

Rejecting Arguments Based on Doubt in Structured Bipolar Argumentation

Michael A. Müller
CRIL, CNRS, Univ Artois
Lens, France
Université de Fribourg
Fribourg, Switzerland
muller@cril.fr

Srdjan Vesic
CRIL, CNRS, Univ Artois
Lens, France
vesic@cril.fr

Bruno Yun
Univ Lyon, UCBL, CNRS
INSA Lyon, LIRIS, UMR5205
Villeurbanne, France
bruno.yun@univ-lyon1.fr

ABSTRACT

This paper develops a new approach to computational argumentation that is informed by philosophical and linguistic views. Namely, it takes into account two ideas that have received little attention in the literature on computational argumentation: First, an agent may rationally reject an argument based on mere doubt, thus not all arguments they could defend must be accepted; and, second, that it is sometimes more natural to think in terms of which individual sentences or claims an agent accepts in a debate, rather than which arguments. In order to incorporate these two ideas into a computational approach, we first define the notion of structured bipolar argumentation frameworks (SBAFs), where arguments consist of sentences and we have both an attack and a support relation between them. Then, we provide semantics for SBAFs with two features: (1) Unlike with completeness-based semantics, our semantics do not force agents to accept all defended arguments. (2) In addition to argument extensions, which give acceptable sets of arguments, we also provide semantics for language extensions that specify acceptable sets of sentences. These semantics represent reasonable positions an agent might have in a debate. Our semantics lie between the admissible and complete semantics of abstract argumentation. Further, our approach can be used to provide a new perspective on existing approaches. For instance, we can specify the conditions under which an agent can ignore support between arguments (i.e. under which the use of abstract argumentation is warranted) and we show that deductive support semantics is a special case of our approach.

KEYWORDS

Abstract Argumentation; Doubt; Bipolar Argumentation; Structured Argumentation; Interpreting Semantics

ACM Reference Format:

Michael A. Müller, Srdjan Vesic, and Bruno Yun. 2026. Rejecting Arguments Based on Doubt in Structured Bipolar Argumentation. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/WZOU6484>



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/WZOU6484>

1 INTRODUCTION

For most people, the primary contact point with arguments is through discussions and debates, be it in person or online. In these contexts, they are confronted with the following central question: Given all the arguments, what should you believe? One way to answer this question is through computational argumentation [5, 22]. We can build models of debates and use them to evaluate which sets of arguments, called *extensions*, an agent considering the debate should. This can help agents navigate complex and potentially inconsistent information. The literature has produced a wide range of *semantics* that determine the acceptability of arguments (see [4]). Many of these approaches come out of an interest in modelling different ways of formal reasoning based on specific logics, logic programming, or how agents can reason with inconsistent or uncertain knowledge [1, 3, 14, 22, 33, 38, 40].

This paper uses computational approaches such as abstract [22], structured [11], and bipolar [15] argumentation as its starting point, but incorporates aspects of informal approaches to argumentation such as philosophy [25, 27, 42, 49], linguistics [28], and communication studies [45, 46].¹ Amongst these approaches, we can find two general ideas that are not always explicitly taken into account in formal approaches: (1) In formal approaches, rational agents are assumed to automatically accept all defended arguments, while in informal ones they are allowed to reject such an argument, e.g. if it relies on a premise the agents deem very implausible. (2) Formal approaches often require agents to evaluate arguments as a whole, meaning they often take the sentences of an argument (its premises and conclusion) into account only to determine its relations to other arguments. Informal approaches, in contrast, often assume that agents evaluate directly which sentences they should accept. While some computational approaches implement aspects of these ideas [9, 17, 24, 37, 41], we take both of them fully into account.

The idea of rejecting defended arguments comes from the notion of critical reactions [29, 47]. While an agent has to react critically to an argument in order to reject it, a critical reaction does not have to consist of a counterargument. For instance, instead of attacking an argument, the agent can also challenge it. Such a challenge expresses *doubt* about a premise of the argument and if that doubt cannot be overcome, the argument can be rejected. Thus, even in the absence of any attacks, agents can reject arguments if doubted premises are not supported. Our approach is accordingly one of *bipolar argumentation* [15], where arguments can be both attacked and supported. The use of an explicit support relation is also in

¹This is an extended version of a paper that was accepted to ArgXAI@ECAI'25. Proofs for all propositions can be found in <http://arxiv.org/abs/2602.03286>.

line with empirical results on how we intuitively deal with arguments [28, 36]. The following example illustrates the novelty and expressiveness of our approach.

EXAMPLE 1. Consider the following situation, loosely inspired by [2], where an agent considers a debate containing the following arguments: a_1 : “This violin is a Stradivarius since Alex says so”, a_2 : “This violin is expensive since it is a Stradivarius”, a_3 : “We know Clara says that Anne-Sophie owns this violin since she is cited in a newspaper saying so.”, a_4 : “Anne-Sophie owns this violin since Clara says so”, and a_5 : “Hilary owns this violin since Diego says so”. This gives the following situation of bipolar argumentation, where \rightarrow are attacks and \dashrightarrow are supports.

$$a_1 \dashrightarrow a_2 \quad a_3 \dashrightarrow a_4 \longleftrightarrow a_5$$

If the agent uses any form of complete semantics [22, 50], they have to accept all defended arguments. Thus any such agent has to accept at least a_1 and a_2 . However, one might reasonably doubt whether Alex really made the claim in a_1 . The agent might suspect that Alex does not know anything about violins and thus his claim strikes them as odd. This suspicion does not amount to a full counterargument that could be added to the framework, but it is enough for the agent to doubt a_1 and reject it if there is no support to remove their doubts.

The next weaker semantics are based on admissibility [22]. Admissibility allows agents to doubt a_1 , but it also removes any notion of support to a_2 , as the two arguments can be accepted fully independently of each other. Using d -admissibility of deductive support semantics (see [13, 16] and Section 4.2), support is taken into account and a_1 can only be accepted together with a_2 , giving the following extensions: \emptyset , $\{a_1, a_2\}$, $\{a_3, a_4\}$, $\{a_5\}$, $\{a_1, a_2, a_5\}$, and $\{a_1, a_2, a_3, a_4\}$. But now also a_3 can only be accepted together with a_4 . However, as a_4 and a_5 contradict each other, one might want to abstain from accepting either of them while still accepting a_3 . That is, an agent might want to accept that Clara claims that Anne-Sophie owns the violin based on the newspaper while suspending judgment on whether she really owns it. D -admissibility does not allow for this option.

One of the semantics we define, weak coherence (see Section 3.2), captures these dynamics that elude both completeness-based semantics and d -admissibility by allowing agents to accept the extension $\{a_3\}$. A full analysis of this case is provided in Example 10.

The idea of doubting individual sentences also illustrates how informal approaches often evaluate arguments on the level of their sentences. For people confronted with arguments, it is often more intuitive to think about what sentences they accept rather than which arguments [12]. The latter can be complex structures and it can be challenging to analyse them and figure out to what you are committed when accepting them. Taking into account sentences can help explain why some agent accepts (or should accept) an argument extension, namely by referring to the set of sentences they agree with. This requires taking into account the structure of arguments and accordingly our approach takes the perspective of structured argumentation.

