

How we designed winning algorithms for abstract argumentation and which insight we attained

Federico Cerutti^a, Massimiliano Giacomin^b, Mauro Vallati^c

^a*School of Computer Science & Informatics, Queen's Buildings, Cardiff University, CF24 3AA, Cardiff, United Kingdom*

^b*Department of Information Engineering, University of Brescia, via Branze, 38, 25123, Brescia, Italy*

^c*School of Computing and Engineering, University of Huddersfield, Huddersfield, HD1 3DH, United Kingdom*

Abstract

In this paper we illustrate the design choices that led to the development of **ArgSemSAT**, the winner of the preferred semantics track at the 2017 International Competition on Computational Models of Arguments (ICCMA 2017), a biennial contest on problems associated to the Dung's model of abstract argumentation frameworks, widely recognised as a fundamental reference in computational argumentation. The algorithms of **ArgSemSAT** are based on multiple calls to a SAT solver to compute complete labellings, and on encoding constraints to drive the search towards the solution of decision and enumeration problems. In this paper we focus on preferred semantics (and incidentally stable as well), one of the most popular and complex semantics for identifying acceptable arguments. We discuss our design methodology that includes a systematic exploration and empirical evaluation of labelling encodings, algorithmic variations and SAT solver choices. In designing the successful **ArgSemSAT**, we discover that: (1) there is a labelling encoding that appears to be universally better than other, logically equivalent ones; (2) composition of different techniques such as AllSAT and enumerating stable extensions when searching for preferred semantics brings advantages; (3) injecting domain specific knowledge in the algorithm design can lead to significant improvements.

Keywords: Abstract Argumentation, SAT-based Algorithms, Stable and Preferred Semantics

1. Introduction

Computational Argumentation is emerging in Artificial Intelligence as a reasoning paradigm able to handle incomplete and inconsistent information in a

Email addresses: CeruttiF@cardiff.ac.uk (Federico Cerutti), massimiliano.giacomin@unibs.it (Massimiliano Giacomin), m.vallati@hud.ac.uk (Mauro Vallati)

way that fosters the development of robust intelligent systems in e.g. the legal, medical, e-government and debate systems domains [1]. In particular, argument-based debate systems aim not only at supporting human-human argumentation, but also at developing machine-generated arguments.

In this setting, Dung’s model of abstract argumentation frameworks [2] is widely recognised as a fundamental reference in virtue of its simplicity and ability to capture a variety of more specific approaches as special cases, in the areas of non-monotonic reasoning, logic programming and structured argumentation [2, 3, 4, 5, 6]. More specifically, Dung’s model plays a role in the *assessment layer* of argumentation, i.e. to establish the justification of arguments and their conclusions. In argument-based debate technology applications, it supports the assessment of debatewide features such as which arguments are winning [1].

An abstract argumentation framework (*AF*) consists of a set of arguments and an *attack* relation between them. The concept of *extension* plays a key role in this simple setting, where an extension is intuitively a set of arguments which can “survive the conflict together.” Different notions of extension correspond to alternative *argumentation semantics*, whose definitions and properties are actively investigated since more than two decades (see [7, 8] for an introduction).

The three *traditional* argumentation semantics introduced in Dung’s paper [2], namely *grounded*, *stable*, and *preferred* semantics, currently represent the most widely adopted approaches to determine the justification status of arguments in abstract argumentation. They are all based on *complete* semantics, also introduced in [2] as a unifying framework for different semantics, since extensions are selected among the complete ones. While grounded semantics always identifies a single extension which can be computed in polynomial time, stable and preferred semantics allow multiple extensions, yielding in some cases the capability of assigning a committed status to arguments that are left undecided by grounded semantics [9]. Differently from stable semantics, preferred semantics also guarantees the existence of extensions for any argumentation framework, thus providing a justification status of arguments in cases where applying stable semantics collapses [2]. In this respect, the name of preferred semantics reflects a sort of preference w.r.t. the other traditional argumentation semantics.

The main computational problems in abstract argumentation are naturally related to extensions and include decision problems, e.g. determining whether an argument belongs to all extensions, and construction problems, such as enumerating all of the extensions prescribed by a given semantics for a given *AF*. Unfortunately, the complexity of extension-related decision problems turns out to be intractable for most of the semantics [10, 11] and this clearly has a substantial detrimental effect whenever dealing with real-world problems, e.g. [12]. As to stable semantics, credulous and skeptical reasoning are in the first level of the polynomial hierarchy. As to preferred semantics, the advantages of preferred semantics mentioned above come at a price in terms of computational complexity, in particular skeptical reasoning is located at the second level. These complexity issues motivate a recent interest for the investigation on efficient algorithms for abstract argumentation and their empirical assessment. In order to foster the development of such algorithms as well as determining the state-of-the-art

of current implementations, the *International Competition on Computational Models of Argumentation* (ICCMA)¹ has been established on a biennial basis starting in 2015.

While different approaches turn out to be complementary and can be successfully combined in portfolios [13], reduction-based approaches, most of them based on SAT, had the best performance during ICCMA 2017,² the last competition. It is worth noticing that this is not a universal truth as we already discussed at length in [13].

In this paper we illustrate the design choices that led to the development of SAT-based algorithms for preferred semantics, the most challenging semantics—from the complexity perspective—among the traditional ones [2]. These design choices are at the basis of the solver *ArgSemSAT*, the winner of the track at ICCMA 2017 devoted to preferred semantics. We also consider stable semantics since it can play a significant role also in the context of preferred semantics.

As it will be clear in the following, high performance of *ArgSemSAT* has been achieved by applying SAT-reductions in a relatively simple way. In particular, no decomposition techniques on the argumentation framework have been applied, since previous experiences on the application of the SCC-recursive schema [14, 15] seem to indicate that such techniques are efficient only when the number of strongly connected components is large. Moreover, we have exploited classical SAT solvers with hard clauses, leaving apart the use of soft clauses which may play a role in the identification of preferred extensions [16]. It should be pointed out that we are not claiming that more advanced SAT-based techniques have a detrimental effect on performance, rather that one of the main reasons why *ArgSemSAT* achieves efficiency is a careful identification of the features of the reduction to SAT.

In particular, it turns out that modelling labellings (i.e. a counterpart of extensions based on labels) with different boolean formulas, as well as expressing the constraints involving labellings with different logically equivalent encodings, has a dramatic impact on performance. In addition, there might be different options in the exploration of the set of labellings; for instance in the context of preferred semantics one may either begin from scratch or start from stable labellings by exploiting the relevant algorithms. Finally, algorithms can be implemented using different SAT solvers.

Since all of these aspects are interrelated, a systematic analysis is needed to explore all different options, followed by an empirical evaluation to identify the most efficient combinations. This is the reason why in this paper we do not compare *ArgSemSAT* against the other solvers of the competition, rather we focus on the above features that—as experiments show—have a dramatic impact on performance. This way, we describe a methodology that can be applied also to different SAT-based approaches for abstract argumentation, possibly exploiting advanced features based e.g. on graph decomposition and soft clauses.

¹<http://argumentationcompetition.org> (on 15 February 2019).

²<http://argumentationcompetition.org/2017/results.html> (on 15 February 2019).

Our methodology comprises two steps: first correct algorithms and their variations need to be identified. In this paper we introduce at an abstract level the algorithms of `ArgSemSAT` in Section 3. We then explore possible encodings of the labellings and of the relevant constraints in Section 4, and detail algorithm implementations in Section 5.

Then, the various algorithms and variants need to be empirically evaluated, and Section 6 is devoted to the experimental analysis of `ArgSemSAT` configurations and to assess a number of hypotheses concerning the interplay of the different features mentioned above.

One of the main take-home messages of our work is that it is of pivotal importance to promote the flexibility and configurability of solvers. The well-known *no free-lunch theorem* [17] indicates that any two solvers will deliver indistinguishable performance when compared on a sufficiently large and diverse set of benchmarks. However, since usually solvers are exploited on a much smaller set of instances, the ability to specialise a solver for a class of instances of interest can lead to paramount performance improvement. What our work fosters, and our experimental analysis suggests, is that any general (i.e., domain independent) solver dealing with complex problems should incorporate different methods and techniques, and then allow the selection, configuration, and use of the most suitable ones given the instances at hand. This is well-aligned with the research done in the area of algorithm configuration [18, 19]. Furthermore, it is worth emphasising that the specialisation of solvers according to some testing benchmarks can also lead to a better understanding of solvers behaviours, in terms of weaknesses and strengths (see, for instance [20]). Analysing and comparing different configurations can shed some light on the elements that are affecting performance, potentially leading to further improvements.

Such a data-driven methodology can then be applied to analyse advantages and disadvantages of approaches based on advanced features. For instance, in [15, 21] we discussed how SCC-decomposition might play a role when dealing with certain classes of *AFs*. Besides the presence at the local level of similar variants as in a global algorithm, such an advanced feature could potentially be exploited by an algorithm, hence the need to analyse and compare different configurations against a large set of benchmarks.

The validity of such a methodology is ensured by the fact that `ArgSemSAT` is the winner of the preferred semantics track at ICCMA 2017: further analyses with currently prototypical approaches to exploit advanced features are left as future work.

Comparison with relevant work is provided in Section 7, while in Section 8 we have some conclusive remarks. Unless otherwise stated, proofs are listed in Appendix A. Appendix B summarises the parameters used in the algorithms we introduce in Section 5. Finally, Appendix C illustrates a pilot study investigating the difference in performance when using one propositional variable instead of three in the search for stable extensions using a SAT solver.

2. Background

2.1. Argumentation Frameworks and Semantics

An argumentation framework [2] consists of a set of arguments³ and a binary attack relation between them.

Definition 1. An *argumentation framework* (AF) is a pair $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ where \mathcal{A} is a set of arguments and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$. We say that \mathbf{b} *attacks* \mathbf{a} iff $(\mathbf{b}, \mathbf{a}) \in \mathcal{R}$, also denoted as $\mathbf{b} \rightarrow \mathbf{a}$. The set of attackers of an argument \mathbf{a} is denoted as $\mathbf{a}^- \triangleq \{\mathbf{b} : \mathbf{b} \rightarrow \mathbf{a}\}$, the set of arguments attacked by \mathbf{a} is denoted as $\mathbf{a}^+ \triangleq \{\mathbf{b} : \mathbf{a} \rightarrow \mathbf{b}\}$. An argument \mathbf{a} without attackers, i.e. such that $\mathbf{a}^- = \emptyset$, is said *initial*. We also extend attack notations to sets of arguments, i.e. given $E, S \subseteq \mathcal{A}$, $E \rightarrow \mathbf{a}$ iff $\exists \mathbf{b} \in E$ s.t. $\mathbf{b} \rightarrow \mathbf{a}$; $\mathbf{a} \rightarrow E$ iff $\exists \mathbf{b} \in E$ s.t. $\mathbf{a} \rightarrow \mathbf{b}$; $E \rightarrow S$ iff $\exists \mathbf{b} \in E, \mathbf{a} \in S$ s.t. $\mathbf{b} \rightarrow \mathbf{a}$; $E^- \triangleq \{\mathbf{b} \mid \exists \mathbf{a} \in E, \mathbf{b} \rightarrow \mathbf{a}\}$ and $E^+ \triangleq \{\mathbf{b} \mid \exists \mathbf{a} \in E, \mathbf{a} \rightarrow \mathbf{b}\}$.

Each argumentation framework has an associated directed graph where the vertices are the arguments, and the edges are the attacks.

The basic properties of conflict-freeness, acceptability, and admissibility of a set of arguments are fundamental for the definition of argumentation semantics.

Definition 2. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$:

- a set $S \subseteq \mathcal{A}$ is a *conflict-free* set of Γ if $\nexists \mathbf{a}, \mathbf{b} \in S$ s.t. $\mathbf{a} \rightarrow \mathbf{b}$;
- an argument $\mathbf{a} \in \mathcal{A}$ is *acceptable* in Γ with respect to a set $S \subseteq \mathcal{A}$ if $\forall \mathbf{b} \in \mathcal{A}$ s.t. $\mathbf{b} \rightarrow \mathbf{a}$, $\exists \mathbf{c} \in S$ s.t. $\mathbf{c} \rightarrow \mathbf{b}$;
- the function $\mathcal{F}_\Gamma : 2^{\mathcal{A}} \rightarrow 2^{\mathcal{A}}$ such that $\mathcal{F}_\Gamma(S) = \{\mathbf{a} \mid \mathbf{a} \text{ is acceptable in } \Gamma \text{ w.r.t. } S\}$ is called the *characteristic function* of Γ ;
- a set $S \subseteq \mathcal{A}$ is an *admissible* set of Γ if S is a conflict-free set of Γ and every element of S is acceptable in Γ with respect to S , i.e. $S \subseteq \mathcal{F}_\Gamma(S)$.

An argumentation semantics σ prescribes for any AF Γ a set of *extensions*, denoted as $\mathcal{E}_\sigma(\Gamma)$, namely a set of sets of arguments satisfying the conditions dictated by σ . The paper is focused on stable (denoted as ST), and preferred (denoted as PR) semantics, introduced as follows.

Definition 3. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$:

- a set $S \subseteq \mathcal{A}$ is a *stable extension* of Γ , i.e. $S \in \mathcal{E}_{\text{ST}}(\Gamma)$, iff S is a conflict-free set of Γ and $S \cup S^+ = \mathcal{A}$;
- a set $S \subseteq \mathcal{A}$ is a *preferred extension* of Γ , i.e. $S \in \mathcal{E}_{\text{PR}}(\Gamma)$, iff S is a maximal (w.r.t. set inclusion) admissible set of Γ .

³In this paper we consider only *finite* sets of arguments: see [22] for a discussion on infinite sets of arguments.

The notion of complete extension introduced in [2] represents a common ground for both semantics, i.e. both preferred and stable extensions can be defined as complete extensions satisfying specific maximality properties. This is recalled in Proposition 1.

Definition 4. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ is a *complete extension* of Γ iff S is a conflict-free set of Γ and $S = \mathcal{F}_\Gamma(S)$. The set of complete extensions of Γ is denoted as $\mathcal{E}_{\text{CO}}(\Gamma)$.

Proposition 1 (Originally introduced in [2]). Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ is a preferred extension of Γ iff S is a maximal (w.r.t. set inclusion) complete extension, it is a stable extension of Γ iff it is a complete extension such that for all $\mathbf{a} \in \mathcal{A}$, $\mathbf{a} \in (S \cup S^+)$.

From the above proposition it is easy to see that $\mathcal{E}_{\text{ST}}(\Gamma) \subseteq \mathcal{E}_{\text{PR}}(\Gamma) \subseteq \mathcal{E}_{\text{CO}}(\Gamma)$.

Taking into account that the extensions prescribed by any semantics are conflict-free, one can note that each extension S implicitly defines a three-valued *labelling* of arguments: an argument \mathbf{a} is labelled **in** iff $\mathbf{a} \in S$; it is labelled **out** iff $S \rightarrow \mathbf{a}$; it is labelled **undec** if neither of the above conditions holds. In the light of this correspondence, argumentation semantics can be equivalently defined in terms of labellings rather than of extensions [23, 24, 25, 8].

Definition 5. Given a set of arguments S , a *labelling* of S is a total function $\mathcal{L}ab : S \mapsto \{\mathbf{in}, \mathbf{out}, \mathbf{undec}\}$. The set of all *labellings* of S is denoted as \mathcal{L}_S . Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, a *labelling* of Γ is a labelling of \mathcal{A} . The set of all *labellings* of Γ is denoted as $\mathcal{L}(\Gamma)$.

Given a labelling $\mathcal{L}ab$, we write $\mathbf{in}(\mathcal{L}ab)$ for $\{A \mid \mathcal{L}ab(A) = \mathbf{in}\}$, $\mathbf{out}(\mathcal{L}ab)$ for $\{A \mid \mathcal{L}ab(A) = \mathbf{out}\}$ and $\mathbf{undec}(\mathcal{L}ab)$ for $\{A \mid \mathcal{L}ab(A) = \mathbf{undec}\}$.

Labellings can be related by the commitment relation introduced in [26]:

Definition 6. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab_1, \mathcal{L}ab_2 \in \mathcal{L}(\Gamma)$. $\mathcal{L}ab_2$ is more or equally committed than $\mathcal{L}ab_1$ ($\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$) iff $\mathbf{in}(\mathcal{L}ab_1) \subseteq \mathbf{in}(\mathcal{L}ab_2)$ and $\mathbf{out}(\mathcal{L}ab_1) \subseteq \mathbf{out}(\mathcal{L}ab_2)$. Furthermore, $\mathcal{L}ab_1 \sqsubset \mathcal{L}ab_2$ iff $\mathcal{L}ab_1 \sqsubseteq \mathcal{L}ab_2$ and $\mathcal{L}ab_2 \not\sqsubseteq \mathcal{L}ab_1$.

As shown in [26], \sqsubseteq is a partial order between labellings:

Proposition 2 (Originally introduced in [26]). The relation \sqsubseteq between labellings is reflexive, transitive and anti-symmetric.

Complete labellings can be defined⁴ as follows.

Definition 7. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework. A labelling $\mathcal{L}ab \in \mathcal{L}(\Gamma)$ is a *complete labelling* of Γ iff it satisfies the following conditions for any $\mathbf{a} \in \mathcal{A}$:

⁴The reader may notice that the definition is redundant, e.g. the third condition may be dropped since it is entailed by the first two ones. The choice of this definition is for the sake of exposition.

- $\mathcal{L}ab(\mathbf{a}) = \text{in} \Leftrightarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) = \text{out};$
- $\mathcal{L}ab(\mathbf{a}) = \text{out} \Leftrightarrow \exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) = \text{in};$
- $\mathcal{L}ab(\mathbf{a}) = \text{undec} \Leftrightarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) \neq \text{in} \wedge \exists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{c}) = \text{undec}.$

The stable and preferred labellings can then be defined on the basis of complete labellings.

Definition 8. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework. A labelling $\mathcal{L}ab \in \mathfrak{L}(\Gamma)$ is

- a *stable labelling* of Γ iff it is a complete labelling of Γ and there is no argument labelled **undec**;
- a *preferred labelling* of Γ iff it is a complete labelling of Γ maximising the set of arguments labelled **in**.

The sets including the complete, stable and preferred labellings of an argumentation framework Γ are denoted as $\mathfrak{L}_{\text{CO}}(\Gamma)$, $\mathfrak{L}_{\text{ST}}(\Gamma)$ and $\mathfrak{L}_{\text{PR}}(\Gamma)$, respectively. From the above definition, it is easy to see that $\mathfrak{L}_{\text{ST}}(\Gamma) \subseteq \mathfrak{L}_{\text{PR}}(\Gamma) \subseteq \mathfrak{L}_{\text{CO}}(\Gamma)$.

On the basis of the results in [2] and [26], the following relations between complete and preferred labellings can be derived:

Proposition 3 (Originally introduced in [26]). Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ and a labelling $\mathcal{L}ab$, $\mathcal{L}ab$ is a preferred labelling of Γ if and only if $\mathcal{L}ab$ is a maximal complete labelling (w.r.t. \sqsubseteq) of Γ . Moreover, given a complete labelling $\mathcal{L}ab$, there exists a preferred labelling $\mathcal{L}ab'$ such that $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$.

In order to show the connection between extensions and labellings, let us recall from [8] the definition of the function **Ext2Lab**, returning the labelling corresponding to a conflict-free set of arguments S .

Definition 9. Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ and a conflict-free set $S \subseteq \mathcal{A}$, the corresponding labelling **Ext2Lab**(S) is the labelling $\mathcal{L}ab$ such that

- $\mathcal{L}ab(\mathbf{a}) = \text{in} \Leftrightarrow \mathbf{a} \in S$
- $\mathcal{L}ab(\mathbf{a}) = \text{out} \Leftrightarrow \exists \mathbf{b} \in S \text{ s.t. } \mathbf{b} \rightarrow \mathbf{a}$
- $\mathcal{L}ab(\mathbf{a}) = \text{undec} \Leftrightarrow \mathbf{a} \notin S \wedge \nexists \mathbf{b} \in S \text{ s.t. } \mathbf{b} \rightarrow \mathbf{a}$

As shown in [25], there is a bijective correspondence between the complete, stable, preferred extensions and the complete, stable, preferred labellings, respectively.

Proposition 4 (Originally introduced in [25]). Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathcal{L}ab$ is a complete (stable, preferred) labelling of Γ if and only if there is a complete (stable, preferred) extension S of Γ such that $\mathcal{L}ab = \text{Ext2Lab}(S)$.

Semantics σ		
	ST	PR
DC- σ	NP-compl.	NP-compl.
DS- σ	coNP-compl.	Π_2^P -compl.

Table 1: Computational complexity of credulous and skeptical acceptance in finite AFs: DC denotes credulous acceptance, while DS skeptical acceptance.

2.2. Computational Problems in Abstract Argumentation

The main computational problems in abstract argumentation are related to the different semantics and can be partitioned into two classes: *decision* problems and *construction* problems. Decision problems pose yes/no questions like “Does this argument belong to one (all) extensions?” or “Is this set an extension?”, while construction problems require to explicitly produce some of the extensions prescribed by a semantics. The complexity of extension-related decision problems has been deeply investigated and, for most of the semantics proposed in the literature they have been proven to be intractable. Intractability extends directly to construction/enumeration problems, given that their solutions provide direct answers to decision problems.

In this paper we focus on *extension enumeration* (a construction problem), *credulous acceptance* and *skeptical acceptance* of an argument (both of them decision problems).

Extension enumeration is the problem of constructing all of the extensions prescribed by a given semantics for a given AF. Following [27], we denote such a problem as EE- σ .

Credulous and skeptical acceptance refer to a specific argument, where the natural questions arising are related to extension inclusion. In particular, an argument \mathbf{a} is *credulously* accepted with regard to a given semantics σ and a given AF Γ iff \mathbf{a} belongs to at least one extension of Γ under σ , i.e. $\exists E \in \mathcal{E}_\sigma(\Gamma)$ s.t. $\mathbf{a} \in E$. We denote the problem of determining whether an argument of AF is credulously accepted with regard to a semantics σ as DC- σ . An argument \mathbf{a} is *skeptically* accepted with regard to a given semantics σ and a given AF Γ iff \mathbf{a} belongs to each extension of Γ under σ , i.e. $\forall E \in \mathcal{E}_\sigma(\Gamma) \mathbf{a} \in E$. We denote the corresponding decision problem as DS- σ .⁵

The complexity of credulous and skeptical acceptance problems for stable and preferred semantics are shown in Table 1 [10]. We refer the reader to [11] for a recent investigation on the complexity of the extension enumeration problem under different semantics.

Given the bijective correspondence between extensions and labellings proved by Proposition 4, the above problems can be directly formulated as the problem

⁵Please note that if $\mathcal{E}_\sigma(\Gamma) = \emptyset$, it is trivially true that $\forall \mathbf{a} \in \mathcal{A} \text{ DS-}\sigma(\mathbf{a}) = \top$ (vacuous truth).

of computing all labellings prescribed by a given semantics (corresponding to extensions enumeration), determining whether there is a labelling where an argument is labelled **in** (credulous acceptance) and whether an argument is labelled **in** by all labellings (skeptical acceptance).

2.3. Propositional satisfiability and SAT solvers

A propositional formula over a set of boolean variables is satisfiable if and only if there exists a truth assignment of the variables such that the formula evaluates to True. Checking whether such an assignment exists is the satisfiability (SAT) problem, and an algorithm able to solve this problem is called SAT solver.

Most SAT solvers require as input a propositional formula in conjunctive normal form (CNF), i.e. expressed as a conjunction of clauses, where each clause is a disjunction of literals (in the context of propositional logic, a literal is either a propositional variable or its negation). Moreover, if the formula is satisfiable SAT solvers are generally able to return as output a satisfying assignment.

