now see that α must satisfy $\varphi(Z_n)$: in order for $\psi \notin ge(X_\varphi, A_\varphi \setminus B)$ to hold, it must be the case that $\varphi \in ge((X_\varphi, A_\varphi \setminus B)$, i.e. each of the C_j attacks on φ must be counterattacked by one of its constituent literal (arguments) y_i. Noting that v_i is always in $ge(X_\varphi, A_\varphi \setminus B)$, if $a_i = \top$ clauses containing $\neg z_i$ cannot be attacked (since the attack $(v_i, \neg z_i)$ is still present). It follows that the instantiation, α, attacks each clause so that $\varphi \in ge(X_\varphi, A_\varphi \setminus B)$. In sum, if $\langle\langle X_\varphi, A_\varphi, w\rangle, \psi, n+1\rangle$ is not accepted then $\varphi(Z_n)$ is satisfiable, so completing the proof. □

Note that in some cases, considering sceptical grounded extensions is of limited value. Let $unch(X, A)$ denote the set of arguments in X that are unchallenged (have no attackers) according to A. Then we have:

PROPOSITION 4. *Let $\langle X, A, w\rangle$ be a weighted argument system and β be an inconsistency budget. Then $unch(X, A) \neq \emptyset$ iff $(\bigcap_{Y \in wge(X,A,w,\beta)} Y) \neq \emptyset$.*

6. HOW MUCH INCONSISTENCY DO WE NEED?

Another obvious question is now raised. Suppose we have a weighted argument system $\langle X, A, w\rangle$ and a set of arguments S. Then *what is the smallest amount of inconsistency would we need to tolerate in order to make S a solution*? Now, when considering consistency and admissibility, the answer is easy: we know exactly which attacks we would have to disregard to make a set of arguments admissible or consistent — we have no choice in the matter. However, when considering grounded extensions, the answer is not so easy. As we saw above, there may be multiple ways of getting a set of arguments into a weighted extension, each with potentially different costs; we are thus typically interested in solving the problem:

$$\text{minimise } \beta^* \text{ s.t.} \qquad \exists Y \in wge(X, A, w, \beta^*) : S \subseteq Y \qquad (1)$$

What can we say about (1)? First, consider the following problem. We are given a weighted argument system $\langle X, A, w\rangle$ and an inconsistency budget $\beta \in \mathbb{R}_{\geq}$, and asked whether β is *minimal*, i.e., whether $\forall \beta' < \beta$ and $\forall Y \in wge(X, A, w, \beta')$, we have that $S \not\subseteq Y$. (This problem does *not* require that S is contained in a some β-grounded extension of $\langle X, A, w\rangle$.)

PROPOSITION 5. *Given a weighted argument system $\langle X, A, w\rangle$, set of arguments $S \subseteq X$, and inconsistency budget β, checking whether β is minimal w.r.t. $\langle X, A, w\rangle$ and S is co-NP-complete.*

PROOF. Consider the complement problem, i.e., the problem of checking whether $\exists \beta' < \beta$ and $\exists Y \in wge(X, A, w, \beta')$ such that $S \subseteq Y$. Membership in NP is immediate. For NP-hardness, we can reduce SAT, using essentially the same construction for the weighted argument system as Proposition 2; we ask whether $n + 1$ is not minimal for argument set $\{\varphi\}$. □

This leads very naturally to the following question: is β the *smallest* inconsistency budget required to ensure that S is contained in some β-grounded extension. We refer to this problem as *checking whether β is the minimal budget for S*.

PROPOSITION 6. *Given a weighted argument system $\langle X, A, w\rangle$, set of arguments $S \subseteq X$, and inconsistency budget β, checking whether β is the minimal budget for S is D^p-complete.*

PROOF. For membership of D^p, we must exhibit two languages L_1 and L_2 such that $L_1 \in$ NP, $L_2 \in$ co-NP, and $L_1 \cap L_2$ is the set of instances accepted by the minimal budget problem. Language L_1 is given by Proposition 2, while language L_2 is given by Proposition 5. For hardness, we reduce the Critical Variable Problem (CVP) [7, p.66]. An instance of CVP is given by a propositional formula φ in CNF, and a variable z from φ. We are asked if, under the valuation $z = \top$ the formula φ is satisfiable, while under the valuation $z = \perp$ it is unsatisfiable. We proceed to create in instance of the minimal budget problem by using essentially the same construction as Proposition 2, except that the attack $(z, \neg z)$ is given a weight of 0.5. Now, in the resulting system, n is the minimal budget for $\{\varphi\}$ iff z is a critical variable in φ. □