EXAMPLE 2. It might be intuitive for agents to indicate which individual sentences they accept without thinking about the arguments directly. In Example 1, we might accept the sentences Str : “This violin is a Stradivarius”, Exp : “This violin is expensive”, Cl : “Clara says that Anne-Sophie owns this violin”, and Hil : “Hilary owns this

violin” (see Example 3 for a full translation). Is $\{Str, Exp, Cl, Hil\}$ an acceptable set of sentences? In the terms we introduce in Section 3.2, we can understand this as a weakly coherent language extension and it corresponds to the argument extension $\{a_2\}$.

In this paper, we define structured bipolar argumentation frameworks (SBAFs) and provide a range of semantics for them. At the argument level, we define two versions of what we call coherent semantics. They contrast with completeness-based semantics in that they sometimes allow rejection of defended arguments, but they also contrast with (abstract) admissible semantics in that they take support between arguments into account. We additionally provide two versions of adequate semantics that operate on sets of sentences directly. This way, we offer a new perspective on computational argumentation which is informed by informal theories.

The paper is organised as follows. We first sketch the familiar preliminaries from abstract argumentation (Section 2) and then introduce SBAFs in Section 3. Both argument semantics, which evaluate sets of arguments, and language semantics, which evaluate sets of sentences, are defined in Section 3.2. These semantics are then used to interpret existing semantics from abstract and bipolar argumentation in Section 4. Finally, we discuss related approaches in the literature (Section 5).

2 PRELIMINARIES

We briefly recall the definitions of abstract argumentation [4, 22].

DEFINITION 1 (ABSTRACT ARGUMENTATION FRAMEWORK). An abstract argumentation framework (AF) is a tuple $\mathcal{A} = \langle A, \rightarrow \rangle$ where A is a finite set of arguments and $\rightarrow \subseteq A \times A$ an attack relation.

We write $a \rightarrow b$ in case $(a, b) \in \rightarrow$ and generalise to sets of arguments, i.e. $E \rightarrow b$ in case $a \rightarrow b$ for some $a \in E$ and $a \rightarrow E$ in case $a \rightarrow b$ for some $b \in E$. Further, we write $a \not\rightarrow b$ if $(a, b) \notin \rightarrow$. For two sets of arguments, $E \rightarrow E'$ means $a \rightarrow b$ for some $a \in E$ and $b \in E'$. A set of arguments $E \subseteq A$ defends an argument $a \in A$ if $\forall b \in A : b \rightarrow a \implies E \rightarrow b$.

The semantics are as usual. Let $\mathcal{A} = \langle A, \rightarrow \rangle$ be an AF. An extension $E \subseteq A$ is called *conflict-free* if $\forall a, b \in E : a \not\rightarrow b$, *admissible* if it is conflict-free and defends all its arguments, *complete* if it is admissible and it contains all arguments it defends, *preferred* if it is \subseteq -maximal among admissible extensions.

Note the difference between admissible and complete semantics. Whereas admissible extensions are never forced to include an argument, complete semantics requires accepting all defended arguments, including any unattacked ones. Even when it comes to bipolar semantics (see Section 4.2), the focus still lies on variants of complete semantics (cf. [50]). The semantics we define in Section 3.2 fall in the gap between admissible and complete semantics.

3 STRUCTURED BIPOLAR ARGUMENTATION

In this section, we first introduce structured bipolar argumentation frameworks and then go on to define their semantics.

3.1 Frameworks

Now we start introducing the frameworks for structured bipolar argumentation. We take a structured approach and as such we first have to define the language which we use to represent arguments.

DEFINITION 2 (LANGUAGE). A language $\mathcal{L} = \langle L, \bar{\cdot}, n \rangle$ consists of a non-empty set of sentences L , a (partial) incompatibility function $\bar{\cdot} : L \rightarrow 2^L$, and a (partial) naming function $n : 2^L \times L \rightarrow L$.

We assume $\bar{\cdot}$ to be symmetric, i.e. $\forall s, t \in L : s \in \bar{t} \iff t \in \bar{s}$. Additionally, we assume that $n(\langle \{t\}, t \rangle) = \emptyset$.

The set of sentences can be any set of objects. All the structure we need to represent arguments is given by the incompatibility and the naming functions. Incompatibility is used to model conflict between sentences in that it associates each sentence with the set of sentences incompatible with it. It is a symmetric notion of contrariness, cf. [41]. Intuitively, two incompatible sentences should not be accepted together. The naming function allows us to talk about arguments within the language. We represent arguments as a tuple of a set of sentences (the premises) and another sentence (the conclusion), see Definition 3. Thus, n takes arguments and gives them names which express the claim that one can infer the conclusion from the premises, cf. [32]. The main use of this is to define undercutting attacks (Definition 4) using sentences that are incompatible with the name of an argument. Finally, the condition $n(\langle \{t\}, t \rangle) = \emptyset$ states that an argument that uses the same sentence as its single premise and as its conclusion (see Definition 3) cannot be undercut, as there are no circumstances where we cannot infer a sentence from itself. Note that we do not assume any logical structure or consequence relation on the set of sentences.

EXAMPLE 3. Let us continue Example 1 and consider a language with $L = \{Ale, Str, Exp, New, Cla, Ann, Die, Hil\}$, corresponding to: *Ale* : “Alex says this violin is a Stradivarius”, *Str* : “This violin is a Stradivarius”, *Exp* : “This violin is expensive”, *New* : “Clara is cited in a newspaper article mentioning that Anne-Sophie owns this violin”, *Cla* : “Clara says that Anne-Sophie owns this violin”, *Ann* : “Anne-Sophie owns this violin”, *Die* : “Diego says Hilary owns this violin”, and *Hil* : “Hilary owns this violin”, with $Ann \in Hil$ and $Hil \in Ann$. We can also add $n(\langle \{Ale\}, Str \rangle) : \text{“One can infer that this violin is a Stradivarius from Alex saying so”}$.

As in abstract argumentation, we do not construct arguments. Rather, we take them as given, e.g. through a debate. This means that we do not need any logic or inference rules that determine from which sentences we can infer others. Formally, any combination of premises and conclusion could be an argument. In the following definition, we use *Prem* as a function from arguments to sets of sentences in order to indicate the premises of an argument and *Conc* as a function from arguments to sentences in order to indicate the conclusion of an argument. For an argument a , $Prem(a)$ is its set of sentences and $Conc(a)$ is its conclusion.

DEFINITION 3 (ARGUMENTS). An argument in a language $\mathcal{L} = \langle L, \bar{\cdot}, n \rangle$ is a tuple $a = \langle Prem(a), Conc(a) \rangle$ where $Prem(a)$ is a non-empty finite subset of L and $Conc(a) \in L$.

We also define the set of sentences of an argument a as $Sent(a) := Prem(a) \cup \{Conc(a)\}$. The set of sentences generalises to sets of arguments $Sent(E) = \bigcup_{a \in E} Sent(a)$.

We say a is a minimal argument for sentence $s \in L$ if $a = \langle \{s\}, s \rangle$ (see e.g. [32]).