A variation of the SAT problem is the so-called AllSAT problem, namely given a CNF formula as input, determining the set of all its satisfying assignments. A typical AllSAT solver (ALLSatS) iteratively computes satisfying assignments using a traditional SAT solver and adds to the input formula some *blocking clauses* which are the complement of total or partial assignments to exclude previously identified solutions. More specifically, a practical ALLSatS receives a CNF formula as input and solves the AllSAT problem by determining an equivalent formula in disjunctive normal form (DNF), i.e. as a disjunction of clauses where each clause is a conjunction of literals. From [28]—that builds on top of [29] and [30]—the generic structure of an AllSAT solver is presented in Algorithm 1. Line 3 uses a SAT solver (**SatS**) to compute a satisfying assignment for the formula F . Line 4 calls the unspecified function **compute** to obtain the partial assignment (cube) q_i from m_i (total assignment). A blocking clause c_i is found and added to F_i preventing these assignments from being enumerated any further. This loop ends when F_i becomes unsatisfiable.

Algorithm 1 Algorithm template for AllSAT solver, [28, Algorithm 1]

Input: F (a CNF Formula)
Output: Q (a DNF Formula) such that $F \iff Q$

- 1: $i := 1, F_1 := F, Q := \perp$
- 2: **while not unsat**(F_i) **do**
- 3: $m_i := \text{SatS}(F_i)$ ▷ get a satisfying assignment of F_i
- 4: $(q_i, c_i) := \text{compute}(m_i)$ ▷ get cube and blocking clause
- 5: $Q := Q \vee q_i, F_{i+1} := F_i \wedge c_i$ ▷ update Q and F_i
- 6: $i := i + 1$
- 7: **end while**
- 8: **return** Q

As noticed in [28], a very simple implementation of **compute** is to set $q_i := m_i$, and $c_i := \neg q_i$. In [28] the authors discuss at length more efficient ways to

implement such a `compute` function, and achieve speed improvements up to three orders of magnitude.

There are also extensions to the SAT and AllSAT problems where clauses are partitioned into hard and soft clauses, and the solutions are determined among all assignments satisfying all hard clauses and maximising the sum of numerical weights associated to soft clauses. However, as specified in Section 1, in this paper we focus on *classical* SAT and AllSAT solvers: an application of weighted MaxSAT to compute the preferred extensions is described in [16].

3. Abstract Algorithms for Stable and Preferred Semantics

In this section we provide in abstract form the algorithms we consider for the solution to the computational problems under stable and preferred semantics, respectively. These algorithms perform a search in the space of (a subset of) the complete labellings, where each step of the search process requires the invocation of a non-deterministic function which returns a complete labelling satisfying a number of specific constraints. Acceptable solutions of the non-deterministic functions are encoded as a propositional formula, in such a way that each satisfying assignment corresponds to a desired complete labelling. This way, the function can be implemented by means of a SAT solver. Different ways of identifying the corresponding formulas will be introduced in Section 4.

3.1. Stable Semantics

The first problem we consider is the enumeration of stable labellings. Algorithm 2 starts from an empty set of labellings \mathcal{L}_s (initialised at line 1) and simply consists of a loop (lines 2–4) which at each iteration identifies a novel stable labelling (line 2) and records it in \mathcal{L}_s (line 3). In particular, FindST^\neq is a non-deterministic function that, given an argumentation framework Γ and a set of stable labellings $\mathcal{L}_s \subseteq \mathcal{L}_{\text{ST}}(\Gamma)$, returns a stable labelling $\text{Lab} \in \mathcal{L}_{\text{ST}}(\Gamma)$ such that $\text{Lab} \notin \mathcal{L}_s$ if such a labelling exists, and returns \perp in the other case. As shown in Section 4, this function (as well as FindST_a in Algorithms 3 and 4) can identify a stable labelling as a satisfying assignment of a propositional formula which encodes the constraints of complete labellings enriched with the condition of having no `undec` arguments. The issue of encoding complete labellings is dealt with in Section 4.1.

It can be seen that the loop of Algorithm 2 enumerates all the satisfying assignments of a propositional formula, corresponding to the stable labellings of Γ . As a consequence, the loop can be replaced by a single call to an AllSAT solver: this will be available as an option of the implemented algorithm (see Section 5).

Credulous acceptance under stable semantics can simply be checked by searching for a stable labelling which assigns the `in`-label to the input argument. Accordingly, Algorithm 3 consists of a single line exploiting the non-deterministic function FindST_a which receives as input an argumentation framework $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, an argument $a \in \mathcal{A}$ and a label $s \in \{\text{in}, \text{out}\}$, and returns as output a

Algorithm 2 Enumeration of Stable Labellings

SAT-EL-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ **Output:** $\mathcal{L}_s = \mathcal{E}_{\text{ST}}(\Gamma)$

- 1: $\mathcal{L}_s := \emptyset$
 - 2: **while** $\mathcal{L}ab := \text{FindST}^\neq(\Gamma, \mathcal{L}_s)$ **do**
 - 3: $\mathcal{L}_s := \mathcal{L}_s \cup \{\mathcal{L}ab\}$
 - 4: **end while**
 - 5: **return** \mathcal{L}_s
-

stable labelling $\mathcal{L}ab \in \mathcal{E}_{\text{ST}}(\Gamma)$ such that $\mathcal{L}ab(\mathbf{a}) = s$ if such a labelling exists, \perp otherwise.

Algorithm 3 Credulous acceptance under Stable Semantics

SAT-DC-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{a} \in \mathcal{A}$ **Output:** \top if \mathbf{a} is credulously accepted under ST, \perp otherwise

- 1: **return** $\text{FindST}_{\mathbf{a}}(\Gamma, \text{in}) \neq \perp$
-

As to skeptical acceptance, Algorithm 4 checks whether there is a stable labelling such that the argument is out (line 1), returning \perp if this is the case (line 2). If such a labelling does not exist, then either the argumentation framework has no stable labellings or every stable labelling labels the argument in : in both cases the argument is skeptically accepted, and the algorithm returns \top accordingly (line 4).

Algorithm 4 Skeptical acceptance under Stable Semantics

SAT-DS-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{a} \in \mathcal{A}$ **Output:** \top if \mathbf{a} is skeptically accepted under ST, \perp otherwise

- 1: **if** $\text{FindST}_{\mathbf{a}}(\Gamma, \text{out}) \neq \perp$ **then**
 - 2: **return** \perp
 - 3: **end if**
 - 4: **return** \top
-

Correctness of Algorithms 2, 3 and 4 is pointed out in the following theorem, whose proof is obvious and is thus omitted.

Theorem 1. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{a} \in \mathcal{A}$:

- $\{\text{in}(\mathcal{L}ab) \mid \mathcal{L}ab \in \text{SAT-EL-ST}(\Gamma)\} = \mathcal{E}_{\text{ST}}(\Gamma)$
- $\text{SAT-DC-ST}(\Gamma, \mathbf{a}) \equiv (\exists S \in \mathcal{E}_{\text{ST}}(\Gamma), \mathbf{a} \in S)$
- $\text{SAT-DS-ST}(\Gamma, \mathbf{a}) \equiv (\forall S \in \mathcal{E}_{\text{ST}}(\Gamma), \mathbf{a} \in S)$

3.2. Preferred Semantics

Due to the complexity gap between SAT and the preferred labelling enumeration problem, it is not possible to identify in polynomial time a propositional formula (of polynomial size) whose models correspond to preferred labellings. In particular, differently to the case of stable semantics, in order to identify preferred labellings it is necessary to maximise **in**-arguments among complete labellings, thus the constraints corresponding to preferred labellings cannot be obtained by adding a simple condition to the constraints of complete labellings.

Algorithm 5 enumerates the preferred labellings of a framework Γ by means of two nested loops. The outer loop (lines 2-7) starts from an empty set of labellings \mathcal{L}_p (line 1) and identifies at each iteration a novel preferred labelling $\mathcal{L}ab$, which is added to \mathcal{L}_p (line 6).

In particular, FindCL^{\boxplus} is a non-deterministic function that, given an argumentation framework Γ and a set of preferred labellings $\mathcal{L}_p \subseteq \mathfrak{L}_{\text{PR}}(\Gamma)$, picks out a labelling $\mathcal{L}ab$ from a set of complete labellings \mathcal{L} subject to the following constraints:

- $\mathcal{L} \subseteq \mathfrak{L}_{\text{CO}}(\Gamma)$, i.e. \mathcal{L} only includes complete labellings of Γ
- $\forall \mathcal{L}ab' \in \mathcal{L}, \forall \mathcal{L}ab'' \in \mathcal{L}_p, \mathcal{L}ab' \not\boxplus \mathcal{L}ab''$, i.e. \mathcal{L} does not include those labellings that are less or equally committed w.r.t. a labelling of \mathcal{L}_p
- $(\mathfrak{L}_{\text{PR}}(\Gamma) \setminus \mathcal{L}_p) \subseteq \mathcal{L}$, i.e. \mathcal{L} includes the preferred labellings that are not already included in \mathcal{L}_p

If there is no labelling in \mathcal{L} , the function returns \perp .

Note that the above constraints on \mathcal{L} are compatible, in particular the preferred labellings that are not included in \mathcal{L}_p cannot be less committed than any labelling of \mathcal{L}_p , due to Proposition 3.

Ideally, the set \mathcal{L} should only include the preferred labellings of Γ that have not yet been found, i.e. $\mathcal{L} = \mathfrak{L}_{\text{PR}}(\Gamma) \setminus \mathcal{L}_p$, entailing that a preferred labelling is directly identified at line 2. Since this is not possible in general, $\text{FindCL}^{\boxplus}(\Gamma, \mathcal{L}_p)$, line 2, identifies a complete labelling which is not necessary preferred, but due to the second constraint on \mathcal{L} there is a preferred labelling which is more or equally committed w.r.t. it. This preferred labelling is identified by the inner loop (lines 3-5) which starts from the complete labelling $\mathcal{L}ab$ and produces a sequence of more committed complete labellings strictly ordered w.r.t. \boxminus . In particular, $\text{FindCL}^{\boxminus}$ is a non-deterministic function that, given an argumentation framework Γ and a complete labelling $\mathcal{L}ab \in \mathfrak{L}_{\text{CO}}(\Gamma)$, returns a complete labelling $\mathcal{L}ab'$ such that $\mathcal{L}ab \boxminus \mathcal{L}ab'$ if such a labelling exists, \perp otherwise. When the sequence can no more be extended, $\text{FindCL}^{\boxminus}(\Gamma, \mathcal{L}ab)$ returns \perp at line 3, and its last element $\mathcal{L}ab$ correspond to a maximal complete labelling w.r.t. \boxminus , namely to a preferred labelling.

Since stable labellings are also preferred, instead of starting from an empty set of preferred labellings at line 1 it is possible to first compute the stable labellings by Algorithm 2, i.e. line 1 is replaced with $\mathcal{L}_p := \text{SAT-EL-ST}(\Gamma)$. This straightforward variation is reported in Algorithm 6.

Algorithm 5 Enumeration of Preferred Labellings

EL-PR**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ **Output:** $\mathcal{L}_p = \mathfrak{L}_{\text{PR}}(\Gamma)$

```
1:  $\mathcal{L}_p := \emptyset$ 
2: while  $\mathcal{L}ab := \text{FindCL}^{\oplus}(\Gamma, \mathcal{L}_p)$  do
3:   while  $\mathcal{L}ab' := \text{FindCL}^{\overline{\oplus}}(\Gamma, \mathcal{L}ab)$  do
4:      $\mathcal{L}ab := \mathcal{L}ab'$ 
5:   end while
6:    $\mathcal{L}_p := \mathcal{L}_p \cup \{\mathcal{L}ab\}$ 
7: end while
8: return  $\mathcal{L}_p$ 
```

Algorithm 6 Enumeration of Preferred Labellings: variation using the enumeration of Stable Labellings

EL-PR-withST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ **Output:** $\mathcal{L}_p = \mathfrak{L}_{\text{PR}}(\Gamma)$

```
1:  $\mathcal{L}_p := \text{SAT-EL-ST}(\Gamma)$ 
2: while  $\mathcal{L}ab := \text{FindCL}^{\oplus}(\Gamma, \mathcal{L}_p)$  do
3:   while  $\mathcal{L}ab' := \text{FindCL}^{\overline{\oplus}}(\Gamma, \mathcal{L}ab)$  do
4:      $\mathcal{L}ab := \mathcal{L}ab'$ 
5:   end while
6:    $\mathcal{L}_p := \mathcal{L}_p \cup \{\mathcal{L}ab\}$ 
7: end while
8: return  $\mathcal{L}_p$ 
```

Since preferred labellings do not directly correspond to the satisfying assignments of a propositional formula, differently to the case of stable semantics AllSAT solvers are not directly applicable here. An option would be to determine all complete labellings by solving an AllSAT problem and then to identify the maximal ones w.r.t. \sqsubseteq , however this would be inefficient compared to Algorithms 5 and 6, which are able to cut the search space of complete labellings.

Turning to credulous acceptance, similarly to the case of stable semantics it suffices to determine whether a complete labelling exists which assigns the **in** label to the input argument (thus ensuring that there is a preferred labelling, possibly coinciding with the identified complete one, satisfying this condition). This corresponds to Algorithm 7, where $\text{FindCL}_{\mathbf{a}}$ is a function that, given an argumentation framework Γ , an argument $\mathbf{a} \in \mathcal{A}$ and a label $s \in \{\mathbf{in}, \mathbf{out}\}$, returns a complete labelling $\mathcal{L}ab \in \mathcal{L}_{\text{CO}}(\Gamma)$ such that $\mathcal{L}ab(\mathbf{a}) = s$ if such a labelling exists, \perp otherwise.

Algorithm 7 Credulous acceptance under Preferred Semantics

DC-PR

Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{a} \in \mathcal{A}$

Output: \top if \mathbf{a} is credulously accepted under PR, \perp otherwise

1: **return** $\text{FindCL}_{\mathbf{a}}(\Gamma, \mathbf{in}) \neq \perp$

Skeptical acceptance can be determined by enumerating the preferred labellings of the input argumentation framework, and checking whether all of them assign the label **in** to the input argument. This corresponds to Algorithm 8, a straightforward adaptation of Algorithm 5 which iterates over the preferred labellings of Γ and returns \perp whenever it identifies a preferred labelling where the input argument \mathbf{a} is not labelled **in** (lines 6 and 7). If for all preferred labellings this is not the case, then \mathbf{a} is skeptically accepted and the algorithm returns \top at line 11.

Similarly to the case of Algorithm 5, it is also possible to start from the stable labellings to check whether one of them does not assign the label **in** to \mathbf{a} . However Algorithm 9 shows a rather different alternative to Algorithm 8. Since stable extensions by definition do not assign the label **undec**, it is more efficient to check with a single call to the function $\text{FindCL}_{\mathbf{a}}$ whether there is a complete labelling which assigns the label **out** to \mathbf{a} , entailing the existence of a preferred labelling satisfying this condition and thus disproving skeptical acceptance of \mathbf{a} . Moreover, it can also be checked whether there is a complete labelling assigning the label **in** to \mathbf{a} , and if this is not the case then \mathbf{a} is not skeptically accepted either (line 1-3).

If the algorithm enters line 4 then there are no preferred labellings where \mathbf{a} is **out**, and the algorithm looks for a preferred labelling assigning the label **undec** to \mathbf{a} in order to prove that \mathbf{a} is not skeptically accepted. To this purpose, a loop (lines 5–13) starts from an empty set of complete labellings \mathcal{L}_c , and considers at each iteration a maximal complete labelling such that \mathbf{a} is labelled **undec**. For ease of explanation, let us define \mathcal{L}_{mu} as the set including the

Algorithm 8 Skeptical acceptance under Preferred Semantics 1

DS-PR-straightforward

Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{a} \in \mathcal{A}$

Output: \top if \mathbf{a} is skeptically accepted under PR, \perp otherwise

```

1:  $\mathcal{L}_p := \emptyset$ 
2: while  $\mathcal{L}ab := \text{FindCL}^{\boxplus}(\Gamma, \mathcal{L}_p)$  do
3:   while  $\mathcal{L}ab' := \text{FindCL}^{\boxminus}(\Gamma, \mathcal{L}ab)$  do
4:      $\mathcal{L}ab := \mathcal{L}ab'$ 
5:   end while
6:   if  $\mathcal{L}ab(\mathbf{a}) \neq \text{in}$  then
7:     return  $\perp$ 
8:   end if
9:    $\mathcal{L}_p := \mathcal{L}_p \cup \{\mathcal{L}ab\}$ 
10: end while
11: return  $\top$ 

```

labellings of this kind, i.e. $\mathcal{L}_{mu} \triangleq \{\mathcal{L}ab \in \mathfrak{L}_{CO}(\Gamma) \mid \mathcal{L}ab(\mathbf{a}) = \text{undec} \wedge \nexists \mathcal{L}ab' \in \mathfrak{L}_{CO}(\Gamma) \text{ s.t. } (\mathcal{L}ab'(\mathbf{a}) = \text{undec} \wedge \mathcal{L}ab \sqsubseteq \mathcal{L}ab')\}$. $\text{FindCLU}_{\mathbf{a}}^{\boxplus}$ is a non-deterministic function that, given an argument $\mathbf{a} \in \mathcal{A}$, an argumentation framework Γ and a set of complete labellings $\mathcal{L}_c \subseteq \mathcal{L}_{mu}$, picks out a complete labelling $\mathcal{L}ab$ (not necessarily maximal) from a set of complete labellings \mathcal{L} subject to the following constraints:

- $\mathcal{L} \subseteq \{\mathcal{L}ab' \in \mathfrak{L}_{CO}(\Gamma) \mid \mathcal{L}ab'(\mathbf{a}) = \text{undec}\}$, i.e. \mathcal{L} only includes complete labellings of Γ which assign the label **undec** to \mathbf{a} ;
- $\forall \mathcal{L}ab' \in \mathcal{L}, \forall \mathcal{L}ab'' \in \mathcal{L}_c, \mathcal{L}ab' \not\sqsubseteq \mathcal{L}ab''$, i.e. \mathcal{L} does not include those labellings that are less or equally committed w.r.t. a labelling of \mathcal{L}_c ;
- $(\{\mathcal{L}ab' \in \mathfrak{L}_{PR}(\Gamma) \mid \mathcal{L}ab'(\mathbf{a}) = \text{undec}\} \setminus \mathcal{L}_c) \subseteq \mathcal{L}$, i.e. \mathcal{L} includes the preferred labellings where \mathbf{a} is **undec** that are not already included in \mathcal{L}_c .

If there is no such labelling in \mathcal{L} , the function returns \perp .

Note that the definition of $\text{FindCLU}_{\mathbf{a}}^{\boxplus}$ is well founded, i.e. the above constraints are compatible due to the maximality w.r.t. \sqsubseteq of preferred labellings (see Proposition 3).

At each iteration of the loop (lines 5–13), a complete labelling $\mathcal{L}ab$ such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$ is selected (line 5). Then the inner loop (lines 6–8) starts from the labelling $\mathcal{L}ab$ and replaces it with a (not necessarily strictly) more committed labelling in \mathcal{L}_{mu} . In particular, similarly to $\text{FindCL}^{\boxminus}$, $\text{FindCL}_{\mathbf{a}}^{\boxminus}$ is a non-deterministic function that, given an argument \mathbf{a} , an argumentation framework Γ , a complete labelling $\mathcal{L}ab \in \mathfrak{L}_{CO}(\Gamma)$ such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$ and a label $s \in \{\text{in}, \text{undec}\}$, returns a complete labelling $\mathcal{L}ab' \in \mathfrak{L}_{CO}(\Gamma)$ such that $\mathcal{L}ab'(\mathbf{a}) = s$ and $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$ if such a labelling exists, \perp otherwise. At each iteration of the inner loop, $\text{FindCL}_{\mathbf{a}}^{\boxminus}(\Gamma, \mathcal{L}ab, \text{undec})$ produces a strictly more committed labelling such that \mathbf{a} is **undec**, until the resulting sequence can not

more be extended and the loop is exited. This way, at line 9 the variable $\mathcal{L}ab$ includes a labelling of \mathcal{L}_{mu} .

It should be noted that such a labelling $\mathcal{L}ab$ obtained is not necessarily a preferred labelling, since there might exist a strictly more committed labelling such that \mathbf{a} is **in** (recall that from line 4 we are guaranteed that there are no preferred labellings where \mathbf{a} is labelled **out**). The existence of such a labelling is checked at line 9. If $\text{FindCL}_{\mathbf{a}}^{\overline{\perp}}(\Gamma, \mathcal{L}ab, \mathbf{in}) = \perp$, then there is no strictly more committed labelling $\mathcal{L}ab''$ such that $\mathcal{L}ab'' = \mathbf{in}$, i.e. $\mathcal{L}ab$ is a preferred labelling such that $\mathcal{L}ab(\mathbf{a}) = \mathbf{undec}$, thus \mathbf{a} is not skeptically accepted and the algorithm returns \perp . Otherwise, $\mathcal{L}ab$ is a labelling of \mathcal{L}_{mu} which is not preferred (there is a strictly more committed complete labelling) and such that all more committed preferred labellings assign to \mathbf{a} the label **in** (by the maximality of $\mathcal{L}ab$ there is no more committed preferred labelling assigning to \mathbf{a} the label **undec**, and as mentioned above there is no preferred labelling where \mathbf{a} is **out**). As a consequence, $\mathcal{L}ab$ is recorded by adding it to \mathcal{L}_c at line 12, so that at the next iteration of the loop all complete labellings less committed w.r.t. it (including $\mathcal{L}ab$ itself) are excluded from the search, since according to the considerations above this does not exclude the possible preferred labellings where \mathbf{a} is **undec**.

If at line 5 the function $\text{FindCLU}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}_c)$ does not return any labelling, then there are no preferred labellings where \mathbf{a} is **undec**, and the algorithm returns \top (line 14).

Algorithm 9 Skeptical acceptance under Preferred Semantics 2

DS-PR
Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{a} \in \mathcal{A}$
Output: \top if \mathbf{a} is skeptically accepted under PR, \perp otherwise

- 1: **if** $\text{FindCL}_{\mathbf{a}}(\Gamma, \mathbf{in}) = \perp \vee \text{FindCL}_{\mathbf{a}}(\Gamma, \mathbf{out}) \neq \perp$ **then**
- 2: **return** \perp
- 3: **end if**
- 4: $\mathcal{L}_c := \emptyset$
- 5: **while** $\mathcal{L}ab := \text{FindCLU}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}_c)$ **do**
- 6: **while** $\mathcal{L}ab' := \text{FindCL}_{\mathbf{a}}^{\overline{\perp}}(\Gamma, \mathcal{L}ab, \mathbf{undec})$ **do**
- 7: $\mathcal{L}ab := \mathcal{L}ab'$
- 8: **end while**
- 9: **if** $\text{FindCL}_{\mathbf{a}}^{\overline{\perp}}(\Gamma, \mathcal{L}ab, \mathbf{in}) = \perp$ **then**
- 10: **return** \perp
- 11: **end if**
- 12: $\mathcal{L}_c := \mathcal{L}_c \cup \{\mathcal{L}ab\}$
- 13: **end while**
- 14: **return** \top

The correctness of Algorithms 5-9 is proved by the following theorem.

Theorem 2. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{a} \in \mathcal{A}$:

- $\{\mathbf{in}(\mathcal{L}ab) \mid \mathcal{L}ab \in \mathbf{EL-PR}(\Gamma)\} = \mathcal{E}_{\text{PR}}(\Gamma)$ (Algorithm 5);

- $\{\text{in}(\mathcal{L}ab) \mid \mathcal{L}ab \in \mathbf{EL-PR-withST}(\Gamma)\} = \mathcal{E}_{\text{PR}}(\Gamma)$ (Algorithm 6);
- $\mathbf{DC-PR}(\Gamma, \mathbf{a}) \equiv (\exists S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 7);
- $\mathbf{DS-PR-straightforward}(\Gamma, \mathbf{a}) \equiv (\forall S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 8);
- $\mathbf{DS-PR}(\Gamma, \mathbf{a}) \equiv (\forall S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 9).

