Open Problems in Abstract Argumentation

Ringo Baumann and Hannes Strass

Computer Science Institute, Leipzig University

Abstract We give a list of currently unsolved problems in abstract argumentation. For each of the problems, we motivate why it is interesting and what makes it (apparently) hard to solve.

1 Introduction

Formal argumentation has established itself as an active subfield of artificial intelligence [16]. Argumentation is concerned with how conflicts between different pieces of knowledge, possibly involving preferences among them, can be resolved in a principled manner. The further subfield of *abstract argumentation* ignores the potential internal structure of arguments, and instead concentrates on the interaction between different arguments. The predominantly used approach is that by Dung [20], where argumentation scenarios are represented using *argumentation frameworks (AFs)* $F = (A, R)$ consisting of a set A of abstract arguments and a relation R of attacks between these arguments.

This seemingly lightweight formalism has led to a large amount of research around it. Gerd Brewka is among those who had a lasting impact on the field. With this paper, we want to honor his contributions and take the opportunity to point out some avenues for future work.

We do this by collecting together various open problems from different areas and presenting them all in one place.[1] The list we give here is not necessarily complete, nor is it representative. However, we think that it nicely illustrates the breadth of abstract argumentation research, and the various connections to other fields of mathematics and logic that have been discovered. For presentation purposes, we keep the common background to a minimum, and rather introduce the necessary background that is needed for each problem individually.

2 Background

In the following we consider a fixed countably infinite set $\mathcal{U}$ of arguments, called *universe*. Furthermore, we define $\mathcal{A} = \{F \mid F = (A, R), A \subseteq \mathcal{U}, R \subseteq A \times A\}$ containing all AFs w.r.t. $\mathcal{U}$. Instead of $(a, b) \in R$ we write $a \rightarrowtail b$ and say that

[1] Independently, Stefan Woltran had the same idea for his invited talk "Abstract Argumentation: All Problems Solved?" at ECAI 2014 (as part of the *Frontiers of Artificial Intelligence* series). We took up several suggestions for open problems from that talk and subsequent personal communication with Stefan.

a attacks b. For sets $E_1, E_2 \subseteq A$ and arguments $a, b \in A$ we say that $E_1 \rightarrowtail b$ if some $a \in E_1$ attacks b, $a \rightarrowtail E_2$ if a attacks some $b \in E_2$ and $E_1 \rightarrowtail E_2$ if some $a \in E_1$ attacks some $b \in E_2$. An argument $a \in A$ is defended by a set $E \subseteq A$ in F if for each $b \in A$ with $b \rightarrowtail a$, $E \rightarrowtail b$. The range E^+ of a set of arguments E is defined by the extension of E with all arguments attacked by E, i.e. $E^+ = E \cup \{a \in A \mid E \rightarrowtail a\}$.

A semantics σ is a function which assigns to any F a set of sets of arguments denoted by $\sigma(F)$. Each one of them, a so-called σ-extension, is considered to be acceptable with respect to F (for a recent overview see [1]). In the following we define conflict-free and admissible sets as well as complete, preferred, semi-stable, stable, naive, stage, grounded, ideal and eager semantics which will be frequently considered throughout the paper (abbreviated by *cf*, *adm*, *com*, *pr*, *ss*, *st*, *nai*, *stg*, *grd*, *id*, *eg*). Semantics that are used only once will be defined in the corresponding sections.

Definition 1. *Given an AF $F = (A, R)$. We call a set $E \subseteq A$*

1. $E \in cf(F)$ *if for all* $a, b \in E$ *we have* $a \not\rightarrowtail b$,
2. $E \in adm(F)$ *if* $E \in cf(F)$ *and for all* $a \rightarrowtail E$ *also* $E \rightarrowtail a$,
3. $E \in com(F)$ *if* $E \in adm(F)$ *and for any* $a \in A$ *defended by* E *in* F, $a \in E$,
4. $E \in pr(F)$ *if* $E \in adm(F)$ *and there is no* $E' \in adm(F)$ *s.t.* $E \subsetneq E'$,
5. $E \in ss(F)$ *if* $E \in adm(F)$ *and there is no* $E' \in adm(F)$ *s.t.* $E^+ \subsetneq E'^+$,
6. $E \in st(F)$ *if* $E \in cf(F)$ *and* $E^+ = A$,
7. $E \in nai(F)$ *if* $E \in cf(F)$ *and there is no* $E' \in cf(F)$ *s.t.* $E \subsetneq E'$,
8. $E \in stg(F)$ *if* $E \in cf(F)$ *and there is no* $E' \in cf(F)$ *s.t.* $E^+ \subsetneq E'^+$,
9. $E \in grd(F)$ *if* $E \in com(F)$ *and there is no* $E' \in com(F)$ *s.t.* $E' \subsetneq E$.
10. $E \in id(F)$ *if* $E \in adm(F)$, $E \subseteq \bigcap_{P \in pr(F)} P$ *and there is no* $E' \in adm(F)$ *satisfying* $E' \subseteq \bigcap_{P \in pr(F)} P$ *s.t.* $E \subsetneq E'$,
11. $E \in eg(F)$ *if* $E \in adm(F)$, $E \subseteq \bigcap_{P \in ss(F)} P$ *and there is no* $E' \in adm(F)$ *satisfying* $E' \subseteq \bigcap_{P \in ss(F)} P$ *s.t.* $E \subsetneq E'$.