EXAMPLE 4. In Example 3, we already saw one argument: $a_1 : \langle \{Ale\}, Str \rangle$, which corresponds to “This violin is a Stradivarius

since Alex says so”. The other arguments are: $a_2 : \langle \{Str\}, Exp \rangle$, $a_3 : \langle \{New\}, Cla \rangle$, $a_4 : \langle \{Cla\}, Ann \rangle$, and $a_5 : \langle \{Die\}, Hil \rangle$. An example of a minimal argument would be $\langle \{Str\}, Str \rangle$, i.e. “This violin is a Stradivarius since this violin is a Stradivarius”. It is an edge case of an argument as it does not contain any real inference step, but they are useful to represent single sentences in frameworks.

Next, we define supports and attacks between arguments. Attacks are defined such that an argument attacks another argument if its conclusion is incompatible with some part of it [32]. We call an attack on the name of an argument an undercut and we say that a set of arguments E contains *undercutting information* for an argument a if $n(a) \cap Sent(E) \neq \emptyset$. Note that attacks are defined only through the conclusions of arguments. Accordingly, a set of arguments can contain undercutting information for an argument without attacking it and an argument containing a premise incompatible with the premise of another argument does not necessarily attack it. For instance, there are no attacks between the arguments $a : \langle \{r\}, x \rangle$ and $b : \langle \{y\}, z \rangle$ with $r \in \bar{n(b)}$ and $y \in \bar{r}$, since the conflicts occur only between their premises. However, a contains undercutting information for b . While there are more fine-grained notions of attack in the literature [19], they are often variations of the attacks defined here.

The support relation requires more explanation. We want to capture situations where accepting some set of arguments commits you to accepting another: If you accept an argument $a_1 : \langle \{s\}, t \rangle$ and there is an argument $a_2 : \langle \{t\}, u \rangle$, you should also accept a_2 since you accept all its premises. When it comes to arguments with more than one premise, only a set of arguments can potentially force its acceptance. For instance, accepting a_1 does not force acceptance of $a_3 : \langle \{t, v\}, w \rangle$. But if you also accept $a_4 : \langle \{v\}, v \rangle$, then you should accept a_3 . Thus, we define support not as a binary relation between arguments but as a relation between sets of arguments and arguments. Hence, we capture a notion of *premise support* [18].

But there are more situations where accepting some arguments can commit you to accepting another. Suppose there is also an argument $a_5 : \langle \{s\}, r \rangle$. Then, accepting a_1 means accepting all premises of a_5 and thus a_5 should be accepted as well. This is a kind of *premise-sharing-support*, sometimes called *exhaustion* [43]. In that sense, our notion of support combines premise-support and premise-sharing-support in order to capture all situations where accepting a set of arguments commits you to accept another one as well. This leads to the following definition.

DEFINITION 4 (SUPPORT AND ATTACK). Let a, b be arguments in language \mathcal{L} .

We say a set of arguments E supports a if $Prem(a) \subseteq Sent(E)$.

We say that a attacks b if $Conc(a) \in \bar{s}$ for some $s \in Sent(b)$ or if $Conc(a) \in n(b)$.

EXAMPLE 5. Using the setting of Examples 1, 3 and 4, we have a support from $a_1 : \langle \{Ale\}, Str \rangle$ to $a_2 : \langle \{Str\}, Exp \rangle$ and a mutual rebut between $a_4 : \langle \{Cla\}, Ann \rangle$ and $a_5 : \langle \{Die\}, Hil \rangle$.


The following definition collects all of this together.

DEFINITION 5 (STRUCTURED BIPOLAR ARGUMENTATION FRAMEWORK). A structured bipolar argumentation framework (SBAF) is a tuple $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashv \rangle$ where A is a finite set of arguments

in language \mathcal{L} and $\rightarrow, \dashrightarrow$ are the corresponding attack and support relations.

We write $a \dashrightarrow b$ in case $(a, b) \in \dashrightarrow$ and use the same notational conventions as with attacks. The intuitive running example is revisited in Example 10. Here, we introduce a more technical example.

EXAMPLE 6. The following example is based on $t \in \bar{r}$, $r \in \bar{t}$, $r \in \overline{n(a_6)}$, and $z \in \bar{p}$ (and accordingly $p \in \bar{z}$). Note that a_5 undercuts a_6 as it attacks its inference claim.



It is useful to define (strongly) saturated frameworks. These are such that they contain minimal arguments for sentences that are incompatible with others. This ensures that all important relations between sentences are visible as relations between arguments.

DEFINITION 6 (SATURATED SBAFs). An SBAF \mathcal{SB} is called *saturated* (resp. *strongly saturated*) if $\forall s \in \text{Sent}(A)$ s.t. $\exists t \in \text{Sent}(A) \cap \bar{s}$, there is a minimal argument for s or (resp. and) for t in A , and $\forall u \in \text{Sent}(A)$ s.t. $u \in \overline{n(a)}$ for some $a \in A$, there is a minimal argument for u in A .

The SBAF in Example 6 is not saturated, but could be made so by adding, for instance, $\langle\{r\}, r\rangle$ and $\langle\{z\}, z\rangle$. To make it strongly saturated, $\langle\{t\}, t\rangle$ and $\langle\{p\}, p\rangle$ would also have to be added.

3.2 Semantics

We now develop semantics for SBAFs with two goals in mind: to allow rejecting certain defended arguments, and to evaluate both acceptable sets of arguments and acceptable sets of sentences. This dual perspective assists agents with different capabilities to communicate and resolve potential disagreements. That is, it allows agents that stores their knowledge in form of sentences to communicate and interact with agents that only see arguments.

For argument semantics, we find a middle ground between admissible and complete semantics. One is not forced to accept all defended arguments, but their rejection is limited by support. A simple way of doing this is to start with admissible semantics and to add a condition that makes sure that extensions are closed under support (cf. [16, 43]). You should accept an argument not if it is defended, but if you already accept all its premises.

However, there are cases where this condition may be relaxed. A first exception occurs in case you already accept undercutting information for the supported argument. In such a case, you accept that the inference in the argument does not hold and as such you can reject it, even if you accept all its premises. Another exception can be argued for in case you accept all premises of an undefended argument. Depending on how strongly one interprets the support relation, you might be forced to accept it or not.

We give two semantics, based on whether the second exception is allowed. This gives a *strong* and a *weak* interpretation of support.

DEFINITION 7 (COHERENT ARGUMENT EXTENSIONS). A *strongly coherent argument extension* in an SBAF $\mathcal{SB} = \langle\mathcal{L}, A, \rightarrow, \dashrightarrow\rangle$ is an admissible extension $E \subseteq A$ that satisfies:

Strong Support-Closure: $\forall a \in A$: if E supports a and E does not contain undercutting information for a , then $a \in E$.

A *weakly coherent argument extension* is an admissible extension E that satisfies:

Weak Support-Closure: $\forall a \in A$: if E supports a , E does not contain undercutting information for a , and E defends a , then $a \in E$.

OBSERVATION 1. Strongly coherent extensions are weakly coherent.

EXAMPLE 7. Consider the SBAF of Example 6. The extension $\{a_1, a_2, a_3, a_4, a_6\}$ is weakly coherent, but not strongly so. According to strong coherence, it is not possible to accept all of a_1, a_2 , and a_3 together since strong support-closure forces acceptance of a_4 and a_5 , violating admissibility.

The difference between strong and weak coherence lies in different answers to the question whether support can force you to accept undefended arguments. Strong coherence says “yes”, weak coherence says “no”. Weak coherence allows you to reject an argument if it would violate defence, even if you accept all its premises. In some sense, it allows you to infer that the inference of the argument must be faulty without having explicit undercutting information. Strong coherence, in contrast, requires explicit undercutting information to reject an inference. Without it, the inference is assumed to work.

