Investigating the Relationship between Argumentation Semantics via Signatures

Paul E. Dunne and **Christof Spanring**
Department of Computer Science
University of Liverpool, UK
{p.e.dunne,c.spanring}@liverpool.ac.uk

Thomas Linsbichler and **Stefan Woltran**
Institute of Information Systems
TU Wien, Vienna, Austria
{linsbich,woltran}@dbai.tuwien.ac.at

Abstract

Understanding the relation between different semantics in abstract argumentation is an important issue, not least since such semantics capture the basic ingredients of different approaches to nonmonotonic reasoning. The question we are interested in relates two semantics as follows: What are the necessary and sufficient conditions, such that we can decide, for any two sets of extensions, whether there exists an argumentation framework which has exactly the first extension set under one semantics, and the second extension set under the other semantics. We investigate in total nine argumentation semantics and give a nearly complete landscape of exact characterizations. As we shall argue, such results not only give an account on the independency between semantics, but might also prove useful in argumentation systems by providing guidelines for how to prune the search space.

1 Introduction

Within Artificial Intelligence argumentation has become one of the major fields over the last two decades (Rahwan and Simari, 2009; Bench-Capon and Dunne, 2007). In particular, abstract argumentation frameworks (AFs) introduced by Dung (1995) are a simple, yet powerful formalism for modeling and deciding argumentation problems that are integral to many advanced argumentation systems, see e.g. (Caminada and Amgoud, 2007). Evaluating AFs is done via so-called semantics (cf. (Baroni *et al.*, 2011a) for an overview) that deliver subsets of jointly acceptable arguments.

Several semantics have been introduced over the years (Dung, 1995; Verheij, 1996; Caminada, 2007; Caminada *et al.*, 2012; Baroni *et al.*, 2011b). It is thus important to understand the behaviour of different semantics when applied to an AF. For instance, it is known that for any AF, its set of stable extensions is a subset of its set of preferred extensions (already proven by Dung (1995)). More recently, Dunne *et al.* (2015) have shown that for any AF F, its set of stable extensions (if not empty) can be realized via preferred semantics (i.e. there exists an AF F' such that the preferred extensions of F' equal the stable extensions of F). However, there is one aspect which has not been addressed yet. In fact, in this paper we are interested in questions of the following kind: Given sets $\mathbb{S}, \mathbb{T}$ of extensions, does there exist an AF F such that its stable extensions are given by $\mathbb{S}$ and its preferred extensions are given by $\mathbb{T}$. More formally, we are interested in characterizing the following concepts for semantics σ, τ, which we call two-dimensional signatures:

$$\Sigma_{\sigma,\tau} = \{\langle \sigma(F), \tau(F)\rangle \mid F \text{ is an AF }\}$$

The motivation for such work is manyfold. First it tells us about the independence between semantics. Let us again consider stable (*sb*) and preferred (*pr*) extensions and suppose we have two AFs F, F' with $sb(F) \subseteq pr(F')$. Is there also an AF F'' with $sb(F'') = sb(F)$ and $pr(F'') = pr(F')$? This might not always be possible since there are certain dependencies between the two semantics which can make the existence of such an F'' impossible, and – as we will show – this is indeed the case for this particular pair of semantics, i.e.

$$\Sigma_{sb,pr} \neq \{\langle sb(F), pr(F')\rangle \mid sb(F) \subseteq pr(F'); F, F' \text{ AFs}\} \quad (1)$$

However, for certain other pairs of semantics we shall prove such a strong form of independence; for instance, for naive (*na*) semantics, it is known that $sb(F) \subseteq na(F)$, and this is sufficient for the corresponding two-dimensional signature:

$$\Sigma_{sb,na} = \{\langle sb(F), na(F')\rangle \mid sb(F) \subseteq na(F'); F, F' \text{ AFs}\}$$

The second motivation for our work is that it helps to prune the search space for systems designed to enumerate all extensions of a given semantics τ. This is of particular interest when the complexity for some other semantics σ is milder than the one for τ. Again consider stable and preferred semantics, the latter being more complex (Dunne and Bench-Capon, 2002). Results like (1) indicate that for enumerating all preferred extensions, starting with the computation of all stable extensions not only yields a subset of the desired preferred extensions but ultimately rules out certain candidates to become preferred extensions, besides those being comparable (wrt. subset inclusion) to already obtained ones.

For an example in a more concrete application, consider examination of evidence and facts in a legal case. The nature of which subsets are argued may depend on the environment in which the case is argued, e.g. standards of proof differ between civil and criminal cases so that based on the *same* evidence different conclusions may be reached. One view of stable semantics is that for a fixed stable extension, S, every argument has a definite status: either in S or attacked by S. In a legal setting, stable semantics is appropriate to cases that have to be demonstrated "beyond reasonable doubt". In contrast preferred semantics allows some arguments to have an undetermined state with respect to a fixed S: neither in S nor attacked by S. Our results shed light on how much undeterminism preferred semantics may add in contrast to stable semantics in such a situation.

Related work. There has been thorough research on translations (Dvořák and Woltran, 2011; Dvořák and Spanring, 2012) where mappings θ are studied such that, for any AF F, $\sigma(\theta(F))$ is in a certain relation to $\tau(F)$. Naturally, these results are concerned with two different AFs; we on the other hand explore the range of pairs of extensions a single AF is able to express via two types of semantics. The already mentioned work by Dunne *et al.* (2015) has initiated this kind of research but treated semantics separately. Other results on dependencies between semantics occur in the work on spectra (Baumann and Brewka, 2013), which is concerned with enforcing single arguments.

Main Contributions. We aim for exactly characterizing all two-dimensional signatures for 9 prominent semantics. Due to symmetry this requires in total 36 results, from which we succeed to show 32. We also discuss the particular issue for the four open problems (which are all depending on the two-dimensional signature for preferred and semi-stable semantics). Hereby, we provide an interesting observation concerning the role of implicit conflicts (Linsbichler *et al.*, 2015).

2 Background

We first recall basic notions of Dung's abstract frameworks (the reader is referred to (Dung, 1995; Baroni *et al.*, 2011a) for further background).

An *Argumentation Framework* (AF) is a pair $F = (A, R)$, where $A \subset \mathfrak{A}$ is a finite set of *arguments* for $\mathfrak{A}$ being the (countably infinite) universe of all arguments available, and $R \subseteq A \times A$ is its *attack* relation. The collection of all AFs is given by $AF_{\mathfrak{A}}$. For $(a, b) \in R$ we say that a attacks b (in F), accordingly a set $S \subseteq A$ attacks an argument $a \in A$ (in F) if $\exists b \in S : (b, a) \in R$. The *range* in F of a set of arguments $S \subseteq A$ is given as $S^+_F = S \cup \{a \in A \mid S \textit{ attacks } a\}$. Subscript F may be dropped if clear from the context. A set $S \subseteq A$ *defends* argument $a \in A$ (in F) if S attacks all attackers of a.

