Selecting Extensions in Weighted Argumentation Frameworks

Sylvie Coste-Marquis, Sébastien Konieczny, Pierre Marquis
and Mohand Akli Ouali
CRIL - CNRS, Université d'Artois, Lens, France

Abstract. Recently, Dunne et al. [9,10] introduced the concept of WAF (Weighted Argumentation Framework). Such frameworks extend standard Dung's ones for abstract argumentation by associating weights with attacks. In the WAF setting, weights are used for relaxing extensions, which proves useful when there are too few extensions. In this paper, we exploit weights in a different perspective. We show how to take advantage of attacks weights within an argumentation process for selecting some extensions among Dung's ones, which proves useful when there are too many extensions, in order to improve the inferential power of the argumentation framework.

Keywords. Argumentation. Weighted argumentation frameworks.

Introduction

Recently Dunne et al. [10] introduced the notion of weighted argumentation frameworks (WAFs), which extend Dung's argumentation frameworks by associating a weight (i.e., a positive real number) with each attack. Among other things, Dunne et al. [10] provided motivations for extending Dung's setting with such weights, explained how weights can be interpreted, and how they might be derived. They considered general algorithmic and combinatorial properties of the WAF setting, and showed that it is more expressive than four existing settings generalizing Dung's one.

There can be several interpretations of these weights:

Explicit strength of the attacks Weights can reflect the evaluation of the strengths of the attacks between arguments. In some applications, it proves sensible to split the attack relation into a set of different types of attacks. For instance weak attacks, medium ones, and strong ones, weak attacks being attacks that are not completely reliable, and strong ones being attacks that cannot be ignored. Different attack strengths can also be derived when attacks come from different sources, which can be more or less reliable or significant. For instance, attacks coming from legal laws must be considered as more important than attacks coming from some cultural habits. We term this as "explicit strength" since it requires some additional information which have to be given explicitly when the weighted argumentation framework is built.

Implicit strength of the attacks Often arguments have more structure than abstract ones in Dung's setting. Thus, when arguments are built from some logic, the notion of attack can be refined and different attack strengths can be defined by determining "how much" an argument attacks another one. This can be done for instance by evaluating how much the conclusion of an argument is inconsistent with the premisses of another one, using an inconsistency measure [11,12,13]. See e.g. [3] for discussions on the possible definitions of degree of undercuts. We call this case "implicit strength" since it does not depend on any additional information, but just on the structure of the arguments.

Votes on the attacks Perhaps the most natural interpretation of attack weights is when weights indicate how many agents from a given group agree with the attack under consideration. In this case a WAF can be considered as an aggregated view of the (classical) argumentation frameworks of the group of agents. This is a natural way to aggregate the points of view of a group of agents. Furthermore, under this interpretation, it also makes sense to ignore some attacks *a posteriori* (i.e., those which did not receive sufficient supports for being significant). This interpretation has already been evoked in [6].

Dunne et al. [10] used WAF for expanding Dung's extensions, which is particularly useful when extensions are trivial ones (i.e., empty sets) since in this case inference trivializes. The process goes through a relaxation of the usual notion of conflict-free sets of arguments: some inconsistencies are tolerated in sets S of arguments, provided that the sum of the weights of attacks between arguments of S does not exceed a given inconsistency budget β. Admissibility is defined in the standard way, and standard semantics are considered leading to various notions of β-extensions which echo Dung's ones (i.e., grounded, preferred, stable extensions are defined). See [7] for a generalization of this definition with other aggregation functions than sum.

In this paper, we exploit weights in a different perspective. Whereas the original use of WAF was when there are too few extensions, we use them now in the case where there are too many extensions. We show how to take advantage of attacks weights within an argumentation process for selecting some extensions among Dung's ones, which proves useful when there are too many extensions, or when extensions share too few arguments; in the latter case more precise skeptical conclusions can be drawn. Since weights can also be used to derive a refined notion of defence [7], they play with this respect a role similar to those achieved by a preference relation over arguments [1]. Note that [7] introduces a more general notion of refined defence than [5]. A comparison with other related works as [16,17] can also be found in [7].