A more formal way to bring out the difference between strong and weak coherence is the property of directionality [6, 50]: Arguments should only affect the acceptability of other arguments if they are connected via a directed path through support and attack. This is true for weak coherence, but not for strong coherence.

Directionality: Let \mathcal{SB} be an SBAF, $U \subseteq A$ be such that $A \setminus U \not\rightarrow U$ and $A \setminus U \not\rightarrow U$ and define $\mathcal{SB}|_U = \langle\mathcal{L}, U, \rightarrow \cap (U \times U), \dashrightarrow \cap (U \times U)\rangle$. A semantics σ satisfies *directionality* if an extension $E \subseteq U$ is acceptable according to σ in $\mathcal{SB}|_U$ iff there exists a σ -acceptable extension $E' \subseteq A$ in \mathcal{SB} such that $E = E' \cap U$.

PROPOSITION 2. Strong coherence fails directionality. Weak coherence satisfies directionality.

What does a semantics from the language perspective look like? If we are interested in acceptable sets of sentences, we first require that they do not contain incompatible sentences. We call a set of sentences S *compatible* if $\forall s, t \in S$: $s \notin \bar{t}$. But apart from this condition, it is the arguments that limit the choice of sentences we can accept—this is the point of arguing. For this we need to know to which arguments you are committed to when accepting some sentences and vice-versa.

Given an argument extension, we can simply say that the corresponding language extension consists of all sentences of accepted arguments. When we start with a language extension, however, it is less clear how the accepted sentences translate to accepted arguments. For instance, perhaps it is possible to accept all premises and the conclusion of an argument without accepting the argument itself—one can simply disbelief that one can infer the conclusion from the premises. We define two ways of translating from language extensions to argument extensions that are analogous to strong and weak coherence.

The first way, similar to strong coherence, is to count all arguments as accepted of which the premises are accepted and there is no undercutting information. This gives the *strong argument set*. Starting with a set of accepted sentences, it collects together all arguments that should be accepted according to the strong interpretation of support.

DEFINITION 8 (STRONG ARGUMENT SET). *Given a set of sentences S in $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashv \rangle$, we define its strong argument set $Arg_s(S) := \{a \in A \mid Prem(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$.*

Analogously, we can define a *weak argument set*. Given a set of sentences S in an SBAF \mathcal{SB} , we define its *weak argument set* which collects together all the arguments in \mathcal{SB} which should be accepted according to the weak interpretation of support. As we only require accepting defended arguments, we define it using a fixpoint construction analogous to that in abstract argumentation for complete extensions [4, 22]. We first define the analogue of the characteristic function in abstract argumentation.

DEFINITION 9 (CHARACTERISTIC FUNCTION). *Given an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashv \rangle$ and a set of sentences $S \subseteq Sent(A)$, we define the characteristic function $R_{\mathcal{SB}}^S : 2^A \rightarrow 2^A$ as*

$$R_{\mathcal{SB}}^S(E) := \{a \in A \mid a \in Arg_s(S) \text{ and } E \text{ defends } a\}.$$

With this, we can define the *weak argument set*.

DEFINITION 10 (WEAK ARGUMENT SET). *If S is compatible, then its initial set, $Init(S)$, is defined as the largest admissible subset of $\{a \in A \mid Sent(a) \subseteq S \text{ and } \overline{n(a)} \cap S = \emptyset\}$. The weak argument set of a compatible S , $Arg_w(S)$, is the least fixpoint of $R_{\mathcal{SB}}^S$ containing $Init(S)$.*

We need to make sure that the weak argument set contains $Init(S)$, as otherwise the least fixpoint of $R_{\mathcal{SB}}^S$ would sometimes be too small. For instance, in case of two arguments $a_1 : \langle \{s\}, t \rangle$ and $a_2 : \langle \{u\}, u \rangle$ with $u \in \bar{t}$ and $t \in \bar{u}$, the least fixpoint for the language extension $S = \{s, t\}$ would be empty. However, we want it to contain a_1 as all its sentences are accepted and it defends itself.

PROPOSITION 3. *The weak argument set is well-defined.*

One motivation for introducing language extensions is that agents may find it easier to judge which sentences they accept than which arguments. The formal complexity of defining the weak argument set confirms this intuition. Further, it allows us explain why certain arguments are accepted or rejected by reference to their underlying sentences.

EXAMPLE 8. *Consider again the SBAF of Example 6 and take the language extension $S = \{s, t, u, v, w, x, y\}$. Its strong argument set is $Arg_s(S) = \{a_1, a_2, a_3, a_4, a_5, a_6\}$, which is not conflict-free. For the weak argument set, we have: $Init(S) = \{a_1, a_2, a_3, a_4\}$ and $Arg_w(S) = \{a_1, a_2, a_3, a_4, a_6\}$.*

Consider also $S' = \{s, u, v, w, x, r\}$ with $Arg_w(S') = \{a_1, a_2, a_3, a_5\}$. While both S and S' end up accepting all of a_1, a_2 , and a_3 , thus supporting both a_4 and a_5 , they differ in which of them they accept. On the argument level, the situation looks symmetrical, but using the sentence perspective, we can explain this difference through the prior commitment of S to the conclusion of a_4 and the prior commitment of S' to that of a_5 .

PROPOSITION 4. *For a compatible language extension S , $Arg_w(S)$ is admissible.*

We can now specify which sets of sentences should be accepted in an SBAF. We again get a strong and a weak semantics, depending on whether we take language extensions to commit to their strong or their weak argument set. For both, we require language extensions to include conclusions of accepted arguments. Further, we know from Proposition 4 that weak argument sets are defended, but for strong argument sets, we have to require it explicitly.

DEFINITION 11 (ADEQUATE LANGUAGE EXTENSIONS). *A strongly adequate language extension in an SBAF $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashv \rangle$ is a compatible set of sentences $S \subseteq Sent(A)$ such that $Arg_s(S)$ defends all its arguments and it satisfies:*

Sentence-Closure: $\forall a \in Arg_s(S) : Sent(a) \subseteq S$.

A weakly adequate language extension $S \subseteq Sent(A)$ is a compatible set of sentences that satisfies sentence-closure w.r.t. $Arg_w(S)$.

EXAMPLE 9. *As seen in Example 8, $S = \{s, t, u, v, w, x, y\}$ is not strongly adequate, as sentence-closure would require adding r to it, violating compatibility. However, it is weakly adequate, as $Arg_w(S)$ does not include a_5 . In contrast, $S' = \{u, v, w, x, r, y, z\}$ has $Arg_s(S') = \{a_2, a_3, a_5\}$ and is strongly adequate. Note that S' contains undercutting information for a_6 (namely r) and does not have to accept it.*

PROPOSITION 5. *Strongly adequate language extensions are also weakly adequate.*

EXAMPLE 10. *We can now fully analyse Example 1. The language is provided in Example 3. We get the following SBAF.*

$$a_1 : \langle \{Ale\}, Str \rangle \dashv a_2 : \langle \{Str\}, Exp \rangle$$

$$a_3 : \langle \{New\}, Cla \rangle \dashv a_4 : \langle \{Cla\}, Ann \rangle \leftrightarrow a_5 : \langle \{Die\}, Hil \rangle$$

We can see that $\{a_3\}$ is weakly, but not strongly coherent, and that $\{Str, Exp, Cla, Hil\}$ is weakly but not strongly adequate. This is because a strongly adequate extension would have to add Ann , since a_4 is in its strong argument set, but then compatibility would be violated. In this example, strong coherence coincides with d -admissibility (Section 4.2 expands on this observation). But weak coherence captures novel dynamics of doubt in argumentation.