A *semantics* σ is a mapping from AFs to sets of arguments. For a given AF $F = (A, R)$ the members of $\sigma(F)$ are called (σ-)*extensions*. A set $S \subseteq A$ is *conflict-free* in F ($S \in \mathit{cf}(F)$) if it does not contain any attacks, i.e. $(S \times S) \cap R = \emptyset$; $S \in \mathit{cf}(F)$ is *admissible* in F ($S \in \mathit{ad}(F)$) if each $a \in S$ is defended by S; $S \in \mathit{ad}(F)$ is *complete* ($S \in \mathit{co}(F)$) if S contains all $a \in A$ it defends. We define the *naive*, *stable*, *preferred* and *semi-stable* extensions:

- $S \in \mathit{na}(F)$, if $S \in \mathit{cf}(F)$ and $\nexists T \in \mathit{cf}(F)$ s.t. $S \subset T$;
- $S \in \mathit{sb}(F)$, if $S \in \mathit{cf}(F)$ and $S^+_F = A$;
- $S \in \mathit{pr}(F)$, if $S \in \mathit{ad}(F)$ and $\nexists T \in \mathit{ad}(F)$ s.t. $S \subset T$;
- $S \in \mathit{sm}(F)$, if $S \in \mathit{ad}(F)$ and $\nexists T \in \mathit{ad}(F)$ s.t. $S^+_F \subset T^+_F$.

Finally, for semantics σ, τ we define the *ideal reasoning* semantics (see e.g. (Dunne *et al.*, 2013)) for σ under τ ($id_{\sigma,\tau}$) as sets $S \in \sigma(F)$ being $\subseteq$-maximal in satisfying $S \subseteq T$ for each $T \in \tau(F)$. In this paper, we use the *grounded* ($\mathit{gr}(F) = id_{\mathit{ad},\mathit{co}}(F)$), *ideal* ($\mathit{id}(F) = id_{\mathit{ad},\mathit{pr}}(F)$) and *eager* ($\mathit{eg}(F) = id_{\mathit{ad},\mathit{sm}}(F)$) semantics. Since *gr*, *id* and *eg* always provide exactly one extension, we sometimes refer to this by *Gr*, *Id* and *Eg*; thus, for instance, $\mathit{gr}(F) = \{\mathit{Gr}(F)\}$.

Towards the characterization of signatures, we require a few more concepts, mostly taken from (Dunne *et al.*, 2015; Baumann *et al.*, 2014) (however, written in a slightly different way). A set of sets of arguments $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is called *extension-set* if $\bigcup \mathbb{S}$ is finite. Given an extension-set $\mathbb{S}$, we denote the $\subseteq$-maximal elements of $\mathbb{S}$ by $\mathit{max}(\mathbb{S})$. Moreover, we define the *conflicts* in $\mathbb{S}$ ($\mathit{Confs}_{\mathbb{S}}$) and the *borders* of $\mathbb{S}$ ($\mathit{bd}(\mathbb{S})$) as

$$\mathit{Confs}_{\mathbb{S}} = \{(a, b) \in \bigcup \mathbb{S} \times \bigcup \mathbb{S} \mid \nexists S \in \mathbb{S} : a, b \in S\}, \text{ and}$$

$$\mathit{bd}(\mathbb{S}) = \{T \subseteq \bigcup \mathbb{S} \mid b \in \bigcup \mathbb{S} \setminus T \textit{ iff } \exists a \in T : (a, b) \in \mathit{Confs}_{\mathbb{S}}\}.$$

Finally, given $\mathbb{S}, \mathbb{T} \subseteq 2^{\mathfrak{A}}$, $\mathbb{S}$ is called *conflict-sensitive* wrt. $\mathbb{T}$, or $\mathbb{S} \rtimes \mathbb{T}$ for short, if for all $A, B \in \mathbb{S}$ such that $A \cup B \notin \mathbb{S}$ there are $a \in A$, $b \in B$ with $(a, b) \in \mathit{Confs}_{\mathbb{T}}$.

Example 1. Let $\mathbb{S} = \{\{a, b\}, \{a, c, e\}, \{b, d, e\}\}$. First, $\mathbb{S}$ is an extension-set. Moreover, we have $\mathit{max}(\mathbb{S}) = \mathbb{S}$, $\mathit{Confs}_{\mathbb{S}} = \{(a, d), (d, a), (b, c), (c, b), (c, d), (d, c)\}$, and $\mathit{bd}(\mathbb{S}) = \{\{a, b, e\}, \{a, c, e\}, \{b, d, e\}\}$. Finally, $\mathbb{S} \rtimes \mathbb{S}$ as, for instance, $(b, c) \in \mathit{Confs}_{\mathbb{S}}$ for $\{a, b\}$ and $\{a, c, e\}$.

3 Signatures

We first recall the results from (Dunne *et al.*, 2015) (similar in style of presentation to (Baumann *et al.*, 2014)) on signatures and then generalize this concept to multiple semantics.

Definition 1. Given a semantics σ, a set $\mathbb{S} \subseteq 2^{\mathfrak{A}}$ is *realizable* under σ if there is an AF F with $\sigma(F) = \mathbb{S}$ (F realizes $\mathbb{S}$ under σ). The *signature* of σ is defined as $\Sigma_\sigma = \{\sigma(F) \mid F \in AF_{\mathfrak{A}}\}$.

Proposition 1. *The following collections of extension-sets $\mathbb{S}$ yield the signatures of the semantics under consideration.*

- $\Sigma_{\mathit{gr}} = \Sigma_{\mathit{id}} = \Sigma_{\mathit{eg}} = \{\mathbb{S} \mid |\mathbb{S}| = 1\}$;
- $\Sigma_{\mathit{cf}} = \{\mathbb{S} \neq \emptyset \mid \mathit{max}(\mathbb{S}) = \mathit{bd}(\mathbb{S}), \forall S \in \mathbb{S} \forall S' \subseteq S : S' \in \mathbb{S}\}$;
- $\Sigma_{\mathit{na}} = \{\mathbb{S} \neq \emptyset \mid \mathbb{S} = \mathit{bd}(\mathbb{S})\}$;
- $\Sigma_{\mathit{sb}} = \{\mathbb{S} \mid \mathbb{S} \subseteq \mathit{bd}(\mathbb{S})\}$;
- $\Sigma_{\mathit{ad}} = \{\mathbb{S} \neq \emptyset \mid \emptyset \in \mathbb{S}, \mathbb{S} \rtimes \mathbb{S}\}$;
- $\Sigma_{\mathit{pr}} = \Sigma_{\mathit{sm}} = \{\mathbb{S} \neq \emptyset \mid \mathbb{S} = \mathit{max}(\mathbb{S}), \mathbb{S} \rtimes \mathbb{S}\}$.

The signatures for $\sigma \in \{\mathit{id}, \mathit{eg}\}$ have not been given explicitly in (Dunne *et al.*, 2015), but they directly follow from the realization of any extension-set $\mathbb{S} \in \Sigma_\sigma$ by the AF $(\bigcup \mathbb{S}, \emptyset)$.

We now give (slight modifications of) the canonical AFs used in (Dunne *et al.*, 2015) to realize certain extension-sets; they are the basis for some constructions used in Section 4.