We noted above that one way of deriving a preference order over sets of arguments is to consider the minimal inconsistency budget required to make a set of arguments a solution. A related idea is to *count* the number of weighted extensions that an argument set appears in, for a given budget: we prefer argument sets that appear in more weighted grounded extensions. Formally, we denote the *rank* of an argument set S, given a weighted argument system $\langle X, A, w\rangle$ and inconsistency budget β, by $\rho(S, X, A, w, \beta)$:

$$\rho(S, X, A, w, \beta) = |\{Y \in wge(X, A, w, \beta) : S \subseteq Y\}|.$$

PROPOSITION 7. *Given weighted argument system $\langle X, A, w\rangle$, argument set $S \subseteq X$, and inconsistency budget β, computing $\rho(S, X, A, w, \beta)$ is #P-complete.*

PROOF. (Outline) For membership, consider a non-deterministic Turing machine that guesses some subset R of A, and verifies that both $\sum_{e \in R} w(e) \leq \beta$ and $S \subseteq ge(X, A \setminus R)$. The number of accepting computations of this machine will be $\rho(S, X, A, w, \beta)$. For hardness, we can reduce #SAT [23, p.439], using the construction of Proposition 2. It is straightforward to see that the reduction is parsimonious. □

7. RELATED WORK

We have already described some of the work that is most closely related to ours in the brief survey of Section 3 but there is additional work that should be mentioned and which does not fit into the broad historical sweep we were describing there.

To begin, there are other interesting developments of abstract argumentation such as a framework for defeasible reasoning about preferences that provides a context dependent mechanism for determining which argument is preferred to which [20, 21]. This also offers a valuable solution to dealing with multiple extensions, but conceptually and formally the proposal is complementary to ours. Also of interest are the proposals for introducing information about how the audience views each argument [6].

The framework we present is also clearly related to preference-based argument systems such as that described in [1]. However, while our approach disregards attacks whose combined weight is less than the inconsistency budget, systems such as that in [1] disregard all attacks whose individual weight is below that of the argument being attacked. This is broadly equivalent, in our terms, to setting the inconsistency budget to the weight of the argument being attacked, and taking the combined weights of the attacking arguments to be the maximum of the weights rather than the sum. Our work is also related to work on possibilistic truth-maintenance systems [12] where assumptions are weighted, conclusions based on the assumptions inherit the weights, and consistent "environments" are established. What is particularly reminiscent about the work in [12] is that, again in our terms, it makes use of inconsistency budget — this is exactly the weight with which the inconsistency $\perp$ can be inferred. Anything that can be inferred with a greater weight than $\perp$ is then taken to hold, anything with a lesser weight is taken to be unreliable, which is broadly the effect of the inconsistency budget in our work.

Finally, we should point out that there has been a good deal of work on incorporating numerical and non-numerical strengths (though not strengths of attack) into argumentation systems that are not based on Dung's work. [16], to take the earliest example, describes the use of probability measures and beliefs in the sense of Shafer's theory of evidence [29]. [25] presents an argumentation system that uses weights which are qualitative abstractions of probability values, while in [31] the weights are infinitesimal probabilities in the sense of [30]. There is also much work on combinations of logic and probability such as [4], [22] and [26], which, while they don't explictly take the form of argumentation, have much in common with it.

8. DISCUSSION AND CONCLUSIONS

Our proposal in this paper, namely weighted argument systems (WAS), is a further contribution to the development of formalisms for abstract argumentation that started with the seminal work by Dung. The WAS approach uses a numerical weight on the attacks between arguments, as do the proposals based on game theory [19] and bipolar argumentation [10], but those proposals are restricted to determining the strength based on the other arguments available in the graph, together with connectively, and so the strength of an argument is a function solely of the graph. In contrast, our proposal allows for the weight to be given as an extra piece of information. There are other proposals that allow for extra information to be given about the strength of arguments in a constellation, in particular the VAF approach [5] and the VSA approach [18], but they are restricted to using explicit ranking information over the arguments or the attacks by arguments, rather than numerical information. By introducing numerical weights, we can simplify and generalize some of the underlying conceptualization of handling the strength of attackss, and furthermore, we can introduce the interesting and potentially valuable idea of inconsistency budgets for finer grained analysis of inconsistent information.

Several possibilities suggest themselves for future research: to investigate specific interpretations for weights; another is to investigate the framework experimentally, to obtain a better understanding of the way the approach behaves. One obvious issue here is to look for "discontinuities" as the inconsistency budget grows, i.e. points where large increases in the number of accepted arguments occur for only a small increase in the inconsistency budget. A third avenue is to investigate the question of the exact relationship between the argument strength and the strength of attacks.

