

On Supported Inference and Extension Selection in Abstract Argumentation Frameworks

Sébastien Konieczny Pierre Marquis Srdjan Vesic

CRIL, CNRS and Université d'Artois, France
{konieczny,marquis,vesic}@cril.fr

Abstract. We present two approaches for deriving more arguments from an abstract argumentation framework than the ones obtained using sceptical inference, that is often too cautious. The first approach consists in selecting only some of the extensions. We point out several choice criteria to achieve such a selection process. Choices are based either on the attack relation between extensions or on the support of the arguments in each extension. The second approach consists of the definition of a new inference policy, between sceptical and credulous inference, and based as well on the support of the arguments. We illustrate the two approaches on examples, study their properties, and formally compare their inferential powers.

1 Introduction

An abstract argumentation system is often represented as an *oriented graph*, where nodes correspond to *arguments* and arcs correspond to *attacks* between them [15]. Different *semantics* are used to calculate *extensions* (sets of arguments that can be accepted together). From the extensions, a *status*, accepted or rejected, is assigned to each argument, using some *acceptance policy*. They are two main acceptance policies. In the first one, the *sceptical* policy, an argument is accepted if (there are extensions and) it appears in each extension. For the second one, the *credulous* policy, an argument is accepted if it belongs to (at least) one extension.

When the number of extensions is large, using a sceptical / credulous approach can be sub-optimal. Namely, if there is a lot of extensions, only few (if any) arguments are in all of them. Thus, using sceptical inference gives almost no information. Conversely, the credulous approach may result in too many arguments.

There exist settings for abstract argumentation where preferences, weighted attacks or similar extra information are considered [20,21,8,16,11,2]. Those additional data can be exploited to reduce the number of extensions. Contrastingly, the problem addressed in this paper is to increase the number of accepted arguments when there is no further data, i.e., other data except the arguments and the attacks between them.

We investigate this problem and present two approaches for dealing with it. The first one consists in selecting only some of the extensions (the “best” ones, for a given semantics). The idea is to discriminate the extensions by taking advantage of the attack relation. The selection achieved in this way leads to increase the number of sceptically accepted arguments. Two methods for selecting extensions are pointed out. The first

one is based on a pairwise comparison of extensions. The second method is based on a global evaluation of each extension, followed by a selection of the best evaluated ones. The second approach we developed goes through the definition of a new policy for accepting arguments. We introduce a third acceptance policy, which can be viewed as a trade-off between the credulous and the sceptical policy. The very idea is to consider the number of times an argument appears in the extensions. For the sceptical policy a "good" argument is one that appears in all extensions. If no such argument exists, then it makes sense to consider that arguments that appears in every extension but one are "quite good", and better than the ones that appear in less extensions.

A technical report [19] containing all the proofs, more explanations, more examples and more figures is available online at <http://www.cril.fr/~vesic>.

2 Formal Setting

This section introduces basic definitions and notations we use throughout the paper. An *argumentation system* (AS) is a pair $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ where $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$. \mathcal{A} is called the set of *arguments* and \mathcal{R} is called the *attack relation*. We restrict ourselves to the case when \mathcal{A} is finite.

Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an AS, and let $\mathcal{E}, \mathcal{E}', \mathcal{E}'' \subseteq \mathcal{A}$ and $a \in \mathcal{A}$. \mathcal{E} is *conflict-free* if and only if there exist no arguments $a, b \in \mathcal{E}$ such that $a \mathcal{R} b$. \mathcal{E} *defends* a if and only if for every $b \in \mathcal{A}$ we have that if $b \mathcal{R} a$ then there exists $c \in \mathcal{E}$ such that $c \mathcal{R} b$. Argument a is *strongly defended* from \mathcal{E}' by \mathcal{E}'' (written $sd(a, \mathcal{E}', \mathcal{E}'')$) if and only if $(\forall b \in \mathcal{E}')$ if $(b \mathcal{R} a)$ then $(\exists c \in \mathcal{E}'' \setminus \{a\})((c \mathcal{R} b) \wedge sd(c, \mathcal{E}', \mathcal{E}'' \setminus \{a\}))$.

Usual semantics for Dung's AS are considered, especially the complete, preferred, grounded [15], semi-stable [7] and ideal semantics [14]. A semantics σ is said to return conflict-free sets iff for every AS \mathcal{F} , every extension of \mathcal{F} is conflict-free. For an argumentation system $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ we denote $\text{Ext}_\sigma(\mathcal{F})$; or, by a slight abuse of notation, $\text{Ext}_\sigma(\mathcal{A}, \mathcal{R})$ the set of its extensions with respect to semantics σ . We use abbreviations c, p, s, ss, g and i for respectively complete, preferred, stable, semi-stable, grounded and ideal semantics. For example, $\text{Ext}_p(\mathcal{F})$ denotes the set of preferred extensions of \mathcal{F} .

