

Abstract argumentation frameworks with strong and weak constraints [☆]

Gianvincenzo Alfano ^{*}, Sergio Greco, Domenico Mandaglio, Francesco Parisi, Irina Trubitsyna

Department of Informatics, Modeling, Electronics and System Engineering, University of Calabria, Italy

ARTICLE INFO

Keywords:

Formal argumentation
Hard constraints
Soft constraints

ABSTRACT

Dealing with controversial information is an important issue in several application contexts. Formal argumentation enables reasoning on arguments for and against a claim to decide on an outcome. Dung's abstract Argumentation Framework (AF) has emerged as a central formalism in argument-based reasoning. Key aspects of the success and popularity of Dung's framework include its simplicity and expressiveness. Integrity constraints help to express domain knowledge in a compact and natural way, thus keeping easy the modeling task even for problems that otherwise would be hard to encode within an AF. In this paper, we first explore two intuitive semantics based on Kleene and Lukasiewicz logics, respectively, for AF augmented with (strong) constraints—the resulting argumentation framework is called Constrained AF (CAF). Then, we propose a new argumentation framework called Weak constrained AF (WAF) that enhances CAF with weak constraints. Intuitively, these constraints can be used to find “optimal” solutions to problems defined through CAF. We provide a detailed complexity analysis of CAF and WAF, showing that strong constraints do not increase the expressive power of AF in most cases, while weak constraints systematically increase the expressive power of CAF (and AF) under several well-known argumentation semantics.

1. Introduction

Argumentation is a well-known human process used in our daily life to explain something, persuade people, derive conclusions, and in general it is fundamental during debates. Most of the situations where argumentation takes place are inherently characterized by the presence of controversial information. Enabling automated systems to process such kind of information, much in the same way as organized human discussions are carried out, is an important challenge that has deserved increasing attention from the Artificial Intelligence community in the last decades. This has led to the development of an important and active research area called *formal argumentation* [22,63], that has been explored in several application contexts, e.g., legal reasoning [19], decision support systems [12], E-Democracy [26], healthcare [91], medical applications [72], financial analysis [84], explanation of results [30], as well as multi-agent systems and social networks [73].

[☆] This paper is a substantially revised and expanded version of [5].

^{*} Corresponding author.

E-mail addresses: g.alfano@dimes.unical.it (G. Alfano), greco@dimes.unical.it (S. Greco), d.mandaglio@dimes.unical.it (D. Mandaglio), fparisi@dimes.unical.it (F. Parisi), i.trubitsyna@dimes.unical.it (I. Trubitsyna).

<https://doi.org/10.1016/j.artint.2024.104205>

Received 22 December 2023; Received in revised form 30 July 2024; Accepted 12 August 2024

Available online 20 August 2024

0004-3702/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. (a) AF Λ of Example 1; (b) AF Λ' of Example 3.

Dung's abstract Argumentation Framework (AF) has emerged as a central formalism for modelling disputes between two or more agents [48]. An AF consists of a set of *arguments* and a binary *attack* relation over the set of arguments that specifies conflicts between arguments (if argument a attacks argument b , then b is acceptable only if a is not). Hence, arguments are abstract entities whose role is determined by attacks. We can think of an AF as a directed graph whose nodes represent arguments and edges represent attacks. As for graph theory, an important aspect of the success of Dung's framework is that it is a simple yet powerful formalism. The meaning of an AF is given in terms of argumentation *semantics*, which intuitively tell us the sets of arguments (called *extensions*) that can collectively be accepted to support a point of view in a dispute.

Despite the expressive power and generality of AFs, in some cases it is difficult to accurately model domain knowledge by an AF in a natural and easy-to-understand way. For this reason, Dung's framework has been extended by the introduction of further constructs, such as preferences [10,79,69,6,7] and *integrity constraints* [41,17], to achieve more comprehensive, natural, and compact ways of representing useful relationships among arguments. In particular, enhancing AF with constraints allows us to naturally and compactly express domain conditions that need to be taken into account to filter out unfeasible solutions, as illustrated in the following example.

Example 1. Albert, Betty and Charlie wish to attend a basketball game on Saturday evening, but only two tickets are available. In an attempt to model this situation by an AF Λ , the following six arguments can be used: a (resp., b , c) states that Albert (resp., Betty, Charlie) attends the game, whereas \bar{a} (resp., \bar{b} , \bar{c}) states that Albert (resp., Betty, Charlie) does not attend the game. The direct graph encoding Λ is shown in Fig. 1(a), where double arrows are used to represent mutually attacks between arguments. Specifically, argument a (resp., b , c) attacks and is attacked by argument \bar{a} (resp., \bar{b} , \bar{c}), i.e., only one of them can be accepted. Moreover, argument \bar{a} (resp., \bar{b} , \bar{c}) is attacked by the other two arguments \bar{b} and \bar{c} (resp., \bar{a} and \bar{c} ; \bar{a} and \bar{b}) since the argument that Albert (resp., Betty, Charlie) attends the game can be accepted only if one of the arguments stating that Betty or Charlie (resp., Albert or Charlie; Albert or Betty) do not attend the game is accepted. Thus, the set of attacks between every pair in $\{\bar{a}, \bar{b}, \bar{c}\}$ models the fact that at most one argument among \bar{a} , \bar{b} and \bar{c} can be accepted and then, as a consequence, at least two arguments among a , b and c can be accepted, i.e., all available tickets are sold.