Definition 2. Given extension-set $\mathbb{S}$, define $F_{\mathit{sb}}(\emptyset) = (\{a\}, \{(a, a)\})$ and, for $\mathbb{S} \neq \emptyset$, $F_{\mathit{sb}}(\mathbb{S}) = (A, \mathit{Confs}_{\mathbb{S}} \cup R)$ with

$$A = \bigcup \mathbb{S} \cup \{x_S \mid S \in \mathit{bd}(\mathbb{S}) \setminus \mathbb{S}\};$$

$$R = \{(x_S, x_S), (a, x_S) \mid S \in \mathit{bd}(\mathbb{S}) \setminus \mathbb{S}, a \in \bigcup \mathbb{S} \setminus S\}.$$

Table 1: Summary of results.

	id	*eg*	*na*	*sb*	*sm*	*pr*	*cf*	*ad*
gr	Th.1	Th.1	Th.2	Th.3	Th.3	Th.3	Th.8	Th.10
id	-	Th.1	Th.2	Th.3	Th.3	Th.3	Th.8	Th.9
eg		-	Th.2	Th.3	Th.3	?	Th.8	?
na			-	Th.5	Th.5	Th.5	Th.7	Th.9
sb				-	Th.4	Th.6	Th.8	Th.9
sm					-	?	Th.8	?
pr						-	Th.8	Th.7
cf							-	Th.8

Definition 3. Given extension-set $\mathbb{S}$, let $\mathbb{S}^{ns} = \bigcup\mathbb{S}\setminus\bigcap\mathbb{S}$, and $\mathbb{S}_a = \{S{\in}\mathbb{S} \mid a{\in}S\}$ for $a \in \bigcup\mathbb{S}$. The defense formula for a in $\mathbb{S}$ is $\bigvee_{S\in\mathbb{S}_a}\bigwedge_{b\in S\setminus\{a\}} b$; $\Gamma_{\mathbb{S},a}$ is a logically equivalent formula in conjunctive normal form. Define $F_{pr}(\mathbb{S}) = F_{ad}(\mathbb{S}) = (A, \mathit{Confs}_{\mathbb{S}} \cup R)$, $F_{sm}(\mathbb{S}) = (A \cup \bar{A}, \mathit{Confs}_{\mathbb{S}} \cup R \cup \bar{R})$ with

$$A = \bigcup\mathbb{S} \cup \{C \in \Gamma_{\mathbb{S},a} \mid a \in \mathbb{S}^{ns}\};$$
$$R = \{(C,C),(C,a),(b,C) \mid a \in \mathbb{S}^{ns}, C \in \Gamma_{\mathbb{S},a}, b \in C\};$$
$$\bar{A} = \{\bar{a} \mid a \in \bigcup\mathbb{S}\}; \quad \bar{R} = \{(a,\bar{a}),(\bar{a},\bar{a}) \mid a \in \bigcup\mathbb{S}\}.$$

Proposition 2. *For $\sigma \in \{sb, pr, sm, ad\}$, given $\mathbb{S} \in \Sigma_\sigma$ it holds that $\sigma(F_\sigma(\mathbb{S})) = \mathbb{S}$.*

The natural generalization of signatures is now defined as follows. It captures the capabilities of AFs with respect to different sets of semantics.

Definition 4. Given semantics $\sigma_1, \ldots, \sigma_n$, their *(n-dimensional) signature* is defined as

$$\Sigma_{\sigma_1,\ldots,\sigma_n} = \{\langle\sigma_1(F), \ldots, \sigma_n(F)\rangle \mid F \in AF_{\mathfrak{A}}\}.$$

We say that AF F realizes $\langle\mathbb{S}_1, \ldots, \mathbb{S}_n\rangle$ under $(\sigma_1, \ldots, \sigma_n)$ if $\sigma_i(F) = \mathbb{S}_i$ for all $i \in \{1, \ldots, n\}$.

In this paper, we will restrict to two-dimensional signatures. The following observation is crucial. Given arbitrary semantics σ and τ it always holds for members $\langle\mathbb{S},\mathbb{T}\rangle$ of $\Sigma_{\sigma,\tau}$ that $\mathbb{S} \in \Sigma_\sigma$ and $\mathbb{T} \in \Sigma_\tau$. When characterizing the two-dimensional signatures we will omit this necessary condition by using the following abbreviation:

$$\langle\mathbb{S},\mathbb{T}\rangle_{\sigma,\tau} := \langle\mathbb{S},\mathbb{T}\rangle \in \Sigma_\sigma \times \Sigma_\tau$$

4 Characterizing 2-Dimensional Signatures

In this section, we give characterizations for the nine semantics we consider in this paper. Exploiting the obvious symmetry $\Sigma_{\sigma,\tau} = \{\langle\mathbb{S},\mathbb{T}\rangle \mid \langle\mathbb{T},\mathbb{S}\rangle \in \Sigma_{\tau,\sigma}\}$, characterizing in total 36 signatures is still required. We proceed as follows. First, we consider signatures involving only unique status semantics (that is *gr*, *id*, and *eg*); second, we characterize signatures involving a unique status semantics and a multiple status semantics (that is *na*, *sb*, *sm*, and *pr*); third, we consider the combinations of multiple status semantics; then, we characterize the signatures involving *cf* and *ad*; Table 1 summarizes our results, leaving four open questions, all being related to the issue of *sm* and *pr* which we discuss in Section 4.5.

4.1 Unique Status Semantics

It is well known that $Gr(F) \subseteq Id(F) \subseteq Eg(F)$ for any AF F. In order to characterize the respective signatures, the question occurs whether this condition is also sufficient. In other words, can we find, for any given finite sets S, T with $S \subseteq T$, AFs F, F' and F'' such that $Gr(F) = S$, $Id(F) = T$, $Gr(F') = S$, $Eg(F') = T$, $Id(F'') = S$, and $Eg(F'') = T$. The following example answers this question positively.

Example 2. For index sets I, J, K take into account the modular AF $F = (A, R)$ with $A = \{a_i, b_j, \bar{b}_j, c_k, c'_k, \bar{c}_k \mid i \in I, j \in J, k \in K\}$, and $R = \{(b_j, \bar{b}_j), (\bar{b}_j, b_j), (\bar{b}_j, \bar{b}_j) \mid j \in J\} \cup \{(c_k, c'_k), (c'_k, c_k), (c_k, \bar{c}_k), (\bar{c}_k, \bar{c}_k) \mid k \in K\}$. We have that $Gr(F) = \{a_i \mid i \in I\}$, $Id(F) = Gr(F) \cup \{b_j \mid j \in J\}$, and $Eg(F) = Id(F) \cup \{c_k \mid k \in K\}$. Hence, with any of I, J, K possibly empty, it is possible for *gr* and *id*, *id* and *eg*, or *gr* and *eg* to be in arbitrarily extending relationships.

The exact relations now immediately follow.

Theorem 1. *For $(\sigma,\tau) \in \{(gr, id), (gr, eg), (id, eg)\}$,*

$$\Sigma_{\sigma,\tau} = \{\langle\{S\},\{T\}\rangle_{\sigma,\tau} \mid S \subseteq T \subseteq \mathfrak{A}\}.$$

4.2 Unique vs. Multiple Status Semantics

We begin with the signatures involving naive semantics.

Theorem 2. *For $\sigma \in \{gr, id, eg\}$, it holds that*

$$\Sigma_{\sigma,na} = \{\langle\{S\},\mathbb{T}\rangle_{\sigma,na} \mid \exists T \in \mathbb{T} : S \subseteq T\}.$$

Proof. First observe that clearly for any AF the σ-extension is conflict-free and thus also contained in some naive extension.

Now take an extension-set $\mathbb{T} \in \Sigma_{na}$, some $T \in \mathbb{T}$ and some $S \subseteq T$ as given. Consider the AF $F = (A, R)$ with $A = \bigcup\mathbb{T} \cup \{x\}$ and $R = (\mathit{Confs}_{\mathbb{T}} \setminus \{(b,a) \mid a \in S, b \in \bigcup\mathbb{T} \setminus S\}) \cup \{(x,x),(x,b) \mid b \in \bigcup\mathbb{T} \setminus S\}$. It can be shown that $na(F) = \mathbb{T}$ and $\sigma(F) = \{S\}$. $\square$

For the two-dimensional signatures with $\tau \in \{pr, sm, sb\}$ we have the following condition which is more restrictive than with naive semantics (note that $\bigcap\emptyset = \mathfrak{A}$).