4. Encoding Complete Labellings and their Constraints

Each of the non-deterministic functions exploited by the algorithms described in Section 3 returns a single complete labelling subject to specific constraints (e.g. being related by \sqsubseteq w.r.t. a given complete labelling). Taking into account that a stable labelling is a complete labelling, this also holds for the functions returning stable labellings that appear in the algorithms of Section 3.1.

Since a complete labelling can be expressed by means of local conditions relating the label of each node with those of its attackers (see Definition 7), using a SAT solver is a natural way to implement these functions. According to these considerations, two problems have to be faced in this respect:

- How to encode the constraints corresponding to complete labellings as a propositional formula?
- How to encode as propositional clauses, to be conjoined with the formula above, the specific constraints on the complete labelling returned by each function?

While these might seem clear-cut tasks, several syntactically different encodings can be devised which, although logically equivalent, can significantly affect the performance of the overall process of searching a satisfying assignment. For instance, adding some “redundant” clauses to a formula may speed up the search process, thanks to the additional information provided by the additional constraints. On the other hand, increasing syntactic complexity might lead to worse performances, thus a careful selection of the encoding is needed. These aspects will be the main focus of Section 6.

The aim of the following two subsections instead is to explore possible encodings for complete labellings and the relevant constraints.

4.1. Logical Encodings of Complete Labellings

Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ we are interested in identifying a boolean formula such that its satisfying assignments have a one-to-one correspondence with the complete labellings of Γ .

As mentioned above, several ways to encode complete labellings can be envisaged. In order to explore alternative encodings, let us consider again the requirements of Definition 7. They can be expressed as a conjunction of 6 terms, i.e. $C_{\text{in}}^{\rightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow}$, where

- $C_{\text{in}}^{\rightarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{in} \Rightarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) = \text{out});$




Figure 1: Partial Hasse Diagram of 64 possible encodings of complete labellings.

- $C_{\text{in}}^{\leftarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{in} \Leftarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) = \text{out})$;
- $C_{\text{out}}^{\rightarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{out} \Rightarrow \exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) = \text{in})$;
- $C_{\text{out}}^{\leftarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{out} \Leftarrow \exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) = \text{in})$;
- $C_{\text{undec}}^{\rightarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{undec} \Rightarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) \neq \text{in} \wedge \exists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{c}) = \text{undec})$;
- $C_{\text{undec}}^{\leftarrow} \equiv (\mathcal{L}ab(\mathbf{a}) = \text{undec} \Leftarrow \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) \neq \text{in} \wedge \exists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{c}) = \text{undec})$.

Let us also define $C_{\text{in}}^{\leftrightarrow} \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{in}}^{\leftarrow}$, $C_{\text{out}}^{\leftrightarrow} \equiv C_{\text{out}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow}$, $C_{\text{undec}}^{\leftrightarrow} \equiv C_{\text{undec}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow}$.

We aim at exploring all the encodings corresponding to the 64 possible subsets of the 6 terms above, characterised by a cardinality (i.e. the number of terms) between 0 and 6 and partially ordered according to the \subseteq -relation. We will denote these encodings as $C_{\mathbf{b}}^i$, with $i, o, u \in \{\rightarrow, \leftarrow, \leftrightarrow, \sim\}$, i.e. $C_{\mathbf{b}}^i \equiv C_{\text{in}}^i \wedge C_{\text{out}}^o \wedge C_{\text{undec}}^u$, where $C_{\text{in}}^{\sim} = C_{\text{out}}^{\sim} = C_{\text{undec}}^{\sim} = \top$ (in other words, the symbol

\sim denotes that the corresponding term is not present in the encoding). For instance, $C^{\leftrightarrow} \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow}$.

Using basic combinatorial calculus it can be seen that the 64 subsets can be partitioned as follows: one encoding with cardinality 0 (i.e. the empty encoding C^{\approx}), 6 encodings with cardinality 1, 15 encodings with cardinality 2, 20 encodings with cardinality 3, 15 encodings with cardinality 4, 6 encodings with cardinality 5 and one encoding with cardinality 6 (i.e. C^{Ξ} corresponding to Definition 7). Figure 1 depicts these encodings by arranging them according to their cardinality (indicated at the left), and partially shows the inclusion relationships (focusing on those that will be relevant in the following).

The encodings can be partitioned into three classes:

1. *weak* encodings, i.e. such that there is an argumentation framework and a labelling satisfying all their terms which however is not a complete labelling;
2. *correct and non-redundant* encodings, i.e. able to correctly identify complete labellings and such that any strict subset of their terms is weak;
3. *correct and redundant* encodings, i.e. able to correctly identify complete labellings and such that there is a strict subset of their terms which is correct.

The corresponding characterisation of the 64 encodings is illustrated in Figure 1: the dotted line separates the weak encodings (under the line) from the correct ones (over the line), and among the correct ones those that are non-redundant are depicted as squares, redundant encodings appear as circles⁶. This characterisation is shown to be correct by Theorem 3, which requires two preliminary propositions.

Proposition 5. Let $\langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework. A total function $\mathcal{L}ab : \mathcal{A} \mapsto \{\text{in}, \text{out}, \text{undec}\}$ is a complete labelling iff it satisfies any of the following conjunctive constraints for any $\mathbf{a} \in \mathcal{A}$:

$$\begin{aligned}
C^{\Xi} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftrightarrow} \\
C^{\approx} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow} \\
C^{\approx} &= C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow} \\
C^{\Xi} &= C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow} \\
C^{\Xi} &= C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}
\end{aligned}$$

⁶ $C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftrightarrow}$ and $C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ correspond to the alternative definitions of complete labellings in [31], where a proof of their equivalence is provided.




Figure 2: Hasse-like diagram for the 18 correct encodings of complete labellings.

Proposition 6. The following 6 encodings are weak:

$$\begin{aligned}
C_{\exists\exists} &= C_{\text{undec}}^{\leftrightarrow} \wedge C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow} \\
C_{\exists\exists} &= C_{\text{undec}}^{\leftrightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \\
C_{\exists\exists} &= C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{in}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow} \\
C_{\exists\exists} &= C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow} \\
C_{\exists\exists} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow} \\
C_{\exists\exists} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow}
\end{aligned}$$

In particular, the proof of Proposition 6 identifies for each encoding an argumentation framework admitting a unique complete labelling which drastically differs from the one satisfying the weak encoding, i.e. there are arguments labelled **in** that should be labelled **undec** or there are arguments labelled **out** or **undec** that should be labelled **in** (see the proof in Appendix A for details).

Theorem 3. All the encodings of cardinality 0, 1, and 2 are weak. Among the encodings of cardinality 3, $C_{\exists\exists} = C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ and $C_{\exists\exists} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$ are correct and non-redundant, the other 18 encodings are weak. Among the encodings of cardinality 4, $C_{\exists\exists} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftrightarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$ and $C_{\exists\exists} = C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$ are correct and non-redundant, the 6 encodings $C_{\exists\exists} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow}$, $C_{\exists\exists} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow}$ are correct and redundant, the other 6 encodings are weak. All the encodings of cardinality 5 and 6 are correct and redundant.

Let $\mathfrak{C} = \{C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}, C_{\exists\exists}\}$ be the set of all the correct encodings of complete labellings as identified by Theorem 3. The Hasse diagram for these 18 correct encodings is

reported in Figure 2. Correct and non-redundant encodings are the minimal ones under set inclusion and correspond to the five encodings identified in Proposition 5.

4.2. Logical Encodings of Complete Labelling Constraints

We then need to encode the specific constraints that characterise the complete labellings returned as output of the various functions considered in the algorithms presented in Section 3. In this subsection we identify different ways such constraints can be encoded, thus completing the analysis at the base of the implementation of the algorithms detailed in Section 5.

Let us first consider the algorithms described in Section 3.1 relevant to the stable semantics, which exploit the functions FindST^\neq and FindST_a (where \mathbf{a} is an argument). Both of them return a stable labelling $\mathcal{L}ab$, which is defined as a complete labelling without arguments labelled **undec** (see Section 2) and can thus be obtained by constraining the encoding of a complete labelling with the condition $C_{st} \equiv \forall \mathbf{b} \in \mathcal{A}, \mathcal{L}ab(\mathbf{b}) \neq \text{undec}$.

Focusing on the function FindST^\neq , used in Algorithm 2, $\text{FindST}^\neq(\Gamma, \mathcal{L}_s)$ returns a stable labelling $\mathcal{L}ab$ of Γ which is not included in \mathcal{L}_s , where \mathcal{L}_s is a set of stable labellings. For each $\mathcal{L}ab' \in \mathcal{L}_s$, the condition $\mathcal{L}ab \neq \mathcal{L}ab'$ must thus be encoded. Since there are no arguments labelled **undec** by stable labellings, a natural way to do it amounts to requiring that an argument labelled **in** by $\mathcal{L}ab'$ is labelled **out** by $\mathcal{L}ab$, or vice versa. While this can be encoded as a disjunction over the whole set of arguments, Proposition 7, based on Lemma 1, shows that such disjunction can be restricted to the arguments labelled **in** by $\mathcal{L}ab'$ (or $\mathcal{L}ab$), or equivalently on arguments labelled **out** by $\mathcal{L}ab'$ (or $\mathcal{L}ab$). Lemma 1 shows that an argument \mathbf{a} is labelled **out** by a complete labelling $\mathcal{L}ab_1$ and **in** or **undec** by another complete labelling $\mathcal{L}ab_2$ if and only if another argument \mathbf{b} is labelled **in** by $\mathcal{L}ab_1$ and **out** or **undec** by $\mathcal{L}ab_2$.

Lemma 1. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab_1, \mathcal{L}ab_2 \in \mathfrak{L}(\Gamma)$ complete labellings. $\exists \mathbf{a} \in \mathcal{A}$ such that $\mathcal{L}ab_1(\mathbf{a}) = \text{out}$ and $\mathcal{L}ab_2(\mathbf{a}) \neq \text{out}$ iff $\exists \mathbf{b} \in \mathcal{A}$ such that $\mathcal{L}ab_1(\mathbf{b}) = \text{in}$ and $\mathcal{L}ab_2(\mathbf{b}) \neq \text{in}$.

Proposition 7. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab, \mathcal{L}ab' \in \mathfrak{L}(\Gamma)$ stable labellings. All of the following statements are equivalent to $\mathcal{L}ab \neq \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$
3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

Summing up, the encoding corresponding to $\mathcal{L}ab = \text{FindST}^\neq(\Gamma, \mathcal{L}_s)$ can be expressed as $C \wedge C_{st} \wedge C_\neq$, where C is any encoding belonging to \mathfrak{C} and C_\neq is the conjunction of any of the three conditions of Proposition 7 applied to

all labellings of \mathcal{L}_s . It can be noted that the third condition of Proposition 7 corresponds to the conjunction of the first two conditions: As with redundant encodings for complete labellings, we also consider redundant constraints that may improve the efficiency of SAT-based algorithms, introduced in Section 3.

As to the function $\text{FindST}_{\mathbf{a}}$, exploited in Algorithms 3 and 4, it returns a stable labelling $\mathcal{L}ab$ of Γ such that $\mathcal{L}ab(\mathbf{a})$ is a specific label $s \in \{\text{in}, \text{out}\}$. Obviously, the corresponding encoding is $C \wedge C_{st} \wedge \mathcal{L}ab(\mathbf{a}) = s$.

Let us now turn to preferred semantics and first consider the function FindCL^{\boxplus} used in Algorithms 5 and 6 for the enumeration of preferred labellings and in Algorithm 8 for skeptical acceptance under preferred semantics. In particular, given an argumentation framework Γ and a set of preferred labellings $\mathcal{L}_p \subseteq \mathfrak{L}_{\text{PR}}(\Gamma)$, $\text{FindCL}^{\boxplus}(\Gamma, \mathcal{L}_p)$ returns a complete labelling $\mathcal{L}ab$ such that $\forall \mathcal{L}ab' \in \mathcal{L}_p, \mathcal{L}ab \boxplus \mathcal{L}ab'$. Proposition 8, in the following, shows that this constraint is equivalent to require an argument labelled *out* by $\mathcal{L}ab'$ to be labelled *in* by $\mathcal{L}ab$, or vice versa.

Proposition 8. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab' \in \mathfrak{L}_{\text{PR}}(\Gamma)$ a preferred labelling. Then, $\forall \mathcal{L}ab \in \mathfrak{L}_{\text{CO}}(\Gamma)$, all of the following statements are equivalent to $\mathcal{L}ab \boxplus \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$
3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

According to Proposition 8, the constraint $\forall \mathcal{L}ab' \in \mathcal{L}_p, \mathcal{L}ab \boxplus \mathcal{L}ab'$ can be expressed in the same way as in the case of the function FindST^{\neq} , i.e. the encoding corresponding to $\mathcal{L}ab = \text{FindCL}^{\boxplus}(\Gamma, \mathcal{L}_p)$ is $C \wedge C_{\neq}$, where $C \in \mathfrak{C}$ and C_{\neq} is the conjunction of any of the three conditions of Proposition 8 applied to all labellings of \mathcal{L}_p . Note that $C \wedge C_{\neq}$ captures all preferred labellings that do not belong to \mathcal{L}_p , complying with the definition of FindCL^{\boxplus} (see Section 3.2).

As to the function $\text{FindCL}^{\supseteq}$, used again in Algorithms 5, 6 and 8, given an argumentation framework Γ and a complete labelling $\mathcal{L}ab'$, $\text{FindCL}^{\supseteq}(\Gamma, \mathcal{L}ab')$ returns a complete labelling $\mathcal{L}ab$ such that $\mathcal{L}ab' \supseteq \mathcal{L}ab$. This constraint amounts to require that $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab)$, $\text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab)$, and that there is an argument $\mathbf{a} \in \text{undec}(\mathcal{L}ab')$ such that $\mathcal{L}ab(\mathbf{a}) \neq \text{undec}$. Proposition 9, based on Lemma 1, shows that this constraint can be strengthened by replacing $\mathcal{L}ab(\mathbf{a}) \neq \text{undec}$ with either $\mathcal{L}ab(\mathbf{a}) = \text{in}$ or $\mathcal{L}ab(\mathbf{a}) = \text{out}$.

Proposition 9. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab, \mathcal{L}ab' \in \mathfrak{L}_{\text{CO}}(\Gamma)$. All of the following statements are equivalent to $\mathcal{L}ab' \supseteq \mathcal{L}ab$:

1. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{b} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) = \text{out}$
2. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{c} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{c}) = \text{in}$




Figure 3: A counterexample to the direct use of Proposition 8.

3. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{b} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) = \text{out} \wedge \exists \mathbf{c} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{c}) = \text{in}$

According to the proposition above, the encoding corresponding to $\mathcal{L}ab = \text{FindCL}_{\mathbf{a}}^{\exists}(\Gamma, \mathcal{L}ab')$ is $C \wedge C_{\exists}$, where $C \in \mathfrak{C}$ and C_{\exists} is any of the three conditions of Proposition 9.

We now consider the function $\text{FindCL}_{\mathbf{a}}$, which is exploited by Algorithms 7 and 9 to identify a complete labelling where \mathbf{a} is **in** and by Algorithm 9 to identify a complete labelling where \mathbf{a} is **out**. Accordingly, the encoding corresponding to the labelling $\mathcal{L}ab = \text{FindCL}_{\mathbf{a}}(\Gamma, s)$, where $s \in \{\text{in}, \text{out}\}$, is simply $C \wedge \mathcal{L}ab(\mathbf{a}) = s$, where $C \in \mathfrak{C}$.

Let us finally turn to the functions appearing in Algorithm 9, namely $\text{FindCLU}_{\mathbf{a}}^{\exists}$ and $\text{FindCL}_{\mathbf{a}}^{\exists}$, which roughly correspond to the functions FindCL^{\exists} and FindCL^{\exists} , respectively, but with an additional constraint on the label assigned to \mathbf{a} by the output labelling $\mathcal{L}ab$.

In particular, given an argumentation framework Γ and a set of complete labellings $\mathcal{L}_c \subseteq \mathcal{L}_{mu}$, $\text{FindCLU}_{\mathbf{a}}^{\exists}(\Gamma, \mathcal{L}_c)$ returns a complete labelling $\mathcal{L}ab$ such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$ and $\forall \mathcal{L}ab' \in \mathcal{L}_c, \mathcal{L}ab \not\sqsubseteq \mathcal{L}ab'$. It would be tempting to directly apply Proposition 8 to this case, however the example depicted in Figure 3 shows that the corresponding result does not hold in general under the hypothesis that $\mathcal{L}_c \subseteq \mathcal{L}_{mu}$. Indeed, there are only three complete labellings such that \mathbf{a} is **undec**, namely the grounded labelling where all arguments are labelled **undec**, $\mathcal{L}ab = \{(e_1, \text{undec}), (d_1, \text{undec}), (c_1, \text{undec}), (e_2, \text{out}), (d_2, \text{in}), (c_2, \text{out}), (\mathbf{b}, \text{undec}), (\mathbf{a}, \text{undec})\}$ and $\mathcal{L}ab' = \{(e_1, \text{out}), (d_1, \text{in}), (c_1, \text{out}), (e_2, \text{undec}), (d_2, \text{undec}), (c_2, \text{undec}), (\mathbf{b}, \text{undec}), (\mathbf{a}, \text{undec})\}$. The grounded labelling is not maximal w.r.t. \sqsubseteq , thus $\mathcal{L}_{mu} = \{\mathcal{L}ab, \mathcal{L}ab'\}$. It can be easily checked that, despite $\mathcal{L}ab' \in \mathcal{L}_{mu}$ and $\mathcal{L}ab \not\sqsubseteq \mathcal{L}ab'$, none of the conditions of Proposition 8 holds, i.e. there is no argument which is labelled **out** by $\mathcal{L}ab'$ and **in** by $\mathcal{L}ab$, or vice versa. On the other hand, the following proposition shows that these conditions are satisfied under the hypothesis that $\mathcal{L}ab$ is a preferred labelling.

Proposition 10. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework, \mathbf{a} an argument in \mathcal{A} and $\mathcal{L}ab' \in \mathfrak{L}(\Gamma)$ a maximal complete labelling such that $\mathcal{L}ab'(\mathbf{a}) = \text{undec}$, i.e. $\nexists \mathcal{L}ab^* \in \mathfrak{L}(\Gamma)$ such that $\mathcal{L}ab^*(\mathbf{a}) = \text{undec}$ and $\mathcal{L}ab' \sqsubset \mathcal{L}ab^*$. Then, $\forall \mathcal{L}ab \in \mathfrak{L}_{PR}(\Gamma)$ such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$, all of the following statements are

equivalent to $\mathcal{L}ab \neq \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$
3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

According to the proposition above, the encoding corresponding to $\mathcal{L}ab = \text{FindCLU}_{\mathbf{a}}^{\neq}(\Gamma, \mathcal{L}_c)$ is $C \wedge C_{\neq} \wedge \mathcal{L}ab(\mathbf{a}) = \text{undec}$, where $C \in \mathfrak{C}$ and C_{\neq} is the conjunction of any of the three conditions of Proposition 10 applied to all labellings of \mathcal{L}_c . Indeed, the set of labellings satisfying these constraints complies with the set \mathcal{L} in the definition of $\text{FindCLU}_{\mathbf{a}}^{\neq}$ (see Section 3.2): it is a subset of the complete labellings (due to C) where \mathbf{a} is labelled **undec** (due to $\mathcal{L}ab(\mathbf{a}) = \text{undec}$) which does not include those labellings that are less or equally committed w.r.t. a labelling of \mathcal{L}_c (due to the conditions of C_{\neq}), while including all preferred labellings where \mathbf{a} is **undec** that are outside \mathcal{L}_c (this is guaranteed by Proposition 10).

We have finally to consider the function $\text{FindCL}_{\mathbf{a}}^{\supseteq}$, which given an argument \mathbf{a} , an argumentation framework Γ and a complete labelling $\mathcal{L}ab' \in \mathfrak{L}_{\text{CO}}(\Gamma)$ returns a complete labelling $\mathcal{L}ab$ such that $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab$ and $\mathcal{L}ab(\mathbf{a}) = s$, with $s \in \{\text{in}, \text{undec}\}$. Taking into account Proposition 9, it is easy to see that the encoding corresponding to $\mathcal{L}ab$ is $C \wedge C_{\supseteq} \wedge \mathcal{L}ab(\mathbf{a}) = s$, where again $C \in \mathfrak{C}$ and C_{\supseteq} is any of the three conditions of Proposition 9.

5. Algorithms Implementation

In this section we present the detailed pseudo-codes⁷ implementing the abstract algorithms presented in Section 3.

In order to allow us to explore all the different encodings of complete labellings and the relevant constraints, we will exploit a number of parameters, denoted as π_{*}^{*} , and select algorithmic variants with conditional statements based on them: Table B.5 summarises them.

Before presenting the detailed algorithms, we provide some details on the encoding of complete labellings as required by existing SAT solvers.

5.1. SAT Encodings of Complete Labellings

To exploit existing SAT solvers, we need first to express the correct constraints of \mathfrak{C} identified in Theorem 3 as a formula in conjunctive normal form. To this

⁷The actual C++ code is available at <https://sourceforge.net/projects/argsemsat/> (on 11 May 2018)

purpose, for each argument $\mathbf{a} \in \mathcal{A}$ we define three boolean variables,⁸ $I_{\mathbf{a}}$, $O_{\mathbf{a}}$, and $U_{\mathbf{a}}$, with the intended meaning that $I_{\mathbf{a}}$ is true when argument \mathbf{a} is labelled **in**, false otherwise, and analogously $O_{\mathbf{a}}$ and $U_{\mathbf{a}}$ correspond to labels **out** and **undec**. Formally, given $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ we define the corresponding set of variables as $\mathcal{V}(\Gamma) \triangleq \bigcup_{\mathbf{a} \in \mathcal{A}} \{I_{\mathbf{a}}, O_{\mathbf{a}}, U_{\mathbf{a}}\}$. The constraints of Definition 7 can then be expressed in terms of the variables $\mathcal{V}(\Gamma)$. The detail of the resulting CNF formula is given in Proposition 11: proof is omitted as it follows from straightforward manipulations.