An *acceptance policy* is a function $\text{Inf}_\sigma : \text{Ext}_\sigma(\mathcal{F}) \rightarrow 2^{\mathcal{A}}$. The two main acceptance policies are sceptical and credulous policies. We say that x is *sceptically accepted* under semantics σ (or in short *s-sceptically accepted*) iff $\text{Ext}_\sigma(\mathcal{F}) \neq \emptyset$ and $x \in \bigcap_{\mathcal{E} \in \text{Ext}_\sigma(\mathcal{F})} \mathcal{E}$. x is *credulously accepted* under semantics σ iff $x \in \bigcup_{\mathcal{E} \in \text{Ext}_\sigma(\mathcal{F})} \mathcal{E}$. We denote the set of sceptically accepted arguments by $\text{Sc}_\sigma(\mathcal{F})$ and the set of credulously accepted arguments by $\text{Cr}_\sigma(\mathcal{F})$. We denote by $\mathcal{R}_{\downarrow \mathcal{E}}$ the restriction of attack relation \mathcal{R} on set \mathcal{E} .

3 Comparing Extensions by Pairwise Comparison

This section studies the way to select the "best" extensions based on the following process:

1. Compare all pairs of extensions based on a given criterion (e.g. the number of arguments in one extension not attacked by the other extension)

2. Choose the "best" extension(s) given the winners of pairwise comparisons

Definition 1 (Pairwise comparison criteria). Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an AS, σ a semantics and $\text{Ext}_\sigma(\mathcal{F})$ the set of extensions of \mathcal{F} . Let $\mathcal{E}, \mathcal{E}' \in \text{Ext}_\sigma(\mathcal{F})$. Then:

1. $\mathcal{E} \succeq_{\text{nonatt}} \mathcal{E}'$ if the number of arguments in \mathcal{E} non attacked by \mathcal{E}' is greater than or equal to the number of arguments in \mathcal{E}' non attacked by arguments of \mathcal{E}
2. $\mathcal{E} \succeq_{\text{strdef}} \mathcal{E}'$ if the number of arguments in \mathcal{E} strongly defended from \mathcal{E}' by \mathcal{E} is greater than or equal to the number of arguments in \mathcal{E}' strongly defended from \mathcal{E} by \mathcal{E}'
3. $\mathcal{E} \succeq_{\text{delarg}} \mathcal{E}'$ if the cardinality of any largest subset S of \mathcal{E} such that if all the attacks from S to \mathcal{E}' are deleted then \mathcal{E} is an extension of $(\mathcal{E} \cup \mathcal{E}', \mathcal{R}_{\downarrow \mathcal{E} \cup \mathcal{E}'})$ is greater than or equal to the cardinality of any largest subset S' of \mathcal{E}' such that if all the attacks from S' to \mathcal{E} are deleted then \mathcal{E} is an extension of $(\mathcal{E} \cup \mathcal{E}', \mathcal{R}_{\downarrow \mathcal{E} \cup \mathcal{E}'})$
4. $\mathcal{E}' \succeq_{\text{delatt}} \mathcal{E}$ if the maximal number of attacks from \mathcal{E} to \mathcal{E}' that can be deleted such that \mathcal{E} is still an extension of $(\mathcal{E} \cup \mathcal{E}', \mathcal{R}_{\downarrow \mathcal{E} \cup \mathcal{E}'})$ is greater than or equal to the maximal number of attacks from \mathcal{E}' to \mathcal{E} that can be deleted such that \mathcal{E}' is still an extension of $(\mathcal{E} \cup \mathcal{E}', \mathcal{R}_{\downarrow \mathcal{E} \cup \mathcal{E}'})$

The two first criteria are based on the number of non attacked or (strongly) defended arguments. The last two ones are based on a notion of robustness from attacks stemming from the other extension. One could also consider other criteria, for example by comparing the total number of attacks from \mathcal{E} to \mathcal{E}' and the total number of attacks from \mathcal{E}' to \mathcal{E} . For a criterion γ , we write $\mathcal{E} \succ_\gamma \mathcal{E}'$ iff $\mathcal{E} \succeq_\gamma \mathcal{E}'$ and it is not the case that $\mathcal{E}' \succeq_\gamma \mathcal{E}$. We also write $\mathcal{E} \sim_\gamma \mathcal{E}'$ iff $\mathcal{E} \succeq_\gamma \mathcal{E}'$ and $\mathcal{E}' \succeq_\gamma \mathcal{E}$.