The rest of the paper is organized as follows. In the following section, classical definitions from Dung's setting are first recalled. The second section indicates how weights can be used to select some extensions (for a given semantics). Section 3 explore how to define best defended extensions, using a balance between attacks on and attacks from an extension. There are two ways to achieve this, either globally or locally. Whereas in this setting the extensions are evaluated independently one another, Section 4 shows how to evaluate extensions by comparing them in a pairwise fashion. We first define when an extension is better than another one. Then it is possible to take advantage of methods coming from voting theory in order to define the best extensions. The last section concludes the paper and gives some perspectives for further research.

1. Preliminaries

Let us start by presenting some basic definitions at work in Dung's theory of abstract argumentation [8]. A (finite) argumentation framework is a pair $\text{AF} = \langle A, R\rangle$ where A is a (finite) set of so-called arguments and R is a binary relation over A (a subset of $A \times A$), the attack relation. An argument a is *acceptable* with respect to a set of arguments S whenever it is defended by the set, i.e., for every $b \in A$ s.t. $(b,a) \in R$, there exists $c \in S$ such that $(c,b) \in R$. We say that a subset S of A is *conflict-free* if and only if for every $a,b \in S$, we have $(a,b) \notin R$. A subset S of A is *admissible* for AF if and only if S is conflict-free and acceptable with respect to S.

"Solutions" of an argumentation systems are sets of arguments that can be accepted together given the attacks. This gives rise to the notion of extensions. Several definitions are possible. For instance:

- S is *a preferred extension* of AF if and only if it is maximal (with respect to set inclusion) among the set of admissible sets for AF .
- S is *a stable extension* of AF if and only if S is conflict-free and $\forall a \in A \setminus S$, $\exists b \in S$ such that $(b,a) \in R$.

Now given a set $\mathscr{E}_\sigma(\text{AF})$ of extensions for a given semantics σ, one has to make precise the arguments which can be inferred. This calls for an inference relation. In this work we focus on skeptical inference (w.r.t. a semantics σ): $\text{AF} \mid\!\sim^{\forall,\sigma} S$ if S is included into every σ-extension, i.e., $\text{AF} \mid\!\sim^{\forall,\sigma} S$ iff $\forall E \in \mathscr{E}_\sigma(\text{AF}),\ S \subseteq E$.

Let us now turn to the Weighted Argumentation Frameworks:

Definition 1 (weighted argumentation framework) *A* Weighted Argumentation Framework *(*WAF*) is a triple* $\text{WAF} = \langle A, R, w\rangle$ *where* $\langle A, R\rangle$ *is a Dung-style abstract argumentation framework, and* $w : A \times A \to \mathbb{N}$ *is a function assigning a natural number*[1] *to each attack (i.e.* $w(a,b) > 0$ *iff* $(a,b) \in R$*), and a null value otherwise (*$w(a,b) = 0$ *iff* $(a,b) \notin R$*).*

In [10] the weight function is defined as a real value function. In most situations we think that natural numbers are enough, and this simplifies some forthcoming definitions. Another difference is that in [10] the weight function is defined only for attacks (i.e., $w : R \to \mathbb{R}_*^+$). We extend its definition also to non-attacks, also for simplifying some incoming definitions. [10] discuss the possibility to assign a 0 weight to an attack, but conclude that *"0-weight attacks is perhaps counter-intuitive"*. We adhere to the view consisting in considering 0 weight for an ordered pair of arguments (a,b) as meaning "no attack" from a to b.

Let $\text{WAF} = \langle A, R, w\rangle$ be a weighted argumentation framework, we denote by $\widehat{\text{WAF}}$ the corresponding standard argumentation framework, obtained by forgetting the weights, i.e., $\widehat{\text{WAF}} = \langle A, R\rangle$.

[1] We let $\mathbb{N}$ denote the natural numbers greater than or equal to 0, $\mathbb{N}_*$ denote the natural numbers strictly greater than 0, and $\mathbb{R}_*^+$ denote the real numbers stricIty greater than 0.

2. Selecting Extensions

In the general case, an argumentation framework may admit a large number of extensions for some semantics (including the preferred one and the stable one). This leads to very weak skeptical inferences. Selecting some extensions, the best ones for some criterion, is thus a way to get more significant inferences.