3 Open Problems

1. Given an AF, can all implicit conflicts (pairs of arguments that do not occur jointly in any extension) be made explicit (by adding one or two attacks between them)?
2. What are the signatures (sets of extension-sets that can be realized by AFs under a semantics) of complete, cf2 and resolution-based grounded semantics?
3. What is the precise computational complexity of credulous acceptance with respect to ideal semantics?
4. What is the maximal number of complete extensions in an AF with n arguments?
5. Is there a closed-form expression for the average number of stable extensions of AFs with n arguments and x attacks?
6. What is the (σ, Φ)-characteristic of semi-stable semantics?
7. What is the (stable, semi-stable, preferred)-spectrum?
8. How can normal deletion equivalence in case of stage, semi-stable, eager, preferred, ideal and naive semantics be characterized?

3.1 Explicit-conflict conjecture

The fundamental building blocks of Dung's AFs are arguments. The fundamental means of *expression*, however, are attacks between arguments, as these ultimately influence which arguments can be accepted together. An attack between two arguments a and b is an explicit manifestation of a conflict between the two. But in addition to such syntactic, *explicit* conflicts, incompatibilities between arguments may also arise on the semantical level, that is, whenever two arguments never occur in an extension together. In such a case, we will speak about an *implicit* conflict. Clearly, for semantics based on conflict-freeness, each explicit conflict leads to an implicit conflict. But it is also possible to have implicit conflicts that are not explicit, as we show below in Figure 1. To make matters more formal, consider the following definition. Roughly, for a set X of sets of arguments (say, extensions), $Pairs_X$ captures which arguments co-occur in at least one of the elements of X. This relation directly yields implicit conflicts, and can be used to figure out whether there are implicit conflicts that are not explicit.

Definition 2. *Let* $X \subseteq 2^{\mathcal{U}}$ *and* $Pairs_X = \{(a, b) \mid \textit{exists } E \in X \textit{ s.t. } \{a, b\} \subseteq E\}$. *An AF* $F = (A, R)$ *is* conflict-explicit under semantics σ *iff for each* $a, b \in A$ *such that* $(a, b) \notin Pairs_{\sigma(F)}$*, we find* $(a, b) \in R$ *or* $(b, a) \in R$ *(or both).*

In words, a framework is conflict-explicit under σ if any two arguments of the framework that do not occur jointly in a σ-extension are explicitly conflicting, that is, there is an attack one way or the other.

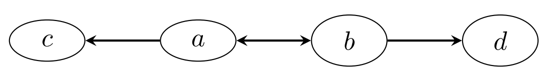

Figure 1: An argumentation framework that is not conflict-explicit under stable semantics. Observe that $st(F) = \{\{a, d\}, \{b, c\}\}$ *and* $(c, d) \notin Pairs_{\mathbb{S}}$ *but* $(c, d) \notin R$ *as well as* $(d, c) \notin R$*. If we add attacks* (c, d) *or* (d, c) *we obtain an equivalent (under stable semantics) conflict-explicit (under stable semantics) AF.*

The open problem now consists of proving or disproving whether every AF F has a conflict-explicit AF F' over the same arguments with the same stable extensions.

Conjecture 1 *For each AF* $F = (A, R)$ *there exists an AF* $F' = (A, R')$ *which is conflict-explicit under the stable semantics such that* $st(F) = st(F')$.

While formulating this conjecture is reasonably straightforward (it is perhaps the "easiest" conjecture of this paper, in terms of required background), Baumann et al. [13] have illustrated in a series of examples that the problem itself is far from easy. Clearly, given an argumentation framework F that is not conflict-explicit, our first try at making it conflict-explicit would be to add, for each conflict that is implicit but not explicit, an attack (or two). However, as Figure 2 shows, we cannot choose attacks to add at random. This creates a combinatorial problem, since for each of k non-explicit implicit conflicts, we have three possibilities of how to deal with it, thus 3^k possibilities in total. Even worse, just adding attacks

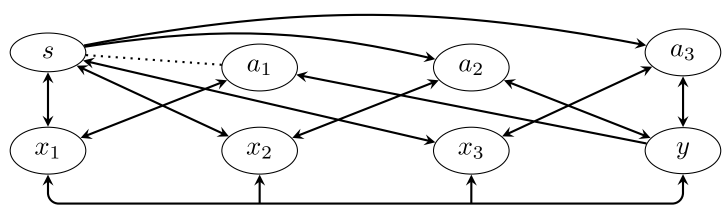

Figure 2: Orientation of attacks due to previously non-explicit conflicts matters: First observe that $st(F) = \{\{a_1, a_2, x_3\}, \{a_1, a_3, x_2\}, \{a_2, a_3, x_1\}, \{s, y\}\}$. *Next,* $Pairs_{st(F)}$ *yields one pair of arguments* a_1 *and* s *whose conflict is not explicit by* F, *that is,* $(a_1, s) \notin Pairs_{st(F)}$, *but* $(a_1, s), (s, a_1) \notin R_F$. *Now adding the attack* (a_1, s) *to* F *would create the additional stable extension* $\{a_1, a_2, a_3\} \notin st(F)$. *On the other hand, by adding the attack* (s, a_1), *we get the conflict-explicit AF* F' *with* $st(F) = st(F')$.

does not suffice in the general case. In an example that is too large to reproduce here, Baumann et al. [13] show that there are cases where one has to modify parts of the framework that are not directly involved in the implicit conflicts.