Example 1. Consider the AS $\mathcal{F}_1 = (\mathcal{A}_1, \mathcal{R}_1)$ with $\mathcal{A}_1 = \{a, b, c, d\}$ and $\mathcal{R}_1 = \{(a, c), (a, d), (b, c), (c, a), (d, b)\}$. $\text{Ext}_p(\mathcal{F}_1) = \{\mathcal{E}, \mathcal{E}'\}$ with $\mathcal{E} = \{a, b\}$, $\mathcal{E}' = \{c, d\}$. All the arguments are attacked, so $\mathcal{E} \sim_{\text{nonatt}} \mathcal{E}'$. No argument is strongly defended, so $\mathcal{E} \sim_{\text{strdef}} \mathcal{E}'$. We also have $\mathcal{E} \succ_{\text{delarg}} \mathcal{E}'$ since for $S = \{b\}$ \mathcal{E} is still an extension even if all the attacks from S are deleted; whereas there are no $S' \subseteq \mathcal{E}'$ with $S' \neq \emptyset$ such that \mathcal{E}' is still a preferred extension even after deleting all the attacks from S' . Finally, $\mathcal{E} \succ_{\text{delatt}} \mathcal{E}'$ since even if the attack from a to c is deleted, \mathcal{E} is still a preferred extension, whereas as soon as one attack from \mathcal{E}' is deleted, \mathcal{E}' is no longer a preferred extension.

Definition 2 (Copeland-based extensions). Let $\gamma \in \{\text{nonatt}, \text{strdef}, \text{delarg}, \text{delatt}\}$ be one of the criteria from Definition 1. Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an argumentation system, σ a semantics and $\text{Ext}_\sigma(\mathcal{F})$ the set of extensions of \mathcal{F} with respect to σ . We define the set of Copeland-based extensions (CBE) as follows

$$\text{CBE}_{\sigma, \gamma}(\mathcal{F}) = \arg \max_{\mathcal{E} \in \text{Ext}_\sigma(\mathcal{F})} |\{\mathcal{E}' \in \text{Ext}_\sigma(\mathcal{F}) \mid \mathcal{E} \succeq_\gamma \mathcal{E}'\}| - |\{\mathcal{E}'' \in \text{Ext}_\sigma(\mathcal{F}) \mid \mathcal{E}'' \succeq_\gamma \mathcal{E}\}|$$

We call this selection "Copeland-based" since it is inspired by the Copeland's rule from voting theory [22]. Of course, one can envisage other ways to select the extensions given criterion γ , for instance all voting methods based on the majority graph (such as Miller, Fishburn, Schwartz, Banks or Slater's methods [6]). Clearly, selecting some extensions is a way to increase the number of sceptically accepted arguments (and to decrease the number of credulously accepted arguments):

Fact 1 For every $\gamma \in \{\text{nonatt}, \text{strdef}, \text{delarg}, \text{delatt}\}$, for every semantics σ , for every AS $\mathcal{F} = (\mathcal{A}, \mathcal{R})$, for every $x \in \mathcal{A}$:

- $\text{CBE}_{\sigma, \gamma}(\mathcal{F}) \subseteq \text{Ext}_{\sigma}(\mathcal{F})$
- if x is σ -sceptically accepted then x is $\text{CBE}_{\sigma, \gamma}$ -sceptically accepted
- if x is $\text{CBE}_{\sigma, \gamma}$ -credulously accepted then it is σ -credulously accepted.

Example 2. Consider the argumentation system from Example 1. For example, we have that $\text{CBE}_{\sigma, \text{delarg}}(\mathcal{F}_1) = \text{CBE}_{\sigma, \text{delatt}}(\mathcal{F}_1) = \{\mathcal{E}\}$.

Baroni and Giacomin [4] pointed out a set of extension evaluation criteria that can be seen as properties for characterizing good semantics. We now show that the semantics defined in this section satisfy the same properties as the underlying semantics they are built from, with the exception of directionality.

Proposition 1. *Let x be any property among I-maximality, Admissibility, Strong Admissibility, Reinstatement, Weak Reinstatement, CF-Reinstatement [4]. If the semantics σ satisfies property x , then the semantics $\text{CBE}_{\sigma, \gamma}$ satisfies property x .*

Note that the relations among different semantics do not carry over in case of CBE approach. For instance, it is not guaranteed that each CBE-stable extension is also a CBE-preferred extension.