Adequate language extensions and coherent argument extensions capture similar ideas. But, as the following example illustrates, there are SBAFs where the two notions come apart.

EXAMPLE 11. *Consider an SBAF with two arguments $a_1 : \langle \{s\}, t \rangle$ and $a_2 : \langle \{u\}, v \rangle$ with $s \in \bar{u}$. Then, $\{a_1, a_2\}$ is a strongly (and thus weakly) coherent argument extension, but its set of sentences is not compatible and hence neither strongly nor weakly adequate.*

Nevertheless, we can show that for a large class of SBAFs, there is a direct correspondence between adequate language extensions and coherent argument extensions. This confirms that even when we are interested in whether an agent accepts an acceptable argument extension, we can evaluate directly their accepted set of sentences.

PROPOSITION 6. *Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashv \rangle$ be a saturated SBAF. Then for every strongly adequate language extension S , its strong argument set $Arg_s(S)$ is a strongly coherent argument extension.*

Also, for every strongly coherent argument extension E , its set of sentences $Sent(E)$ is a strongly adequate language extension.

PROPOSITION 7. Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then for every weakly adequate language extension S , its weak argument set $Arg_w(S)$ is a weakly coherent argument extension.

Also, for every weakly coherent argument extension E , its set of sentences $Sent(E)$ is a weakly adequate language extension.

This concludes the semantics for SBAF, implementing both the critical reaction of doubt and the language perspective.

4 INTERPRETING ABSTRACT AND BIPOLAR SEMANTICS IN SBAFS

We now give new interpretations of preferred semantics and deductive support semantics [13, 16], as they both describe special cases of weakly, resp. strongly coherent semantics. Thus, in some circumstances, SBAFs can be seen as instantiating these semantics.

4.1 SBAFs and Preferred Semantics

We now compare our semantics with those of Dung-style abstract argumentation. Those semantics take neither support between arguments nor argument structure into account, so the following results show the circumstances under which we can ignore these aspects of SBAFs. The following observation shows that Dung-style semantics are related to weak coherence.

OBSERVATION 8. Complete extensions are weakly coherent and a weakly coherent extension is \subseteq -maximal iff it is preferred.

But we can use the language perspective to reach a more interesting result. While preferred extensions correspond to an agent who accepts as many arguments as possible, what kind of extensions do we get from an agent who accepts as many sentences as possible? We first define *confident* extensions to capture this idea.

DEFINITION 12 (CONFIDENT EXTENSIONS). Let \mathcal{SB} be an SBAF. A language extension $S \subseteq Sent(A)$ is called a *confident strongly (resp. weakly) adequate language extension* if it is \subseteq -maximal amongst strongly (resp. weakly) adequate language extensions.

An argument extensions $E \subseteq A$ is called a *confident strongly (resp. weakly) coherent argument extension* if it is strongly (resp. weakly) coherent and there exists a confident strongly (resp. weakly) adequate language extension $S \subseteq Sent(A)$ such that $E = Arg_s(S)$ (resp. $E = Arg_w(S)$).

EXAMPLE 12. In the SBAF of Example 6, we have the following confident strongly coherent extensions: $\{a_1, a_2, a_4, a_6\}$, $\{a_1, a_3, a_4, a_6\}$, and $\{a_2, a_3, a_5, a_7\}$. And the following confident weakly coherent extensions: $\{a_1, a_2, a_3, a_5, a_7\}$, and $\{a_1, a_2, a_3, a_4, a_6\}$.

We can again note correspondence between argument and language extensions. This time, the direction from arguments to sentences works only indirectly, as the set of sentences of a confident coherent argument extension might not itself be confident adequate.

PROPOSITION 9. Let \mathcal{SB} be a saturated SBAF. For every confident strongly (resp. weakly) adequate language extension, its strong (resp. weak) argument set is confident strongly (resp. weakly) coherent.

Also, for every confident strongly (resp. weakly) coherent argument extension E , there exists a confident strongly (resp. weakly) adequate language extension S such that $Arg_s(S) = E$ (resp. $Arg_w(S) = E$).




Figure 1: SCA: strongly coherent argument extensions, WCA: weakly coherent argument extensions, WAL: weakly adequate language extensions, SAL: strongly adequate language extensions. $SCA \subseteq WCA$ indicates that each strongly coherent extension is also weakly coherent. “ \subseteq^* ” relates argument extensions with the argument sets of language extensions. Dashed relations only hold in saturated SBAFs.

In Example 12, confident weakly coherent extensions correspond to preferred extensions. Thus, there, maximising arguments is the same as maximising sentences. But this is not always the case.

EXAMPLE 13. In the following SBAF (with $s \in \bar{u}$), \emptyset is confident weakly coherent, based on $\{t, u\}$ being confident weakly adequate, but it is not preferred.

$$a_1 : \langle \{s\}, s \rangle \longrightarrow a_2 : \langle \{t, u\}, t \rangle$$

Nevertheless, we find that in saturated SBAFs, preferred extensions are confident weakly coherent. For the other direction, we need strongly saturated frameworks. Thus, it is only in strongly saturated frameworks, where many minimal arguments are added, that maximising arguments and maximising sentences coincide. The correspondence between confident weakly coherent extensions and preferred extensions also shows that, in strongly saturated SBAFs, the former does not take support into account, as the latter can be calculated in pure attack-frameworks.

PROPOSITION 10. Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a saturated SBAF. Then any preferred extension $E \subseteq A$ is confident weakly coherent.

PROPOSITION 11. Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be a strongly saturated SBAF. Then any confident weakly coherent argument extension $E \subseteq A$ is preferred.

These two propositions show that at least in a large class of SBAFs, an agent who wants to maximise either arguments or sentences can safely disregard support between arguments and solely focus on attacks. In some sense, then, support becomes redundant if we use preferred-style semantics. While it is an open question whether we can find a full correspondence between weakly coherent and complete semantics, it indicates that we should indeed use admissible-style semantics if we want to account for supports. Figure 1 shows all our semantics and their relations.

4.2 SBAFs and Deductive Support Semantics

We base our discussion of bipolar semantics on the notion of *deductive support* [13, 16]. This approach implements the idea that accepting a supporting argument entails accepting the supported argument as well. We show that it is related to strong coherence.


DEFINITION 13 (BIPOLAR ARGUMENTATION FRAMEWORK). A bipolar argumentation framework (BAF) is a tuple $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ where A is a finite set of arguments, $\rightarrow \subseteq A \times A$ an attack relation, and $\dashrightarrow \subseteq A \times A$ a support relation.

DEFINITION 14 (SUPPORTED AND MEDIATED ATTACKS). Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ be a bipolar argumentation framework and take $a, b \in A$.

We say that there is a supported attack from a to b if there exists $c \in A$ such that $a \dashrightarrow c$ and $c \rightarrow b$.