Proposition 11. Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$:

$$\begin{aligned}
C_{\text{in}}^{\leftarrow} &\equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(I_{\mathbf{a}} \vee \left(\bigvee_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} (\neg O_{\mathbf{b}}) \right) \right) \\
C_{\text{in}}^{\rightarrow} &\equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} \neg I_{\mathbf{a}} \vee O_{\mathbf{b}} \right) \\
C_{\text{out}}^{\leftarrow} &\equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} \neg I_{\mathbf{b}} \vee O_{\mathbf{a}} \right) \\
C_{\text{out}}^{\rightarrow} &\equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\neg O_{\mathbf{a}} \vee \left(\bigvee_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} I_{\mathbf{b}} \right) \right) \\
C_{\text{undec}}^{\leftarrow} &\equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} \left(U_{\mathbf{a}} \vee \neg U_{\mathbf{b}} \vee \left(\bigvee_{\{\mathbf{c} \mid \mathbf{c} \rightarrow \mathbf{a}\}} I_{\mathbf{c}} \right) \right) \right)
\end{aligned}$$

⁸It is worth noticing that variations of the proposed algorithm can be used with fewer variables. For instance, there are algorithms for preferred semantics using only one variable per argument, cf. [32, 33]. In ICCMA-15 an earlier version of ArgSemSAT still using three boolean variables per argument resulted faster than such algorithms when computing the skeptical acceptance of arguments, despite relying on an external SAT solver as a new process. For what concerns stable semantics, whose definition does not make use of the **undec** label, in Appendix C an interested reader can find the results of a pilot study aimed at quantifying an alleged improvement in using one variable instead of three when enumerating stable extensions. Our preliminary results show cases where there is not a statistical significant change; and when there is one, it might be that the version using three variables has a lower median runtime. It seems to us that there is evidence that performance does not necessary improve when considering fewer propositional variables when transforming the constraints imposed by semantics into CNFs.

$$C_{\text{undec}}^{\rightarrow} \equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} (-U_{\mathbf{a}} \vee \neg I_{\mathbf{b}}) \right) \wedge \left(-U_{\mathbf{a}} \vee \left(\bigvee_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} U_{\mathbf{b}} \right) \right) \right)$$

Furthermore, since for each argument \mathbf{a} the variables $I_{\mathbf{a}}$, $O_{\mathbf{a}}$, and $U_{\mathbf{a}}$ are independent, an additional condition is needed to enforce an assignment corresponding to a total function, i.e. to guarantee that exactly one of the variables is assigned the value True. Moreover, as a small engineering improvement we explicitly constrain unattacked arguments to be labelled **in**, and for technical reasons we restrict to “non-empty” labellings (in the sense that at least one of the arguments is labelled **in**). In order to express the resulting CNF identifying non-empty complete labellings of a given argumentation framework, let us define a function `cnf` that receives as input an element of \mathfrak{C} and an argumentation framework Γ , and returns the equivalent CNF representation as specified by Proposition 11, in conjunction with the following formulæ:

$$\bigwedge_{\mathbf{a} \in \mathcal{A}} \left((I_{\mathbf{a}} \vee O_{\mathbf{a}} \vee U_{\mathbf{a}}) \wedge (\neg I_{\mathbf{a}} \vee \neg O_{\mathbf{a}}) \wedge (\neg I_{\mathbf{a}} \vee \neg U_{\mathbf{a}}) \wedge (\neg O_{\mathbf{a}} \vee \neg U_{\mathbf{a}}) \right) \quad (1)$$

$$\bigwedge_{\{\mathbf{a} \mid \mathbf{a}^- = \emptyset\}} (I_{\mathbf{a}} \wedge \neg O_{\mathbf{a}} \wedge \neg U_{\mathbf{a}}) \quad (2)$$

$$\bigvee_{\mathbf{a} \in \mathcal{A}} I_{\mathbf{a}} \quad (3)$$

In particular, (1) enforces a total function, (2) constrains an unattacked argument to be labelled **in**, and (3) enforces that at least one argument is labelled **in**, thus forcing the search for a non-empty labelling.

5.2. Algorithms for Stable Semantics

In this section we detail the implementation of the algorithms presented in Section 3.1 on the basis of the encodings identified in Section 4.

To this purpose, we first need to introduce some auxiliary functions referring to a generic argumentation framework $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$.

The call to a SAT solver is expressed by means of the function `SatS`, which receives as input a CNF formula and returns a labelling over \mathcal{A} satisfying the constraints expressed by the formula, or \emptyset if the formula is not satisfiable.

Similarly, `ALLSatS` represents a call to an AllSAT solver. It receives as input a CNF formula as well as two parameters $\pi_{\text{All};0}^S$ and $\pi_{\text{All};1}^S$, and returns two elements, where the first is the (possibly empty) set of labellings over \mathcal{A} satisfying the constraints expressed by the formula, and the second element is the set of blocking formulæ generated in the process. This second element is

sometimes needed to exclude from the subsequent search the set of labellings returned by **ALLSatS**. Specifically, the implementation exploits the **AllSAT** solver developed by [28] with the only modification to enable it to return the set of blocking clauses along with the set of the found models. The input parameters $\pi_{\text{All};\text{O}}^{\text{S}}$ and $\pi_{\text{All};\text{I}}^{\text{S}}$ allow one to choose the form of the blocking clauses according to Proposition 7. In particular, $\pi_{\text{All};\text{O}}^{\text{S}}$ generates for each found labelling a blocking clause based on arguments labelled **in**, requiring one of them to be labelled **out**, while $\pi_{\text{All};\text{I}}^{\text{S}}$ generates a blocking clause that considers only arguments labelled **out**. As shown below, they are in essence analogous to the parameters $\pi_{\text{O}}^{\text{S}}$ and $\pi_{\text{I}}^{\text{S}}$ used at lines 8–13 of Algorithm 10.

Finally, **I-ARGS** (resp. **O-ARGS**, **U-ARGS**) receives as input a labelling and returns the set of arguments labelled **in** (resp. **out**, **undec**).

Algorithm 10 details the implementation of Algorithm 2 for the enumeration of stable labellings. Besides $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, the algorithm receives as input the specific encoding $C^* \in \mathfrak{C}$ of complete labellings to be exploited (see Section 4.1) as well as five parameters, i.e. $\pi_{\text{All}}^{\text{S}}$, $\pi_{\text{All};\text{O}}^{\text{S}}$, $\pi_{\text{All};\text{I}}^{\text{S}}$, $\pi_{\text{O}}^{\text{S}}$, and $\pi_{\text{I}}^{\text{S}}$. The first one enables the algorithm to directly exploit an **AllSAT** solver to identify all the stable labellings, using in this case as input parameters $\pi_{\text{All};\text{O}}^{\text{S}}$ and $\pi_{\text{All};\text{I}}^{\text{S}}$, while $\pi_{\text{O}}^{\text{S}}$ and $\pi_{\text{I}}^{\text{S}}$, in case the **AllSAT** solver is not exploited, allow one to select the way to encode the constraint relevant to the function FindST^{\neq} .

In particular, if $\pi_{\text{All}}^{\text{S}} = \top$ then the **AllSAT** solver is called at line 2 on a CNF which encodes the complete labellings without undecided arguments, i.e. satisfying the additional requirement $\bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}}$. In the case $\pi_{\text{All}}^{\text{S}} = \perp$, the function $\text{FindST}^{\neq}(\cdot, \cdot)$, used at line 2 of Algorithm 2, is implemented within lines 6–15. Line 6 identifies a stable labelling (if this exists) that has not been identified in an earlier execution of the cycle. Proposition 7 shows that there are three possibilities to ensure that: an argument labelled **in** should then become labelled **out**, implemented at line 9 if $\pi_{\text{O}}^{\text{S}} = \top$; an argument labelled **out** should then become labelled **in**, implemented at line 12 if $\pi_{\text{I}}^{\text{S}} = \top$; or both of the previous ones, in case $\pi_{\text{O}}^{\text{S}} = \top \wedge \pi_{\text{I}}^{\text{S}} = \top$. As discussed in Section 4.2, we leave room for all the three possibilities, including the redundant constraint, since logically equivalent encodings may yield different behaviours from the efficiency point of view.

Note that the algorithm returns as output also the blocking formulæ used to exclude the generated solutions from the search. These are useful in case the algorithm is exploited as a first step of the enumeration of preferred labellings (cf. Algorithm 13).

Moving to credulous and skeptical acceptance, Algorithm 11 uses a call to a SAT solver (line 1) to implement in a rather straightforward fashion the function $\text{FindST}_{\mathbf{x}}(\cdot, \cdot)$ called at line 1 of Algorithm 3. Similarly, at line 1 of Algorithm 12 the function $\text{FindST}_{\mathbf{x}}(\cdot, \cdot)$ used at line 1 of Algorithm 4 is implemented with a simple call to a SAT Solver.

Algorithm 10 Enumeration of Stable Labellings using SAT solvers

SAT-EL-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle, C^* \in \mathcal{C}, \langle \pi_{\text{All}}^S, \pi_{\text{All};\text{O}}^S, \pi_{\text{All};\text{I}}^S, \pi_{\text{O}}^S, \pi_{\text{I}}^S \rangle \in \{\top, \perp\}^5$ **Output:** $\langle L_s \subseteq \mathcal{L}_{\text{ST}}(\Gamma), \text{blocking} \rangle$

```
1: if  $\pi_{\text{All}}^S$  then
2:   return ALLSatS  $\left( \text{cnf}(C^*, \Gamma) \wedge \bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}}, \pi_{\text{All};\text{I}}^S, \pi_{\text{All};\text{O}}^S \right)$ 
3: end if
4:  $L_s := \emptyset, \text{blocking} := \top$ 
5: do
6:    $stb := \text{SatS} \left( \text{cnf}(C^*, \Gamma) \wedge \bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}} \wedge \text{blocking} \right)$ 
7:   if  $stb \neq \emptyset$  then
8:     if  $\pi_{\text{O}}^S$  then
9:        $\text{blocking} := \text{blocking} \wedge \bigvee_{\mathbf{a} \in \text{I-ARGS}(stb)} O_{\mathbf{a}}$ 
10:    end if
11:    if  $\pi_{\text{I}}^S$  then
12:       $\text{blocking} := \text{blocking} \wedge \bigvee_{\mathbf{a} \in \text{O-ARGS}(stb)} I_{\mathbf{a}}$ 
13:    end if
14:     $L_s := L_s \cup \{stb\}$ 
15:  end if
16: while  $stb \neq \emptyset$ 
17: return  $\langle L_s, \text{blocking} \rangle$ 
```

Algorithm 11 Credulous acceptance under Stable Semantics using SAT solvers

SAT-DC-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle, \mathbf{x} \in \mathcal{A}, C^* \in \mathcal{C}$ **Output:** $\{\top, \perp\}$

```
1: return  $\left( \text{SatS} \left( \text{cnf}(C^*, \Gamma) \wedge \bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}} \wedge I_{\mathbf{x}} \right) \neq \emptyset \right)$ 
```

Algorithm 12 Skeptical acceptance under Stable Semantics using SAT solvers

SAT-DS-ST**Input:** $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle, \mathbf{x} \in \mathcal{A}, C^* \in \mathcal{C}$ **Output:** $\{\top, \perp\}$

```
1: if  $\text{SatS} \left( \text{cnf}(C^*, \Gamma) \wedge \bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}} \wedge O_{\mathbf{x}} \right) \neq \emptyset$  then
2:   return  $\perp$ 
3: end if
4: return  $\top$ 
```

Algorithm 13 Enumeration of Preferred Labellings using SAT solvers

SAT-EL-PR

Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $C^* \in \mathfrak{C}$, $\langle \pi_S^P, \pi_{All}^S, \pi_{All;O}^S, \pi_{All;l}^S, \pi_O^S, \pi_l^S, \pi_{iO}^P, \pi_{il}^P, \pi_{eO}^P, \pi_{el}^P \rangle \in \{\top, \perp\}^{10}$

Output: $L_s \subseteq \mathfrak{L}_{ST}(\Gamma)$

```

1:  $L_p := \emptyset$ ,  $blocking := \top$ 
2: if  $\pi_S^P$  then
3:    $\langle L_p, blocking \rangle := \mathbf{SAT-EL-ST}(\Gamma, C^*, \langle \pi_{All}^S, \pi_{All;O}^S, \pi_{All;l}^S, \pi_O^S, \pi_l^S \rangle)$ 
4: end if
5: do
6:    $iblock := \top$ ,  $prf := \emptyset$ 
7:   do
8:      $cmp := \mathbf{SatS}(\mathit{cnf}(C^*, \Gamma) \wedge blocking \wedge iblock)$ 
9:     if  $cmp \neq \emptyset$  then
10:       $prf := cmp$ 
11:       $iblock := \bigwedge_{a \in \mathbf{I-ARGS}(cmp)} I_a \wedge \bigwedge_{a \in \mathbf{O-ARGS}(cmp)} O_a$ 
12:      if  $\pi_{iO}^P$  then
13:         $iblock := iblock \wedge \bigvee_{a \in \mathbf{U-ARGS}(cmp)} O_a$ 
14:      end if
15:      if  $\pi_{il}^P$  then
16:         $iblock := iblock \wedge \bigvee_{a \in \mathbf{U-ARGS}(cmp)} I_a$ 
17:      end if
18:    end if
19:    while  $cmp \neq \emptyset \wedge \mathbf{U-ARGS}(cmp) \neq \emptyset$ 
20:    if  $prf \neq \emptyset$  then
21:       $L_p := L_p \cup \{prf\}$ 
22:      if  $\pi_{eO}^P$  then
23:         $blocking := blocking \wedge \bigvee_{a \in \mathbf{I-ARGS}(prf)} O_a$ 
24:      end if
25:      if  $\pi_{el}^P$  then
26:         $blocking := blocking \wedge \bigvee_{a \in \mathbf{O-ARGS}(prf)} I_a$ 
27:      end if
28:    end if
29:  while  $prf \neq \emptyset$ 
30:  if  $L_p = \emptyset$  then
31:     $L_p = \{\mathcal{LU}\}$ 
32:  end if
33:  return  $L_p$ 

```

5.3. Algorithms for Preferred Semantics

Algorithm 13 details the implementation of Algorithms 5 and 6 for the enumeration of preferred labellings. It receives in input an argumentation framework $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, the encoding $C^* \in \mathfrak{C}$ (with the same meaning as in the case of the previous algorithms), and 10 configuration parameters. The first parameter $\pi_{\mathfrak{S}}^{\mathfrak{P}}$ is the selector between Algorithms 5 and 6, while the other 9 parameters specify how to encode the constraints relevant to the non-deterministic functions corresponding to the SAT solver invocations. In particular, $\langle \pi_{\text{All}}^{\mathfrak{S}}, \pi_{\text{All;O}}^{\mathfrak{S}}, \pi_{\text{All;I}}^{\mathfrak{S}}, \pi_{\text{O}}^{\mathfrak{S}}, \pi_{\text{I}}^{\mathfrak{S}} \rangle$ are related to the enumeration of stable labellings and have the same meaning as detailed in the previous section.

If $\pi_{\mathfrak{S}}^{\mathfrak{P}} = \top$ then the algorithm **SAT-EL-ST** is called at line 3, corresponding to line 1 of Algorithm 6, passing it the parameters $\langle \pi_{\text{All}}^{\mathfrak{S}}, \pi_{\text{All;O}}^{\mathfrak{S}}, \pi_{\text{All;I}}^{\mathfrak{S}}, \pi_{\text{O}}^{\mathfrak{S}}, \pi_{\text{I}}^{\mathfrak{S}} \rangle$ it expects. Then the outer loop of Algorithms 5 and 6 correspond to lines 5-29 of Algorithm 13, the inner loop to lines 7-19. The function FindCL^{\oplus} , used at line 2 of Algorithms 5 and 6, is implemented within lines 8 and 22–27. Line 8 identifies a non-empty complete labelling that is not less committed w.r.t. any found preferred labelling. Proposition 8 shows that there are three ways to ensure that: an argument labelled **in** should then become labelled **out** (cf. line 23); an argument labelled **out** should then become labelled **in** (cf. line 26); or both of the previous ones. Those options are controlled by the parameters $\pi_{\text{eO}}^{\mathfrak{P}}$ and $\pi_{\text{el}}^{\mathfrak{P}}$.

Moreover, the function $\text{FindCL}^{\supseteq}$, used at line 3 of Algorithm 6, is implemented within lines 8–18 of Algorithm 13. In particular, as from the second iteration of the inner loop, line 8 identifies a new complete labelling which is strictly more committed w.r.t. a previously found complete labelling. Proposition 9 shows that there are three ways to ensure that: besides maintaining the labels of arguments labelled **in** and **out** (line 11), an argument labelled **undec** should then become **out** (line 13); or an argument labelled **undec** should then become **in** (line 16); or both conditions can be enforced. Those options are controlled by the parameters $\pi_{\text{iO}}^{\mathfrak{P}}$ and $\pi_{\text{il}}^{\mathfrak{P}}$.

It should be noted that if in the argumentation framework there is only a preferred extension which is empty, then due to the non-emptiness condition in the encodings of \mathfrak{C} no labelling is identified in the main loop. In this case, the algorithm correctly returns a unique labelling, denoted as \mathcal{LU} , assigning the label **undec** to all arguments (line 31).

While the credulous acceptance w.r.t. preferred semantics is rather straightforward—see Algorithms 7 and 14—and we report it only for completeness, skeptical acceptance needs a little further explanation.

First of all, Algorithm 15 implements Algorithm 8, which as shown in Section 3.2 is a direct extension of the algorithm for enumerating preferred labellings to the case of checking skeptical acceptance. As such, the input parameters are exactly those of Algorithm 13.

Finally, Algorithm 16 implements Algorithm 9 which represents an improved algorithm for skeptical acceptance. As shown in Section 4.2, the constraints on the complete labelling \mathcal{Lab} returned by the functions $\text{FindCLU}_{\mathfrak{a}}^{\oplus}$ and $\text{FindCL}_{\mathfrak{a}}^{\supseteq}$ are the same as those relevant to the functions FindCL^{\oplus} and $\text{FindCL}^{\supseteq}$, respectively,

with the addition of the condition $\mathcal{L}ab(\mathbf{a}) = \text{undec}$ or $\mathcal{L}ab(\mathbf{a}) = \text{in}$. As a consequence, the input parameters $\langle \pi_{\text{O}}^{\text{P}}, \pi_{\text{I}}^{\text{P}}, \pi_{\text{eO}}^{\text{P}}, \pi_{\text{eI}}^{\text{P}} \rangle$ have the same meaning as in Algorithm 13. In particular, it is easy to see that $\text{FindCLU}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}_c)$ is implemented within lines 8-28, with the relevant constraint enforced in lines 23-28, while $\text{FindCL}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}ab, \text{undec})$ is implemented within lines 8-16, with the relevant constraint enforced in lines 10-16.

Algorithm 14 Credulous acceptance under Preferred Semantics using SAT solvers

SAT-DC-PR

Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{x} \in \mathcal{A}$, $C^* \in \mathfrak{C}$

Output: $\{\top, \perp\}$

1: **return** $(\text{SatS}(\text{cnf}(C^*, \Gamma) \wedge I_{\mathbf{x}}) \neq \emptyset)$

Algorithm 15 Skeptical acceptance under Preferred Semantics using SAT solvers 1

SAT-DS-PR-straightforward
Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle, \mathbf{x} \in \mathcal{A}, C* \in \mathcal{C}, \langle \pi_S^P, \pi_{All}^S, \pi_{All;O}^S, \pi_{All;l}^S, \pi_O^S, \pi_l^S, \pi_{iO}^P, \pi_{il}^P, \pi_{eO}^P, \pi_{el}^P \rangle \in \{\top, \perp\}^{10}$
Output: $\{\top, \perp\}$

- 1: $L_p := \emptyset, blocking := \top$
- 2: **if** π_S^P **then**
- 3: $\langle L_p, blocking \rangle := \text{SAT-EL-ST}(\Gamma, C*, \langle \pi_{All}^S, \pi_{All;O}^S, \pi_{All;l}^S, \pi_O^S, \pi_l^S \rangle)$
- 4: **for all** $aL_p \in L_p$ **do**
- 5: **if** $\mathbf{x} \notin \text{I-ARGS}(aL_p)$ **then**
- 6: **return** \perp
- 7: **end if**
- 8: **end for**
- 9: **end if**
- 10: **do**
- 11: $iblock := \top, prf := \emptyset$
- 12: **do**
- 13: $cmp := \text{SatS}(\text{cnf}(C*, \Gamma) \wedge blocking \wedge iblock)$
- 14: **if** $cmp \neq \emptyset$ **then**
- 15: $prf := cmp, iblock := \bigwedge_{a \in \text{I-ARGS}(cmp)} I_a \wedge \bigwedge_{a \in \text{O-ARGS}(cmp)} O_a$
- 16: **if** π_{iO}^P **then**
- 17: $iblock := iblock \wedge \bigvee_{a \in \text{U-ARGS}(cmp)} O_a$
- 18: **end if**
- 19: **if** π_{il}^P **then**
- 20: $iblock := iblock \wedge \bigvee_{a \in \text{U-ARGS}(cmp)} I_a$
- 21: **end if**
- 22: **end if**
- 23: **while** $cmp \neq \emptyset \wedge \text{U-ARGS}(cmp) \neq \emptyset$
- 24: **if** $prf \neq \emptyset$ **then**
- 25: **if** $\mathbf{x} \notin \text{I-ARGS}(prf)$ **then**
- 26: **return** \perp
- 27: **end if**
- 28: $L_p := L_p \cup \{prf\}$
- 29: **if** π_{eO}^P **then**
- 30: $blocking := blocking \wedge \bigvee_{a \in \text{I-ARGS}(prf)} O_a$
- 31: **end if**
- 32: **if** π_{el}^P **then**
- 33: $blocking := blocking \wedge \bigvee_{a \in \text{O-ARGS}(prf)} I_a$
- 34: **end if**
- 35: **end if**
- 36: **while** $prf \neq \emptyset$
- 37: **if** $L_p = \emptyset$ **then**
- 38: **return** \perp
- 39: **end if**
- 40: **return** \top

Algorithm 16 Skeptical acceptance under Preferred Semantics using SAT solvers 2

SAT-DS-PR
Input: $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$, $\mathbf{x} \in \mathcal{A}$, $C^* \in \mathfrak{C}$, $\langle \pi_{iO}^P, \pi_{iI}^P, \pi_{eO}^P, \pi_{eI}^P \rangle \in \{\top, \perp\}^4$
Output: $\{\top, \perp\}$

- 1: **if** $\text{SatS}(\text{cnf}(C^*, \Gamma) \wedge I_{\mathbf{x}}) = \emptyset \vee \text{SatS}(\text{cnf}(C^*, \Gamma) \wedge O_{\mathbf{x}}) \neq \emptyset$ **then**
- 2: **return** \perp
- 3: **end if**
- 4: $\text{blocking} := \top$
- 5: **do**
- 6: $\text{iblock} := \top$, $\text{prf} := \emptyset$
- 7: **do**
- 8: $\text{cmp} := \text{SatS}(\text{cnf}(C^*, \Gamma) \wedge \text{blocking} \wedge \text{iblock} \wedge U_{\mathbf{x}})$
- 9: **if** $\text{cmp} \neq \emptyset$ **then**
- 10: $\text{prf} := \text{cmp}$, $\text{iblock} := \bigwedge_{\mathbf{a} \in \text{I-ARGS}(\text{cmp})} I_{\mathbf{a}} \wedge \bigwedge_{\mathbf{a} \in \text{O-ARGS}(\text{cmp})} O_{\mathbf{a}}$
- 11: **if** π_{iO}^P **then**
- 12: $\text{iblock} := \text{iblock} \wedge \bigvee_{\mathbf{a} \in \text{U-ARGS}(\text{cmp})} O_{\mathbf{a}}$
- 13: **end if**
- 14: **if** π_{iI}^P **then**
- 15: $\text{iblock} := \text{iblock} \wedge \bigvee_{\mathbf{a} \in \text{U-ARGS}(\text{cmp})} I_{\mathbf{a}}$
- 16: **end if**
- 17: **end if**
- 18: **while** $\text{cmp} \neq \emptyset$
- 19: **if** $\text{prf} \neq \emptyset$ **then**
- 20: **if** $\text{SatS} \left(\text{cnf}(C^*, \Gamma) \wedge \text{blocking} \wedge \bigwedge_{\mathbf{a} \in \text{I-ARGS}(\text{cmp})} I_{\mathbf{a}} \wedge \bigwedge_{\mathbf{a} \in \text{O-ARGS}(\text{cmp})} O_{\mathbf{a}} \wedge I_{\mathbf{x}} \right) = \emptyset$ **then**
- 21: **return** \perp
- 22: **end if**
- 23: **if** π_{eO}^P **then**
- 24: $\text{blocking} := \text{blocking} \wedge \bigvee_{\mathbf{a} \in \text{I-ARGS}(\text{prf})} O_{\mathbf{a}}$
- 25: **end if**
- 26: **if** π_{eI}^P **then**
- 27: $\text{blocking} := \text{blocking} \wedge \bigvee_{\mathbf{a} \in \text{O-ARGS}(\text{prf})} I_{\mathbf{a}}$
- 28: **end if**
- 29: **end if**
- 30: **while** $\text{prf} \neq \emptyset$
- 31: **return** \top

6. Experimental Evaluation

6.1. Experimental Settings

We ground our experimental analysis with the ICCMA 2015⁹ [27] and ICCMA 2017 benchmarks.¹⁰ Experiments have been run on a cluster with computing nodes equipped with 2.5 Ghz Intel Core 2 Quad Processors, 4 GB of RAM and Linux operating system. A cutoff of 900 seconds was imposed to compute the labellings—either preferred or stable—for each AF as well as for addressing each decision problem. Decision problems have been run three times on randomly selected arguments and results averaged, following the methodology designed for ICCMA 2015 [27]. For each ArgSemSAT-configuration we recorded the overall result: success (if it solved the considered problem), crashed, timed-out or ran out of memory.