9. ACKNOWLEDGMENTS

This work was partially supported by the US Army Research Laboratory (USARL) and the UK Ministry of Defence (MoD) under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the USARL, the US Government, the UK MoD, or the UK Government. The US and UK Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

10. REFERENCES

[1] L. Amgoud and C. Cayrol. A model of reasoning based on the production of acceptable arguments. *Annals of Mathematics and Artificial Intelligence*, 34:197–215, 2002.

[2] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proc. 4th International Conf. on Multi-Agent Systems*, pages 31–38, Boston, MA, USA, 2000. IEEE Press.

[3] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare Volume 1*. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 2002.

[4] F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87:75–143, 1996.

[5] T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *J. of Logic and Computation*, 13(3):429–448, 2003.

[6] T. Bench-Capon, S. Doutre, and P. E. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.

[7] T. J.M. Bench-Capon, S. Doutre, and P. E. Dunne. Audiences in argumentation frameworks. *Artificial Intelligence*, 171(1):42–71, 2007.

[8] L. Bertossi, A. Hunter, and T. Schaub, editors. *Inconsistency Tolerance* volume 3300 of *LNCS* Springer-Verlag, 2004.

[9] Ph. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.

[10] C. Cayrol and M.-C. Lagasquie-Schiex. Gradual valuation for bipolar argumentation frameworks. In *Proc. 8th ECSQARU*, volume 3571 of *LNCS*, pages 366–377. Springer-Verlag, 2005.

[11] Y. Dimopolous and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.

[12] D. Dubois, J. Lang, and H. Prade. A possibilistic truth-maintenance system with uncertain justifications, and its application to belief revision. In J. P. Martins and M. Reinfrank, editors, *Truth Maintenance Systems*, volume 515 of *LNCS*, pages 87–106. Springer-Verlag, 1991.

[13] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

[14] P. E. Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artificial Intelligence*, 171:701–729, 2007.

[15] P. E. Dunne and T. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141:187–203, 2002.

[16] P. Krause, S. Ambler, M. Elvang-Gørannson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11 (1):113–131, 1995.

[17] P. Krause, J. Fox, P. Judson, and M. Patel. Qualitative risk assessment fulfils a need. In *Applications of Uncertainty Formalisms*, volume 1455 of *LNCS*, pages 138–156. Springer, 1998.

[18] D. Martinez, A. Garcia, and G. Simari. An abstract argumentation framework with varied-strength attacks. In *Proc. KR'08*, 2008.

[19] P. Matt and F. Toni. A game-theoretic measure of argument strength for abstract argumentation. In *Proc. JELIA'08* volume 5293 of *LNAI*, pages 285–297. Springer, 2008.

[20] S. Modgil. An abstract theory of argumentation that accommodates defeasible reasoning about preferences. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, volume 4724 of *LNCS*. Springer, 2007.

[21] S. Modgil and T. Bench-Capon. Integrating object and meta-level value based argumentation. In *Proc. COMMA 2008*, pages 240–251. IOS Press, 2008.

[22] E. Neufeld and D. Poole. Towards solving the mutiple extension problem: combining defaults and probability. In L. N. Kanal, T. S. Levitt, and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence 3*, pages 35–44. Elsevier Science Publishers, Amsterdam, The Netherlands, 1989.

[23] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley: Reading, MA, 1994.

[24] S. Parsons. *Qualitative methods for reasoning under uncertainty*. MIT Press, Cambridge, MA, 2001.

[25] S. Parsons. On precise and correct qualitative probabilistic reasoning. *International Journal of Approximate Reasoning*, 35:111–135, 2004.

[26] D. Poole. Probabilistic horn abduction and Bayesian networks. *Artificial Intelligence*, 64:81–129, 1993.

[27] H. Prakken and G. Sartor. Argument-based logic programming with defeasible priorities. *Journal of Applied Non-classical Logics*, 7:25–75, 1997.

[28] I. Rahwan and K. Larson. Pareto optimality in abstract argumentation. In *Proc. AAAI'08*. AAAI Press, 2008.

[29] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, NJ, 1976.

[30] W. Spohn. A general non-probabilistic theory of inductive reasoning. In R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, editors, *Uncertainty in Artificial Intelligence 4*, pages 149–158. Elsevier Science Publishers, Amsterdam, The Netherlands, 1990.

[31] V. Tamma and S. Parsons. Argumentation and qualitative reasoning with kappa calculus. In *Proc. 6th ECSQARU*, volume 2143 of *LNCS*, pages 680–691. Springer, 2001.

[32] U.S. Environmental Protection Agency. Guidelines for carcinogen risk assessment. *U.S. Federal Register*, 51:33991–34003, 24 September 1986.