Within the WAF setting, it is possible to take advantage of the available weights, in order to select the extensions which best defend themselves. This selection process goes through a comparison of the extension scores, expressing intuitively how strong they are. The computation of such scores requires attacks weights to be somehow aggregated, and this calls for an aggregation function:

Definition 2 (aggregation function) *A* symmetric aggregation function $\oplus$ *is a mapping from* $\mathbb{N}^n$ *to* $\mathbb{N}$ *such that:*[2]

- *if* $x \leq y$, *then* $\oplus(x_1,...,x,...,x_n) \leq \oplus(x_1,...,y,...,x_n)$ *(monotony)*
- $\oplus(0,x_1,\ldots,x_n) = \oplus(x_1,\ldots,x_n)$ *(neutral element)*
- $\oplus(x_1,\ldots,x_n) = 0$ *iff* $x_1 = \ldots = x_n = 0$ *(minimality)*
- $\oplus(x) = x$ *(identity)*
- $\oplus(x_1,\ldots,x_n) = \oplus(x_{\pi(1)},\ldots,x_{\pi(n)})$, *for any permutation* π *(symmetry)*

Usual aggregation functions are Σ and *max*. We focus on these functions fin the following, but many other choices are possible (including leximax, leximin, Σ^n (sum of the n^{th} powers), etc.).

One also needs to consider the attackers of an extension. Formally, given a WAF $= \langle A,R,w\rangle$ and two subsets A_1, A_2 of A, we note by $R(A_1,A_2) = \{b \in A_2 \mid \exists a \in A_1, (b,a) \in R\}$. $R(A_1,A_2)$ contains all the arguments from A_2 that attack at least one argument from A_1. Finally, the attack scores can be defined as follows:

Definition 3 ($\oplus$-attack) *Let* WAF $= \langle A,R,w\rangle$ *be a weighted argumentation framework. Let* A_1, A_2 *be two subsets of* A. *Let* $\oplus$ *be an aggregation function. The* $\oplus$-attack *from* A_1 *on* A_2 *is :* $S_\oplus(A_1 \rightarrow A_2) = \oplus_{a\in A_1, b\in A_2} w(a,b)$.

Now, there are several ways to select extensions based on their scores. We explore these ways in the following. Let us start with some basic definitions.

Definition 4 ($\oplus$-most attacking extensions) *Let* WAF $= \langle A,R,w\rangle$ *be a weighted argumentation framework. Let* $\mathscr{E}$ *be the set of extensions of* $\hat{\text{WAF}} = \langle A,R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. We define* $out_\oplus(E) = S_\oplus(E \rightarrow R(E,A))$. *The* $\oplus$-most attacking extensions *of* $\mathscr{E}$ *are then given by:* $ma_\oplus(\mathscr{E}) = argmax_{E\in\mathscr{E}}(out_\oplus(E))$. [3]

$\oplus$-most attacking extensions are those for which the aggregated weight of outgoing attacks is maximal. Such extensions have some flavor of semi-stable extensions [4]. Note however that there are no links between semi-stable extensions and Σ-most attacking ex-

[2] More formally, it is a family of such mappings, one for each positive integer n.

[3] Given a mapping g returning a real number, $argmax_{E\in\mathscr{E}}(g(E))$ (resp. $argmin_{E\in\mathscr{E}}(g(E))$) denotes the subset of elements of $\mathscr{E}$ maximizing (resp. minimizing) the value of g.

tensions for grounded or preferred semantics, even when attack weights have a uniform, positive value since maximality is considered with respect to cardinality for these two semantics, while it is considered with respect to set-inclusion for semi-stable extensions.

Dually, one can also focus on $\oplus$-least attacked extensions:

Definition 5 ($\oplus$-least attacked extensions) *Let* WAF $= \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* $\mathscr{E}$ *be the set of extensions of* $\hat{\text{WAF}} = \langle A, R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. We define* $in_{\oplus}(E) = S_{\oplus}(R(E,A) \rightarrow E)$. *The* $\oplus$-least attacked extensions *of* $\mathscr{E}$ *are then given by:* $la_{\oplus}(\mathscr{E}) = argmin_{E\in\mathscr{E}}(in_{\oplus}(E))$.

According to this definition, the best extensions are the ones receiving the least (amount of) attacks.