To analyse the results we considered two metrics: IPC (the higher the better) and PAR10 (the lower the better). IPC is the International Planning Competition score, as defined for the Agile track of the 2014 edition of the competition [34, 35], and represents a relative performance measure among a set of compared solvers: for a solver \mathcal{C} and a problem p , $Score(\mathcal{C}, p)$ is 0 if p is unsolved, and $1/(1 + \log_{10}(T_p(\mathcal{C})/T_p^*))$ otherwise (where $T_p(\mathcal{C})$ is the amount of time required by \mathcal{C} to solve the problem p , and T_p^* is the minimum amount of time required by any compared system). The IPC score on a set of instances is given by the sum of the scores achieved on each considered instance. PAR10 is the Penalized Average Runtime with a penalty equal to ten times the cutoff time. It represents an absolute performance measure which trades off coverage—i.e. the percentage of the AF s that are correctly solved below the cutoff time—and runtime for successfully analysed AF s: runs that do not solve the given problem get ten times the cutoff time, other runs get the actual runtime. The PAR10 score of a solver on a set of AF s is the average of the associated runtimes.

It should be noted that IPC and PAR10 are independent measures, thus considering both of them yields a detailed picture of solver performance. For instance, a configuration of ArgSemSAT may have a good IPC but a worse PAR10, allowing one to infer that the coverage is low but the solver is fast in solving those instances it is able to solve (the opposite holds in the reverse case).

Algorithms have been implemented using two different SAT solvers, Minisat [36] as a library, and Glucose [37] invoked through PIPE communication.¹¹

In the following we present the results of the experimental analysis, describing in particular four findings that have been achieved therefrom.




Figure 4: IPC scores of enumerating stable labellings using Minisat varying the chosen encoding on the ICCMA 2015 benchmarks.



Figure 5: IPC scores of enumerating stable labellings using Minisat varying the chosen encoding on the ICCMA 2017 benchmarks.




Figure 6: IPC scores of enumerating stable labellings using Glucose varying the chosen encoding on the ICCMA 2015 benchmarks.




Figure 7: PAR10 scores of enumerating stable labellings using both Minisat and Glucose varying the chosen encoding on the ICCMA 2015 benchmarks.




Figure 8: PAR10 scores of enumerating stable labellings using Minisat varying the chosen encoding on the ICCMA 2017 benchmarks.

6.2. Experimental Results

6.2.1. $C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\rightarrow}$ is the most efficient encoding of complete labellings

In order to determine the most efficient encoding of complete labellings in the context of the problems tackled by ArgSemSAT, we computed the enumeration of stable and preferred labellings varying the chosen encoding among the 18 correct ones identified in Section 4.1 (see Figure 2). Moreover, we also considered credulous and skeptical acceptance w.r.t. preferred semantics, once again varying the chosen encoding. Results for credulous and skeptical acceptance w.r.t. stable semantics are omitted as they do not provide further insights.

Algorithms for stable and preferred semantics have been configured respectively with options: $\pi_1^S \wedge \pi_0^S$; and $P_{|O}^O \equiv \pi_{|i}^P \wedge \pi_{|o}^P \wedge \pi_{e|}^P \wedge \pi_{eO}^P$ (cf. Table B.5). We sampled 50 data points considering other possible configurations: the results we gathered do not provide evidence that different configurations could lead to rejecting our experimental hypothesis.

Figures 4 and 5 depict the IPC scores obtained by each encoding in enumerating stable labellings using Minisat. The same dynamics can be appreciated in Figure 6, where Glucose is used instead (on ICCMA 2015 benchmarks only). We chose not to depict the two graphs on the same figure because IPC is a relative measure—i.e. the value of configuration depends on the pool of configurations considered—therefore the resulting figure would have been misleading, as the

⁹http://argumentationcompetition.org/2015/ICCMA2015_benchmarks.zip (on 17th July 2017).

¹⁰<http://argumentationcompetition.org/2017/results.html> (on 6th February 2019).

¹¹We refer any reader interested in knowing the effect of different SAT solvers configuration parameters when addressing argumentation problems to [13].




Figure 9: IPC scores of enumerating preferred labellings using Minisat varying the chosen encoding on the ICCMA 2015 benchmarks.

numeric values of the results gathered with a SAT solver are not directly comparable with the results gathered with the second SAT solver. Figures 7 and 8 depict the PAR10 scores obtained in the enumeration of stable labellings. As PAR10 is not a relative measure, i.e. the PAR10 score of a configuration does not depend on the pool of configurations considered, the results can be safely displayed on the same figure without risk of confusion. It is worth noticing that $C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\rightarrow}$ has slightly higher PAR10 values than $C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\leftrightarrow}$ and $C_{out}^{\rightarrow} \wedge C_{undec}^{\leftrightarrow}$ when considering the ICCMA 2017 benchmarks (Figure 8), while still being the one with highest IPC score (Figure 5). This seems to suggest that for a few instances—the actual difference between the PAR10 values is very small—the algorithm running with $C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\leftrightarrow}$ and $C_{out}^{\rightarrow} \wedge C_{undec}^{\leftrightarrow}$ managed to find the solution in the assigned time, while the algorithm running with $C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\rightarrow}$ failed to do so hence incurring in the penalty.

Results obtained by enumerating preferred labellings are similarly shown in Figures 9 and 10, which depict the IPC scores obtained by each encoding using Minisat, Figure 11, which evidences a substantially similar dynamics using Glucose (on ICCMA 2015 benchmarks only), and Figures 12 and 13, which directly compare w.r.t. PAR10 score the implementations.

Overall, five encodings emerge as *optimal* according to IPC score (Figures 4, 5, 6, 9, 10 and 11), namely $C_{in}^{\leftrightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\rightarrow}$, $C_{in}^{\rightarrow} \wedge C_{out}^{\leftrightarrow} \wedge C_{undec}^{\rightarrow}$, $C_{in}^{\leftrightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\leftrightarrow}$, $C_{in}^{\rightarrow} \wedge C_{out}^{\leftrightarrow} \wedge C_{undec}^{\leftrightarrow}$ and $C_{in}^{\leftrightarrow} \wedge C_{out}^{\leftrightarrow} \wedge C_{undec}^{\rightarrow}$. The same observation substantially holds for PAR10 in the case of preferred labellings (Figures 12 and 13) and in the case of stable labellings using Glucose (Figure 7), while the resulting scores are more flat with Minisat, which is used as a library. It is worth noting that all these five configurations have a single commonality, i.e. they lack $C_{undec}^{\leftrightarrow}$, thus with a small abuse of notation we




Figure 10: IPC scores of enumerating preferred labellings using Minisat varying the chosen encoding on the ICCMA 2017 benchmarks.




Figure 11: IPC scores of enumerating preferred labellings using Glucose varying the chosen encoding on the ICCMA 2015 benchmarks.




Figure 12: PAR10 scores of enumerating preferred labellings using both Minisat and Glucose varying the chosen encoding on the ICCMA 2015 benchmarks.




Figure 13: PAR10 scores of enumerating preferred labellings using Minisat varying the chosen encoding on the ICCMA 2017 benchmarks.




Figure 15: IPC (left axis, red) and PAR10 (right axis, blue) scores of solving credulous acceptance w.r.t. preferred semantics using Minisat only and varying the chosen encoding, cf. Figure 2 , on the ICCMA 2017 benchmarks




Figure 16: IPC (left axis, red) and PAR10 (right axis, blue) scores of solving skeptical acceptance w.r.t. preferred semantics using Minisat only and varying the chosen encoding, cf. Figure 2, on the ICCMA 2015 benchmarks.




Figure 17: IPC (left axis, red) and PAR10 (right axis, blue) scores of solving skeptical acceptance w.r.t. preferred semantics using Minisat only and varying the chosen encoding, cf. Figure 2 , on the ICCMA 2017 benchmarks

we consider Minisat or Glucose, in the interest of time and space we limit our following results to Minisat only.

Figures 14, 15, 16, and 17 show the results of the experimental analysis when considering, respectively, credulous and skeptical acceptance w.r.t. preferred semantics using Minisat. Results are substantially in line with those highlighted for the enumeration of preferred labellings.

Summing up, this large experimentation counting 97,560 executions¹⁴ of our algorithms shows that $C_{\exists}^{\rightarrow} \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ is consistently the encoding leading to the best performance.

We believe that the above results can be explained by looking at the CNF formulæ of the encodings of complete labellings presented in Proposition 11.

Indeed, it is known [38] that the 2-SAT problem—i.e. the SAT problem where all the formulæ are binary—is polynomial. Moreover in [39] the authors show how reducing the length of disjunctive formulæ leads to significant computational benefits. Modern SAT solvers rely on conflict-driven clause learning [40], which is an improvement built around the DPLL unit-propagation algorithm—a typical conflict-driven clause learning based SAT solver spends at least 80% of the time running the unit propagation procedure [41]—and thus it will also benefit from short—best with 2-SAT—formulæ.

According to Proposition 11, $C_{\text{undec}}^{\leftarrow}$ maximises the number of disjunctions per clause. This seems to explain why the best configurations belong to $\neg C_{\text{undec}}^{\leftarrow}$.

¹⁴192 AFs in the ICCMA 2015 benchmark, 350 in the ICCMA 2017 benchmark, each of those analysed in 36 configurations both for enumeration of stable and of preferred labellings, on top of 3 runs for credulous and skeptical acceptance w.r.t. preferred semantics against 18 different configurations.

Moreover, there are only two non-redundant encodings in $\neg C_{\text{undec}}^{\leftarrow}$, namely $C_{\exists}^{\rightarrow}$ and C_{\exists}^{\leftarrow} , and both of them minimise the number of clauses. The fact that $C_{\exists}^{\rightarrow}$ achieves a better performance can be explained in the case of the enumeration of stable labellings by taking into account that, for each argument \mathbf{a} , the condition $\neg U_{\mathbf{a}}$ is enforced, thus the constraint $C_{\text{undec}}^{\rightarrow}$ does not play any role. As a consequence, differently from $C_{\exists}^{\rightarrow}$, C_{\exists}^{\leftarrow} becomes redundant, and the presence of redundant clauses is known to have a negative impact on the performance of SAT solvers [42]. This may also explain the relatively good performance achieved by $C_{\exists}^{\rightarrow}$. We conjecture that similar reasons underlie the performance results in the case of preferred semantics, though the performance difference is less noticeable.

As to the comparison between Glucose and Minisat, the first seems to be more sensible to the presence of $C_{\text{undec}}^{\leftarrow}$, and its decrease of performance is possibly particularly accentuated by the fact that it has been not implemented as an internal library but called as an external process using PIPE communications. On the other hand, from Figure 12 Glucose emerges as the *winner* against Minisat according to PAR10 for the configurations in $\neg C_{\text{undec}}^{\leftarrow}$. It must also be acknowledged that the differences between the two are quite minimal so that they do not justify further investigations: Glucose is the result of almost 10 years of improvements from the SAT community, and enumerating preferred labellings is significantly more complex than enumerating stable labellings. It seems reasonable that a better solver would have the opportunity to demonstrate its value.

In this section we have experimentally shown that $C_{\exists}^{\rightarrow} \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ overall is the encoding leading to the best performance for all the considered algorithms and different configurations of the algorithm parameters. In the subsequent subsections we discuss how such parameters exhibit positive interactions that can be exploited to gain a better performance.

6.2.2. Using an AllSAT solver leads to a significant improvement for enumerating stable labellings

We fixed the encoding to be the most efficient, i.e. $C_{\exists}^{\rightarrow} \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$, and we enumerated stable labellings on the benchmark *AFs* varying the configuration of the algorithm.

Tables 2 and 3 summarise the results in terms of IPC score, PAR10, and coverage. Only the parameters set to \top are indicated, while false parameters are omitted for brevity. It is immediate to see that the use of AllSAT leads to a significant improvement for enumerating stable labellings. It is also worth noticing that varying the other parameters, i.e. $\pi_{\text{O}}^{\text{S}}$ and $\pi_{\text{I}}^{\text{S}}$, leads to statistically significant results with $p < 0.05$ but not with $p < 0.01$ (Friedman test, $\chi^2(2) = 6.776, p = 0.034$) on the ICCMA 2015 benchmarks. In the case of the ICCMA 2017, it is self-evident from the magnitude of values that the results are also statistically significant. Varying parameters $\pi_{\text{All;I}}^{\text{S}}$ and $\pi_{\text{All;O}}^{\text{S}}$ do not lead to statistically significant variations on the ICCMA 2015 benchmarks, hence in Tables 2 and 3 we considered only the combination $\pi_{\text{All}}^{\text{S}} \wedge \pi_{\text{All;I}}^{\text{S}} \wedge \pi_{\text{All;O}}^{\text{S}}$.

Parameters	IPC Score	PAR10	Coverage
$\pi_{\text{O}}^{\text{S}}$	177.57	8.20	100.0
$\pi_{\text{T}}^{\text{S}}$	177.26	8.24	100.0
$\pi_{\text{O}}^{\text{S}} \wedge \pi_{\text{T}}^{\text{S}}$	177.15	8.49	100.0
$\pi_{\text{All}}^{\text{S}} \wedge \pi_{\text{All;l}}^{\text{S}} \wedge \pi_{\text{All;o}}^{\text{S}}$	191.87	5.40	100.0

Table 2: IPC score, PAR10 and coverage (percentage of *AF*s successfully analysed) when considering different algorithmic parameters for enumerating stable labellings on the ICCMA 2014 benchmarks. In bold the best results.

Parameters	IPC Score	PAR10	Coverage
$\pi_{\text{O}}^{\text{S}}$	214.45	4460.49	73.43
$\pi_{\text{T}}^{\text{S}}$	215.02	4503.22	73.14
$\pi_{\text{O}}^{\text{S}} \wedge \pi_{\text{T}}^{\text{S}}$	210.31	4834.84	71.14
$\pi_{\text{All}}^{\text{S}} \wedge \pi_{\text{All;l}}^{\text{S}} \wedge \pi_{\text{All;o}}^{\text{S}}$	303.58	2263.87	86.86

Table 3: IPC score, PAR10 and coverage (percentage of *AF*s successfully analysed) when considering different algorithmic parameters for enumerating stable labellings on the ICCMA 2017 benchmarks. In bold the best results.

The improvement yielded by the AllSAT solver is hardly a surprise, given that our algorithm for enumerating stable labellings is less engineered than the AllSAT algorithm that uses only Minisat data structures.

6.2.3. Enumerating stable labellings first can lead to a significant improvement for enumerating preferred labellings

Also in this case we fixed the encoding to be the most efficient one, i.e. $C \rightleftharpoons \equiv C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$, and we enumerated the preferred labellings on the benchmark *AF*s varying the configuration of the algorithm.

Figures 18, 19, 20, and 21 depict the box-and-whisker plot of three groups of results. As above, false parameters are omitted for brevity, unless necessary to avoid confusion.

Group (A) considers performance results collected when running the algorithm with the option $\pi_{\text{S}}^{\text{P}} = \perp$, i.e. without searching for stable labellings first, with combinations of options $\pi_{\text{il}}^{\text{P}}$, $\pi_{\text{IO}}^{\text{P}}$, $\pi_{\text{el}}^{\text{P}}$, and $\pi_{\text{eO}}^{\text{P}}$.

Group (B) considers the case where the algorithm is run with $\pi_{\text{S}}^{\text{P}} = \top$ and $\pi_{\text{All}}^{\text{S}} = \perp$, i.e. stable labellings are enumerated first, but without using an AllSAT solver.

Finally, group (C) considers the usual combinations of options $\pi_{\text{il}}^{\text{P}}$, $\pi_{\text{IO}}^{\text{P}}$, $\pi_{\text{el}}^{\text{P}}$, $\pi_{\text{eO}}^{\text{P}}$ with $\pi_{\text{S}}^{\text{P}} = \pi_{\text{All}}^{\text{S}} = \top$.




Figure 18: IPC scores (higher better) when enumerating preferred labellings on ICCMA 2015 benchmarks varying algorithm configurations: (A) computing without searching for stable labellings first; (B) searching for stable labellings first (π_S^P) but without using an AllSAT solver; (C) searching for stable labellings first (π_S^P) using an AllSAT solver (π_{All}^S , with $\pi_{All;l}^S \wedge \pi_{All;o}^S$).

Overall, our empirical evaluation supports the claim that searching for stable labellings first leads to a significant improvement for enumerating preferred labellings, but only when an AllSAT solver is exploited. If this is not the case, Figures 18, 19, 20, and 21 suggest that starting from stable labellings improves the coverage of the algorithm, while slowing down the overall search. We believe this is due to the additional time required for proving that the search for stable labellings is completed, as well as to the cost of parameters passing, particularly in cases where the stable labellings are a few percentage of the preferred labellings. These disadvantages appear as completely compensated by the use of the engineered AllSAT algorithm.

Post-hoc analysis: on the positive interactions between options when enumerating stable labellings first. We also considered the interactions between choosing different options for enumerating stable labellings at the beginning using an AllSAT solver, together with the acceptable combination of parameters for computing blocking clauses when exploring the space of complete labellings, i.e. parameters π_{il}^P , π_{iO}^P , π_{el}^P , and π_{eO}^P . Figures 22 and 23 depict the results of this analysis. The best performance is obtained by combining $\pi_S^P \wedge \pi_{All}^S \wedge \pi_{All;o}^S$ with $P_{iO}^I \equiv \pi_{il}^P \wedge \pi_{iO}^P \wedge \pi_{el}^P \wedge \pi_{eO}^P$ over the ICCMA 2015 benchmarks, and by combining $\pi_S^P \wedge \pi_{All}^S \wedge \pi_{All;o}^S$ with $P_l^I \equiv \pi_{il}^P \wedge \pi_{el}^P$ over the ICCMA 2017 benchmarks. As noticed earlier, the ICCMA 2017 benchmark set is substantially more complex




Figure 19: IPC scores (higher better) when enumerating preferred labellings on ICCMA 2017 benchmarks varying algorithm configurations: (A) computing without searching for stable labellings first; (B) searching for stable labellings first (π_S^P) but without using an AllSAT solver; (C) searching for stable labellings first (π_S^P) using an AllSAT solver (π_{All}^S , with $\pi_{All;I}^S \wedge \pi_{All;O}^S$).

Version	IPC Score	PAR10	Coverage
Straightforward	186.49	2.22	100.0
Improved	188.30	1.39	100.0

Table 4: IPC score, PAR10 and coverage (percentage of *AF*s successfully analysed) when considering straightforward and improved algorithms for computing skeptical acceptability w.r.t. preferred semantics on ICCMA 2015 benchmarks. In bold the best results.

when dealing with preferred semantics, hence this difference suggests that such parameter optimisation is heavily benchmark dependent consistently with the *no free-lunch theorem* [17].

6.2.4. The improved algorithm for skeptical acceptance w.r.t. preferred semantics leads to significant improvements

We fixed the encoding to be the most efficient, i.e. $C\exists \equiv C_{in}^{\rightarrow} \wedge C_{out}^{\rightarrow} \wedge C_{undec}^{\rightarrow}$, and we computed skeptical acceptance of arguments w.r.t. preferred semantics on the ICCMA 2015 benchmark *AF*s varying the configuration of the **straightforward** algorithm, i.e. Algorithm 15, and of the **improved** algorithm, i.e. Algorithm 16. In order to provide a fair comparison where neither algorithm is allowed to exploit an AllSAT solver, we did not consider the option of computing




Figure 20: PAR10 scores (lower better) when enumerating preferred labellings on ICCMA 2015 benchmarks varying algorithm configurations: (A) computing without searching for stable labellings first; (B) searching for stable labellings first (π_S^P) but without using an AllSAT solver; (C) searching for stable labellings first (π_S^P) using an AllSAT solver (π_{All}^S , with $\pi_{All;I}^S \wedge \pi_{All;O}^S$).

stable extensions first in the straightforward algorithm. As a matter of fact, considering this option would not lead to different conclusions, as discussed in the post-hoc analysis below.

Table 4 compares the best configuration for the straightforward algorithm against the best configuration for the improved algorithm. According to the results we achieved, the best configuration of the straightforward algorithm is with options $\pi_{il}^P \wedge \pi_{el}^P$ while for the improved version is with $\pi_{il}^P \wedge \pi_{iO}^P \wedge \pi_{el}^P \wedge \pi_{eO}^P$. The results of Table 4 are statistically significant (Wilcoxon, $Z=-2.487$, $p = 0.013$) and show that there is an improvement of performance using the improved algorithm.

This analysis, first performed before submitting ArgSemSAT to the ICCMA 2017, has been confirmed in the results of that competition. Indeed Algorithm 16 resulted to be the winner of the sub-track of ICCMA 2017 for skeptical acceptance of preferred semantics by a large margin.¹⁵

Post-hoc analysis: on the positive interactions between options when considering skeptical acceptance w.r.t. preferred semantics. We also analysed the interactions between choosing to use different options for considering (or not) stable labellings

¹⁵<http://argumentationcompetition.org/2017/ICCMA2017-slides-tafa.pdf> (on 6th February 2019), page 26.




Figure 21: PAR10 scores (lower better) when enumerating preferred labellings on ICCMA 2017 benchmarks varying algorithm configurations: (A) computing without searching for stable labellings first; (B) searching for stable labellings first (π_S^P) but without using an ALLSAT solver; (C) searching for stable labellings first (π_S^P) using an ALLSAT solver (π_{All}^S , with $\pi_{All;i}^S \wedge \pi_{All;o}^S$).

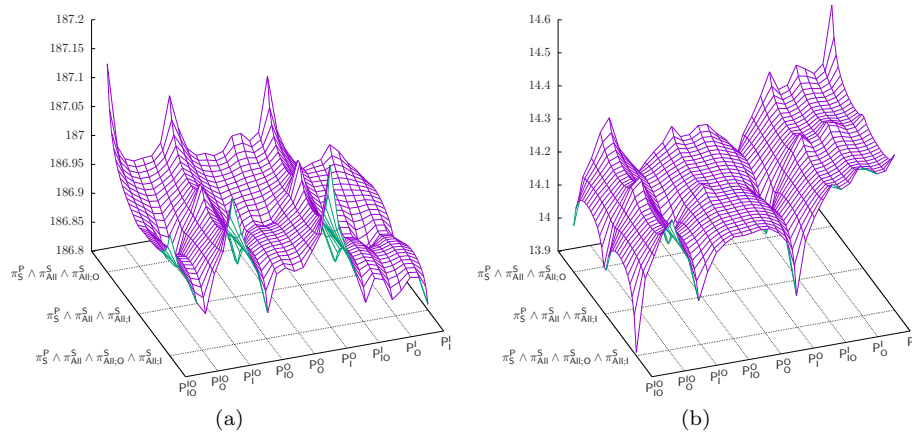


Figure 22: IPC scores (a) and PAR10 (b) when enumerating preferred labellings varying algorithm parameters on ICCMA 2015 benchmarks. Here, $P_i^I \equiv \pi_{il}^P \wedge \pi_{el}^P$, $P_o^I \equiv \pi_{il}^P \wedge \pi_{eo}^P$, $P_{io}^I \equiv \pi_{il}^P \wedge \pi_{el}^P \wedge \pi_{eo}^P$, $P_i^O \equiv \pi_{io}^P \wedge \pi_{el}^P$, $P_o^O \equiv \pi_{io}^P \wedge \pi_{eo}^P$, $P_{io}^O \equiv \pi_{io}^P \wedge \pi_{el}^P \wedge \pi_{eo}^P$, $P_i^{IO} \equiv \pi_{io}^P \wedge \pi_{eo}^P$, $P_o^{IO} \equiv \pi_{io}^P \wedge \pi_{el}^P \wedge \pi_{eo}^P$.