4 Comparing Extensions by Global Evaluation

In Section 3 we considered different criteria for *pairwise comparison* of extensions. In this section we define the score of an argument as the number of extensions it appears in. One may justify this choice of score as some kind of generalization of the principles behind sceptical acceptance. For sceptical acceptance a "good" argument is an argument that appears in all extensions. But, if no such argument exists, it could make sense to consider that arguments that appears in every extension but one are "good", and typically better than the ones that appears in less extensions. Note that one can use other scores in the construction and obtain similar results.

Definition 3 (Scores and support vectors). *Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an argumentation system, σ a semantics, x be an argument, and $\text{Ext}_{\sigma}(\mathcal{F})$ the set of extensions of \mathcal{F} with respect to σ . We define ne as the number of extensions x appears in. Formally, $\text{ne}_{\sigma}(x, \mathcal{F}) = |\{\mathcal{E} \in \text{Ext}_{\sigma}(\mathcal{F}) \mid x \in \mathcal{E}\}|$. For an extension $\mathcal{E} \in \text{Ext}_{\sigma}(\mathcal{F})$, with $\mathcal{E} = \{a_1, \dots, a_n\}$ we define its support as $\text{vsupp}_{\sigma}(\mathcal{E}, \mathcal{F}) = (\text{ne}_{\sigma}(a_1, \mathcal{F}), \dots, \text{ne}_{\sigma}(a_n, \mathcal{F}))$.*

When \mathcal{F} and σ are clear from the context, we write $\text{ne}(x)$ and $\text{vsupp}(\mathcal{E})$ instead of $\text{ne}_{\sigma}(x, \mathcal{F})$ and $\text{vsupp}_{\sigma}(\mathcal{E}, \mathcal{F})$.

Definition 4 (Aggregation functions). *Let $v = (v_1, \dots, v_n)$ be a vector of natural numbers. We denote by $\text{sum}(v)$ the sum of all elements of v , by $\text{max}(v)$ the maximal element of v , by $\text{min}(v)$ the minimal element of v , by $\text{leximax}(v)$ the re-arranged version of v where v_1, \dots, v_n are put in decreasing order, by $\text{leximin}(v)$ the re-arranged version of v where v_1, \dots, v_n are put in increasing order.*

For example, if $v = (2, 1, 4, 2, 5)$, then we have $sum(v) = 14$ and $leximin(v) = (1, 2, 2, 4, 5)$. Note that there exist other ways to aggregate vectors [13].

For the next definition we need the notion of lexicographic order $<_{lex}$ (for *leximin* and *leximax*). Let $v = (v_1, \dots, v_n)$ and $v' = (v'_1, \dots, v'_n)$ be two vectors of natural numbers. We have $v <_{lex} v'$ iff $\exists j \in 1, \dots, n (\forall i \in 1, \dots, j-1, v_i = v'_i)$ and $v_j < v'_j$. We also have $v <_{leximin} v'$ iff $leximin(v) <_{lex} leximin(v')$ and $v <_{leximax} v'$ iff $leximax(v) <_{lex} leximax(v')$.

Definition 5 (Order-based extensions). Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an argumentation system, σ a semantics, $Ext_\sigma(\mathcal{F})$ be the set of extensions of \mathcal{F} with respect to σ , and γ be an aggregation function. We have $OBE_{\sigma, \gamma}(\mathcal{F}) = \arg \max_{\mathcal{E} \in Ext_\sigma(\mathcal{F})} \gamma(vsupp_\sigma(\mathcal{E}, \mathcal{F}))$.

The idea of the previous definition is to calculate the popularity of an extension by taking into account the popularity of the arguments it contains.

Example 3. Let $\mathcal{F}_3 = (\mathcal{A}_3, \mathcal{R}_3)$ be AS $\mathcal{F}_3 = (\mathcal{A}_3, \mathcal{R}_3)$ with $\mathcal{A}_3 = \{a, b, c, d, e, f, g, h\}$ and $\mathcal{R}_3 = \{(a, b), (b, a), (e, f), (f, e), (b, g), (f, g), (g, h), (h, d), (d, c), (c, d)\}$. There are five preferred extensions: $\{a, e, g, c\}$, $\{a, e, g, d\}$, $\{a, f, h, c\}$, $\{b, h, c, e\}$, $\{b, h, c, f\}$. So $ne_p(a, \mathcal{F}_3) = 3$, $ne_p(b, \mathcal{F}_3) = 2$, $ne_p(c, \mathcal{F}_3) = 4$, $ne_p(d, \mathcal{F}_3) = 1$, $ne_p(e, \mathcal{F}_3) = 3$, $ne_p(f, \mathcal{F}_3) = 2$, $ne_p(g, \mathcal{F}_3) = 2$, $ne_p(h, \mathcal{F}_3) = 3$.