3. Best Defended Extensions

The previous definitions take into account only the attacks from or towards the considered extension. Of course it is easy to combine the corresponding scores $out_{\oplus}(E)$ and $in_{\oplus}(E)$ (for instance using difference or any lexicographic combination of them) in order to define additional or still refined further notions of extensions.

Furthermore, the defence of an extension can be measured by considering the extension as a whole (global defence) or in an argument-wise fashion (local defence). We are now ready to define the corresponding notions of globally $\oplus$-best defended extensions and locally $\oplus$-best defended extensions.

Definition 6 (globally $\oplus$-best defended extensions) *Let* WAF $= \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* $\mathscr{E}$ *be the set of extensions of* $\hat{\text{WAF}} = \langle A, R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. For any extension* E *of* $\mathscr{E}$*, one defines the score* $def^{g}_{\oplus}(E)$ *of global defence of* E *by :* $def^{g}_{\oplus}(E) = out_{\oplus}(E) - in_{\oplus}(E)$. *The globally* $\oplus$-best defended extensions *of* $\mathscr{E}$ *are then given by:* $gbd_{\oplus}(\mathscr{E}) = argmax_{E\in\mathscr{E}}(def^{g}_{\oplus}(E))$.

It can be the case that there exists an argument of a globally $\oplus$-best defended extension E such that the aggregated weights of its attackers exceeds the aggregated weights of its defenders. Accordingly, E would not be considered as an acceptable set with respect to the refined notion of defence introduced in [7]. Thus, a more demanding selection criterion also makes sense, based on the local score of defence:

Definition 7 (locally $\oplus$-best defended extensions) *Let* WAF $= \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* $\mathscr{E}$ *be the set of extensions of* $\hat{\text{WAF}} = \langle A, R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. For any extension* E *of* $\mathscr{E}$*, one defines the score* $def^{l}_{\oplus}(E)$ *of local defence of* E *by* $def^{l}_{\oplus}(E) = min_{a\in E} def^{l}_{\oplus}(a,E)$*, where* $def^{l}_{\oplus}(a,E) = min_{b\in R(\{a\},A)} \oplus_{c\in E} w(c,b) - w(b,a)$. *The locally* $\oplus$-best defended extensions *of* $\mathscr{E}$ *are then given by:* $lbd_{\oplus}(\mathscr{E}) = argmax_{E\in\mathscr{E}}(def^{l}_{\oplus}(E))$.

The local score of defence of an extension E is equal to the defence score of the ”least defended” argument of E.

4. Adversarial Extension Selection

All the definitions above select extensions by evaluating them independently from others, and then by choosing the best ones.

A more sensible way to select extensions seems to be to compare them in pairwise contests. It is then possible to take advantage of methods coming from voting theory in order to select the best extensions with respect to these contests.

An extension E_1 can be considered as better than another one E_2 when the attack of E_1 against E_2 is stronger than the attack from E_2 against E_1:

Definition 8 ($>_\oplus$) *Let* WAF $= \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* E_1 *and* E_2 *be two extensions of* $\widehat{\text{WAF}} = \langle A, R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. Then* $E_1 >_\oplus E_2$ *iff* $S_\oplus(E_1 \to E_2) > S_\oplus(E_2 \to E_1)$.

There are at least four natural ways to exploit the ordering $>_\oplus$ in order to select extensions:

Definition 9 ($\mathbf{best}_1^\oplus$) *Let* WAF $= \langle A, R, w\rangle$ *be a weighted argumentation framework. Let* $\mathscr{E}$ *be the set of extensions of* $\widehat{\text{WAF}} = \langle A, R\rangle$ *for a given semantics (preferred, stable, etc.). Let* $\oplus$ *be an aggregation function. Then*

- $best_1^\oplus(\mathscr{E}) = \{E \in \mathscr{E} \mid \nexists E' \in \mathscr{E}, E' >_\oplus E\}$.
- $best_2^\oplus(\mathscr{E}) = argmax_{E\in\mathscr{E}}|\{E' \in \mathscr{E} \mid E >_\oplus E'\}|$.
- $best_3^\oplus(\mathscr{E}) = argmax_{E\in\mathscr{E}}|\{E' \in \mathscr{E} \mid E >_\oplus E'\}| - |\{E' \in \mathscr{E} \mid E' >_\oplus E\}|$.
- $best_4^\oplus(\mathscr{E}) = argmax_{E\in\mathscr{E}} KS_\oplus(E)$, *where* $KS_\oplus(E) = \min_{E'\in\mathscr{E}, E'\neq E}(S_\oplus(E \to E'))$.