3.2 Signatures of complete, cf2 and resolution-based grounded semantics

Given an argumentation semantics σ, the *signature of* σ is the set

$$\Sigma_\sigma = \{\sigma(F) \mid F \text{ is an AF } \}.$$

That is, the signature of a semantics collects all sets of sets of arguments that can possibly arise as extension-set of some argumentation framework. This is a quite fundamental concept, since it provides a bird's eye view on capabilities and limitations of the semantics. For example, the signature of the grounded semantics clearly contains only (and all) singleton sets, since the grounded semantics is unique for any given AF, but an arbitrary singleton $\{E\}$ is realized by the AF $(E, \emptyset)$.

The notion of signature has been defined and studied by Dunne et al. [24,25], who also provide characterizations of the signatures for conflict-free, naive, stage, admissible, preferred and stable semantics. A characterization of Σ_σ consists of necessary and sufficient conditions that allow to decide (in a more sophisticated way than using brute force), given a set X of desired extensions, whether there exists an AF F such that $\sigma(F) = X$. For example, for the grounded semantics, the property of X being a singleton is both necessary and sufficient; therefore, the easily checkable singleton property precisely characterizes Σ_{grd}. For stable semantics, it is a necessary condition that X is a $\subseteq$-antichain, but this condition is not sufficient as the extension-set $X = \{\{a,b\},\{a,c\},\{b,c\}\}$ is not stable-realizable [25] (while being a $\subseteq$-antichain).

However, for several semantics, precise characterizations of their signatures are as yet unknown. Among these are the complete, cf2 and resolution-based grounded semantics. We will first recall some additional necessary technical prerequisites to formulate the open problems. However, for a lack of space, we have to refer the reader to [4] for details on the *cf2* semantics.[2] The resolution-based family of semantics is defined as follows [2]: for an AF $F = (A, R)$, a *resolution* of F is any AF $F' = (A, R')$ such that $R' \subseteq R$, $(a, a) \in R$ implies $(a, a) \in R'$, $(a, b) \in R$ with $a \neq b$ implies either $(a, b) \in R'$ or $(b, a) \in R'$ (but not both). Denoting the set of all resolutions of F by $\gamma(F)$, for a semantics σ, its resolution-based version σ^* is defined by

$$\sigma^*(F) = \min_{\subseteq} \left(\bigcup_{G \in \gamma(F)} \sigma(G) \right)$$

The resolution-based grounded semantics is then the grounded instance of this general scheme, that is, $rbg = grd^*$.

Now we can sketch the current state of knowledge and formulate the open problems: For complete semantics, we have $\Sigma_{adm} \subsetneq \Sigma_{com}$ [25]. For cf2, the current knowledge only says that $\Sigma_{nai} \subsetneq \Sigma_{cf2} \subsetneq \Sigma_{stg}$.[3] For the resolution-based grounded semantics, we know that $\Sigma_{rbg} \subsetneq \Sigma_{pr}$ and that Σ_{rbg} is incomparable to the signatures of naive, stage and stable semantics [26]. Thus the open problem is this:

Open Problem 2 *What are exact characterizations of Σ_{com}, Σ_{cf2}, Σ_{rbg}?*

3.3 Computational complexity of ideal semantics

The ideal semantics was introduced by Dung, Mancarella and Toni [21]. It covers an important middle ground between the grounded semantics (that is sometimes too restrictive) and sceptical reasoning over the preferred semantics (that is sometimes too permissive). As an illustration, consider Figure 3, an example taken from [22]. Recall that formally, for an argumentation framework $F = (A, R)$, a set $S \subseteq A$ is an *ideal set* if it is admissible and a subset of each preferred extension. Furthermore, S is the *ideal extension* if it is the $\subseteq$-maximal ideal set. Thus arises the question of the computational complexity of ideal semantics, that is,

[2] Roughly, the computation of *cf2* semantics proceeds along the strongly connected components of AFs. Naive extensions are determined in all components in the order of their dependence on one another, and statuses of arguments in previous SCCs are propagated to subsequent SCCs.

[3] Stefan Woltran, personal communication.

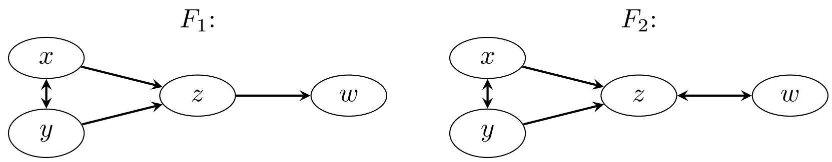

Figure 3: Two argumentation frameworks F_1, F_2. In both, the grounded extension is empty, and argument w is contained in every preferred extension. The ideal semantics can distinguish between the two, since in F_1 argument w cannot defend itself (ideal extension $\emptyset$) while in F_2 it can (ideal extension $\{w\}$).

whether its attractive properties (from a semantical standpoint) are (somewhat negatively) reflected in a high computational cost.