We obtain $OBE_{\sigma, max}(\mathcal{F}_3) = OBE_{\sigma, min}(\mathcal{F}_3) = \{\{a, e, g, c\}, \{a, f, h, c\}, \{b, h, c, e\}, \{b, h, c, f\}\}$. So, whereas $Sc_p(\mathcal{F}_3) = \emptyset$, we have $Sc_{OBE_{p, min}}(\mathcal{F}_3) = \{c\}$. Similarly, we have $OBE_{\sigma, sum}(\mathcal{F}_3) = OBE_{\sigma, leximin}(\mathcal{F}_3) = OBE_{\sigma, leximax}(\mathcal{F}_3) = \{\{a, e, g, c\}, \{a, f, h, c\}, \{b, h, c, e\}\}$.

Fact 2 For every $\gamma \in \{sum, max, min, leximin, leximax\}$, for every semantics σ , for every AS $\mathcal{F} = (\mathcal{A}, \mathcal{R})$, for every $x \in \mathcal{A}$:

- $OBE_{\sigma, \gamma}(\mathcal{F}) \subseteq Ext_\sigma(\mathcal{F})$
- if x is σ -sceptically accepted then x is $OBE_{\sigma, \gamma}$ -sceptically accepted
- if x is $OBE_{\sigma, \gamma}$ -credulously accepted then it is σ -credulously accepted.

Proposition 2. Let x be any property among I-maximality, Admissibility, Strong Admissibility, Reinstatement, Weak Reinstatement, CF-Reinstatement [4]. If the semantics σ satisfies property x , then the semantics $OBE_{\sigma, \gamma}$ satisfies property x .

Like in Section 3, directionality is not always satisfied by the OBE approach.

A natural issue is to determine how the proposed criteria are connected. Do some of the rules coincide? Are some of them refinements of others? In the rest of this section we provide the answer to this question. Essentially, all the criteria give different results; the exceptions come from the obvious fact that *leximin* (resp. *leximax*) refines *min* (resp. *max*). We used the preferred semantics to construct the counter-examples; a similar study can be conducted for the other semantics.

Definition 6. Let Γ and Γ' be two functions. We write $\Gamma \sqsubseteq \Gamma'$ iff for every \mathcal{F} , $\Gamma(\mathcal{F}) \subseteq \Gamma'(\mathcal{F})$. The relation \sqsubseteq is a pre-order. Let us denote its strict part by \sqsubset , its symmetric part by \doteq and its negation by $\not\sqsubseteq$. We write $\Gamma \text{ ind } \Gamma'$ iff $\Gamma \not\sqsubseteq \Gamma'$ and $\Gamma' \not\sqsubseteq \Gamma$.

Proposition 3. For every acceptability semantics σ ,

$$\text{OBE}_{\sigma, \text{leximin}} \sqsubseteq \text{OBE}_{\sigma, \text{min}} \text{ and } \text{OBE}_{\sigma, \text{leximax}} \sqsubseteq \text{OBE}_{\sigma, \text{max}}$$

We now provide a complete comparison between pairs of criteria under preferred semantics.

Proposition 4. It holds that $\text{OBE}_{p, \text{leximin}} \sqsubseteq \text{OBE}_{p, \text{min}}$ and $\text{OBE}_{p, \text{leximax}} \sqsubseteq \text{OBE}_{p, \text{max}}$. The other pairs of rules (x, y) with $x, y \in \{\text{OBE}_{\text{sum}}, \text{OBE}_{\text{min}}, \text{OBE}_{\text{max}}, \text{OBE}_{\text{leximin}}, \text{OBE}_{\text{leximax}}\}$, $x \neq y$ are incomparable, i.e., $x \text{ ind } y$.

5 Support-Based Acceptance Policy

This section presents a completely different approach for selecting arguments. We focus on arguments that have the greatest supports among extensions to construct what we call “candidate sets”. Then, an argument is called *supportedly accepted* if it is in all the candidate sets.

Definition 7 (Candidate sets). Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an AS and let σ be a semantics. Let \succeq be any pre-order defined on \mathcal{A} . Let $|\mathcal{A}| = m$. For a permutation θ of $\{1, \dots, m\}$, let $>_{\theta}$ be the linear order on \mathcal{A} defined by $a_{\theta(1)} >_{\theta} \dots >_{\theta} a_{\theta(m)}$. $>_{\theta}$ is said to be compatible with \succeq iff $a_{\theta(1)} \succeq \dots \succeq a_{\theta(m)}$. A set $\mathcal{E} \subseteq \mathcal{A}$ is a candidate set of \mathcal{F} under semantics σ w.r.t. \succeq iff there exists a permutation θ of $\{1, \dots, m\}$ such that $>_{\theta}$ is compatible with \succeq and \mathcal{E} is obtained by the following greedy procedure:

```

S := ∅;
for j = 1, ..., m do
    if (neσ(aθ(j),  $\mathcal{F}$ ) ≥ 1) and (S ∪ {aθ(j)} is conflict-free)
        then S := S ∪ {aθ(j)}
end for;
 $\mathcal{E}$  := S.