The first rule is the most natural one, where the selected extensions are the ones that are not beaten by any other extension. The problem with this method is that it can provide an empty set as answer. This is not the case with the next ones (provided that the input set of extensions is non-empty, of course). The second approach for selecting extensions consists in counting how many other extensions an extension defeats. Indeed, voting theory [2] provides us with interesting methods for deciding which are the best extensions, such as the Copeland rule [18] or the Kramer-Simpson (also called maximin) rule [19,14]. In fact the third rule is similar to the Copeland voting rule [18], which takes as score of an extension E the difference between the number of extensions beaten by E and the number of extensions which beats E. The fourth method is related, but instead of just counting wins and losses, it counts in some sense how much an extension wins. This can be related to the Kramer-Simpson voting rule [19,14].

5. An Example

We give in this section an example of WAF in order to illustrate the results obtained with the different methods. It is not possible to discriminate all the methods on a single example, but one can observe than four of the five preferred extensions can be selected as best extensions, depending of the chosen method.

Consider WAF $_1 = \langle A_1, R_1, w_1\rangle$ illustrated on the digraph of Figure 1 with $A_1 = \{a,b,c,d,e,f,g\}$, $R_1 = \{(a,d),(a,e),(a,f),(b,a),(c,b),(c,g),(d,c),(d,e),(e,c),(e,d),$

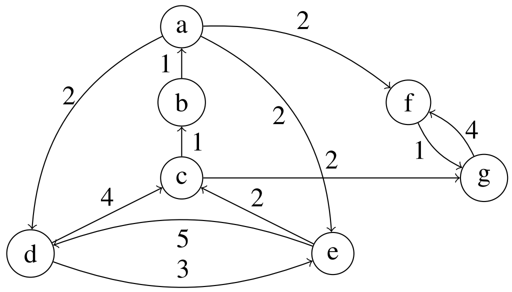

Figure 1. The digraph of WAF $_1$

$(f,g),(g,f)\}$ and w_1 : $(a,d) \to 2$; $(a,e) \to 2$; $(a,f) \to 2$; $(b,a) \to 1$; $(c,b) \to 1$; $(c,g) \to 2$; $(d,c) \to 4$; $(d,e) \to 3$; $(e,c) \to 2$; $(e,d) \to 5$; $(f,g) \to 1$; $(g,f) \to 4$.

WÂF$_1$ has five preferred extensions : $\mathscr{E} = \{E_1,E_2,E_3,E_4,E_5\}$, with $E_1 = \{a,c\}$, $E_2 = \{b,d,f\}$, $E_3 = \{b,e,f\}$, $E_4 = \{b,d,g\}$, $E_5 = \{b,e,g\}$.

We focus on two aggregation functions only: Σ and max.

- The most attacking and least attacked extensions are:
  - The Σ-most attacking extensions are: $ma_\Sigma(\mathscr{E}) = \{E_4,E_5\}$.
  - The max-most attacking extensions are: $ma_{\max}(\mathscr{E}) = \{E_3,E_5\}$.
  - The Σ-least attacked extension is: $la_\Sigma(\mathscr{E}) = \{E_1\}$.
  - The max-least attacked extension is: $la_{\max}(\mathscr{E}) = \{E_5\}$.
- The globally and locally best defended extensions are:
  - The globally Σ-best defended extensions coincide with the globally *max*-best defended extensions: $gbd_\Sigma(\mathscr{E}) = gbd_{max}(\mathscr{E}) = \{E_5\}$.
  - As to local defence, for any aggregation function $\oplus$ the locally $\oplus$-best defended extension is $lbd_\oplus(\mathscr{E}) = \{E_5\}$.
- The adversarial methods gives:
  - $best_1^\Sigma(\mathscr{E}) = \{E_1\}$.
  - $best_2^\Sigma(\mathscr{E}) = \{E_5\}$.
  - $best_3^\Sigma(\mathscr{E}) = \{E_1,E_5\}$.
  - $best_4^\Sigma(\mathscr{E}) = \{E_1\}$.
  - $best_1^{\max}(\mathscr{E}) = \{E_5\}$.
  - $best_2^{\max}(\mathscr{E}) = \{E_5\}$.
  - $best_3^{\max}(\mathscr{E}) = \{E_5\}$.
  - $best_4^{\max}(\mathscr{E}) = \{E_4\}$.