As a quick recapitulation [31], recall that the complexity class NP contains all problems L that have polytime-computable witness relation; that is, $L \in \mathsf{NP}$ iff there are $W_L \in \mathsf{P}$ and $k \in \mathbb{N}$ such that: $x \in L$ iff there is a y such that $(x, y) \in W_L$ and $|y| \leq |x|^k$. (Intuitively y is the polynomial-size witness proving that $x \in L$.) The class coNP contains all languages L whose complement $\overline{L}$ is in NP. The complexity class $\Theta_2^P = \mathsf{P}_{\|}^{\mathsf{NP}}$ contains all problems that are decidable in deterministic polynomial time using polynomially many non-adaptive calls to an NP oracle. An NP oracle call can be understood as having a constant-time decision subroutine for NP problems. Non-adaptive means that the oracle calls are independent of each other, that is, the answer to one oracle call may not influence a latter query to the oracle. (In the class Δ_2^P, on the other hand, oracle calls can build upon one another.) There is a special subclass of Θ_2^P, the class $\mathsf{DP} = \mathrm{D}_2^P$, where the number of oracle calls is exactly two. We clearly find that $\mathsf{NP} \subseteq \mathrm{D}_2^P \subseteq \Theta_2^P = \mathsf{P}_{\|}^{\mathsf{NP}} \subseteq \mathsf{P}^{\mathsf{NP}} = \Delta_2^P$.

Dunne [22] studies the following decision problems for ideal semantics:[4]

CAI Given $F = (A, R)$ and $a \in A$, is a contained in the ideal extension of F?
INE Given $F = (A, R)$, is its ideal extension non-empty?
VIE Given $F = (A, R)$ and $S \subseteq A$, is S the ideal extension?

Theorem 1 ([22, Theorem 1]). CAI *is* coNP*-hard;* **INE** *is* NP*-hard;* **VIE** *is* DP*-hard.*

Dunne [22] later provides conditional completeness results, dependent on knowing the exact complexity of **CAI**:

Theorem 2 ([22, Theorem 6]).

- *If* **CAI** *is* NP*-hard, then* **CAI** *is* $\mathsf{P}_{\|}^{\mathsf{NP}}$*-complete.*
- *If* **CAI** *is in* coNP*, then* **INE** *is* NP*-complete.*
- *If* **CAI** *is in* coNP*, then* **VIE** *is* DP*-complete.*

[4] The paper contains many more results, but for the purpose of this paper we are only interested in the open problems.

Thus many of the open problems rest on resolving whether **CAI** is NP-hard or **CAI** is in coNP. Currently, there is strong evidence that **CAI** is not in coNP. This evidence rests on the (open) complexity of the unique satisfiability problem (given a propositional formula φ, is there *exactly one* model for φ?), and randomised reductions [22]. Dunne [22] shows that with high probability:

Conjecture 1. **CAI**, **INE** and **VIE** are $\mathsf{P}_{\|}^{\mathsf{NP}}$-complete.

Dunne et al. [23] observed that the ideal semantics can be parameterized with respect to base semantics. They also conjecture the gap between the complexity of credulous and skeptical acceptance for preferred extensions to be a major influence on the difficulty in determining the precise complexity of ideal semantics.

3.4 Maximal number of complete extensions

In [14] the authors presented a first analytical and empirical study of the maximal and average numbers of extensions in case of abstract argumentation frameworks. The study was restricted to the case of stable semantics. In particular, it was shown that for any AF possessing n arguments the maximal number of stable extensions does not exceed $3^{\frac{n}{3}}$. Interestingly, the authors reduced the problem of determining the maximal number of stable extensions in argumentation frameworks to the problem of determining the maximal number of maximal independent sets in undirected graphs. The latter was already solved by John W. Moon and Leo Moser in 1965 [29].

We recapitulate the main theorem. The upper bound is presented as a function in the number of arguments denoted by $\sigma_{\max}(n)$.

Theorem 3 ([14, Theorem 1]). *In the case of stable sematics, the function* $\sigma_{\max} : \mathbb{N} \to \mathbb{N}$ *is given by*

$$\sigma_{\max}(n) = \begin{cases} 1, & \text{if } n = 0 \text{ or } n = 1, \\ 3^s, & \text{if } n \geq 2 \text{ and } n = 3s, \\ 4 \cdot 3^{s-1}, & \text{if } n \geq 2 \text{ and } n = 3s+1, \\ 2 \cdot 3^s, & \text{if } n \geq 2 \text{ and } n = 3s+2. \end{cases}$$

Recently, it was shown that $\sigma_{\max}(n)$ also serves as the maximal number of semi-stable, preferred, stage as well as naive extensions [25].

Why is it interesting to study the maximal number of extensions? The obtained results can be used to provide lower bounds for the minimal realizability of certain sets of extensions (cf. [13] for a detailed analysis). Furthermore, the results may yield upper bounds for algorithms computing extensions. Last but not least, the maximal number of extensions is simply a further criterion (or better, fundamental property) which helps to classify the plethora of argumentation semantics. This line of research was motivated and initiated by Pietro Baroni and Massimiliano Giacomin [3].

In case of admissible and conflict-free sets we may only state the naive bound 2^n in case of n arguments. This is due to the fact that for any set A and its associated AF $F_A = (A, \emptyset)$ we have $cf(F) = adm(F) = 2^A$. Up to now we were not able to find a proof for the maximal number of complete extensions.

Open Problem 3 *What is $\sigma_{\max}$ in case of complete semantics?*

We as well as our colleagues from Vienna, Thomas Linsbichler and Stefan Woltran, conjecture the following.