Figure 23: IPC scores (a) and PAR10 (b) when enumerating preferred labellings varying algorithm parameters on ICCMA 2017 benchmarks. Here, $P_I^I \equiv \pi_{il}^P \wedge \pi_{el}^P$, $P_I^O \equiv \pi_{il}^P \wedge \pi_{eO}^P$, $P_{IO}^I \equiv \pi_{il}^P \wedge \pi_{eO}^P \wedge \pi_{eO}^P$, $P_I^O \equiv \pi_{iO}^P \wedge \pi_{el}^P$, $P_O^O \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^O \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^P \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^S \equiv \pi_{il}^P \wedge \pi_{iO}^P \wedge \pi_{eO}^P$.




Figure 24: IPC scores (a) and PAR10 (b) when computing skeptical acceptance w.r.t. preferred semantics varying algorithm parameters on ICCMA 2015 benchmarks. Here, $P_I^I \equiv \pi_{il}^P \wedge \pi_{el}^P$, $P_I^O \equiv \pi_{il}^P \wedge \pi_{eO}^P$, $P_{IO}^I \equiv \pi_{il}^P \wedge \pi_{eO}^P \wedge \pi_{eO}^P$, $P_I^O \equiv \pi_{iO}^P \wedge \pi_{el}^P$, $P_O^O \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^O \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^P \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_{IO}^S \equiv \pi_{il}^P \wedge \pi_{iO}^P \wedge \pi_{eO}^P$.




Figure 25: IPC scores (a) and PAR10 (b) when computing skeptical acceptance w.r.t. preferred semantics varying algorithm parameters on ICCMA 2017 benchmarks. Here, $P_1^I \equiv \pi_{il}^P \wedge \pi_{el}^P$, $P_0^I \equiv \pi_{il}^P \wedge \pi_{eO}^P$, $P_1^O \equiv \pi_{il}^P \wedge \pi_{el}^P \wedge \pi_{eO}^P$, $P_0^O \equiv \pi_{iO}^P \wedge \pi_{el}^P$, $P_0^S \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_1^S \equiv \pi_{iO}^P \wedge \pi_{eO}^P$, $P_0^S \equiv \pi_{iI}^P \wedge \pi_{iO}^P \wedge \pi_{eO}^P$, $P_1^O \equiv \pi_{iI}^P \wedge \pi_{iO}^P \wedge \pi_{el}^P \wedge \pi_{eO}^P$.

at the beginning of the skeptical acceptance w.r.t. preferred semantics in the **straightforward** algorithm, together with the acceptable combination of parameters for computing blocking clauses when exploring the space of complete extensions, i.e. parameters π_{il}^P , π_{iO}^P , π_{el}^P , and π_{eO}^P . Figures 24 and 25 depict the results of this analysis: although the magnitude of the effect in the performance of varying the parameters is minimal—sometimes even negligible, in particular according to the PAR10 measures, cf. Figure 24b—according to the IPC (Figure 24a and 25a) it appears that the best option is not to exploit the inclusion relation between the sets of stable and preferred labellings, as highlighted by the peak in Figure 24a and confirmed in Figure 25a.

7. Related Work

The idea of exploiting satisfiability techniques for solving abstract argumentation problems (even if no specific algorithm is devised) can be traced back to the work in [43], where encodings of stable and complete extensions are introduced in terms of variables associated to arguments indicating whether each argument belongs to the extension, rather than in terms of labels as in our approach. A similar approach has been used also more recently in a probabilistic setting [44], but with the aim of estimating the probability for a fixed set of arguments of being conflict-free or admissible, rather than to identify a formula whose models characterize the extensions under a given semantics.

In [32, 33] various NP fragments of abstract argumentation reasoning problems located at the second level of the polynomial hierarchy are identified. Driven by this analysis, complexity-sensitive algorithms for preferred semantics (as well as semi-stable and stage semantics) are proposed which require a polynomial

number of calls to a SAT solver in the identified NP fragments. In particular, the proof of Theorem 3 in [33] contains a procedure for the enumeration of preferred extensions resembling Algorithm 5, and a procedure for skeptical acceptance under preferred semantics is reported which shares the basic schema underlying our Algorithms 8 and 9. There are however substantial differences w.r.t. [32, 33]. From a strictly technical point of view, we make use of labelling-based encodings, while in [33] similarly to [43] a single variable is associated to each argument. Moreover, our algorithms are somewhat generic w.r.t. the specific functions that identify the complete labellings at the basis of the search for preferred labellings (see in particular the definitions of $\text{FindCL}^{\#}$ and $\text{FindCLU}_a^{\#}$ in Section 3.2). This feature as well as the use of labellings are at the basis of the novel strategy underlying Algorithm 9, which at each iteration looks for a maximal complete labelling where \mathbf{a} is `undec`, and then checks with a single call to the SAT solver whether it can be extended to a preferred labelling where \mathbf{a} is `in`, whereas the algorithm in [33] does not use this constraint in maximising complete extensions. From a more general point of view, our focus is different w.r.t. [32, 33], since we are concerned with a systematic exploration and evaluation of the possible encodings and algorithmic variants underlying the above mentioned methodology.

The complexity of the problem of enumerating extensions is investigated in [11] under different semantics and by considering classes of argumentation frameworks that lead to tractability. In the case of symmetric argumentation frameworks, preferred extensions coincide with maximal conflict-free sets, that can be computed with polynomial delay (between one solution and the subsequent one) by means of e.g. the algorithm proposed in [45] to enumerate maximal independent sets in undirected graphs. This also applies to stable semantics for symmetric argumentation frameworks without self-defeating arguments, since in this case stable semantics coincide with preferred semantics. Another tractable class considered in [45] includes argumentation frameworks without even cycles, and in particular acyclic graphs. In this case, there is a unique complete and preferred extension coinciding with the grounded extension, which can be computed in polynomial time. As to stable semantics, it can also be easily checked if the grounded extension is stable.

Since it is possible to determine in polynomial time whether a graph belongs to one of the above classes, it would be easy to use a dedicated polynomial algorithm to solve tractable input instances (in the above sense). Dedicated algorithms can then be used in a portfolio fashion: a pre-processing step can be put in place to identify whether a framework is a member of one of the above mentioned classes, and select the appropriate algorithm to use. There is an extensive literature on portfolio approaches, and techniques to select online promising algorithms. While the exploration of portfolio-based approaches for dealing with argumentation problems is outside the scope of this paper, we refer the interested reader to [46, 47] for an overview of the area, and to [48, 49, 50] for more general information about portfolio approaches in AI. On the other hand, it has to be remarked that the tractable classes mentioned above, while interesting from a computational complexity point of view, represent very specific instances from a knowledge representation perspective if one considers the structure of

original arguments. For instance, symmetric argumentation frameworks require any attacker to be counterattacked, a condition which is violated whenever two rebutting arguments have different strength, or whenever an argument undercuts another argument by denying the application of a defeasible rule used in its construction. As to the class where even-length cycles are not present, this condition prevents e.g. two equally plausible arguments to attack each other, a situation which is very common in structured argumentation.

In [21] we discussed a parallel algorithm for computing preferred extensions based on the SCC-decomposability. Since `ArgSemSAT` exploits a SAT solver as a black box, any solver that supports the DIMACS format—the standard language used in SAT competitions—can be easily integrated and tested. Therefore, if `ArgSemSAT` is running on a multi-core machine, this black-box approach allows to take advantage of parallel SAT solvers as well. In particular, since the most effective approach to exploit parallelisation is to run a portfolio of solvers [51], this will result in having multiple SAT solvers invoked by `ArgSemSAT`, and maximising the use of available cores.

Finally, it has to be remarked that some of the results of this paper are based on a preliminary approach we presented in [52], where in particular the logically equivalent encodings of complete labellings exploited in this paper has been first identified. In [52] we only tackle the problem of enumerating preferred extensions, with a preliminary algorithm which is not parameterised w.r.t. all the different variants considered in this paper. Such variants include e.g. different ways to encode the constraints on complete labellings, different SAT solvers with a variety of configurations, and the exploitation of stable labellings. More generally, we believe that the main advancement of the present paper is the identification and application of the general methodology introduced in Section 1, which led to the winning solver of the preferred semantics track at ICCMA 2017.

We refrain from further comparison with state of the art approaches as this has been the subject of ICCMA 2015 [27, 53]; a re-run of the competition to further analyse the results [13]; and a survey of various approaches [54], which in turn has been superseded by a chapter in the Handbook of Formal Argumentation [55], also published as a journal paper in [56]. We refer an interested reader to those references for further analysis on other techniques for solving problems in abstract argumentation.

8. Conclusion

In this paper we illustrated the design choices that led to the development of `ArgSemSAT`, which managed the first place at the track devoted to preferred semantics of the most recent edition of the international competition on abstract argumentation (ICCMA 2017). The comprehensive experimental analysis we performed informed not only the development of `ArgSemSAT` but provided us with useful insight we desire to share with the community.

First of all, as discussed in Section 6.2.1, although there are 18 correct logically equivalent SAT encodings of complete labellings, one of them leads

to the highest performance independently of the problem and the semantics considered. Moreover, as discussed in Sections 6.2.2 and 6.2.3, composing different techniques such as AllSAT and enumerating stable extensions when searching for preferred extensions, brings advantages. Finally, as we discussed in Section 6.2.4 for the case of skeptical acceptance of arguments w.r.t. preferred semantics, injecting domain specific knowledge in the algorithm design can lead to statistically significant improvements.

Such insight has been attained through the systematic analysis that we devised of different labelling-based encodings and of the ways the relevant constraints can be expressed, and the empirical evaluation of a number of algorithm variations as well as the use of different SAT solvers. While relatively simple algorithms, in particular without resorting to graph decomposition, seem to achieve the best performance at the current stage of the research, the experimental analysis shows that our methodology is a key factor in boosting solver coverage and speed. In fact, the choices concerning encodings, constraints and SAT solvers are interdependent, thus a systematic analysis and experimentation is needed to find a good combination.

As to future work, we believe our approach can be extended in several respects. First, there is a significant interest towards graph-decomposition techniques to lower the computational effort [57]. Experiments using the SCC-based schema [58] show that advantages can be achieved in the case the number of strongly connected components is high [15] or the density is low [14]. A related issue is the development of incremental algorithms in presence of addition/removal of arguments and attacks [59, 60, 61]. Another interesting extension is the use of stochastic local search in incomplete SAT-based algorithms, which has recently been proposed in [62] for enumerating preferred extensions. Moreover, SAT-based techniques can be exploited for several other semantics, such as semi-stable, ideal and eager semantics [63]. Again, we remark that the methodology described in this paper may play a key role in boosting performance also in these extended contexts.

Acknowledgements

The authors wish to thank Yinlei Yu, Pramod Subramanyan, Nestan Tsiskaridze, and Sharad Malik, for having shared the implementation of the AllSAT solver they developed for [28]; and Alessandro Cherubini and Tobia Zanetti for their contribution in coding ArgSemSAT.

References

- [1] K. Atkinson, P. Baroni, M. Giacomin, A. Hunter, H. Prakken, C. Reed, G. Simari, M. Thimm, S. Villata, Toward artificial argumentation, *AI Magazine* 38 (3) (2017) 25–36.
- [2] P. M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games, *Artificial Intelligence* 77 (2) (1995) 321–357.

- [3] C. Cayrol, On the relation between argumentation and non-monotonic coherence based entailment, in: Proc. of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95), Morgan Kaufmann Publishers Inc., 1995, pp. 1443–1448.
- [4] G. Governatori, M. Maher, G. Antoniou, D. Billington, Argumentation semantics for defeasible logic, *Journal of Logic and Computation* 14 (5) (2004) 675–702.
- [5] M. Caminada, S. Sá, J. Alcântara, W. Dvůrák, On the equivalence between logic programming semantics and argumentation semantics, *International Journal of Approximate Reasoning* 58 (2015) 87–111.
- [6] H. Prakken, An abstract framework for argumentation with structured arguments, *Argument and Computation* 1 (2) (2010) 93–124.
- [7] P. Baroni, M. Giacomin, Semantics of abstract argumentation systems, in: *Argumentation in Artificial Intelligence*, Springer, 2009, pp. 25–44.
- [8] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, *Knowledge Engineering Review* 26 (4) (2011) 365–410.
- [9] M. David, S. Karl, Floating conclusions and zombie paths: Two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets, *Artificial Intelligence* 48 (2) (1991) 199–209.
- [10] P. E. Dunne, M. Wooldridge, Complexity of abstract argumentation, in: *Argumentation in Artificial Intelligence*, Springer, 2009, pp. 85–104.
- [11] M. Kröll, R. Pichler, S. Woltran, On the complexity of enumerating the extensions of abstract argumentation frameworks, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 1145–1152. doi:10.24963/ijcai.2017/159.
- [12] A. Toniolo, T. Norman, A. Etuk, F. Cerutti, R. Ouyang, M. Srivastava, N. Oren, T. Dropps, J. Allen, P. Sullivan, Supporting reasoning with different types of evidence in intelligence analysis, in: *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, Vol. 2, 2015, pp. 781–789.
- [13] F. Cerutti, M. Vallati, M. Giacomin, On the impact of configuration on abstract argumentation automated reasoning, *International Journal of Approximate Reasoning* 92 (2018) 120 – 138. doi:<https://doi.org/10.1016/j.ijar.2017.10.002>.
URL <http://www.sciencedirect.com/science/article/pii/S0888613X16303085>
- [14] B. Liao, L. Lei, J. Dai, Computing Preferred Labellings by Exploiting SCCs and Most Sceptically Rejected Arguments, in: *Second International Workshop on Theory and Applications of Formal Argumentation (TAFAs-13)*, 2013, pp. 194–208.

- [15] F. Cerutti, M. Giacomin, M. Vallati, M. Zanella, An SCC recursive meta-algorithm for computing preferred labellings in abstract argumentation, in: Proc. of the 14th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2014), 2014, pp. 42–51.
- [16] W. Faber, M. Vallati, F. Cerutti, M. Giacomin, Solving set optimization problems by cardinality optimization with an application to argumentation, in: ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016), 2016, pp. 966–973.
- [17] D. H. Wolpert, W. G. Macready, No free lunch theorems for optimization, IEEE transactions on evolutionary computation 1 (1) (1997) 67–82.
- [18] F. Hutter, M. Lindauer, A. Balint, S. Bayless, H. H. Hoos, K. Leyton-Brown, The configurable SAT solver challenge (CSSC), Artif. Intell. 243 (2017) 1–25.
- [19] S. Kadioglu, Y. Malitsky, M. Sellmann, K. Tierney, Isac–instance-specific algorithm configuration, in: Proceedings of the 2010 conference on ECAI 2010: 19th European Conference on Artificial Intelligence, 2010, pp. 751–756.
- [20] C. Fawcett, H. H. Hoos, Analysing differences between algorithm configurations through ablation, J. Heuristics 22 (4) (2016) 431–458.
- [21] F. Cerutti, I. Tachmazidis, M. Vallati, S. Batsakis, M. Giacomin, G. Antoniou, Exploiting parallelism for hard problems in abstract argumentation, in: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA., 2015, pp. 1475–1481.
- [22] P. Baroni, F. Cerutti, P. E. Dunne, M. Giacomin, Automata for infinite argumentation structures, Artificial Intelligence 203 (0) (2013) 104 – 150.
- [23] B. Verheij, Two approaches to dialectical argumentation: admissible sets and argumentation stages, in: J.-J. C. Meyer, L. C. van der Gaag (Eds.), NAIC’96. Proceedings of the Eighth Dutch Conference on Artificial Intelligence, Vol. 96, 1996, pp. 357–368.
- [24] B. Verheij, Artificial argument assistants for defeasible argumentation, Artificial Intelligence 150 (1) (2003) 291–324. doi:10.1016/S0004-3702(03)00107-3.
- [25] M. Caminada, On the issue of reinstatement in argumentation, in: Proceedings of JELIA 2006, 2006, pp. 111–123.
- [26] M. Caminada, G. Pigozzi, On judgment aggregation in abstract argumentation, Autonomous Agents and Multi-Agent Systems 22 (1) (2011) 64–102.

- [27] M. Thimm, S. Villata, F. Cerutti, N. Oren, H. Strass, M. Vallati, Summary Report of The First International Competition on Computational Models of Argumentation, *AI Magazine*.
URL <http://eprints.hud.ac.uk/26063/>
- [28] Y. Yu, P. Subramanyan, N. Tsiskaridze, S. Malik, All-SAT Using Minimal Blocking Clauses, in: 2014 27th International Conference on VLSI Design and 2014 13th International Conference on Embedded Systems, IEEE, 2014, pp. 86–91. doi:10.1109/VLSID.2014.22.
URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6733111>
- [29] K. L. McMillan, Applying SAT Methods in Unbounded Symbolic Model Checking, in: *Proceedings of the 14th International Conference on Computer Aided Verification*, Springer, Berlin, Heidelberg, 2002, pp. 250–264.
URL http://link.springer.com/10.1007/3-540-45657-0{_}19
- [30] J. Brauer, A. King, J. Kriener, Existential Quantification As Incremental SAT, in: *Proceedings of the 23rd International Conference on Computer Aided Verification*, CAV’11, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 191–207.
- [31] M. Caminada, D. M. Gabbay, A logical account of formal argumentation, *Studia Logica* (Special issue: new ideas in argumentation theory) 93 (2–3) (2009) 109–145.
- [32] W. Dvůrák, M. Järvisalo, J. P. Wallner, S. Woltran, Complexity-sensitive decision procedures for abstract argumentation, in: *Proceedings of KR 2012*, AAAI Press, 2012, pp. 54–64.
- [33] W. Dvůrák, M. Järvisalo, J. P. Wallner, S. Woltran, Complexity-sensitive decision procedures for abstract argumentation, *Artificial Intelligence* 206 (2014) 53–78.
URL <http://www.sciencedirect.com/science/article/pii/S0004370213001069>
- [34] M. Vallati, L. Chrupa, M. Grzes, T. L. McCluskey, M. Roberts, S. Sanner, The 2014 international planning competition: Progress and trends, *AI Magazine* 36 (3) (2015) 90–98.
URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/2571>
- [35] M. Vallati, L. Chrupa, T. L. McCluskey, What you always wanted to know about the deterministic part of the international planning competition (IPC) 2014 (but were too afraid to ask), *Knowledge Eng. Review* 33 (2018) e3.
- [36] N. Eén, N. Sörensson, *An Extensible SAT-solver*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 502–518.
URL http://dx.doi.org/10.1007/978-3-540-24605-3{_}37

- [37] G. Audemard, L. Simon, Lazy clause exchange policy for parallel sat solvers, in: Theory and Applications of Satisfiability Testing–SAT 2014, 2014, pp. 197–205.
- [38] M. R. Krom, The decision problem for a class of first-order formulas in which all disjunctions are binary, *Mathematical Logic Quarterly* 13 (1-2) (1967) 15–20. doi:10.1002/malq.19670130104.
URL <http://dx.doi.org/10.1002/malq.19670130104>
- [39] N. Eén, A. Biere, Effective preprocessing in SAT through variable and clause elimination, in: Theory and Applications of Satisfiability Testing, 8th International Conference, SAT 2005, St. Andrews, UK, June 19-23, 2005, Proceedings, 2005, pp. 61–75.
- [40] A. Biere, M. Heule, H. van Maaren, T. Walsh, Conflict-driven clause learning sat solvers, *Handbook of Satisfiability, Frontiers in Artificial Intelligence and Applications* (2009) 131–153.
- [41] S. Hölldobler, N. Manthey, A. Saptawijaya, Improving resource-unaware sat solvers, in: International Conference on Logic for Programming Artificial Intelligence and Reasoning, Springer, 2010, pp. 519–534.
- [42] O. Fourdrinoy, É. Grégoire, B. Mazure, L. Saïs, Eliminating redundant clauses in sat instances, in: Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 2007, pp. 71–83.
- [43] P. Besnard, S. Doutre, Checking the acceptability of a set of arguments, in: Proceedings of 10th International Workshop on Non-Monotonic Reasoning (NMR 2004), 2004, pp. 59–64.
- [44] B. Fazzinga, S. Flesca, F. Parisi, On efficiently estimating the probability of extensions in abstract argumentation frameworks, *International Journal of Approximate Reasoning* 69 (2016) 106 – 132. doi:<https://doi.org/10.1016/j.ijar.2015.11.009>.
URL <http://www.sciencedirect.com/science/article/pii/S0888613X15001760>
- [45] D. S. Johnson, C. H. Papadimitriou, M. Yannakakis, On generating all maximal independent sets, *Information Processing Letters* 27 (3) (1988) 119–123.
- [46] F. Cerutti, M. Vallati, M. Giacomini, Where are we now? state of the art and future trends of solvers for hard argumentation problems, in: P. Baroni, T. Gordon, T. Scheffler (Eds.), Proceedings of 6th International Conference on Computational Models of Argument, COMMA 2016, Vol. 287 of Frontiers in Artificial Intelligence and Applications, IOS Press, 2016, pp. 207–218. doi:10.3233/978-1-61499-686-6-207.