Proposition 3. *For $\sigma \in \{Gr, Id, Eg\}$, $\tau \in \{pr, sm, sb\}$, $(\sigma,\tau) \neq (Eg, pr)$, $\sigma(F) \subseteq \bigcap\tau(F)$ holds for any $F \in AF_{\mathfrak{A}}$.*

Example 3. Note that (Eg, pr) steps out of line here. This is, for instance, witnessed by $\mathbb{S} = \{\{a\}\}$ and $\mathbb{T} = \{\{a\},\{b\}\}$. Despite $\mathbb{S} \not\subseteq \bigcap\mathbb{T} = \emptyset$ we can realize $\langle\mathbb{S},\mathbb{T}\rangle$ under (eg, pr) by the AF $(\{a,b,x\}, \{(a,b),(b,a),(a,x),(x,x)\})$. In fact, the relaxed condition $\exists T \in pr(F) : Eg(F) \subseteq T$ obviously holds for each AF F. However, as we will see in Section 4.5, this is far from a sufficient condition for realizability.

The question whether we can select, for σ, an arbitrary subset of the intersection of all τ-extensions has to be tackled differently, depending on the number of τ-extensions. We first consider the case when τ provides exactly one extension.

Proposition 4. *For $(\sigma,\tau) \in \{(id, pr), (eg, sb), (eg, sm), (eg, pr)\}$ and $F \in AF_{\mathfrak{A}}$, if $|\tau(F)| = 1$ then $\sigma(F) = \tau(F)$.*

This property does not hold for other σ-τ-combinations. The following can be derived from Example 2.

Proposition 5. *For $(\sigma,\tau) \in \{(gr, sb), (gr, sm), (gr, pr), (id, sb), (id, sm)\}$ and arbitrary $S \subseteq T \subseteq \mathfrak{A}$ there is an AF F with $\sigma(F) = \{S\}$ and $\tau(F) = \{T\}$.*

If τ provides more than one extension, σ can select an arbitrary subset of the arguments occurring in every τ-extension.

Lemma 1. *Let $\sigma \in \{Gr, Id, Eg\}$, $\tau \in \{pr, sm, sb\}$, $F = (A, R)$ be an AF such that $\sigma(F) = \bigcap \tau(F)$ and $S \subseteq \sigma(F)$. If $|\tau(F)| > 1$ and each $a \in \sigma(F)$ defends itself in F, then there exists an AF F' with $\sigma(F') = S$ and $\tau(F') = \tau(F)$.*

Proof. The AF $F' = (A \cup \{x\}, R')$ with $R' = R \cup \{(x, x), (x, c), (d, x) \mid c \in \sigma(F) \setminus S, d \in A \setminus \sigma(F)\}$ provides the required properties. $\square$

We are now ready to give the exact signatures with the exception of (eg, pr). In order to realize pairs of extension-sets we build upon the canonical frameworks from Section 3.

Theorem 3. *Let $(\sigma_1, \tau_1) \in \{(id, sm), (gr, sm), (gr, pr)\}$, $\sigma_2 \in \{id, gr\}$, $(\sigma_3, \tau_3) \in \{(eg, sm), (id, pr)\}$. It holds that*

$$\Sigma_{\sigma_1,\tau_1} = \{\langle\{S\}, \mathbb{T}\rangle_{\sigma_1,\tau_1} \mid S \subseteq \bigcap \mathbb{T}\};$$

$$\Sigma_{\sigma_2,sb} = \{\langle\{S\}, \mathbb{T}\rangle_{\sigma_2,sb} \mid \mathbb{T} \neq \emptyset, S \subseteq \bigcap \mathbb{T}\} \cup \{\langle\{S\}, \emptyset\rangle_{\sigma_2,sb}\};$$

$$\Sigma_{\sigma_3,\tau_3} = \{\langle\{S\}, \mathbb{T}\rangle_{\sigma_3,\tau_3} \mid |\mathbb{T}| > 1, S \subseteq \bigcap \mathbb{T}\} \cup \{\langle\{S\}, \{S\}\rangle_{\sigma_3,\tau_3}\};$$

$$\Sigma_{eg,sb} = \{\langle\{S\}, \mathbb{T}\rangle_{eg,sb} \mid |\mathbb{T}| > 1, S \subseteq \bigcap \mathbb{T}\} \cup \{\langle\{S\}, \{S\}\rangle_{eg,sb}, \langle\{S\}, \emptyset\rangle_{eg,sb}\}.$$

Proof. The $\subseteq$-direction is by Propositions 3 and 4. The proof of the $\supseteq$-direction proceeds as follows: If $|\mathbb{T}| > 1$ we can take $F_\tau(\mathbb{T})$, which has $\tau(F_\tau(\mathbb{T})) = \mathbb{T}$ by Proposition 2. Moreover, each argument $a \in \bigcap \tau(F_\tau(\mathbb{T}))$ is unattacked, meaning $\sigma(F_{pr}(\mathbb{T})) = \bigcap \tau(F_\tau(\mathbb{T}))$ and a defends itself in $F_\tau(\mathbb{T})$. Therefore we know by Lemma 1 that there is an AF $F'_\tau(\mathbb{T})$ with $\sigma(F'_\tau(\mathbb{T})) = \mathbb{S}$ and $\tau(F'_\tau(\mathbb{T})) = \mathbb{T}$. If, on the other hand, $|\mathbb{T}| = 1$, we can realize $\langle\{S\}, \{S\}\rangle_{\sigma,\tau}$ under (σ_3, τ_3) and (eg, sb) by the AF $(S, \emptyset)$ and get the result for the other σ-τ-combinations by Proposition 5. Finally, we can realize $\langle\{S\}, \emptyset\rangle_{\sigma,sb}$ under (σ, sb) by the AF $(S \cup \{x\}, \{(x, x)\})$. $\square$

4.3 Multiple Status Semantics

Coming to signatures involving two multiple status semantics we begin with the following result, which is immediate by the basic property that stable and semi-stable extensions coincide in case a stable extension exists.

Theorem 4. $\Sigma_{sb,sm} = \{\langle\mathbb{T}, \mathbb{T}\rangle_{sb,sm}\} \cup \{\langle\emptyset, \mathbb{T}\rangle_{sb,sm}\}$.

Stable extensions must always be maximal conflict-free sets, while preferred (resp. semi-stable) extensions can be arbitrary conflict-free sets as long as they are conflict-sensitive.