Conjecture 4 *In case of complete semantics $\sigma_{\max} : \mathbb{N} \to \mathbb{N}$ is given by*

$$\sigma_{\max}(n) = \begin{cases} 1, & \text{if } n = 0 \text{ or } n = 1, \\ 3^{\frac{n}{2}}, & \text{if } n \geq 2 \text{ and } n \text{ even}, \\ 4 \cdot 3^{\frac{n-3}{2}}, & \text{otherwise.} \end{cases}$$

To see that the maximal number is at least as large as conjectured consider the AFs $\mathcal{E}_n$ and $\mathcal{O}_n$ for even or odd n, respectively:

$$\mathcal{E}_n = \left(\left\{a_i, b_i \mid 1 \leq i \leq \frac{n}{2}\right\}, \left\{(a_i, b_i), (b_i, a_i) \mid 1 \leq i \leq \frac{n}{2}\right\}\right)$$
$$\mathcal{C}_3 = (\{a, b, c\}, \{(a, b), (b, a), (a, c), (c, a), (b, c), (c, b)\})$$
$$\mathcal{O}_n = \mathcal{C}_3 \cup \left(\left\{a_i, b_i \mid 1 \leq i \leq \frac{n-3}{2}\right\}, \left\{(a_i, b_i), (b_i, a_i) \mid 1 \leq i \leq \frac{n}{2}\right\}\right)$$

Obviously, for even n, $com(\mathcal{E}_n) = 3^{\frac{n}{2}}$ and for odd $n \geq 3$, $com(\mathcal{O}_n) = 4 \cdot 3^{\frac{n-3}{2}}$. To prove Conjecture 4, it would thus suffice to show that the given values are also upper bounds for the maximal number of complete extensions.

3.5 Average number of stable extensions

What is the average number of extensions for an AF possessing n arguments and k attacks? This means, we are interested in an expectation value *without actually inspecting* the AF except for determining the parameters n and k, which can be done in linear time. This problem was firstly tackled in [14] for the case of stable semantics. The authors presented some precise values, denoted by $\overline{\sigma}(n, k)$, given that the number of attacks k is close to 0 or close to n^2.

Proposition 1 ([14, Proposition 3]). *For any suitable*[5] *$n \in \mathbb{N}$, we have*

$$\overline{\sigma}(n, 0) = 1 \qquad \overline{\sigma}(n, n^2 - 3) = \frac{3 \cdot (n^2 - n - 1)}{(n + 1) \cdot (n^2 - 2)}$$
$$\overline{\sigma}(n, 1) = 1 - \frac{1}{n} \qquad \overline{\sigma}(n, n^2 - 2) = \frac{2}{n + 1}$$
$$\overline{\sigma}(n, 2) = 1 - \frac{2n - 2}{n^2 + n} \qquad \overline{\sigma}(n, n^2 - 1) = \frac{1}{n}$$

[5] Note that AFs do not possess negative numbers of attacks. Consequently, the considered n's have to ensure that the second argument of $\overline{\sigma}$ is non-negative.

The reason why the authors did not present a closed-form function is the enourmus combinatorial blowup which has to be handled efficently. Nevertheless, the achieved results can be used to show that the average number of stable extensions in the case of very small numbers of attacks approaches from below to 1. In the case of very large numbers of attacks we have a convergence to 0 from above. What happens in the middle ground? With an increasing number of attacks, does the average number of stable extensions just decrease in a monotone fashion? It turns out that while the number of attacks linearly increases, the average number of extensions first decreases, then increases and then decreases again. This observation is not restricted to a specific number of arguments (cf. [14, Figures 1 and 2, Table 1]). The main open problem of this section is a sufficiently precise specification of the function $\overline{\sigma}(n,k)$.

Open Problem 5 *What is* $\overline{\sigma}(n,k)$*?*

In this regard we present two conjectures supported by the analytical and empirical results in [14]. The first conjecture claims that the average number of stable extensions of AFs is always located in between 0 and 1.

Conjecture 6 *For any natural numbers* n *and* k *with* $0 < k < n^2$ *we have:*

$$0 < \overline{\sigma}(n,k) < 1.$$

The second conjecture claims that the local maximum always coincides with $n^2 - n$. This conjecture is precisely verified for AFs possessing at most 10 arguments (cf. [14, Table 1]).

Conjecture 7 *Let* $n \in \mathbb{N}$ *and define* $\overline{\sigma}^n(k) : \mathbb{N} \to \mathbb{R}$ *where* $\overline{\sigma}^n(k) = \overline{\sigma}(n,k)$*. Then,*

$$\overline{\sigma}^n(k) \text{ possesses a local maximum at } k_{max} = n^2 - n.$$

3.6 Minimal change problem for semi-stable semantics

More recently several problems regarding *dynamic* aspects of abstract argumentation have been addressed in the literature [18,19,17,27]. One much cited problem among these concerns the acceptability of certain arguments and is called *enforcing problem* [10]. This is, in brief, the question whether it is possible, given a specific set of allowed operations, to modify a given AF such that a desired set of arguments becomes an extension or a subset of an extension of the modified AF. Several sufficient conditions under which enforcements are (im)possible were identified.

Consider the following snapshot of a dialogue among agents A and B depicted in Figure 4. Assume it is A's turn and her desired set of arguments is $E = \{a_1, a_2, a_3\}$. Furthermore, A and B are discussing under preferred semantics.