- [47] M. Vallati, F. Cerutti, M. Giacomin, On the combination of argumentation solvers into parallel portfolios, in: *AI 2017: Advances in Artificial Intelligence - 30th Australasian Joint Conference*, 2017, pp. 315–327.
- [48] B. Hurley, L. Kotthoff, Y. Malitsky, D. Mehta, B. O’Sullivan, Advanced portfolio techniques, in: *Data Mining and Constraint Programming - Foundations of a Cross-Disciplinary Approach*, 2016, pp. 191–225.
- [49] F. Hutter, L. Xu, H. H. Hoos, K. Leyton-Brown, Algorithm runtime prediction: Methods & evaluation, *Artif. Intell.* 206 (2014) 79–111.
- [50] M. Lindauer, H. H. Hoos, K. Leyton-Brown, T. Schaub, Automatic construction of parallel portfolios via algorithm configuration, *Artif. Intell.* 244 (2017) 272–290.
- [51] T. Balyo, C. Sinz, Parallel satisfiability, in: *Handbook of Parallel Constraint Reasoning.*, 2018, pp. 3–29.
- [52] F. Cerutti, P. E. Dunne, M. Giacomin, M. Vallati, Computing preferred extensions in abstract argumentation: A sat-based approach, in: E. Black, S. Modgil, N. Oren (Eds.), *Theory and Applications of Formal Argumentation*, Vol. 8306 LNAI, Springer Verlag, Berlin, Heidelberg, 2014, pp. 176–193.
- [53] M. Thimm, S. Villata, The first international competition on computational models of argumentation: Results and analysis, *Artificial Intelligence* 252 (2017) 267–294.
- [54] G. Charwat, W. Dvůrák, S. A. Gaggl, J. P. Wallner, S. Woltran, Methods for solving reasoning problems in abstract argumentation: A survey, *Artificial Intelligence* 220 (2015) 28 – 63. doi:<https://doi.org/10.1016/j.artint.2014.11.008>.
URL <http://www.sciencedirect.com/science/article/pii/S0004370214001404>
- [55] F. Cerutti, S. A. Gaggl, M. Thimm, J. P. Wallner, Foundations of implementations for formal argumentation, in: P. Baroni, D. Gabbay, M. Giacomin, L. van der Torre (Eds.), *Handbook of Formal Argumentation*, College Publications, 2018, Ch. 15.
- [56] F. Cerutti, S. A. Gaggl, M. Thimm, J. P. Wallner, Foundations of implementations for formal argumentation, *IfCoLog Journal of Logics and their Applications* 4 (8) (2017) 2623–2706.
- [57] W. Dvůrák, R. Pichler, S. Woltran, Towards fixed-parameter tractable algorithms for abstract argumentation, *Artificial Intelligence* 186 (2012) 1 – 37. doi:<https://doi.org/10.1016/j.artint.2012.03.005>.
URL <http://www.sciencedirect.com/science/article/pii/S0004370212000264>

- [58] P. Baroni, M. Giacomin, G. Guida, SCC-recursiveness: a general schema for argumentation semantics, *Artificial Intelligence* 168 (1-2) (2005) 165–210.
- [59] S. Greco, F. Parisi, Efficient computation of deterministic extensions for dynamic abstract argumentation frameworks, in: *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, 2016, pp. 1668–1669. doi:10.3233/978-1-61499-672-9-1668.
URL <https://doi.org/10.3233/978-1-61499-672-9-1668>
- [60] S. Greco, F. Parisi, Incremental computation of deterministic extensions for dynamic argumentation frameworks, in: *Logics in Artificial Intelligence - 15th European Conference, JELIA 2016, Larnaca, Cyprus, November 9-11, 2016, Proceedings*, 2016, pp. 288–304.
- [61] G. Alfano, S. Greco, F. Parisi, Efficient computation of extensions for dynamic abstract argumentation frameworks: An incremental approach, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017*, pp. 49–55. doi:10.24963/ijcai.2017/8.
URL <https://doi.org/10.24963/ijcai.2017/8>
- [62] D. Niu, L. Liu, S. Lü, A new stochastic local search approach for computing preferred extensions of abstract argumentation, in: *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, 2016, pp. 1652–1653. doi:10.3233/978-1-61499-672-9-1652.
URL <https://doi.org/10.3233/978-1-61499-672-9-1652>
- [63] J. Wallner, G. Weissenbacher, S. Woltran, Advanced sat techniques for abstract argumentation, in: *Proceedings of CLIMA2013*, 2013, pp. 138–154.

Appendix A. Proofs

Theorem 2. Given an *AF* $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ and $\mathbf{a} \in \mathcal{A}$:

- $\{\text{in}(\mathcal{L}ab) \mid \mathcal{L}ab \in \mathbf{EL-PR}(\Gamma)\} = \mathcal{E}_{\text{PR}}(\Gamma)$ (Algorithm 5);
- $\{\text{in}(\mathcal{L}ab) \mid \mathcal{L}ab \in \mathbf{EL-PR-withST}(\Gamma)\} = \mathcal{E}_{\text{PR}}(\Gamma)$ (Algorithm 6);
- $\mathbf{DC-PR}(\Gamma, \mathbf{a}) \equiv (\exists S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 7);
- $\mathbf{DS-PR-straightforward}(\Gamma, \mathbf{a}) \equiv (\forall S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 8);
- $\mathbf{DS-PR}(\Gamma, \mathbf{a}) \equiv (\forall S \in \mathcal{E}_{\text{PR}}(\Gamma), \mathbf{a} \in S)$ (Algorithm 9).

Proof. Considering Algorithm 5, we have to show that $\mathbf{EL-PR}(\Gamma) = \mathfrak{L}_{\text{PR}}(\Gamma)$ (the desired result then follows by Proposition 4). Let us first focus on the inner loop (lines 3-5) and let us indicate as $\mathcal{L}ab^0$ the complete labelling assigned to the variable $\mathcal{L}ab$ in line 2 before the inner loop execution. According to the definition of $\mathbf{FindCL}^{\supseteq}$, the labellings assigned to the variable $\mathcal{L}ab$ (line 4) at each iteration of the inner loop satisfy the conditions $\mathcal{L}ab^0 \sqsubseteq \mathcal{L}ab^1 \sqsubseteq \dots$, where each $\mathcal{L}ab^i$ is a complete labelling. By transitivity and anti-symmetry of \sqsubseteq , all labellings $\mathcal{L}ab^i$ are distinct, and taking into account that Γ is finite (and so is the number of its possible labellings) this entails that the corresponding sequence is finite as well and the inner loop terminates. In particular, the loop terminates when $\mathbf{FindCL}^{\supseteq}(\Gamma, \mathcal{L}ab) = \perp$, which by the definition of $\mathbf{FindCL}^{\supseteq}$ entails that the labelling in $\mathcal{L}ab$ is maximal w.r.t. \sqsubseteq , i.e. a preferred labelling on the basis of Proposition 3. Summing up, the execution of lines 3-6 starts from a complete labelling $\mathcal{L}ab$ and includes in \mathcal{L}_p (line 6) a preferred labelling $\mathcal{L}ab'$ such that $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$. Let us now consider the outer loop (lines 2-7). First, the loop terminates, since as shown below the labelling $\mathcal{L}ab$ added to \mathcal{L}_p in line 6 does not belong to \mathcal{L}_p : taking into account that Γ is finite, \mathcal{L}_p cannot grow indefinitely. To see that $\mathcal{L}ab \notin \mathcal{L}_p$ before the execution of line 6, note that $\mathbf{FindCL}^{\sqsupseteq}(\Gamma, \mathcal{L}_p)$ returns in line 6 a complete labelling which is not less committed than any of the labellings in \mathcal{L}_p , thus by transitivity of \sqsubseteq the labelling $\mathcal{L}ab$ added to \mathcal{L}_p in line 6 does not belong to \mathcal{L}_p . Now, according to the considerations above, line 6 adds to \mathcal{L}_p a preferred labelling, thus in line 2 it always holds that $\mathcal{L}_p \subseteq \mathfrak{L}_{\text{PR}}(\Gamma)$. Since the loop terminates when $\mathbf{FindCL}^{\sqsupseteq}(\Gamma, \mathcal{L}_p)$ returns \perp in line 2, by the definition of $\mathbf{FindCL}^{\sqsupseteq}$ it holds that $(\mathfrak{L}_{\text{PR}}(\Gamma) \setminus \mathcal{L}_p) \subseteq \emptyset$, i.e. $\mathfrak{L}_{\text{PR}}(\Gamma) \subseteq \mathcal{L}_p$. As a consequence, $\mathcal{L}_p = \mathfrak{L}_{\text{PR}}(\Gamma)$ in line 8, as desired.

Turning to Algorithm 6, its correctness follows from that of Algorithm 5 taking into account that $\mathfrak{L}_{\text{ST}}(\Gamma) \subseteq \mathfrak{L}_{\text{PR}}(\Gamma)$.

To show the correctness of Algorithm 7, we distinguish two cases. If \mathbf{a} is credulously accepted then there is a preferred labelling where \mathbf{a} is \mathbf{in} , and since a preferred labelling is by definition a complete labelling $\mathbf{FindCL}_{\mathbf{a}}(\Gamma, \mathbf{in})$ is able to return a labelling in line 1, and the algorithm correctly returns \top . In the case where \mathbf{a} is not credulously accepted, there is no complete labelling where \mathbf{a} is \mathbf{in} , since according to Proposition 3 this would entail the presence of a preferred labelling where \mathbf{a} is \mathbf{in} . As a consequence, the algorithm correctly returns \perp .

The correctness of Algorithm 8 follows from the above proof concerning Algorithm 5. In particular, we know that the algorithm terminates and that during the execution of the loop of lines 2-10 the variable $\mathcal{L}ab$ in lines 6-9 iterates over the set of preferred labellings. Thus, if \mathbf{a} is not skeptically accepted then there is an iteration of the loop such that the algorithm returns \perp in line 7, otherwise all preferred labellings assign to \mathbf{a} the label \mathbf{in} and the loop terminates, thus the algorithm correctly returns \top in line 11.

Let us finally turn to Algorithm 9. It is easy to see that if the algorithm terminates then it either returns \top or \perp . We thus show that the algorithm always terminates and whenever it returns a value, this is correct.

As to termination, there are two nested loops in the algorithm. Focusing first

on the inner loop (lines 6-8) and denoting as $\mathcal{L}ab^0$ the labelling in the variable $\mathcal{L}ab$ before the execution of the loop, by the definition of $\text{FindCL}_{\mathbf{a}}^{\overline{\perp}}$ the values in variable $\mathcal{L}ab$ at line 6 identify a sequence of labellings $\mathcal{L}ab^0 \sqsubset \mathcal{L}ab^1 \sqsubset \dots$. As in the proof concerning Algorithm 5, transitivity and anti-symmetry of \sqsubset guarantee that the sequence includes distinct elements, and since the argumentation framework is finite then the sequence is finite as well. As a consequence, after a finite number of iterations the inner loop terminates with a labelling in $\mathcal{L}ab$ such that $\mathcal{L}ab^0 \sqsubseteq \mathcal{L}ab$ and $\mathcal{L}ab \in \mathcal{L}_{mu}$ (i.e. $\mathcal{L}ab$ is a maximal complete labelling such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$). As to the loop of lines 5-13, we prove that at each iteration line 12 increases the cardinality of the set in \mathcal{L}_c , and again, since Γ is finite, this entails termination of the loop. To prove that $\mathcal{L}ab$ before the execution of line 12 does not belong to \mathcal{L}_c , let us consider the case where $\text{FindCLU}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}_c)$ in line 5 returns a labelling $\mathcal{L}ab^0$ (i.e. different than \perp) which is assigned to the variable $\mathcal{L}ab$. By the definition of $\text{FindCLU}_{\mathbf{a}}^{\#}$, it holds that $\forall \mathcal{L}ab' \in \mathcal{L}_c, \mathcal{L}ab^0 \not\sqsubseteq \mathcal{L}ab'$, i.e. $\mathcal{L}ab^0$ is not less or equally committed w.r.t. a labelling of \mathcal{L}_c . As shown above, the inner loop obtains a labelling $\mathcal{L}ab$ such that $\mathcal{L}ab^0 \sqsubseteq \mathcal{L}ab$, thus it can not be the case that $\mathcal{L}ab \in \mathcal{L}_c$.

We now show that the value returned by the algorithm is always correct. Focusing first on lines 1-3, according to the definition of $\text{FindCL}_{\mathbf{a}}$ line 2 is executed if there is no complete labelling where \mathbf{a} is **in** or there is a complete labelling where \mathbf{a} is **out**. It is easy to see that in both cases \mathbf{a} is not skeptically justified, thus the returned value \perp is correct. In particular, since preferred labellings are also complete labellings, in the first case there is no preferred labelling where \mathbf{a} is **in**, and taking into account that a preferred labelling always exists (see e.g. [2]) this entails that \mathbf{a} is not skeptically justified. In the second case, according to Proposition 3 there is a preferred labelling where \mathbf{a} is **out**, and again \mathbf{a} is not skeptically justified.

Let us now consider the case where the algorithm enters line 4. Note that by the considerations above there is no preferred labelling where \mathbf{a} is **out**. Focusing on lines 9-10, let us show that if line 10 is executed then \mathbf{a} is not skeptically justified. According to the definition of $\text{FindCL}_{\mathbf{a}}^{\overline{\perp}}$, this happens if there is no complete labelling $\mathcal{L}ab'$ such that $\mathcal{L}ab'(\mathbf{a}) = \text{in}$ and $\mathcal{L}ab \sqsubset \mathcal{L}ab'$. Since $\mathcal{L}ab \in \mathcal{L}_{mu}$ and there are no preferred labellings where \mathbf{a} is **out**, in this case $\mathcal{L}ab$ is a maximal complete labelling, i.e. a preferred labelling by Proposition 3. As a consequence, there is a preferred labelling where \mathbf{a} is **undec** entailing that it is not skeptically justified.

To complete the proof, we first show that, in line 5, $\mathcal{L}_c \subseteq \mathcal{L}_{mu}$ and $\mathcal{L}_c \cap \mathfrak{L}_{PR}(\Gamma) = \emptyset$. This is obviously true before the first iteration of the outer loop, since $\mathcal{L}_c = \emptyset$. It is also true after each iteration of the loop, since we have shown above that the labelling $\mathcal{L}ab$ added to \mathcal{L}_c in line 12 belongs to \mathcal{L}_{mu} , and taking into account lines 9-11 there is a complete labelling $\mathcal{L}ab'$ such that $\mathcal{L}ab \sqsubset \mathcal{L}ab'$, i.e. $\mathcal{L}ab$ is not maximal w.r.t. \sqsubseteq and thus it is not a preferred labelling. Now, the algorithm enters line 14 only if $\text{FindCLU}_{\mathbf{a}}^{\#}(\Gamma, \mathcal{L}_c)$ in line 5 returns \perp . According to the definition of $\text{FindCLU}_{\mathbf{a}}^{\#}$, this entails that $\{\mathcal{L}ab' \in \mathfrak{L}_{PR}(\Gamma) \mid \mathcal{L}ab'(\mathbf{a}) = \text{undec}\} \subseteq \mathcal{L}_c$, where $\mathcal{L}_c \cap \mathfrak{L}_{PR}(\Gamma) = \emptyset$. Then $\{\mathcal{L}ab' \in \mathfrak{L}_{PR}(\Gamma) \mid \mathcal{L}ab'(\mathbf{a}) = \text{undec}\}$ is empty, and taking into account that there

are no preferred labellings where \mathbf{a} is **out**, we conclude that \mathbf{a} is skeptically justified, and the value returned in line 14 is correct. \square

Proposition 5. Let $\langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework. A total function $\mathcal{L}ab : \mathcal{A} \mapsto \{\mathbf{in}, \mathbf{out}, \mathbf{undec}\}$ is a complete labelling iff it satisfies any of the following conjunctive constraints for any $\mathbf{a} \in \mathcal{A}$:

$$\begin{aligned} C_{\exists}^{\leftrightarrow} &= C_{\mathbf{in}}^{\leftrightarrow} \wedge C_{\mathbf{out}}^{\leftrightarrow} \\ C_{\exists}^{\rightarrow} &= C_{\mathbf{in}}^{\leftrightarrow} \wedge C_{\mathbf{undec}}^{\leftrightarrow} \\ C_{\exists}^{\leftarrow} &= C_{\mathbf{out}}^{\leftrightarrow} \wedge C_{\mathbf{undec}}^{\leftrightarrow} \\ C_{\exists}^{\rightarrow} &= C_{\mathbf{in}}^{\rightarrow} \wedge C_{\mathbf{out}}^{\rightarrow} \wedge C_{\mathbf{undec}}^{\rightarrow} \\ C_{\exists}^{\leftarrow} &= C_{\mathbf{in}}^{\leftarrow} \wedge C_{\mathbf{out}}^{\leftarrow} \wedge C_{\mathbf{undec}}^{\leftarrow} \end{aligned}$$

Proof. We prove that each of the 5 conjunctive constraints is equivalent to $C_{\exists}^{\leftrightarrow} = C_{\mathbf{in}}^{\leftrightarrow} \wedge C_{\mathbf{out}}^{\leftrightarrow} \wedge C_{\mathbf{undec}}^{\leftrightarrow}$, i.e. the constraint expressed in Definition 7. All of them are obviously entailed by $C_{\exists}^{\leftrightarrow}$, thus we have to prove that each of them entails $C_{\exists}^{\leftrightarrow}$.

As to $C_{\exists}^{\rightarrow}$, C_{\exists}^{\leftarrow} and $C_{\exists}^{\leftrightarrow}$, the result derives from the fact that $\mathcal{L}ab$ is a function. In particular, let us consider an argument \mathbf{a} which satisfies $C_{\exists}^{\rightarrow}$. Let us first prove $C_{\mathbf{undec}}^{\leftarrow}$. If $\forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) \neq \mathbf{in} \wedge \exists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{c}) = \mathbf{undec}$, then $\mathcal{L}ab(\mathbf{a}) \neq \mathbf{in}$ because of $C_{\mathbf{in}}^{\rightarrow}$, and $\mathcal{L}ab(\mathbf{a}) \neq \mathbf{out}$ because of $C_{\mathbf{out}}^{\rightarrow}$, thus the only possibility is that $\mathcal{L}ab(\mathbf{a}) = \mathbf{undec}$. As to $C_{\mathbf{undec}}^{\rightarrow}$, if $\mathcal{L}ab(\mathbf{a}) = \mathbf{undec}$ then by $C_{\mathbf{out}}^{\leftarrow} \forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) \neq \mathbf{in}$, thus by $C_{\mathbf{in}}^{\leftarrow} \exists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{c}) = \mathbf{undec}$. Overall, $C_{\mathbf{undec}}^{\leftrightarrow}$ is satisfied, entailing $C_{\exists}^{\leftrightarrow}$. The proofs for C_{\exists}^{\leftarrow} and $C_{\exists}^{\leftrightarrow}$ are analogous.

As to $C_{\exists}^{\rightarrow}$, with respect to $C_{\exists}^{\leftrightarrow}$ this constraint does not include the terms $C_{\mathbf{in}}^{\leftarrow}$, $C_{\mathbf{out}}^{\leftarrow}$ and $C_{\mathbf{undec}}^{\leftarrow}$. Here we prove that $C_{\mathbf{in}}^{\leftarrow}$ (and, similarly, $C_{\mathbf{out}}^{\leftarrow}$ and $C_{\mathbf{undec}}^{\leftarrow}$) is indeed satisfied. Let us consider an argument \mathbf{a} such that $\forall \mathbf{b} \in \mathbf{a}^- \mathcal{L}ab(\mathbf{b}) = \mathbf{out}$, and let us reason by contradiction by assuming that $\mathcal{L}ab(\mathbf{a}) \neq \mathbf{in}$. Since $\mathcal{L}ab$ is a function, if $\mathcal{L}ab(\mathbf{a}) \neq \mathbf{in}$ then either $\mathcal{L}ab(\mathbf{a}) = \mathbf{out}$ or $\mathcal{L}ab(\mathbf{a}) = \mathbf{undec}$. If $\mathcal{L}ab(\mathbf{a}) = \mathbf{out}$, from $C_{\mathbf{out}}^{\rightarrow} \exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) = \mathbf{in} \neq \mathbf{out}$. If $\mathcal{L}ab(\mathbf{a}) = \mathbf{undec}$, from $C_{\mathbf{undec}}^{\rightarrow} \exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) = \mathbf{undec} \neq \mathbf{out}$. The proof for $C_{\mathbf{out}}^{\leftarrow}$ and $C_{\mathbf{undec}}^{\leftarrow}$ is similar.

As to C_{\exists}^{\leftarrow} , the proof follows the same line. We prove that $C_{\mathbf{in}}^{\rightarrow}$ (and, similarly, $C_{\mathbf{out}}^{\rightarrow}$ and $C_{\mathbf{undec}}^{\rightarrow}$) is indeed satisfied. Given an argument \mathbf{a} such that $\mathcal{L}ab(\mathbf{a}) = \mathbf{in}$, assume by contradiction that $\exists \mathbf{b} \in \mathbf{a}^- : \mathcal{L}ab(\mathbf{b}) \neq \mathbf{out}$. Since $\mathcal{L}ab$ is a function, either $\mathcal{L}ab(\mathbf{b}) = \mathbf{in}$ or $\mathcal{L}ab(\mathbf{b}) = \mathbf{undec}$. In the first case, $C_{\mathbf{out}}^{\leftarrow}$ entails that $\mathcal{L}ab(\mathbf{a}) = \mathbf{out} \neq \mathbf{in}$. In the second case, either $C_{\mathbf{out}}^{\leftarrow}$ or $C_{\mathbf{undec}}^{\leftarrow}$ applies, i.e. $\mathcal{L}ab(\mathbf{a}) \in \{\mathbf{out}, \mathbf{undec}\}$ thus $\mathcal{L}ab(\mathbf{a}) \neq \mathbf{in}$. Following the same reasoning line, we can prove that also $C_{\mathbf{out}}^{\rightarrow}$ and $C_{\mathbf{undec}}^{\rightarrow}$ hold. \square




Figure A.26: Identifying some weak encodings.

Proposition 6. The following 6 encodings are weak:

$$\begin{aligned}
C_{\exists}^{\rightarrow} &= C_{\text{undec}}^{\leftrightarrow} \wedge C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow} \\
C_{\exists}^{\leftarrow} &= C_{\text{undec}}^{\leftrightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \\
C_{\exists}^{\rightarrow} &= C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{in}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow} \\
C_{\exists}^{\leftarrow} &= C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow} \\
C_{\exists}^{\rightarrow} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\leftarrow} \\
C_{\exists}^{\leftarrow} &= C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\rightarrow}
\end{aligned}$$

Proof. For each encoding, we identify an argumentation framework and a non-complete labelling which satisfies the relevant constraint. In particular, referring to Figure A.26: for $C_{\exists}^{\rightarrow}$, see the labelling $\{(a, \text{out})\}$ of AF_1 ; for C_{\exists}^{\leftarrow} , see the labelling $\{(a, \text{in})\}$ of AF_2 ; for $C_{\exists}^{\rightarrow}$, see the labelling $\{(a, \text{undec})\}$ of AF_1 ; for C_{\exists}^{\leftarrow} , see the labelling $\{(a, \text{undec}), (b, \text{in}), (c, \text{out})\}$ of AF_3 ; for $C_{\exists}^{\rightarrow}$, see the labelling $\{(a, \text{in}), (b, \text{undec}), (c, \text{undec})\}$ of AF_4 ; for C_{\exists}^{\leftarrow} , see the labelling $\{(a, \text{undec}), (b, \text{out}), (c, \text{in})\}$ of AF_3 . \square

Theorem 3. All the encodings of cardinality 0, 1, and 2 are weak. Among the encodings of cardinality 3, $C_{\exists}^{\rightarrow} = C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ and $C_{\exists}^{\leftarrow} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$ are correct and non-redundant, the other 18 encodings are weak. Among the encodings of cardinality 4, $C_{\exists}^{\rightarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftrightarrow}$, $C_{\exists}^{\leftarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$ and $C_{\exists}^{\rightarrow} = C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow}$ are correct and non-redundant, the 6 encodings $C_{\exists}^{\rightarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$, $C_{\exists}^{\leftarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$, $C_{\exists}^{\rightarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$, $C_{\exists}^{\leftarrow} = C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$, $C_{\exists}^{\rightarrow} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow}$, $C_{\exists}^{\leftarrow} = C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow}$ are correct and redundant, the other 6 encodings are weak. All the encodings of cardinality 5 and 6 are correct and redundant.