Theorem 5. *Let $\sigma \in \{sm, pr\}$. It holds that*

$$\Sigma_{na,sb} = \{\langle\mathbb{S}, \mathbb{T}\rangle_{na,sb} \mid \mathbb{T} \subseteq \mathbb{S}\};$$
$$\Sigma_{na,\sigma} = \{\langle\mathbb{S}, \mathbb{T}\rangle_{na,\sigma} \mid \forall T \in \mathbb{T}\, \exists S \in \mathbb{S} : T \subseteq S, \mathbb{T} \rtimes \mathbb{S}\}.$$

Proof. (na, sb): The $\subseteq$-direction follows immediately by the fact that for any AF F always $sb(F) \subseteq na(F)$. For the $\supseteq$-direction let $\langle\mathbb{S}, \mathbb{T}\rangle_{na,sb}$ with $\mathbb{T} \subseteq \mathbb{S}$. We define $F_{na,sb}(\mathbb{S}, \mathbb{T}) = (A, R)$ with $A = \bigcup \mathbb{S} \cup \{x_S \mid S \in \mathbb{S} \setminus \mathbb{T}\}$ and $R = Confs_{\mathbb{S}} \cup \{(x_S, x_S), (a, x_S) \mid S \in \mathbb{S} \setminus \mathbb{T}, a \in \bigcup \mathbb{S} \setminus S\}$. As $\mathbb{S} = bd(\mathbb{S})$ by $\mathbb{S} \in \Sigma_{na}$, it is clear that $na(F_{na,sb}(\mathbb{S}, \mathbb{T})) = \mathbb{S}$ and each $S \in \mathbb{S}$ attacks each $a \in \bigcup \mathbb{S} \setminus S$. Moreover, all elements of $\mathbb{S}$ are incomparable wrt. $\subseteq$, therefore each $T \in \mathbb{T}$ attacks each x_S with $S \in \mathbb{S} \setminus \mathbb{T}$, hence $T \in sb(F_{na,sb}(\mathbb{S}, \mathbb{T}))$. Finally, each $S \in \mathbb{S} \setminus \mathbb{T}$ does not attack x_S, hence $S \notin sb(F_{na,sb}(\mathbb{S}, \mathbb{T}))$. We conclude that $sb(F_{na,sb}(\mathbb{S}, \mathbb{T})) = \mathbb{T}$.

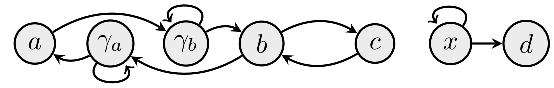

Figure 1: $F_{na,pr}(\{\{a, b, d\}, \{a, c, d\}\}, \{\{a, b\}, \{c\}\})$.

(na, σ): $\subseteq$: Let $F \in AF_{\mathfrak{A}}$. First observe that each $S \in \sigma(F)$ is conflict-free, i.e. $\exists T \in na(F)$ s.t. $S \subseteq T$. Now assume $\sigma(F) \rtimes na(F)$ does not hold, witnessed by $A, B \in \sigma(F)$ with $\forall a \in A, b \in B : (a, b) \notin Confs_{na(F)}$. But as A and B are admissible then also $A \cup B$ must be, a contradiction to A and B being (range-)maximal admissible sets. $\supseteq$: A pair $\langle\mathbb{S}, \mathbb{T}\rangle_{na,\sigma}$ is realized under (na, σ) by the AFs $F_{na,pr}(\mathbb{S}, \mathbb{T}) = (A, R)$ with $A = \bigcup \mathbb{S} \cup \{C \in \Gamma_{\mathbb{T},a} \mid a \in \mathbb{T}^{ns}\} \cup \{x\}$ and $R = Confs_{\mathbb{S}} \cup \{(C, C), (C, a), (b, C) \mid C \in \Gamma_{\mathbb{T},a}, b \in C\} \cup \{(x, x), (x, s) \mid s \in \bigcup \mathbb{S} \setminus \bigcup \mathbb{T}\}$, and $F_{na,sm}(\mathbb{S}, \mathbb{T}) = (A \cup \bar{A}, R \cup \bar{R})$ with $\bar{A} = \{\bar{a} \mid a \in \bigcup \mathbb{T}\}$ and $\bar{R} = \{(a, \bar{a}), (\bar{a}, \bar{a}) \mid a \in \bigcup \mathbb{T}\}$. These are slight modifications of the AF $F_{pr}(\mathbb{T})$ and $F_{sm}(F)$ from Definition 3 with the main difference that the attacks among arguments $\bigcup \mathbb{T}$ are not given by $Confs_{\mathbb{T}}$ but by $Confs_{\mathbb{S}}$. Thus, we immediately get $na(F_{na,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{S}$. For *pr* and *sm* note the role of the arguments $C \in \Gamma_{\mathbb{T},a}$, as shown in (Dunne *et al.*, 2015): the only admissible sets can be (unions of) elements of $\mathbb{T}$. As $Confs_{\mathbb{S}} \subseteq Confs_{\mathbb{T}}$ we get that each $T \in \mathbb{T}$ is admissible in $F_{na,\sigma}(\mathbb{S}, \mathbb{T})$. Moreover, for $T_1, T_2 \in \mathbb{T}$ we know by $\mathbb{T} \rtimes \mathbb{S}$ that there are arguments $a_1 \in T_1$ and $a_2 \in T_2$ with $(a_1, a_2) \in Confs_{\mathbb{S}}$. Hence $T_1 \cup T_2 \notin ad(F_{na,\sigma}(\mathbb{S}, \mathbb{T}))$. Finally, arguments $s \in \bigcup \mathbb{S} \setminus \bigcup \mathbb{T}$ cannot be defended in $F_{na,\sigma}(\mathbb{S}, \mathbb{T})$ by construction. Therefore the result for *pr* and *sm* follows. $\square$

Example 4. We illustrate the proof for (na, pr) by showing a concrete realization by the AF $F_{na,pr}$. Let $\mathbb{S} = \{\{a, b, d\}, \{a, c, d\}\}$ and $\mathbb{T} = \{\{a, b\}, \{c\}\}$. First note that $\mathbb{S} \in \Sigma_{na}$, $\mathbb{T} \in \Sigma_{pr}$, $\forall T \in \mathbb{T} \exists S \in \mathbb{S} : T \subseteq S$, and $\mathbb{T} \rtimes \mathbb{S}$ hold, hence $\langle\mathbb{S}, \mathbb{T}\rangle$ is (na, pr)-realizable. We get CNF defense formulas $\Gamma_{\mathbb{T},a} = \{\{b\}\}$, $\Gamma_{\mathbb{T},b} = \{\{a\}\}$, and $\Gamma_{\mathbb{T},c} = \emptyset$. The realizing AF $F_{na,pr}(\mathbb{S}, \mathbb{T})$ is depicted in Figure 1, with arguments $C \in \Gamma_{\mathbb{T},z}$ denoted by γ_z. It is easy to verify that $na(F_{na,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{S}$ and $pr(F_{na,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{T}$.

For the two-dimensional signature of stable and preferred semantics we first observe that for any pair $\langle\mathbb{S}, \mathbb{T}\rangle \in \Sigma_{sb,pr}$ it holds that $\mathbb{S} \subseteq \mathbb{T}$. However, already the fact that $\Sigma_{sb} \subset \Sigma_{pr}$ (Dunne *et al.*, 2015) implies that this condition cannot be sufficient as not all pairs with $\mathbb{S} = \mathbb{T}$ are realizable under (sb, pr). The following example illustrates that stable extensions may only be certain subsets of the preferred extensions.

Example 5. Consider the extension-sets $\mathbb{S} = \{\{a, d, e\}, \{b, c, e\}, \{c, d, e\}\}$ and $\mathbb{T} = \mathbb{S} \cup \{\{a, b\}\}$. The pair $\langle\mathbb{S}, \mathbb{T}\rangle$ is realized under (sb, pr) by the AF F depicted in Figure 2 (without the dotted or dashed parts). However, observe that $\mathbb{T} \not\subseteq bd(\mathbb{T})$ (since $\{a, b, e\} \in bd(\mathbb{T})$) and therefore $\langle\mathbb{T}, \mathbb{T}\rangle$ is not realizable under (sb, pr). In fact, no AF with $\mathbb{T}$ as preferred extensions can contain $\{a, b\}$ as stable extension since

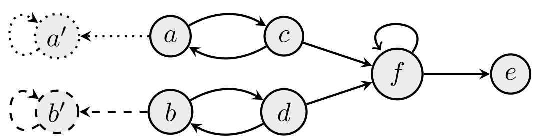

Figure 2: AFs F, F' and F'' used in Examples 5 and 6.