In order to enforce E agent A may come up with new arguments (for example through introducing an argument which attacks b_2 and b_3) and/or question old arguments or attacks between them, respectively (for example through questioning the self-attack of c). Please note that firstly, in this scenario enforcing is possible

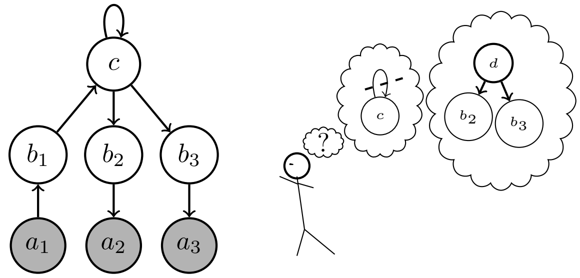

Figure 4: Snapshot of a Dialogue

and secondly, there are at least two different possibilities to achieve that. This observation leads to the more general problem of *minimal change* [7]. That is, in brief, i) is it possible to enforce a desired set of arguments, and if so, ii) what is the minimal number of modifications (additions or removals of attacks) to reach such an enforcement. This value, called (σ, Φ)-characteristic, depends on the underlying semantics σ and type of allowed modifications Φ. Here is the precise definition taken from [7].

Definition 3. *Given a semantics σ, an AF $F = (A, R)$ and a relation $\Phi \subseteq \mathcal{A} \times \mathcal{A}$. The (σ, Φ)-characteristic of a set $C \subseteq A$ is a natural number or infinity defined by the following function*

$$N^F_{\sigma,\Phi} : 2^A \to \mathbb{N}_\infty$$

$$C \mapsto \begin{cases} 0, & \exists C' : C \subseteq C' \text{ and } C' \in \sigma(F) \\ k, & k = \min\{d(F,G) \mid (F,G) \in \Phi, N^G_{\sigma,\Phi}(C) = 0\} \\ \infty, & \text{otherwise.} \end{cases}$$

The distance function $d(F, G)$ is defined as the number of added or removed attacks needed to transform F to G.

Quite surprisingly, it was shown that, in case of stable, preferred, complete and admissible semantics there are local criteria to determine the characteristic, although infinitely many possibilities to modify a given AF exist (see [9] for detailled explanations including all proofs). Let us consider again the dialouge depicted in Figure 4. Using the results in [7] one may show that the characteristic equals 1 if we allow arbitrary modifications, 2 if the deletion of former attacks is forbidden and ∞ (i.e. it is impossible to enforce $\{a_1, a_2, a_3\}$) if A only can come up with weaker arguments. These are fresh arguments which do not attack previous arguments.

Let F be an AF and Φ be a certain modification type. Due to the fact that any stable extension is a semi-stable one and furthermore, any semi-stable extension is preferred we have, $N^F_{st,\Phi} \geq N^F_{ss,\Phi} \geq N^F_{pr,\Phi}$ ([7, Corollary 3]). Whereas $N^F_{st,\Phi}$ and $N^F_{pr,\Phi}$ are already computable a characterization in case of semi-stable semantics remains an open problem.

Open Problem 8 *Are there local criteria determining* $N^F_{ss,\Phi}$*?*

The main reason why semi-stable semantics has defied any attempt of solving is the requirement of range-maximization which cannot be decided by looking at the candidate set only. On a final note, we want to mention that neither $N^F_{st,\Phi}$ nor $N^F_{pr,\Phi}$ coincide with $N^F_{ss,\Phi}$ in general (cf. examples at the end of Sections 4.1, 4.2 and 4.3 in [7]).

3.7 Spectra and Fibres

An at first glance more theoretical problem is the so-called *spectrum problem* [11]. The name was chosen because of its similarity with the famous *Spektralproblem* in model theory [32].[6] Given a certain semantics σ and a modification type Φ, the question is whether there is, for a given natural number n, an AF F and a set of arguments E such that n is the (σ, Φ)-characteristic of E with respect to F. In other words, we ask for the set of natural numbers which may occur as (σ, Φ)-characteristics, the so-called (σ, Φ)-*spectrum*. More generally, one may ask for n-tuples of characteristics representing several semantics simultaneously. Here is the definition of the (st, ss, pr, Φ)-spectrum.[7]

Definition 4. *Let* $\Phi \subseteq \mathcal{A} \times \mathcal{A}$*. The* (st, ss, pr, Φ)*-spectrum is a set of triples (so-called fibres) defined as follows:*

$$\mathcal{S}_{(st,ss,pr,\Phi)} = \{(k,l,m) \mid \exists \text{ } AF \text{ } F = (A,R) \; \exists \; C \subseteq A, \;\; s.t. \; N^F_{st,\Phi} = k, N^F_{ss,\Phi} = l \text{ and } N^F_{pr,\Phi} = m\}.$$

The first open problem is related to the spectrum w.r.t. to weak expansions, denoted by $\mathcal{S}_{(st,ss,pr,W)}$. In case of weak expansion the addition of weaker arguments, i.e. arguments which do not attack previous arguments, is allowed. In case of stable and preferred semantics it is already shown [7, Theorem 6] that there are only two possibilities, namely either a desired set is already contained in an extension, i.e. the characteristic equals zero, or the set is unenforceable, i.e. the characteristic equals infinity. Interestingly, semi-stable semantics does possess values between zero and infinity. Here is an example.

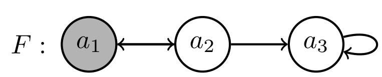

Figure 5: $N^F_{ss,W}(\{a_1\}) = 2$

In [11, Section 3.2] it is formally shown that $N^F_{ss,W}(\{a_1\}) = 2$ holds, indeed. Morever, the AF F and the set $\{a_1\}$ justify $(\infty, 2, 0) \in \mathcal{S}_{(st,ss,pr,W)}$. Unfortunately, up to now, there are no characterization theorems for semi-stable semantics

[6] Roughly speaking, Scholz investigated the possible sizes finite models of a first-order sentence may have.