```

In the following, we consider the pre-order \succeq on \mathcal{A} defined by for all $x, y \in \mathcal{A}$, $x \succeq y$ iff $\text{ne}_{\sigma}(x, \mathcal{F}) \geq \text{ne}_{\sigma}(y, \mathcal{F})$. We denote the set of candidate sets of \mathcal{F} under σ w.r.t. this pre-order by $\text{CS}_{\sigma}(\mathcal{F})$.

Note that, in general, neither each candidate set is an extension nor each extension is a candidate set. Observe also that the construction of candidate sets is reminiscent to the one of preferred subbases from a stratified belief base with respect to the inclusion-based ordering [5]; here the belief base consists of all the arguments and the stratification is based on the $\text{ne}_{\sigma}(\cdot, \mathcal{F})$ score.

Definition 8 (Supported acceptance). Let $\mathcal{F} = (\mathcal{A}, \mathcal{R})$ be an AS, σ be a semantics and let $x \in \mathcal{A}$. We say that x is supportedly accepted under semantics σ iff $x \in \bigcap_{\mathcal{E} \in \text{CS}_{\sigma}(\mathcal{F})} \mathcal{E}$. We denote the set of supportedly accepted arguments $\text{Sp}_{\sigma}(\mathcal{F})$.

We can show that supported inference is “between” sceptical and credulous inference.

Proposition 5. For every AS $\mathcal{F} = (\mathcal{A}, \mathcal{R})$, for every semantics σ returning conflict-free extensions:

$$\text{Sc}_\sigma(\mathcal{F}) \subseteq \text{Sp}_\sigma(\mathcal{F}) \subseteq \text{Cr}_\sigma(\mathcal{F}).$$

Note that the condition telling that σ returns conflict-free extensions is necessary to ensure the link between sceptical and supported acceptance. However, this is not an issue, since all the well-known semantics return conflict-free sets.

Example 4. Let $\mathcal{F}_4 = (\mathcal{A}_4, \mathcal{R}_4)$ be an AS with $\mathcal{A}_4 = \{a, b, c, d, e, f, g, h\}$ and $\mathcal{R}_4 = \{(a, b), (b, a), (b, g), (c, d), (d, c), (d, g), (e, f), (f, e), (f, g), (g, h)\}$. There are eight preferred extensions: $\{a, c, e, g\}$, $\{a, d, e, h\}$, $\{a, c, f, h\}$, $\{a, d, f, h\}$, $\{b, c, e, h\}$, $\{b, d, e, h\}$, $\{b, c, f, h\}$, $\{b, d, f, h\}$. There are no sceptically accepted arguments, i.e. $\text{Sc}_p(\mathcal{F}_4) = \emptyset$. But h is accepted by seven out of the eight extensions, and it is supportedly accepted, i.e., $\text{Sp}_p(\mathcal{F}_4) = \{h\}$.

In the above example the set of candidates is a subset of the set of extensions, but this is not always the case. Consider for instance the AS from Example 3, where there is only one candidate set $\{c, a, e, h\}$, that is *not* an extension. So it is interesting to note that on this example there are four supportedly inferred arguments, whereas with the OBE methods only c is inferred.

A major drawback of credulous inference is that the set of inferred arguments is not always conflict-free. This is problematic since all these arguments cannot be accepted together in such a case. Sceptical inference does not suffer from this problem since the set of inferred arguments is ensured to be conflict-free. Interestingly, supported inference offers the same important property:

Fact 3 For any \mathcal{F} , the set of supportedly accepted arguments is conflict-free.

Note that this set is not necessarily admissible. This should not be shocking since the same observation can be made for the set of sceptically accepted arguments.

Finally, an interesting issue is to determine whether some connections exist between supported inference and the approaches presented in the previous sections. We provide a systematic study of the links between the two approaches under preferred semantics.

Proposition 6. For every $\gamma \in \{\text{sum}, \text{min}, \text{max}, \text{leximin}, \text{leximax}\}$, OBE_γ and CS are incomparable under preferred semantics, i.e., OBE_γ ind CS.

Let us first show that every $\text{OBE}_{\sigma, \text{max}}$ -sceptically accepted argument is also supportedly accepted, for every semantics that returns conflict-free extensions.