We say that there is a mediated attack from a to b if there exists $c \in A$ such that $b \dashrightarrow c$ and $a \rightarrow c$.

EXAMPLE 14. In the following, BAF, the attack from a_1 to a_3 is a mediated attack and that from a_3 to a_4 is a supported attack. Added attacks are indicated with dotted arrows.



DEFINITION 15 (DEDUCTIVE SUPPORT SEMANTICS). Let $\mathcal{B} = \langle A, \rightarrow, \dashrightarrow \rangle$ be a bipolar argumentation framework. We define the set of complex attacks,

$$\rightarrow^{co} := \bigcup_{i \in \mathbb{N}} \rightarrow^i,$$

where $\rightarrow^0 := \rightarrow$ and $\rightarrow^{i+1} := \rightarrow^i \cup \{(a, b) \in A \times A \mid \text{there is a supported or mediated attack from } a \text{ to } b \text{ w.r.t. } \rightarrow^i\}$.

An extension $E \subseteq A$ is called *d-admissible*, *d-complete*, or *d-preferred* if it is admissible and closed under \dashrightarrow , complete, or preferred in $\mathcal{A} = \langle A, \rightarrow^{co} \rangle$.

EXAMPLE 15. In the BAF of Example 14, we have the unique *d-complete*, and *d-preferred* extension $\{a_2, a_3, a_5\}$. On the *d-admissible* side, we have $\emptyset, \{a_5\}, \{a_2, a_5\}$, and $\{a_2, a_3, a_5\}$. Note that $\{a_3, a_5\}$ is not *d-admissible* since it is not closed under \dashrightarrow .

We can now observe that strong coherence behaves like *d-admissibility* under certain conditions. Thus, SBAFs can specify exactly the circumstances under which the structure of the arguments does not need to be considered.

PROPOSITION 12. Let $\mathcal{SB} = \langle \mathcal{L}, A, \rightarrow, \dashrightarrow \rangle$ be an SBAF where $\forall a \in A : |\text{Prem}(a)| = 1$ and $\neg \exists t \in \text{Sent}(A), \exists a \in A : t \in n(a)$. Then an extension $E \subseteq A$ that is strongly coherent is also *d-admissible*.

If further $\forall a, b \in A$, we have $\text{Prem}(a) \neq \text{Prem}(b)$, then an extension $E \subseteq A$ that is *d-admissible* is also strongly coherent.

This result tells us that strong coherence coincides with *d-admissibility* if each argument has exactly one premise that is unique and there is no undercutting information in the framework. Accordingly, we only need to take information about the structure of the arguments into account if they are more complex than unique single-premise. Strong coherence can then be seen as a generalisation of *d-admissibility* to these more complex cases.

5 RELATED APPROACHES

5.1 Bipolar Argumentation

The semantics presented in this paper rely on the notions of strong, resp. weak support-closure. Both variants find their analogues in the literature, be it under different background assumptions.

A version of strong support-closure can be found in the notion of *exhaustion* in *premise-augmented BAFs* (pBAFs) [43]. pBAFs go some way towards a structured approach to bipolar argumentation by adding information about the premises of the arguments. This allows keeping track of which premises an argument extension is committed to. A *p-admissible* extension is then required to be admissible and exhaustive, meaning that it contains all arguments of which all premises are accepted. As there are no exceptions made for undefended arguments this corresponds, disregarding undercutting information, to the notion of strong support-closure.

Crucially, pBAFs do not take the conclusions of arguments into account, only their premises. This can lead to frameworks that work as pBAFs but are uninstantiable if we think about the conclusions.

EXAMPLE 16. Consider the following pBAF, where π indicates the premises of the arguments.


$$\begin{array}{ccc} a_1, \pi(a_1) = \{s\} & \dashrightarrow & a_2, \pi(a_2) = \{t\} \\ \downarrow & & \\ a_3, \pi(a_3) = \{-t\} & & a_4, \pi(a_4) = \{u\} \end{array}$$

This framework is uninstantiable once we think about the conclusion of a_1 . Namely, it has to support both t and $\neg t$. This might be possible if the conclusion of a_1 is itself contradictory, but then it would have to support everything, including a_4 . Thus, this framework cannot be instantiated, but we only see it if we consider the conclusions of arguments. This is exactly the information SBAFs add.

The notion of weak support-closure occurs in *bipolar set-argumentation frameworks* (BSAFs) [10] in the definition of Γ -admissibility. While BSAFs do not take the structure of arguments into account, they allow for collective supports, as do SBAFs. Further, it can be shown that in saturated BSAFs, weak coherence does not need to take into account undercutting information explicitly, as this is covered by the defence-requirement. This makes weak coherence analogous to Γ -admissibility, which requires an extension to contain all supported arguments that are also defended.

The difference between BSAFs (as well as pBAFs) and SBAFs lies in the notion of defence. Namely, BSAFs and pBAFs require defence only against *closed* sets of arguments, i.e. against sets of arguments that are closed under the support relation. SBAFs, in contrast, require defence against all attacking arguments.

EXAMPLE 17. Consider the following SBAF with $s \in \bar{r}, r \in \bar{s}$, and $t \in n(a_3)$.



Argument a_1 defends itself against the closed set $\{a_2, a_3\}$ because it attacks a_3 , but it is not defended against a_2 . Accordingly, $\{a_1\}$ is not weakly coherent, but it would be Γ -admissible.

We can further find an implementation of the notion of doubt in *evidential argumentation systems* [35, 37]. In these frameworks, an argument is only relevant if it is supported by evidence. Any argument that has no evidential support can be rejected. This allows for rejecting unattacked arguments and thus can be seen as an expression of doubt. However, this implementation of doubt is quite different from ours in at least two respects: *First*, doubt in evidential argumentation systems needs to be justified by reference to a lack of evidence and whether an argument lacks evidence is determined by the argumentation system. There is no room for different agents doubting different arguments in the same system, as is possible in SBAFs. In our approach, it is the agents that decide what to doubt, not the framework. We also make no external assumptions on when doubt is justified. *Second* evidential argumentation systems rely on a notion of *necessary* support [33, 34], whereas we showed in Proposition 12 that our notion of support generalises *deductive* support. With necessary support, an argument can only be accepted if all its supporters are also accepted. In that sense, accepting the supporter is necessary for accepting the supported argument. In SBAFs it is never required to accept a supporting argument, support can only force acceptance of supported arguments. Thus, the direction goes the different way and accordingly there is no natural correspondence to necessary support in SBAFs.

5.2 Structured Argumentation

Our account bears similarities to structured approaches to argumentation such as ASPIC⁺ [31, 32, 39] or ABA [14, 41]. Both approaches structure arguments into premises and conclusions, where this structure then determines the relations between the arguments. However, our conception of arguments differs from theirs. *First*, we take the arguments as given, while both ASPIC⁺ and ABA use knowledge bases (resp. assumptions) and a set of inference rules to construct them. This guarantees some rationality in the inferences of the arguments as they must be licenced by some rule. In contrast, SBAF arguments could in principle combine any premises with conclusions, just like abstract argumentation could in principle count any statement as an argument. *Second*, our account of argument structure is less complex than that of ASPIC⁺ or ABA. Our arguments are simple premise-conclusion structures, whereas arguments in other structured approaches are full inference trees that derive their conclusions by means of potentially many inference rules and intermediary steps from their premises. This difference is brought out by our use of a support relation where ASPIC⁺ and ABA work with attacks only. An argument that concludes with a premise of another one counts as a supporting argument (or an element of a supporting set) in SBAFs, but would count as a sub-argument of a larger, more complex argument in ASPIC⁺. *Third*, ABA allows evaluation of frameworks in terms of acceptable sets of sentences, similarly to our language extensions. The difference is that language extensions take account of all sentences that are involved in arguments, where ABA only gives sets of assumptions, which represent the original premises of their inference tree-arguments.