[7] A more general definition including arbitrary n-tuples is given in [11, Definition 2].

(see Problem 8). Nevertheless, several results are already achieved and it turns out that a complete classification of $\mathcal{S}_{(st,ss,pr,W)}$ can be given provided that the following problem is solved.

Open Problem 9 *For any natural number* $n \geq 2$*,* $(\infty, n, 0) \in \mathcal{S}_{(st,ss,pr,W)}$*?*

Conjecture 10 *Yes!*

The reason why we believe *Yes!* is the following proposition stating that there are infinitely many numbers n between 2 and ∞, i.e. $(\infty, n, 0)$ is a fibre of the (st, ss, pr, W)*-spectrum.*

Proposition 2 ([11, Proposition 6]). *For any natural number* $n \in \mathbb{N}$ *there exists* $k \in \mathbb{N}$*, such that* $n \leq k \leq 2n$ *and* $(\infty, k, 0) \in \mathcal{S}_{(st,ss,pr,W)}$*.*

The second open problem regarding spectra and fibres is with respect to arbitrary modifications, so-called *updates* [8, Definition 5]. More precisely, what are the fibres of the corresponding (st, ss, pr)-spectrum, denoted by $\mathcal{S}_{(st,ss,pr,U)}$.

Open Problem 11 *What is* $\mathcal{S}_{(st,ss,pr,U)}$*?*

It is already shown that $(k, l, m) \in \mathcal{S}_{(st,ss,pr,U)}$ implies $k \geq l \geq m$ [11, Proposition 7]. This property is called *m.d.s.* – standing for "monotonic decreasing sequence". We conjecture that the considered spectrum is even *m.d.s.-complete*, i.e. for any $k \geq l \geq m$ we have $(k, l, m) \in \mathcal{S}_{(st,ss,pr,U)}$.

Conjecture 12 $\mathcal{S}_{(st,ss,pr,U)}$ *is m.d.s.-complete.*

To verify this conjecture one has to present witnessing AFs $F_{k,l,m}$ together with a certain set of arguments C, s.t. $N^{F_{k,l,m}}_{st,U}(C) = k$, $N^{F_{k,l,m}}_{ss,U}(C) = l$ and $N^{F_{k,l,m}}_{pr,U}(C) = m$. Due to the multitude of possibilities to modify a certain argumentation scenario together with the lack of local criteria to determine the semi-stable characteristic (Problem 8) we were unable to find a proof so far.

The determination of spectra yields interesting insights into how particular semantics are related. For instance, the fact that $\mathcal{S}_{(st,ss,pr,U)}$ is m.d.s. simply means that whenever enforcing is possible for all of them it is at least as difficult using stable (semi-stable) semantics as it is using semi-stable (preferred) semantics. If it is indeed m.d.s.-complete we know in addition that it can in fact be arbitrarily more difficult.

3.8 Characterizing Normal Deletion Equivalence

Notions of equivalence which guarantee intersubstitutability w.r.t. further modifications have received considerable interest in nonmonotonic reasoning (see [28,34,33] among others). Quite recently this topic emerged in abstract argumentation. In the following we list the notions considered in the literature in chronological order (see [15,12] for recent overviews).

1. expansion and local expansion equivalence [30][8]

[8] In [30] the authors used the term *strong equivalence* instead of expansion equivalence. Due to the different notions defined later, expansion equivalence became similarly popular since the term precisely characterizes the considered modifications.

2. weak expansion equivalence [5]
3. normal and strong expansion equivalence [6]
4. minimal change equivalence [7]
5. update, deletion, local deletion and normal deletion equivalence [8]

Much work has been done to characterize the mentioned equivalence notions. Many characterization theorems rely on *kernels* which are purely syntactical concepts. Quite surprisingly, so-called *context-sensitive* kernels originally introduced to characterize strong expansion equivalence even serve as parts of the characterizations of normal deletion equivalence w.r.t. admissible, complete and grounded semantics [8, Theorem 16]. Unfortunalety, further standard semantics have defied any characterization attempts.

Open Problem 13 *How to characterize normal deletion equivalence in case of stage, semi-stable, eager, preferred, ideal and naive semantics?*

We proceed with the precise definition togehter with an example.

Definition 5. *Two AFs $F = (A, R)$ and $G = (B, S)$ are normal deletion equivalent w.r.t. σ (denoted by $F \equiv^{\sigma}_{ND} G$) iff for any set of argumens C, $\sigma(F \setminus C) = \sigma(G \setminus C)$. Hereby $F \setminus C \stackrel{\text{def}}{=} \big((A, R)|_{A \setminus C}\big)$.*

Roughly speaking, normal deletion equivalent frameworks cannot be semantically distinguished by forgetting arguments together with their corresponding attacks.

Example 1. Consider the following two AFs F and G.

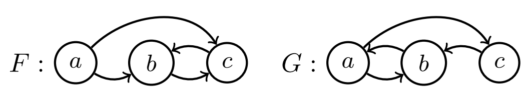

Although both possess the unique preferred extension $\{a\}$ the AFs are not normal deletion equivalent w.r.t. preferred semantics. This can be seen as follows. If we retract the argument c, then $\{b\}$ becomes preferred in $G \setminus \{c\}$ but still not in $F \setminus \{c\}$. Consequently, $F \not\equiv^{pr}_{ND} G$.