$\{a, b\} \notin bd(\mathbb{T})$ cannot achieve full range in such an AF. We can get an arbitrary subset of $\mathbb{S}$ under *sb* though: take AF F' in Figure 2 including the dotted part (argument a' and attacks (a, a), (a, a')). Then, $pr(F') = \mathbb{T}$ and $sb(F') = \{\{a, d, e\}\}$.

In order to make these ideas formal we first extend the AF F_{pr} from Definition 3 using techniques from Definition 2.

Definition 5. Given $\langle \mathbb{S}, \mathbb{T} \rangle_{sb,pr}$, $F_{sb,pr}(\mathbb{S}, \mathbb{T})$ is obtained from $F_{pr}(\mathbb{T})$ with additional arguments $\{x_T \mid T \in \mathbb{T} \setminus \mathbb{S}\}$ and attacks

$$\{(b, C) \mid C \in \Gamma_{\mathbb{T},a}, (a, b) \in Confs_{\mathbb{T}}\} \cup \quad (2)$$

$$\{(x_T, x_T), (a, x_T) \mid T \in \mathbb{T} \setminus \mathbb{S}, a \in \bigcup \mathbb{T} \setminus T\}. \quad (3)$$

Intuitively, taking $F_{pr}(\mathbb{T})$ as a basis ensures that $pr(F_{sb,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{T}$. Moreover, (2) guarantees that each $T \in \mathbb{T}$ also attacks each other argument of $F_{pr}(\mathbb{T})$. Finally, all $\mathbb{T} \setminus \mathbb{S}$ are excluded from $sb(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$ by (3).

Proposition 6. *For each $\langle \mathbb{S}, \mathbb{T} \rangle_{sb,pr}$ such that $\mathbb{S} \subseteq \mathbb{T} \cap bd(\mathbb{T})$ it holds that $sb(F_{sb,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{S}$ and $pr(F_{sb,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{T}$.*

Proof. It can be checked that $ad(F_{pr}(\mathbb{T})) = ad(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$. Therefore, by $pr(F_{pr}(\mathbb{T})) = \mathbb{T}$ (cf. Proposition 2), $pr(F_{sb,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{T}$ follows.

It remains to show that $sb(F_{sb,pr}(\mathbb{S}, \mathbb{T})) = \mathbb{S}$. $\supseteq$: Let $S \in \mathbb{S}$. By assumption $S \in \mathbb{T} \cap bd(\mathbb{T})$, hence $S \in pr(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$ and for all $a \in \bigcup \mathbb{T} \setminus S$ there is some $s \in S$ with $(a, s) \in Confs_{\mathbb{T}}$, meaning that s attacks a in $F_{sb,pr}(\mathbb{S}, \mathbb{T})$. Arguments x_T for $T \in \mathbb{T} \setminus \mathbb{S}$ are attacked by S since $S \in \mathbb{T}$ and elements in $\mathbb{T}$ are pairwise incomparable, hence there is some $a \in S \setminus T$ attacking x_T. Finally let $C \in \Gamma_{\mathbb{T},a}$ for an arbitrary $a \in \bigcup \mathbb{T}$. If there is some $s \in S$ with $s \in C$ then s attacks C. If $S \cap C = \emptyset$ it follows, by construction of $\Gamma_{\mathbb{T},a}$, that for all $s \in S$, $(a, s) \in Confs_{\mathbb{T}}$. By (3) now $(s, C) \in R$. It follows that S has full range in $F_{sb,pr}(\mathbb{S}, \mathbb{T})$, hence $S \in sb(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$. $\subseteq$: Let $E \in sb(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$. We get $E \in pr(F_{sb,pr}(\mathbb{S}, \mathbb{T}))$, hence $E \in \mathbb{T}$. Moreover, as E must have full range, $E \in bd(\mathbb{T})$. Assuming $E \notin \mathbb{S}$ means that there is some $x_E \in A$ which is only attacked by the arguments $\bigcup \mathbb{T} \setminus E$, a contradiction. $\square$

The fact that for stable extensions every argument has a definite status (of in or attacked) together with Propositions 1 and 6 now delivers the following.

Theorem 6. $\Sigma_{sb,pr} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{sb,pr} \mid \mathbb{S} \subseteq \mathbb{T} \cap bd(\mathbb{T})\}$.

4.4 Conflict-Free and Admissible Sets

So far we have disregarded conflict-free and admissible sets from our analysis. This is because of their close connection to naive and preferred semantics:

Theorem 7. *For $(\sigma, \tau) \in \{(cf, na), (ad, pr)\}$, it holds that*

$$\Sigma_{\sigma,\tau} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{\sigma,\tau} \mid max(\mathbb{S}) = \mathbb{T}\}.$$

By transitivity we get the following signatures.

Theorem 8. *For $\sigma \in \{gr, id, eg\}$, $\tau \in \{ad, sm, pr\}$,*

$$\Sigma_{\sigma,cf} = \{\langle \{S\}, \mathbb{T} \rangle_{\sigma,cf} \mid S \in \mathbb{T}\};$$
$$\Sigma_{cf,sb} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{cf,sb} \mid \mathbb{T} \subseteq max(\mathbb{S})\};$$
$$\Sigma_{cf,\tau} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{cf,\tau} \mid \mathbb{T} \subseteq \mathbb{S}, \mathbb{T} \rtimes \mathbb{S}\}.$$

The ideal extension of an AF is determined by its admissible sets. The crucial point for the signatures $\Sigma_{na,ad}$ and $\Sigma_{sb,ad}$ is the fact that changes to F_{pr} made by realizations in Theorem 5 and Proposition 6 do not affect the admissible sets.

Theorem 9. *It holds that*

$$\Sigma_{id,ad} = \{\langle \{S\}, \mathbb{T} \rangle_{id,ad} \mid \{S\} = max(\{T \in \mathbb{T} \mid T \subseteq \bigcap max(\mathbb{T})\})\};$$
$$\Sigma_{na,ad} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{na,ad} \mid \forall T \in \mathbb{T}\, \exists S \in \mathbb{S} : T \subseteq S, \mathbb{T} \rtimes \mathbb{S}\};$$
$$\Sigma_{sb,ad} = \{\langle \mathbb{S}, \mathbb{T} \rangle_{sb,ad} \mid \mathbb{S} \subseteq \mathbb{T} \cap bd(\mathbb{T})\}.$$

Interestingly, $\Sigma_{gr,ad}$ cannot be directly obtained from the case of preferred extensions, since certain subsets of $\bigcap max(\mathbb{T})$ need to lead to the required grounded extension.