Proof. As to the first claim, it is easy to see that any encoding having cardinality 0, 1 and 2 is a strict subset of at least one (weak) encoding introduced in Proposition 6, thus it is obviously weak too. As to the encodings having cardinality 3, $(C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow})$ and $(C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow})$ are correct by Proposition 5, and they are non-redundant since any strict subset has a cardinality strictly lower than 3 (thus it is weak as shown above). The remaining 18 encodings of cardinality 3 can take one of the following two forms: (i) 12 encodings include $C_{\text{in}}^{\leftrightarrow}$, $C_{\text{out}}^{\leftrightarrow}$ or $C_{\text{undec}}^{\leftrightarrow}$ and another single term; (ii) 6 encodings include two “left” and one “right” terms, or vice versa. In both cases, it is easy to check that any of these encodings is a strict subset of one of the weak encodings of Proposition 6. As to the encodings of cardinality 4, $(C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{out}}^{\leftrightarrow})$, $(C_{\text{out}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow})$ and

$(C_{\text{in}}^{\leftrightarrow} \wedge C_{\text{undec}}^{\leftrightarrow})$ are correct by Proposition 5, and they are non-redundant since they do not contain $(C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow})$ nor $(C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow})$, thus any subset is weak according to the considerations above concerning the encodings of cardinality 3. Moreover, the 6 encodings $C_{\exists}^{\rightarrow}, C_{\exists}^{\leftarrow}, C_{\exists}^{\rightarrow}, C_{\exists}^{\leftarrow}, C_{\exists}^{\rightarrow}$ and C_{\exists}^{\leftarrow} are supersets of $(C_{\text{in}}^{\rightarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{undec}}^{\rightarrow})$ or $(C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\leftarrow} \wedge C_{\text{undec}}^{\leftarrow})$, and thus correct and redundant, while the other 6 encodings are the weak ones identified in Proposition 6. Finally, all encodings of cardinality 5 and 6 contain at least one of the correct encodings of Proposition 5, thus they are correct and redundant. \square

Lemma 1. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab_1, \mathcal{L}ab_2 \in \mathfrak{L}(\Gamma)$ complete labellings. $\exists \mathbf{a} \in \mathcal{A}$ such that $\mathcal{L}ab_1(\mathbf{a}) = \text{out}$ and $\mathcal{L}ab_2(\mathbf{a}) \neq \text{out}$ iff $\exists \mathbf{b} \in \mathcal{A}$ such that $\mathcal{L}ab_1(\mathbf{b}) = \text{in}$ and $\mathcal{L}ab_2(\mathbf{b}) \neq \text{in}$.

Proof. As to the first direction of the proof, assume that $\mathcal{L}ab_1(\mathbf{a}) = \text{out}$ and $\mathcal{L}ab_2(\mathbf{a}) \neq \text{out}$. According to Definition 7, $\exists \mathbf{b} \in \mathbf{a}^-$ such that $\mathcal{L}ab_1(\mathbf{b}) = \text{in}$. Since $\mathcal{L}ab_2(\mathbf{a}) \neq \text{out}$, according to Definition 7 $\nexists \mathbf{c} \in \mathbf{a}^- : \mathcal{L}ab_2(\mathbf{c}) = \text{in}$, thus it must be the case that $\mathcal{L}ab_2(\mathbf{b}) \neq \text{in}$.

As to the other direction of the proof, assume that $\mathcal{L}ab_1(\mathbf{b}) = \text{in}$ and $\mathcal{L}ab_2(\mathbf{b}) \neq \text{in}$. We consider two possible cases for $\mathcal{L}ab_2(\mathbf{b})$. If $\mathcal{L}ab_2(\mathbf{b}) = \text{out}$, according to Definition 7 $\exists \mathbf{a} \in \mathbf{b}^-$ such that $\mathcal{L}ab_2(\mathbf{a}) = \text{in}$. Since $\mathcal{L}ab_1(\mathbf{b}) = \text{in}$, according to Definition 7 $\forall \mathbf{c} \in \mathbf{b}^-, \mathcal{L}ab_1(\mathbf{c}) = \text{out}$, and this also holds for \mathbf{a} . If, in the other case, $\mathcal{L}ab_2(\mathbf{b}) = \text{undec}$, according to Definition 7 $\exists \mathbf{a} \in \mathbf{b}^-$ such that $\mathcal{L}ab_2(\mathbf{a}) = \text{undec}$, and again it must also hold that $\mathcal{L}ab_1(\mathbf{a}) = \text{out}$. \square

Proposition 7. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab, \mathcal{L}ab' \in \mathfrak{L}(\Gamma)$ stable labellings. All of the following statements are equivalent to $\mathcal{L}ab \neq \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$
3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

Proof. Taking into account that stable labellings do not assign the label **undec** to any argument, the condition $\mathcal{L}ab \neq \mathcal{L}ab'$ is equivalent to $\exists \mathbf{b} \in \text{out}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) = \text{in}$ or $\exists \mathbf{c} \in \text{in}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{c}) = \text{out}$. According to Lemma 1, this entails all of the three conditions. The fact that each of the three conditions entails $\mathcal{L}ab \neq \mathcal{L}ab'$ is obvious. \square

Proposition 8. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab' \in \mathfrak{L}_{\text{PR}}(\Gamma)$ a preferred labelling. Then, $\forall \mathcal{L}ab \in \mathfrak{L}_{\text{CO}}(\Gamma)$, all of the following statements are equivalent to $\mathcal{L}ab \not\sqsubseteq \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

Proof. Each of the three conditions obviously entails $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$, therefore we can focus on the other direction of the proof.

Let us first prove that if $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$ then the first condition holds.

Letting $S \equiv \text{in}(\mathcal{L}ab) \setminus \text{in}(\mathcal{L}ab')$, it is easy to see that $S \neq \emptyset$. Indeed, $S = \emptyset$ would entail $\text{in}(\mathcal{L}ab) \subseteq \text{in}(\mathcal{L}ab')$ and according to Definition 7 we would also have $\text{out}(\mathcal{L}ab) \subseteq \text{out}(\mathcal{L}ab')$, i.e. $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$.

We now prove that $S \cap \text{out}(\mathcal{L}ab') \neq \emptyset$, i.e. $\exists \mathbf{b} \in \text{out}(\mathcal{L}ab')$ such that $\mathbf{b} \in S \subseteq \text{in}(\mathcal{L}ab)$, entailing the first condition as desired.

For the sake of contradiction, let us assume $S \cap \text{out}(\mathcal{L}ab') = \emptyset$. We prove that $\text{in}(\mathcal{L}ab') \cup S$ is admissible. Let us note that, by Proposition 4, $\text{in}(\mathcal{L}ab)$ is a complete extension and $\text{in}(\mathcal{L}ab')$ is a preferred extension, i.e. maximal admissible. First, we prove that $\text{in}(\mathcal{L}ab') \cup S$ is conflict-free. Indeed, if $S \rightarrow \text{in}(\mathcal{L}ab')$ then $\text{in}(\mathcal{L}ab') \rightarrow S$, since $\text{in}(\mathcal{L}ab')$ is admissible. If $\text{in}(\mathcal{L}ab') \rightarrow S$ then according to Definition 7 we would have $S \cap \text{out}(\mathcal{L}ab') \neq \emptyset$, but this is not the case. Second, we show that $\text{in}(\mathcal{L}ab') \cup S$ defends all of its arguments. In particular, $\forall \mathbf{a} \in \mathcal{A}$ with $\mathbf{a} \rightarrow \text{in}(\mathcal{L}ab') \cup S$ we distinguish two cases:

1. if $\mathbf{a} \rightarrow \text{in}(\mathcal{L}ab')$, then $\text{in}(\mathcal{L}ab') \rightarrow \mathbf{a}$ since $\text{in}(\mathcal{L}ab')$ is admissible.
2. if $\mathbf{a} \rightarrow S$ then $\mathbf{a} \rightarrow \text{in}(\mathcal{L}ab)$, since $S \equiv \text{in}(\mathcal{L}ab) \setminus \text{in}(\mathcal{L}ab')$. This entails that $\text{in}(\mathcal{L}ab) \rightarrow \mathbf{a}$, since $\text{in}(\mathcal{L}ab)$ is a complete extension and thus admissible. Now, if $\mathcal{L}ab'(\mathbf{a}) = \text{out}$ then by definition $\text{in}(\mathcal{L}ab') \rightarrow \mathbf{a}$, otherwise $\text{in}(\mathcal{L}ab) \setminus \text{in}(\mathcal{L}ab') \rightarrow \mathbf{a}$, i.e. $S \rightarrow \mathbf{a}$.

Summing up, $\forall \mathbf{a} \in \mathcal{A}$ with $\mathbf{a} \rightarrow (\text{in}(\mathcal{L}ab') \cup S)$, $(\text{in}(\mathcal{L}ab') \cup S) \rightarrow \mathbf{a}$ holds, i.e. $(\text{in}(\mathcal{L}ab') \cup S)$ is admissible. However, since $S \neq \emptyset$ and by definition $S \cap \text{in}(\mathcal{L}ab') = \emptyset$, it holds that $\text{in}(\mathcal{L}ab') \subsetneq (\text{in}(\mathcal{L}ab') \cup S)$. Quod est absurdum because $\text{in}(\mathcal{L}ab') \cup S$ is admissible contradicting the fact that $\text{in}(\mathcal{L}ab')$ is a preferred extension.

Turning to the second condition, we know from the first one that $\exists \mathbf{b} \in \mathcal{A}$ such that $\mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$. According to Definition 7, $\mathcal{L}ab'(\mathbf{b}) = \text{out}$ entails that $\exists \mathbf{c} \in \mathcal{A}$ (with $\mathbf{c} \rightarrow \mathbf{b}$) such that $\mathcal{L}ab'(\mathbf{c}) = \text{in}$, and $\mathcal{L}ab(\mathbf{b}) = \text{in}$ entails that $\mathcal{L}ab(\mathbf{c}) = \text{out}$.

The third condition obviously follows from the first two conditions. \square

Proposition 9. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework and $\mathcal{L}ab, \mathcal{L}ab' \in \mathcal{L}_{\text{CO}}(\Gamma)$. All of the following statements are equivalent to $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab$:

1. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{b} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) = \text{out}$
2. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{c} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{c}) = \text{in}$
3. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{b} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) = \text{out} \wedge \exists \mathbf{c} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{c}) = \text{in}$

Proof. Each of the three conditions obviously entails $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab$. Let us then assume that $\mathcal{L}ab' \not\sqsubseteq \mathcal{L}ab$, i.e. $\text{in}(\mathcal{L}ab') \subseteq \text{in}(\mathcal{L}ab) \wedge \text{out}(\mathcal{L}ab') \subseteq \text{out}(\mathcal{L}ab) \wedge \exists \mathbf{b} \in \text{undec}(\mathcal{L}ab') : \mathcal{L}ab(\mathbf{b}) \in \{\text{out}, \text{in}\}$. We distinguish two cases for $\mathcal{L}ab(\mathbf{b})$.

If $\mathcal{L}ab(\mathbf{b}) = \text{out}$, then the first condition holds. Moreover, by Lemma 1 (with $\mathcal{L}ab_1 = \mathcal{L}ab$ and $\mathcal{L}ab_2 = \mathcal{L}ab'$) there is an argument \mathbf{c} such that $\mathcal{L}ab(\mathbf{c}) = \text{in}$ and $\mathcal{L}ab'(\mathbf{c}) \neq \text{in}$. Since $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab$, it cannot be the case that $\mathcal{L}ab'(\mathbf{c}) = \text{out}$, thus $\mathcal{L}ab'(\mathbf{c}) = \text{undec}$, entailing the second and third conditions.

In the other case, $\mathcal{L}ab(\mathbf{b}) = \text{in}$ and the second condition holds. Similarly to the previous case, by Lemma 1 there is an argument \mathbf{c} such that $\mathcal{L}ab(\mathbf{c}) = \text{out}$ and $\mathcal{L}ab'(\mathbf{c}) \neq \text{out}$. Taking into account that $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab$, it must be the case that $\mathcal{L}ab'(\mathbf{c}) = \text{undec}$, entailing the first and third conditions. \square

Proposition 10. Let $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$ be an argumentation framework, \mathbf{a} an argument in \mathcal{A} and $\mathcal{L}ab' \in \mathfrak{L}(\Gamma)$ a maximal complete labelling such that $\mathcal{L}ab'(\mathbf{a}) = \text{undec}$, i.e. $\nexists \mathcal{L}ab^* \in \mathfrak{L}(\Gamma)$ such that $\mathcal{L}ab^*(\mathbf{a}) = \text{undec}$ and $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab^*$. Then, $\forall \mathcal{L}ab \in \mathfrak{L}_{\text{PR}}(\Gamma)$ such that $\mathcal{L}ab(\mathbf{a}) = \text{undec}$, all of the following statements are equivalent to $\mathcal{L}ab \neq \mathcal{L}ab'$:

1. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$
2. $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$
3. $\exists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ and $\exists \mathbf{c} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{c}) = \text{in} \wedge \mathcal{L}ab(\mathbf{c}) = \text{out}$

Proof. Each of the three conditions obviously entails $\mathcal{L}ab \neq \mathcal{L}ab'$, therefore we can focus on the other direction of the proof.

Let us first prove that if $\mathcal{L}ab \neq \mathcal{L}ab'$ then the first condition holds.

Since $\mathcal{L}ab$ is a preferred labelling, it must be the case that $\mathcal{L}ab \not\sqsubseteq \mathcal{L}ab'$ (otherwise we would have $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$ contradicting maximality w.r.t. \sqsubseteq). This in turn entails that $\text{in}(\mathcal{L}ab) \not\subseteq \text{in}(\mathcal{L}ab')$. Indeed, according to Definition 7, $\text{in}(\mathcal{L}ab) \subseteq \text{in}(\mathcal{L}ab')$ would entail $\text{out}(\mathcal{L}ab) \subseteq \text{out}(\mathcal{L}ab')$, i.e. $\mathcal{L}ab \sqsubseteq \mathcal{L}ab'$, which is not the case.

Let us now consider the set $\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab)$. Since $\text{in}(\mathcal{L}ab) \not\subseteq \text{in}(\mathcal{L}ab')$, it must be the case that $\text{in}(\mathcal{L}ab') \not\subseteq (\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$. Moreover, since $\mathcal{L}ab'(\mathbf{a}) = \text{undec}$ and $\mathcal{L}ab(\mathbf{a}) = \text{undec}$, $\mathbf{a} \notin (\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$ and it is not the case that $(\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab)) \rightarrow \mathbf{a}$ (otherwise either $\mathcal{L}ab'(\mathbf{a}) = \text{out}$ or $\mathcal{L}ab(\mathbf{a}) = \text{out}$). We reason by contradiction, by proving that if $\nexists \mathbf{b} \in \mathcal{A} : \mathcal{L}ab'(\mathbf{b}) = \text{out} \wedge \mathcal{L}ab(\mathbf{b}) = \text{in}$ then $(\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$ is a complete extension. This leads to a contradiction, since by Proposition 4 the labelling $\mathcal{L}ab^* \equiv \text{Ext2Lab}(\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$ would be a complete labelling and, by the conditions above, $\mathcal{L}ab' \sqsubseteq \mathcal{L}ab^*$ and $\mathcal{L}ab^*(\mathbf{a}) = \text{undec}$, contradicting the maximality of $\mathcal{L}ab'$.

Let us then prove that $(\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$ is a complete extension. Let us note that, by Proposition 4, $\text{in}(\mathcal{L}ab)$ is a preferred extension and $\text{in}(\mathcal{L}ab')$ is a complete extension.

First, $(\text{in}(\mathcal{L}ab') \cup \text{in}(\mathcal{L}ab))$ is conflict-free. Indeed, if $\text{in}(\mathcal{L}ab) \rightarrow \text{in}(\mathcal{L}ab')$ then $\text{in}(\mathcal{L}ab') \rightarrow \text{in}(\mathcal{L}ab)$, since $\text{in}(\mathcal{L}ab')$ is admissible. If $\text{in}(\mathcal{L}ab') \rightarrow \text{in}(\mathcal{L}ab)$

then according to Definition 7 there would be an argument \mathbf{b} such that $\mathcal{L}ab(\mathbf{b}) = \mathbf{in}$ and $\mathcal{L}ab'(\mathbf{a}) = \mathbf{out}$, but this has been excluded by assumption.

Second, we show that $(\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$ defends all of its arguments. For all $\mathbf{b} \in \mathcal{A}$ with $\mathbf{b} \rightarrow (\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$, it holds that $\mathbf{b} \rightarrow \mathbf{in}(\mathcal{L}ab')$ or $\mathbf{b} \rightarrow \mathbf{in}(\mathcal{L}ab)$. Since both $\mathbf{in}(\mathcal{L}ab')$ and $\mathbf{in}(\mathcal{L}ab)$ are admissible, there must be an argument $\mathbf{c} \in (\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$ such that $\mathbf{c} \rightarrow \mathbf{b}$.

Finally, to show that $(\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$ is complete we take into account that $\mathbf{in}(\mathcal{L}ab)$ is a preferred extension. Since $\mathbf{in}(\mathcal{L}ab) \subseteq (\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$ and $(\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$ has been proved to be admissible, it must then be the case that $\mathbf{in}(\mathcal{L}ab) = (\mathbf{in}(\mathcal{L}ab') \cup \mathbf{in}(\mathcal{L}ab))$. Since any preferred extension is a complete extension, the conclusion follows.

Turning to the second condition, as in the proof of Proposition 8 we know from the first one that $\exists \mathbf{b} \in \mathcal{A}$ such that $\mathcal{L}ab'(\mathbf{b}) = \mathbf{out} \wedge \mathcal{L}ab(\mathbf{b}) = \mathbf{in}$. According to Definition 7, $\mathcal{L}ab'(\mathbf{b}) = \mathbf{out}$ entails that $\exists \mathbf{c} \in \mathcal{A}$ (with $\mathbf{c} \rightarrow \mathbf{b}$) such that $\mathcal{L}ab'(\mathbf{c}) = \mathbf{in}$, and $\mathcal{L}ab(\mathbf{b}) = \mathbf{in}$ entails that $\mathcal{L}ab(\mathbf{c}) = \mathbf{out}$.

The third condition obviously follows from the first two conditions. \square

Appendix B. Summary of parameters used in algorithms

Parameter	Domain	Algorithm	Comment
C_*	$\mathfrak{C} = \{C\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists, C\exists\exists\}$	All	The set of all the correct encodings as identified by Theorem 3
π_{All}^S	$\{\top, \perp\}$	Alg. 10, 13, 15	Parameter enabling using a AllSAT solver
$\pi_{\text{All};0}^S$	$\{\top, \perp\}$	Alg. 10, 13, 15	When using a AllSAT solver, this parameter forces, for each found labelling, the creation of a blocking clause based on arguments labelled \mathbf{in} , requiring one of them to be labelled \mathbf{out}

$\pi_{All;l}^S$	$\{\top, \perp\}$	Alg. 10, 13, 15	When using a AllSAT solver, this parameter forces, for each found labelling, the creation of a blocking clause based on arguments labelled out , requiring one of them to be labelled in
π_O^S	$\{\top, \perp\}$	Alg. 10, 13, 15	When not using a AllSAT solver, this parameter forces, for each found labelling, the creation of a blocking clause based on arguments labelled in , requiring one of them to be labelled out
π_I^S	$\{\top, \perp\}$	Alg. 10, 13, 15	When not using a AllSAT solver, this parameter forces, for each found labelling, the creation of a blocking clause based on arguments labelled out , requiring one of them to be labelled in
π_S^P	$\{\top, \perp\}$	Alg. 13, 15	Selector between Algorithms 5 and 6 which exploits the knowledge that stable extensions are also preferred
π_{iO}^P	$\{\top, \perp\}$	Alg. 13, 15, 16	Further to Proposition 9, to identify a complete labelling which is strictly more committed w.r.t. a previously found complete labelling, an argument labelled undec should then become out

π_{ii}^P	$\{\top, \perp\}$	Alg. 13, 15, 16	Further to Proposition 9, to identify a complete labelling which is strictly more committed w.r.t. a previously found complete labelling, an argument labelled undec should then become in
π_{eO}^P	$\{\top, \perp\}$	Alg. 13, 15, 16	Further to Proposition 8, to identify a non-empty complete labelling that is not less committed w.r.t. any found preferred labelling, an argument labelled in should then become labelled out
π_{el}^P	$\{\top, \perp\}$	Alg. 13, 15, 16	Further to Proposition 8, to identify a non-empty complete labelling that is not less committed w.r.t. any found preferred labelling, an argument labelled out should then become labelled in

Table B.5: Table summarising the parameters used in algorithms presented in Section 5.

Appendix C. Pilot study: performance comparison with minimal number of propositional variables

For stable semantics, the CNFs identified in Proposition 11 can be modified as presented in Proposition 12, by taking into consideration the constraint $\bigwedge_{\mathbf{a} \in \mathcal{A}} \neg U_{\mathbf{a}}$ (cf. Algorithm 10 for instance). In Proposition 12, we use one propositional variable $I_{\mathbf{a}}$ for argument $\mathbf{a} \in \mathcal{A}$, implicitly defining two labels: if $I_{\mathbf{a}}$ is assigned \top , \mathbf{a} is **in**, otherwise it is **out**. Proof of Proposition 12 is omitted as it follows from straightforward manipulations.

Proposition 12. Given an AF $\Gamma = \langle \mathcal{A}, \mathcal{R} \rangle$:

$$C_{\text{in}}^{\leftarrow} \equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(I_{\mathbf{a}} \vee \left(\bigvee_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} I_{\mathbf{b}} \right) \right)$$

$$C_{\text{in}}^{\rightarrow} \equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} \neg I_{\mathbf{a}} \vee \neg I_{\mathbf{b}} \right)$$

	Mean runtime (s)		Median runtime (s)		Wilcoxon
	Prop. 12	Prop. 11	Prop. 12	Prop. 11	
<code>st_small</code>	27.91	75.93	7.97	6.02	Yes Z = -1.126 p = 0.260
<code>st_medium</code>	19.33	19.46	13.38	13.85	No Z = -2.057 p = 0.040

Table C.6: Statistics for the pilot study comparing computing EE-ST using three variables per argument (cf. Prop. 11) and using one variable per argument (cf. Prop. 12). The **Wilcoxon** column refers to Wilcoxon signed-rank tests aimed at assessing whether there is a statistical significant change in the execution time.

$$C_{\text{out}}^{\leftarrow} \equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(\bigwedge_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} \neg I_{\mathbf{b}} \vee \neg I_{\mathbf{a}} \right)$$

$$C_{\text{out}}^{\rightarrow} \equiv \bigwedge_{\{\mathbf{a} \in \mathcal{A}\}} \left(I_{\mathbf{a}} \vee \left(\bigvee_{\{\mathbf{b} \mid \mathbf{b} \rightarrow \mathbf{a}\}} I_{\mathbf{b}} \right) \right)$$

In this pilot study, we considered a portion of the ICCMA-15 benchmark, namely the groups `st_small` and `st_medium`, which consist of graphs which feature many complete/preferred/stable extensions.¹⁶ `st_small` features *AFs* with up to 300 arguments, while `st_medium` features *AFs* with 400 arguments.

We ran Algorithm 10 using $C_{\text{in}}^{\rightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow}$ with the CNFs specified in Proposition 11, and we compared that against the same algorithm using $C_{\text{in}}^{\rightarrow} \wedge C_{\text{in}}^{\leftarrow} \wedge C_{\text{out}}^{\rightarrow} \wedge C_{\text{out}}^{\leftarrow}$ with the CNFs specified in Proposition 12.¹⁷ We used a 6 core machine Intel(R) Core(TM) i7-5820K CPU @3.30GHz with 8GB RAM.

Table C.6 summarises the statistics for this study. First of all, according to a Wilcoxon signed-rank test, there is a significant change in the execution time for *AFs* in `st_small`, but there is no significant change for the larger *AFs* in `st_medium`. Moreover, considering the *AFs* in `st_small`, the average runtime for the version using the CNFs of Proposition 11 is higher than for the version using the CNFs of Proposition 12, but this is reversed when considering median values. Indeed, the median runtime of the implementation running on the CNFs specified by Proposition 11 is smaller than the median runtime of the implementation running on the *minimal* CNFs specified by Proposition 12.

¹⁶<http://argumentationcompetition.org/2015/results.html>

¹⁷In this case, references to $U_{\mathbf{a}}$ in Algorithm 10 must be ignored.

Therefore, this pilot study seems to suggest that performance does not necessarily improve when considering fewer propositional variables when transforming the constraints imposed by (stable) semantics into CNFs.