Proposition 7. Let σ be a semantics returning conflict-free extensions. We have

$$\text{Sc}_{\text{OBE}_{\sigma, \text{max}}} \sqsubseteq \text{Sp}_\sigma.$$

Let us now illustrate the indifference between γ -sceptical acceptance and supported acceptance for $\gamma \neq \text{max}$, again on the case of preferred semantics.

Proposition 8. The links between Sc_γ and Sp under preferred semantics are as follows:

1. $\text{Sc}_{\text{OBE}_{p,\max}} \sqsubseteq \text{Sp}_p$.
2. for every $\gamma \in \{\text{sum}, \text{min}, \text{leximin}, \text{leximax}\}$, $\text{Sc}_{\text{OBE}_{p,\gamma}} \text{ind} \text{Sp}_p$.

The two previous propositions show that OBE and supported inference, although both using the scores of arguments defined as the number of extensions they belong to, induce intrinsically different reasoning mechanisms.

6 Conclusion and Related Work

This paper aimed at defining approaches for a better inference from abstract argumentation framework. Indeed, a large number of extensions results in a low number of sceptically accepted arguments. Several approaches have been described for dealing with this problem. First, different criteria for pairwise comparison of extensions and a method for selecting only the best extensions given the winners of pairwise duels have been pointed out. Second, several criteria for ordering the extensions have been presented. Both approaches result in a decrease of the number of extensions; consequently, the number of sceptical arguments increases (and the number of credulous arguments diminishes). The third approach we have put forward does not choose between existing extensions. Instead, it uses extensions to assign a score to every argument (the score of an argument is the number of extensions it belongs to). Then, starting from the arguments having the maximal score, candidate sets can be generated and on this ground supportedly accepted arguments have been defined.

Several papers in the literature are relevant to our work in the sense that their objectives are somehow similar. Thus, some previous work aimed at defining different levels of acceptability for arguments [9,23,18,3]. Such levels can be obtained by attaching numerical scores between 0 and 1 to each argument, or by ranking arguments over an ordinal scale. Contrastingly, the goal of the present paper is not to tackle the problem of gradual acceptance. In this work our objective is not to question the classical binary framework for inference, where an argument is inferred or not, but to define inference relations allowing to infer more arguments than sceptical inference; to make a parallel with logical inference, a similar distinction exists between paraconsistent logics and some weighted logics (such as possibilistic or fuzzy logics).

Settings where argumentation systems are based on preferences or attack weights can also be exploited for reducing the number of extensions. However, those approaches suppose the availability of some extra information such as weights or preferences, whereas our approach is based solely on the argumentation system $\mathcal{F} = (\mathcal{A}, \mathcal{R})$.

Other approaches calculate arguments' scores / statuses without relying on the notion of extension [1,12]. Unlike our approach, semantics (e.g., stable, preferred) are not used at all. Here, we suppose the use of an (arbitrary) semantics to calculate extensions and then point out a way to augment the number of arguments which are accepted. Our criteria are orthogonal to the notion of semantics, so that each criterion can be combined with each semantics.

Another related work is [10] which addresses the problem of defining more prudent inference relations for Dung's argumentation frameworks (i.e., the objective is to derive less arguments). Contrariwise to the present paper, instead of selecting some extensions

or defining a new inference policy, the approach consists in strengthening the usual (direct) conflict-freeness property to indirect conflict-freeness. Thus a prudent extension cannot contain two arguments when there exists an indirect attack among the first one and the second one. When the credulous policy and the preferred semantics (or the stable semantics) are considered, the set of derivable arguments from prudent extensions is included in the set of arguments derivable from the standard extensions.

Baroni et al. [3] show how to define some fine-grained argument justification statuses for abstract argumentation frameworks. For extension-based semantics, the justification status of an argument basically depends on the existence of extensions containing it and the existence of extensions attacking it. Clearly enough, the problem of selecting extensions is orthogonal to the problem of defining argument justification statuses; thus, Baroni's et al. results can be exploited as soon as some extensions exist, even if they come from a selection process. Our notion of supported inference is closer to their proposal since it induces an intermediate argument status, supported acceptance, between sceptical acceptance and credulous acceptance. However, the mechanisms at work for defining this intermediate status and its rationale are quite different from those considered in Baroni's et al. paper: in our work, the support of an argument is based on the number of extensions containing it.

Our approach also departs from the work by Dunne et al. [17] which focusses on ideal semantics. Indeed, ideal acceptance is more demanding than sceptical acceptance. As such, it proves useful when sceptical acceptance is not prudent enough, i.e. when unexpected arguments are sceptically accepted. Contrastingly, our work is motivated by the remaining cases, when sceptical inference is too cautious and discards some expected arguments.