Finally, claim-augmented argumentation frameworks (CAFs) [9, 23] associate each argument with a claim, representing their conclusion, and provide semantics that work purely on the level of claims. They also compare the strategies of maximising accepted

claims and maximising accepted arguments, and they conclude that these strategies are identical in well-formed CAFs. As SBAFs are well-formed by definition, their result differs from ours: Confident weakly coherent extensions and preferred extensions do not coincide in general, as SBAFs also take into account the premises.

6 CONCLUSION

Argumentation in practice has many features and aspects, which in turn lead to many ways of modelling it. In this paper, we followed ideas from informal approaches to argumentation and developed an approach to structured bipolar argumentation. This approach is characterised by semantics that are weaker than typical completeness-based semantics in that not all defended arguments have to be accepted and by the correspondence between the argument and the language perspective. Not only are we able to say which arguments an agent should accept, but we can also directly evaluate which sentences they should accept. As our results show, depending on the properties of the framework (i.e. whether it is saturated), the two perspectives can align or come apart. Further, we can compare the strategy of maximising accepted arguments (i.e. choosing a preferred extension) with that of maximising accepted sentences (i.e. choosing a confident weakly coherent extension). Again, these strategies can lead to different results, but they coincide in strongly saturated frameworks. Amongst the semantics for bipolar argumentation, we find a correspondence to d-admissibility.

Our approach allows agents to doubt any sentence which they cannot infer from already accepted sentences. This implements a very generous notion of doubt that may make it difficult to force some agent to accept an argument, as they could in principle always doubt our premises. Accordingly, it might be desirable to distinguish between sentences that can be freely doubted and other sentences that belong to a form of common ground that should be accepted by all agents. A strong version of this could be captured in our approach by requiring all extensions to contain the sentences representing the common ground. However, it would be useful to implement a weaker common ground where commonly accepted sentences can be rejected if they are attacked, but have to be accepted otherwise.

The comparison between SBAFs and other approaches to structured argumentation should also be explored further. For instance, it would be interesting to examine how a notion of doubt could be added to frameworks such as ASPIC⁺ or ABA. A further point of comparison are approaches to argumentation where it is uncertain which arguments are present. There, doubt could be modelled by removing an argument, thus possibly presenting an alternative to our account. Such approaches can be found in probabilistic argumentation [26, 30] or argumentation with incomplete frameworks [7, 8]. Finally, it would also be interesting to follow empirical approaches on argumentation [20, 21, 44, 48] and examine whether the informal ideas of allowing for doubt and evaluating frameworks on sentences correspond to the intuitions of ordinary reasoners.

ACKNOWLEDGMENTS

Srdjan Vesic benefited from the support of the project AGGREGY ANR-22-CE23-0005 of the French National Research Agency (ANR). Bruno Yun benefited from the University Claude Bernard, Lyon 1 AAP Accueil funding.