Theorem 10. *It holds that*

$$\Sigma_{gr,ad} = \{\langle \{S\}, \mathbb{T} \rangle_{gr,ad} \mid \exists \textit{ strict total order } < \textit{ on } S \textit{ s.t.}$$
$$\forall s \in S : \{s\} \cup \{s' \in S \mid s' < s\} \in \mathbb{T}, S \subseteq \bigcap max(\mathbb{T})\}.$$

Proof. $\subseteq$: Consider AF $F = (A, R)$. We have $Gr(F) \subseteq \bigcap pr(F) = \bigcap max(ad(F))$. $Gr(F)$ is the least fixpoint of the monotone *characteristic function* $\mathcal{F}_F$, mapping each set $S \subseteq A$ to the set of arguments defended by S. Hence, given some $S \in ad(F)$, for all $S' \subseteq \mathcal{F}_F(S) \setminus S$ also $S \cup S' \in ad(F)$. Now let S_i be the set of arguments added by the ith application of $\mathcal{F}_F$ before reaching the fixpoint $Gr(F)$ (beginning with $\emptyset$). Define the strict total order $<$ such that $s < t$ if $s \in S_i$, $t \in S_j$ and $i < j$ and arbitrarily for $s, t \in S_i$. By the observations above, $\{s\} \cup \{s' \in S \mid s' < s\} \in ad(F)$ for all $s \in S$.

$\supseteq$: Let $\langle \{S\}, \mathbb{T} \rangle_{gr,ad}$ such that $S \subseteq \bigcap max(\mathbb{T})$ and there is a strict total order $<$ on S s.t. $\forall s \in S : \{s\} \cup \{s' \in S \mid s' < s\} \in \mathbb{T}$. Define $F_{gr,ad}(S, \mathbb{T})$ as $F_{pr}(\mathbb{T})$ with additional argument x and additional attacks $\{(x, x), (x, c), (c, x) \mid c \in \bigcap max(\mathbb{T}) \setminus S\}$. We get $ad(F_{gr,ad}(S, \mathbb{T})) = \mathbb{T}$ from Proposition 2 and one can show that also $Gr(F_{gr,ad}(S, \mathbb{T})) = S$. $\square$

4.5 Semi-Stable vs. Preferred Semantics

We now turn to the combination of preferred and semi-stable semantics. It is known that $sm(F) \subseteq pr(F)$ for every AF F. Moreover, it follows from Theorem 6 and $sb(F) = sm(F)$ for $sb(F) \neq \emptyset$, that $\mathbb{S} \subseteq \mathbb{T} \cap bd(\mathbb{T})$ is a sufficient condition for realizing $\langle \mathbb{S}, \mathbb{T} \rangle$ under (sm, pr). However, as the next example illustrates, the exact characterization of $\Sigma_{sm,pr}$ must lie somewhere in between these two conditions.

Example 6. Again consider $\mathbb{S}$ and $\mathbb{T}$ from Example 5. Obviously the AF F in Figure 2 realizes $\langle \mathbb{S}, \mathbb{T} \rangle$ also under (sm, pr). But in contrast to before we can now realize $\langle \mathbb{T}, \mathbb{T} \rangle$, namely by the AF F'' including all dotted or dashed arguments and attacks in Figure 2. It is, however, not possible to have $\{a, b\}$ as only extension under *sm*. As it turns out, in any AF G having $pr(G) = \mathbb{T}$, desiring $\{a, b\} \in sm(G)$ requires at least two members of $\mathbb{S}$ to be semi-stable extensions of G as well.

As a side remark, note that this also means that $\langle\{a,b\},\mathbb{T}\rangle \notin \Sigma_{eg,pr}$. Hence, for realizing pairs $\langle\{S\},\mathbb{T}'\rangle_{eg,pr}$ under (eg,pr), the fact that $\exists T \in \mathbb{T}' : S \subseteq T$ is not sufficient, as already suggested in Example 3.

To argue that the exact characterization turns out to be subtle, we define a stricter form of realizability, where each conflict is given explicitly as an attack in the realizing AF.

Definition 6. A pair of extension-sets $\langle\mathbb{S},\mathbb{T}\rangle$ is *naturally realizable* under (σ,τ) if there is an AF $F = (A,R)$ with $\sigma(F) = \mathbb{S}$, $\tau(F) = \mathbb{T}$ and $R \supseteq (\mathit{Confs}_{\mathbb{S}} \cap \mathit{Confs}_{\mathbb{T}})$.

Inspecting the canonical constructions in Definitions 2 and 3 and their modifications in the proofs of the theorems, we see that, so far, we only made use of natural realizations.

Proposition 7. *For* $\sigma,\tau \in \{gr, id, eg, cf, ad, na, sb, sm, pr\}$, $(\sigma,\tau) \notin \{(eg,ad),(eg,pr),(sm,ad),(sm,pr)\}$, $\langle\mathbb{S},\mathbb{T}\rangle$ *is realizable under* (σ,τ) *iff it is naturally realizable under* (σ,τ).

The following example witnesses that this is not the case for (sm,pr). There are pairs of extension-sets such that any simultaneous realization of these extension-sets under *sm* and *pr* must result in some conflicts being implicit, i.e. not represented by an attack. On the one hand this is surprising, as implicit conflicts between non-rejected arguments can be made explicit for preferred and semi-stable semantics on their own. On the other hand it means that for realizing pairs $\langle\mathbb{S},\mathbb{T}\rangle$ under (sm,pr) detection of implicit conflicts is a key ingredient.

Example 7. Consider the AF $F = (A,R)$ as depicted in Figure 3 with arguments b_{xy}, c_y, d_y, symmetric 3-cycles over arguments a^i_{xy}, u^i and directed 3-cycles over arguments e^i_{xy} ($i \in \{1,2,3\}$). Symmetric arrows between regular arguments (or symmetric 3-cycles) and symmetric 3-cycles indicate symmetric attacks between each of the involved arguments. For instance b_{12} as well as a^3_{12} attack and are attacked by u^1, u^2, u^3. Directed arrows between symmetric 3-cycles α^i and directed 3-cycles β^i represent attacks (α^i,β^j) for $i \neq j$. For instance a^2_{12} defends e^2_{12} by attacking e^1_{12} and e^3_{12}. Thus a^i_{xy} is in symmetric attack relationship with b_{xy} and with u^j, while e^i_{xy} can be "activated" by a^i_{1x} as well as by a^i_{2y}. To capture this activation relationship, in what follows we denote, for $S \subseteq A$, by $S^{\uparrow}_e$ the union of S and the arguments e^i_{xy} defended by S in F.

The preferred extensions of F are as follows (with $i,j,k,l \in \{1,2,3\}$ and $x \neq y, x' \neq y' \in \{1,2\}$):

$$E_u = \{u^i, e^i_{11}, e^i_{12}, e^i_{21}, e^i_{22}\},\ E_{ab} = \{a^i_{x1}, a^j_{x2}, a^k_{yx'}, b_{yy'}, d_y\}^{\uparrow}_e,$$
$$E_b = \{a^i_{1x}, a^j_{2x'}, b_{1y}, d_1, b_{2y'}, d_2\}^{\uparrow}_e,\ E_a = \{a^i_{11}, a^j_{12}, a^k_{21}, a^l_{22}\}^{\uparrow}_e.$$

Observe that $E^+_u = E^+_a \subset E^+_{ab}$ and E^+_b is incomparable to E^+_{ab}, hence $sm(F)$ only consists of E_{ab} and E_b. Note that E_b misses some e^i_{xy} in range, while E_{ab} misses some d_x in range; thus any realization of $\langle sm(F), pr(F)\rangle$ has some $e^i_{xy} \notin E^+_b$ and $d_x \notin E^+_{ab}$ as members of $pr(F)$ must be conflict-free.