Caminada and Wu [23] defined different labelling-based justification statuses of arguments. Indeed, they propose to attach to each argument the set of its possible labels (i.e. the collection of all labels it obtains in all complete labellings). Whereas Dung-based approach allows to split the arguments into three classes (sceptically accepted, credulously accepted, rejected), their contribution provides a way for fine-graded classification, by defining six different justification statuses: $\{\text{in}\}$, $\{\text{in}, \text{undec}\}$, $\{\text{undec}\}$, $\{\text{in}, \text{out}, \text{undec}\}$, $\{\text{out}, \text{undec}\}$ and $\{\text{out}\}$. The work of Caminada and Wu is related to our work since it could also be used to reason in cases when there are no (or when there are not enough) accepted arguments. However, the actual way to do it is drastically different from our approach.

References

1. L. Amgoud and J. Ben-Naim. Ranking-based semantics for argumentation frameworks. In V. S. W. Liu and J. Wijsen, editors, *Proceedings of the 7th International Conference on Scalable Uncertainty Management, (SUM'13)*, pages 134–147, 2013.
2. L. Amgoud and S. Vesic. Rich preference-based argumentation frameworks. *International Journal of Approximate Reasoning*, 55:585–606, 2014.
3. P. Baroni, M. Caminada, and M. Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.
4. P. Baroni and M. Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence Journal*, 171:675–700, 2007.

5. S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 640–647, 1993.
6. S. J. Brams and P. C. Fishburn. Voting procedures. In A. K. S. Kenneth J. Arrow and K. Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 1 of *Handbook of Social Choice and Welfare*, pages 173 – 236. Elsevier, 2002.
7. M. Caminada. Semi-stable semantics. In *Proceedings of the 1st International Conference on Computational Models of Argument (COMMA'06)*, pages 121–130, 2006.
8. C. Cayrol, C. Devred, and M.-C. Lagasquie-Schiex. Acceptability semantics accounting for strength of attacks in argumentation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'10)*, pages 995–996, 2010.
9. C. Cayrol and M.-C. Lagasquie-Schiex. Graduality in argumentation. *Journal of Artificial Intelligence Research*, 23:245–297, 2005.
10. S. Coste-Marquis, C. Devred, and P. Marquis. Prudent semantics for argumentation frameworks. In *17th IEEE International Conference on Tools with Artificial Intelligence (IC-TAI'05)*, pages 568–572, 2005.
11. S. Coste-Marquis, S. Konieczny, P. Marquis, and M. A. Ouali. Selecting extensions in weighted argumentation frameworks. In *Proceedings of the Conference on Computational Models of Argument (COMMA'12)*, pages 342–349, 2012.
12. C. da Costa Pereira, A. Tettamanzi, and S. Villata. Changing one’s mind: Erase or rewind? In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'11)*, pages 164–171, 2011.
13. D. Dubois, H. Fargier, and H. Prade. Refinements of the maximin approach to decision-making in fuzzy environment. *Fuzzy Sets and Systems*, 81:103–122, 1996.
14. P. Dung, P. Mancarella, and F. Toni. Computing ideal skeptical argumentation. *Artificial Intelligence Journal*, 171:642–674, 2007.
15. P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence Journal*, 77:321–357, 1995.
16. P. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge. Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artificial Intelligence Journal*, 175(2):457–486, 2011.
17. P. E. Dunne, W. Dvorák, and S. Woltran. Parametric properties of ideal semantics. *Artificial Intelligence Journal*, 202:1–28, 2013.
18. W. Dvorák. On the complexity of computing the justification status of an argument. In *Theory and Applications of Formal Argumentation (TFAFA 2011), Revised Selected Papers*, volume 7132 of *Lecture Notes in Computer Science*, pages 32–49. Springer, 2012.
19. S. Konieczny, P. Marquis, and S. Vesic. On supported inference and extension selection in abstract argumentation frameworks. Technical report, CRIL, CNRS - Univ. Artois, 2015.
20. D. Martínez, A. García, and G. Simari. Strong and weak forms of abstract argument defense. In *Proceedings of the 2nd International Conference on Computational Models of Argument (COMMA'08)*, pages 216–227. IOS Press, 2008.
21. D. C. Martínez, A. García, and G. Simari. An abstract argumentation framework with varied-strength attacks. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, pages 135–144, 2008.
22. H. Moulin. *Axioms of cooperative decision making*. Cambridge University Press, 1988.
23. Y. Wu and M. Caminada. A labelling-based justification status of arguments. *Studies in Logic*, 3(4):12–29, 2010.