6. Conclusion

In this paper we explored how to select extensions of Dung's abstract argumentation frameworks by exploiting weights on the attack relation. We provided a number of criteria for such a selection purpose. We started with simple criteria, that do not make a very fine-grained distinction between extensions, but are quite easy to compute. Then we gave subtler criteria, requiring more complex computations. In particular the criteria $best_3^\oplus$ and $best_4^\oplus$, directly inspired from existing voting rules, are the most promising ones. Other criteria could be easily defined using other voting rules based on the majority graph (see e.g. [15]), that is replaced in our framework by the relation $>_\oplus$. Each criterion we considered is parameterized by an aggregation function which indicates how weights interact. This gives the approach a good level of generality.

As a perspective for further research, we plan to identify the complexity of skeptical inference from selected extensions, for all the criteria introduced in the paper. Another perspective for further work is to compare the inferential powers of skeptical inference for each of the criteria under consideration.

References

[1] L. Amgoud and S. Vesic. Two roles of preferences in argumentation frameworks. In *Proceedings of the Eleventh European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (EECSQARU'11)*, volume 6717 of *Lecture Notes on Artificial Intelligence*, pages 86–97, 2011.

[2] K.J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.

[3] P. Besnard and A. Hunter. *Elements of Argumentation*. MIT Press, 2008.

[4] M. Caminada. Semi-stable semantics. In *Proceedings of the First International Conference on Computational Models of Argument (COMMA'06)*, pages 121–130, 2006.

[5] C. Cayrol, C. Devred, and M.-C. Lagasquie-Schiex. Acceptability semantics accounting for strength of attacks in argumentation. In *Proceedings of the nineteenth European Conference on Artificial Intelligence (ECAI'10)*, pages 995–996, 2010.

[6] C. Cayrol and M.-C. Lagasquie-Schiex. Weighted argumentation systems: A tool for merging argumentation systems. In *Proceedings of the twenty-third IEEE International Conference on Tools with Artificial Intelligence (ICTAI'11)*, pages 629–632, 2011.

[7] S. Coste-Marquis, S. Konieczny, P. Marquis, and M.A. Ouali. Weighted attacks in argumentation frameworks. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR'12)*, 2012.

[8] 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):321–358, 1995.

[9] P. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge. Inconsistency Tolerance in Weighted Argument Systems. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS'09)*, pages 851–858, 2009.

[10] P. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Woolridge. Weighted Argumentation Systems: Basic Definitions, Algorithms and Complexity Results. *Artificial Intelligence*, 175(2):457–486, 2011.

[11] A. Hunter. Measuring inconsistency in knowledge via quasi-classical models. In *Proceedings of the Eighteenth AAAI Conference (AAAI'02)*, pages 68–73, 2002.

[12] A. Hunter and S. Konieczny. *Inconsistency tolerance*, volume 3300 of *Lecture Notes on Computer Science*, chapter Approaches to measuring inconsistent information, pages 189–234. Springer, 2005.

[13] A. Hunter and S. Konieczny. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence*, 174(14):1007–1026, 2010.

[14] G. Kramer. A dynamical model of political equilibrium. *Journal of Economic Theory*, 16:310–334, 1977.

[15] J.F. Laslier. *Tournament solutions and majority voting*. Studies in economic theory. Springer, 1997.

[16] D. Martínez, A. García, and G. Simari. An Abstract Argumentation Framework with Varied Strength Attacks. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, pages 135–144, 2008.

[17] D. Martínez, A. García, and G. Simari. Strong and weak forms of abstract argument defense. In *Proceedings of the Second International Conference on Computational Models of Argument (COMMA'08)*, pages 216–227, 2008.

[18] D. G. Saari and V. R. Merlin. The copeland method: I.: Relationships and the dictionary. *Economic Theory*, 8(1):pp. 51–76, 1996.

[19] PB Simpson. On defining areas of voter choice: Professor tullock on stable voting. *Quarterly Journal of Economics*, 83:478–490, 1969.