REFERENCES

- [1] Leila Amgoud and Claudette Cayrol. 1998. On the Acceptability of Arguments in Preference-based Argumentation. In *Uncertainty in Artificial Intelligence (UAI98)*, Gregory F. Cooper and Serafin Moral (Eds.). Morgan Kaufmann Publishers, 1–7.
- [2] Leila Amgoud and Srdjan Vesic. 2009. Repairing Preference-Based Argumentation Frameworks. In *International Joint Conference on Artificial Intelligence (IJCAI09)*, Craig Boutilier (Ed.). IJCAI Organization, 665–670.
- [3] Ofer Arieli, AnneMarie Borg, Jesse Heyninck, and Christian Straßer. 2021. Logic-Based Approaches to Formal Argumentation. *Journal of Applied Logics* 8, 6 (2021), 1793–1898.
- [4] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. 2018. Abstract Argumentation Frameworks and Their Semantics. In *Handbook of Formal Argumentation. Volume 1*, Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert Van Der Torre (Eds.). College Publications.
- [5] Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre (Eds.). 2018. *Handbook of Formal Argumentation. Volume 1*. College Publications.
- [6] Pietro Baroni and Massimiliano Giacomin. 2007. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence* 171, 10–15 (2007), 675–700.
- [7] Dorothea Baumeister, Matti Järvisalo, Daniel Neugebauer, Andreas Niskanen, and Jörg Rothe. 2021. Acceptance in incomplete argumentation frameworks. *Artificial Intelligence* 295 (2021), 103470.
- [8] Dorothea Baumeister, Daniel Neugebauer, Jörg Rothe, and Hilmar Schadrack. 2018. Complexity of verification in incomplete argumentation frameworks. *Proceedings of the AAI Conference on Artificial Intelligence (AAAI18)* 32, 1 (2018), 1753–1760.
- [9] Michael Bernreiter, Wolfgang Dvořák, Anna Rapberger, and Stefan Woltran. 2024. The Effect of Preferences in Abstract Argumentation under a Claim-Centric View. *Journal of Artificial Intelligence Research* 81 (2024), 203–262.
- [10] Matti Berthold, Anna Rapberger, and Markus Ulbricht. 2024. Capturing Non-flat Assumption-based Argumentation with Bipolar SETAFs. In *Principles of Knowledge Representation and Reasoning (KR24)*, Pierre Marquis, Magdalena Ortiz, and Maurice Pagnucco (Eds.). Association for the Advancement of Artificial Intelligence, 128–133.
- [11] Philippe Besnard, Alejandro J. García, Anthony Hunter, Sanjay Modgil, Henry Prakken, Guillermo R. Simari, and Francesca Toni. 2014. Introduction to structured argumentation. *Argument and Computation* 5, 1 (2014), 1–4.
- [12] Gregor Betz. 2010. *Theorie dialektischer Strukturen*. Klostermann.
- [13] Guido Boella, Dov Gabbay, Leendert van der Torre, and Serena Villata. 2010. Support in Abstract Argumentation. In *Computational Models of Argument (COMMA10)*, Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo R. Simari (Eds.). IOS Press, 111–122.
- [14] Andrei Bondarenko, Francesca Toni, and Robert Anthony Kowalski. 1993. An assumption-based framework for non-monotonic reasoning. In *Proceedings of the second international workshop on Logic programming and non-monotonic reasoning*, Luís Moniz Pereira and Anil Nerode (Eds.). MIT Press.
- [15] Claudette Cayrol and Marie-Christine Lagasque-Schiex. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU05)*, Lluis Godo (Ed.). Springer, 378–389.
- [16] Claudette Cayrol and Marie-Christine Lagasque-Schiex. 2013. Bipolarity in argumentation graphs: Towards a better understanding. *International Journal of Approximate Reasoning* 54 (2013), 876–899.
- [17] Oana Cocarascu and Francesca Toni. 2017. Mining Bipolar ARgumentation Frameworks from natural language text. In *Workshop on Computational Models of Natural Argument (CMNA@ICAL17)*, Floris Bex, Floriana Grasso, and Nancy Green (Eds.). CEUR, 65–70.
- [18] Andrea Cohen, Simon Parsons, Elizabeth I. Sklar, and Peter McBurney. 2018. A characterization of types of support between structured arguments and their relationship with support in abstract argumentation. *International Journal of Approximate Reasoning* 94 (2018), 76–104.
- [19] Esther Anna Corsi. 2025. Attack principles in sequent-based argumentation. *Journal of Logic and Computation* 35, 3 (2025).
- [20] Marcos Cramer and Mathieu Guillaume. 2018. Directionality of attacks in natural language argumentation. *CEUR Workshop Proceedings* 2261 (2018), 40–46.
- [21] Marcos Cramer and Leendert van der Torre. 2019. SCF2 – An argumentation semantics for rational human judgments on argument acceptability. In *Proceedings of the 8th Workshop on Dynamics of Knowledge and Belief (DKB-2019) and the 7th Workshop KI & Kognition (KIK-2019)*, Christoph Beierle, Marco Ragni, Frieder Stolzenburg, and Matthias Thimm (Eds.). 24–35.
- [22] Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77 (1995), 321–357.
- [23] Wolfgang Dvořák, Anna Rapberger, and Stefan Woltran. 2023. A claim-centric perspective on abstract argumentation semantics: Claim-defeat, principles, and expressiveness. *Artificial Intelligence* 324 (7 no 2023), 104011. <https://doi.org/10.1016/j.artint.2023.104011>
- [24] Melisa Esca Ñuela Gonzalez, Maximiliano Budán, Gerardo Simari, and Guillermo Simari. 2021. Labeled Bipolar Argumentation Frameworks. *Journal of Artificial Intelligence Research* 70 (2021), 1557–1636.
- [25] Charles Leonard Hamblin. 1970. *Fallacies*. Methuen and Co.
- [26] Anthony Hunter, Sylwia Polberg, and Matthias Thimm. 2020. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence* 281, 103236 (2020).
- [27] Ralph H. Johnson and J. Anthony Blair. 2006. *Logical Self-Defense*. International Debate Education Association.
- [28] Marcin Koszowy, Katarzyna Budzyska, Barbara Konat, Steve Oswald, and Pascal Gygax. 2022. A Pragmatic Account of Rephrase in Argumentation: Linguistic and Cognitive Evidence. *Informal Logic* 42, 1 (2022), 49–82.
- [29] Erik C.W. Krabbe and Jan Albert van Laar. 2011. The ways of criticism. *Argumentation* 25, 2 (2011), 199–227.
- [30] Hengfei Li, Nir Oren, and Timothy J. Norman. 2012. Probabilistic argumentation frameworks. In *Theorie and Applications of Formal Argumentation (TAF 2011)*, Sanjay Modgil, Nir Oren, and Francesca Toni (Eds.). Springer, 1–16.
- [31] Sanjay Modgil and Henry Prakken. 2013. A general account of argumentation with preferences. *Artificial Intelligence* 195 (2013), 361–397.
- [32] Sanjay Modgil and Henry Prakken. 2014. The ASPIC+ framework for structured argumentation: A tutorial. *Argument and Computation* 5, 1 (2014), 31–62.
- [33] Farid Nouioua. 2013. AFs with Necessities: Further Semantics and Labelling Characterization. In *Scalable Uncertainty Management (SUM13)*, Weiru Liu, V. S. Subrahmanian, and Jef Wijsen (Eds.). Springer, 120–133.
- [34] Farid Nouioua and Vincent Risch. 2011. Argumentation frameworks with necessities. In *Scalable Uncertainty Management (SUM11)*, Salem Benferhat and John Grant (Eds.). Springer, 163–176.
- [35] Nir Oren and Timothy J. Norman. 2008. Semantics for Evidence-Based Argumentation. In *Computational Models of Argument (COMMA08)*, Philippe Besnard, Sylvie Doutre, and Anthony Hunter (Eds.). IOS Press, 276 – 284.
- [36] Sylwia Polberg and Anthony Hunter. 2018. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning* 93 (2018), 487–543.
- [37] Sylwia Polberg and Nir Oren. 2014. Revisiting Support in Abstract Argumentation Systems. In *Computational Models of Argument (COMMA 2014)*, Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti (Eds.). IOS Press, 369–376.
- [38] Nico Potyka. 2018. Continuous Dynamical Systems for Weighted Bipolar Argumentation. In *Principles of Knowledge Representation and Reasoning (KR18)*. Association for the Advancement of Artificial Intelligence, 148–157.
- [39] Henry Prakken. 2010. An abstract framework for argumentation with structured arguments. *Argument and Computation* 1, 2 (2010), 93–124.
- [40] Fuan Pu, Jian Luo, Yulai Zhang, and Guiming Luo. 2015. Attacker and Defender Counting Approach for Abstract Argumentation. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society (CogSci15)*, David C. Noelle, Rick Dale, Anne S. Warlaumont, Jeff Yoshimi, Teenie Matlock, Carolyn D. Jennings, and Paul P. Maglio (Eds.). 1913–1918.
- [41] Francesca Toni. 2014. A tutorial on assumption-based argumentation. *Argument and Computation* 5, 1 (2014), 89–117.
- [42] Stephen E. Toulmin. 2003. *The Uses of Argument: Updated Edition*. Cambridge University Press.
- [43] Markus Ulbricht, Nico Potyka, Anna Rapberger, and Francesca Toni. 2024. Non-flat ABA Is an Instance of Bipolar Argumentation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI24)*, Vol. 38. 10723–10731.
- [44] Frans van Eemeren, Bart Garssen, and Bert Meuffels. 2009. *Fallacies and Judgements of Reasonableness. Empirical Research Concerning the Pragma-Dialectical Discussion Rules*. Springer.
- [45] Frans H. van Eemeren and Rob Grootendorst. 1984. *Speech acts in argumentative discussions. A theoretical model for analysis of discussions directed towards solving conflicts of opinion*. Foris Publications.
- [46] Frans H. van Eemeren and Rob Grootendorst. 2004. *A Systematic Theory of Argumentation. The pragma-dialectical approach*. Cambridge University Press.
- [47] Jan Albert van Laar and Erik C.W. Krabbe. 2013. The Burden of Criticism: Consequences of Taking a Critical Stance. *Argumentation* 27, 2 (2013), 201–224.
- [48] Srdjan Vesic, Bruno Yun, and Predrag Teovanovic. 2022. Graphical Representation Enhances Human Compliance with Principles for Graded Argumentation Semantics. In *Autonomous Agents and Multiagent Systems (AAMAS22)*, Catherine Pelachaud, Matthew E. Taylor, Piotr Faliszewski, and Viviana Mascardi (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 1319 – 1327.
- [49] Douglas Walton, Chris Reed, and Fabrizio Macagno. 2008. *Argumentation Schemes*. Cambridge University Press.
- [50] Liuwen Yu, Caren Al Anaissy, Srdjan Vesic, Xu Li, and Leendert van der Torre. 2023. A Principle-Based Analysis of Bipolar Argumentation Semantics. In *Logics in Artificial Intelligence (JELIA23)*, Sarah Gaggl, Maria Vanina Martinez, and Magdalena Ortiz (Eds.). Springer, 209–224.