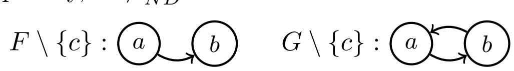

As a final note we mention that it is already checked that none of the existing kernels can contribute anything to solving Open Problem 13. This means, if kernels play a decisive role, then new kernel definitions are required.

4 Conclusion

We presented eight open problems in abstract argumentation, one of Gerd's major research areas in the last decade. For each of the problems, we tried to motivate why the problem is important, gave a formal problem statement and explained why the problem is (or at least seems to be) hard to solve. Some of the problems stem directly from work that Gerd was personally involved in, while others are inspired by his work.

Acknowledgements. We dedicate this paper to Gerhard Brewka, whose support and guidance was of substantial importance in our academic careers. *All the best, Gerd, and thank you for everything!* The authors are also grateful to Stefan Woltran for adding several open problems and acting as a second reader, thus suggesting numerous improvements of the text.

References

1. Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.
2. Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artificial Intelligence*, 175(3):791–813, 2011.
3. Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.
4. Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. SCC-recursiveness: A general schema for argumentation semantics. *Artificial Intelligence*, 168(1–2):162–210, 2005.
5. Ringo Baumann. Splitting an argumentation framework. In *LPNMR*, volume 6645 of *LNCS*, pages 40–53. Springer, 2011.
6. Ringo Baumann. Normal and strong expansion equivalence for argumentation frameworks. *Artificial Intelligence*, 193:18–44, 2012.
7. Ringo Baumann. What does it take to enforce an argument? Minimal change in abstract argumentation. In *ECAI*, pages 127–132, 2012.
8. Ringo Baumann. Context-free and context-sensitive kernels: Update and deletion equivalence in abstract argumentation. In *ECAI*, pages 63–68, 2014.
9. Ringo Baumann. *Metalogical Contributions to the Nonmonotonic Theory of Abstract Argumentation*. College Publications, 2014.
10. Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *COMMA*, volume 216 of *FAIA*, pages 75–86. IOS Press, 2010.
11. Ringo Baumann and Gerhard Brewka. Spectra in abstract argumentation: An analysis of minimal change. In *LPNMR*, pages 174–186. Springer, 2013.
12. Ringo Baumann and Gerhard Brewka. The equivalence zoo for dung-style semantics. *Journal of Logic and Computation: Special Issue*, 2014.
13. Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact Argumentation Frameworks. In *ECAI*, pages 69–74, 2014.
14. Ringo Baumann and Hannes Strass. On the Maximal and Average Numbers of Stable Extensions. In Elizabeth Black, Sanjay Modgil, and Nir Oren, editors, *Proceedings of the Second International Workshop on Theory and Applications of Formal Argumentation (TAFA 2013)*, volume 8306 of *LNAI*, pages 111–126. Springer, August 2014.
15. Ringo Baumann and Stefan Woltran. The role of self-attacking arguments in characterizations of equivalence notions. *Journal of Logic and Computation: Special Issue*, 2014.
16. Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

17. Pierre Bisquert, Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in argumentation systems: Exploring the interest of removing an argument. In Salem Benferhat and Jon Grant, editors, *SUM*, pages 275–288. Springer, 2011.
18. Guido Boella, Souhila Kaci, and Leendert W. N. van der Torre. Dynamics in argumentation with single extensions: Abstraction principles and the grounded extension. In *ECSQARU*, pages 107–118, 2009.
19. Claudette Cayrol, Florence Dupin de Saint-Cyr, and Marie-Christine Lagasquie-Schiex. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research*, pages 49–84, 2010.
20. Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321 – 357, 1995.
21. Phan Minh Dung, Paolo Mancarella, and Francesca Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171(10):642–674, 2007.
22. Paul E. Dunne. The computational complexity of ideal semantics. *Artificial Intelligence*, 173(18):1559–1591, 2009.
23. Paul E Dunne, Wolfgang Dvořák, and Stefan Woltran. Parametric properties of ideal semantics. *Artificial Intelligence*, 202:1–28, 2013.
24. Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In *Proceedings of DKB*, pages 16–30, 2013.
25. Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. In *KR*, pages 72–81, 2014.
26. Wolfgang Dvořák, Thomas Linsbichler, Emilia Oikarinen, and Stefan Woltran. Resolution-based grounded semantics revisited. In *COMMA*. IOS Press, 2014.
27. Bei Shui Liao, Li Jin, and Robert C. Koons. Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, pages 1790–1814, 2011.
28. Vladimir Lifschitz, David Pearce, and Agustín Valverde. Strongly equivalent logic programs. *ACM Transactions on Computational Logic*, pages 526–541, 2001.
29. John W. Moon and Leo Moser. On cliques in graphs. *Israel Journal of Mathematics*, pages 23–28, 1965.
30. Emilia Oikarinen and Stefan Woltran. Characterizing strong equivalence for argumentation frameworks. *Artificial Intelligence*, 175(14-15):1985–2009, 2011.
31. Christos H. Papadimitriou. *Computational complexity*. John Wiley and Sons Ltd., 2003.
32. Heinrich Scholz. Ein ungelöstes Problem in der symbolischen Logik. *Journal of Symbolic Logic*, 17:160, 1952.
33. Mirosław Truszczyński. Strong and uniform equivalence of nonmonotonic theories – an algebraic approach. *Annals of Mathematics and Artificial Intelligence*, pages 245–265, 2006.
34. Hudson Turner. Strong equivalence for causal theories. In *LPNMR*, pages 289–301, 2004.