Now observe that the u^i never occur together with d_x in any semi-stable or preferred extension, a so called implicit conflict. Naturally realizing $\langle sm(F), pr(F)\rangle$ has to make

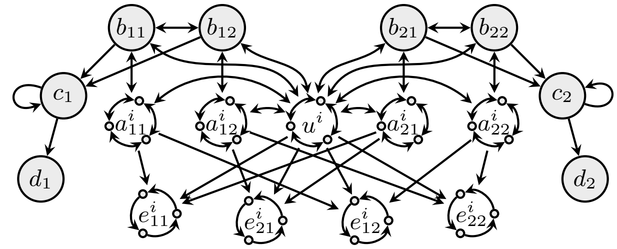

Figure 3: F s.t. $\langle sm(F), pr(F)\rangle$ is not naturally realizable.

all implicit conflicts symmetrically explicit, in this case by adding attacks $(u^i,d_x),(d_x,u^i)$. But then we get $d_x, e^i_{xy} \in E^+_u$ for $x,y \in \{1,2\}, i \in \{1,2,3\}$ and thus E^+_u contains or is at least incomparable to E^+_{ab} and E^+_b, which means that E_u cannot be excluded from the semi-stable extensions.

We can state the following result, leaving the exact characterization of $\Sigma_{sm,pr}$ (as well as for $\Sigma_{eg,pr}$, $\Sigma_{sm,ad}$ and $\Sigma_{eg,ad}$ where we anticipate similar issues) for future work.

Proposition 8. *There exist* $\langle\mathbb{S},\mathbb{T}\rangle \in \Sigma_{sm,pr}$ *which are not naturally realizable under* (sm,pr).

5 Discussion

In this paper, we have given a full characterization of all but four two-dimensional signatures for the semantics of conflict-free, admissible, naive, stable, preferred, semi-stable, grounded, ideal, and eager extensions. Two prominent semantics are missing: complete semantics, as its exact (one-dimensional) signature is not established yet. Stage semantics has been treated in (Dunne *et al.*, 2015), but turns out to behave quite differently when combined with other semantics. We leave them together with investigations on labelling-based semantics (Caminada and Gabbay, 2009) for future work.

Understanding two-dimensional signatures gives further insights about the relationship between semantics, but also yields practical implications. For example, when enumerating preferred extensions, we may start with computing the less complex stable semantics. Assume we have found $\{a,b\}$ and some $S \cup \{a\}$ as stable (and therefore also preferred) extensions. By Theorem 6 we can now exclude any $S' \cup \{b\}$ with $S \cap S' \neq \emptyset$ from the search-space, even if it could still be compatible with Σ_{pr}. We envisage to investigate implications of this kind on a more general level.

Another natural issue for future work is to extend the results to n-dimensional signatures. Some of our results already provide such characterizations; for instance, Example 2 readily delivers the 3-dimensional signature $\Sigma_{gr,id,eg}$, but for other combinations the picture is not clear yet. The research on multi-dimensional signatures is by no ways limited to argumentation. We plan to apply our method also to the world of logic programming in order to compare stable, well-founded, and supported semantics (Gelfond and Lifschitz, 1988; Van Gelder *et al.*, 1988; Clark, 1978) in a similar vein. For instance, understanding to which extent supported and stable models can diverge for a program might lead to shortcuts in the loop-formula approach (Lin and Zhao, 2004) for computing stable models.

Acknowledgements This research has been supported by the Austrian Science Fund (FWF) through projects I1102 and P25521.

References

Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. An introduction to argumentation semantics. *Knowledge Engineering Review*, 26(4):365–410, 2011.

Pietro Baroni, Paul E. Dunne, and Massimiliano Giacomin. On the resolution-based family of abstract argumentation semantics and its grounded instance. *Artificial Intelligence*, 175(3-4):791–813, 2011.

Ringo Baumann and Gerhard Brewka. Spectra in abstract argumentation: An analysis of minimal change. In Pedro Cabalar and Tran Cao Son, editors, *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2013)*, volume 8148 of *LNCS*, pages 174–186. Springer, 2013.

Ringo Baumann, Wolfgang Dvořák, Thomas Linsbichler, Hannes Strass, and Stefan Woltran. Compact argumentation frameworks. In Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan, editors, *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 69–74. IOS Press, 2014.

Trevor J. M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

Martin Caminada and Leila Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171(5-6):286–310, 2007.

Martin Caminada and Dov M. Gabbay. A logical account of formal argumentation. *Studia Logica*, 93(2):109–145, 2009.

Martin Caminada, Walter A. Carnielli, and Paul E. Dunne. Semi-stable semantics. *Journal of Logic and Computation*, 22(5):1207–1254, 2012.

Martin Caminada. Comparing two unique extension semantics for formal argumentation: Ideal and eager. In Mehdi Dastani and Edwin D. de Jong, editors, *Proceedings of the 19th Belgian-Dutch Conference on Artificial Intelligence (BNAIC 2007)*, pages 81–87, 2007.

Keith L. Clark. Negation as failure. In Hervé Gallaire and Jack Minker, editors, *Logic and Data Bases*, pages 293–322. Plenum Press, 1978.

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.

Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1/2):187–203, 2002.

Paul E. Dunne, Wolfgang Dvorák, and Stefan Woltran. Parametric properties of ideal semantics. *Artificial Intelligence*, 202:1–28, 2013.

Paul E. Dunne, Wolfgang Dvořák, Thomas Linsbichler, and Stefan Woltran. Characteristics of multiple viewpoints in abstract argumentation. *Artificial Intelligence*, 228:153–178, 2015.

Wolfgang Dvořák and Christof Spanring. Comparing the expressiveness of argumentation semantics. In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Proceedings of the 4th Conference on Computational Models of Argument (COMMA 2012)*, volume 245 of *Frontiers in Artificial Intelligence and Applications*, pages 261–272. IOS Press, 2012.

Wolfgang Dvořák and Stefan Woltran. On the intertranslatability of argumentation semantics. *J. Artif. Intell. Res. (JAIR)*, 41:445–475, 2011.

Michael Gelfond and Vladimir Lifschitz. The Stable Model Semantics for Logic Programming. In Robert A. Kowalski and Kenneth A. Bowen, editors, *Logic Programming: Proceedings 5th International Conference and Symposium*, pages 1070–1080. MIT Press, 1988.

Fangzhen Lin and Yuting Zhao. ASSAT: computing answer sets of a logic program by SAT solvers. *Artificial Intelligence*, 157(1-2):115–137, 2004.

Thomas Linsbichler, Christof Spanring, and Stefan Woltran. The hidden power of abstract argumentation semantics. In Elizabeth Black, Sanjay Modgil, and Nir Oren, editors, *Theory and Applications of Formal Argumentation (TAFA 2015), Revised Selected Papers*, volume 9524 of *LNCS*, pages 146–162. Springer, 2015.

Iyad Rahwan and Guillermo R. Simari, editors. *Argumentation in Artificial Intelligence*. Springer, 2009.

Allen Van Gelder, Kenneth A. Ross, and John S. Schlipf. Unfounded sets and well-founded semantics for general logic programs. In Chris Edmondson-Yurkanan and Mihalis Yannakakis, editors, *Proceedings of the 7th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 221–230. ACM, 1988.

Bart Verheij. Two approaches to dialectical argumentation: admissible sets and argumentation stages. In John-Jules C. Meyer and Linda C. van der Gaag, editors, *Proceedings of the 8th Dutch Conference on Artificial Intelligence (NAIC 1996)*, pages 357–368, 1996.