The extensions of the AF Λ under the well-known preferred and stable semantics are $E_1 = \{a, b, \bar{c}\}$, $E_2 = \{a, \bar{b}, c\}$, $E_3 = \{\bar{a}, b, c\}$, and $E_4 = \{a, b, c\}$, where the presence of an argument in one of the 4 solutions means that it is accepted. However, the AF Λ fails to capture the knowledge we want to represent due to the presence of extension E_4 , which admits that three people attend the game, while only two tickets are available. \square

With the aim of allowing for a more straightforward and compact encoding of knowledge, several frameworks extending AF have been proposed, such as *Abstract Dialectical Framework* (ADF) [36,92,33,24] and SETAF [80,62,54,55], where the situation of Example 1 can be modeled by using proper *acceptance conditions* over arguments or *collective attacks*, respectively (see Section 9 for a detailed discussion). Moreover, to overcome situations similar to that of Example 1, and thus providing a natural and compact way for expressing such kind of conditions, the use of constraints to filter extensions has been proposed. Considering our example, a constraint κ defined as

$$\kappa = a \wedge b \wedge c \Rightarrow \text{f}$$

can be used. It states that the propositional formula $a \wedge b \wedge c$ must be false. That is, feasible solutions must satisfy the condition that the 3 arguments a , b , and c are not jointly accepted, i.e., Albert, Betty and Charlie cannot attend the game together. The effect of using constraint κ is that E_4 is discarded from the set of solutions of our problem.

The use of constraints in AF has been firstly proposed in [41] and then further investigated in [16–18]. The constrained argumentation frameworks in [16] and [18] are particular cases of those in [17] as the set of constraints is restricted to atomic formulae only. We call an AF with constraints a *Constrained AF* (CAF).

Although constraints in CAF allow restricting the set of feasible solutions, they do not help in finding “best” or preferable solutions. Considering our running example, Albert, Betty and Charlie may agree on the fact that “if there are only two tickets available then Albert and Betty should preferably attend the game”. To express this kind of conditions, in this paper we introduce *weak constraints*, that is, constraints that are required to be satisfied *if possible*. Syntactically, weak constraints have the same form of the above-mentioned kind of constraints, that we call *strong constraints*. Intuitively, weak constraints can be used to find “optimal” solutions to a problem defined by means of an AF or a CAF, that is to filter out, from the set of feasible extensions of a given AF or CAF, the extensions which satisfy a maximal set (or a maximum number) of weak constraints.

A CAF with the addition of weak constraints is said to be a *Weak constrained Argumentation Framework* (WAF).

Example 2. Consider a WAF obtained by adding to the AF of Example 1 the constraint κ and the weak constraint $w = \text{t} \Rightarrow a \wedge b$, stating that is desirable that Albert and Betty attend the game together. Herein, t denotes the truth value true. Then, the extension

$E_1 = \{a, b, \bar{c}\}$ is selected as the “best” one as it is the only one that satisfies the constraint w asking for the presence of arguments a and b . \square

Weak constraints (also called relaxed constraints in some contexts) have been considered in several research areas, including Mathematical Programming with Equilibrium Constraints [86], Answer Set Programming [37,66], (weighted) Max-Sat [74], and for modelling and solving optimization problems [60]. In particular, concerning the field of Answer Set Programming, weak constraints have been implemented in DLV [9], a disjunctive logic programming system with (total) stable models semantics.

The use of strong and weak constraints substantially reduces the effort needed to figure out how to define an AF that models a given problem. In fact, as said before, constraints facilitate to express knowledge in a more compact and easy to understand way. For instance, the problem presented in Example 1, has been represented through an AF which expresses the condition that “at most one argument among \bar{a} , \bar{b} and \bar{c} can be accepted” and then, as a consequence, at least two arguments among a , b and c can be accepted. However, this condition is not easy to be generalized if we have more than three people. Suppose there is a fourth person, David, who wishes to attend the game, and there are again only two available tickets. After adding the arguments d (David attends the game) and \bar{d} (David does not attend the game) to AF Λ of Fig. 1(a), we cannot use the same reasoning as in Example 1 to model the fact that two of the four people attend the game. In fact, having the attacks between every pair in $\{\bar{a}, \bar{b}, \bar{c}, \bar{d}\}$ does not model this situation (it models that at least three of the four people attend the game). Remarkably, using strong and weak constraints allow for using a common reasoning pattern to generalize to this more complex situation, even starting from an AF having a simpler structure.

Example 3. Consider a WAF consisting of AF Λ' of Fig. 1(b) and the following sets, C and \mathcal{W} , of strong and weak constraints, respectively:

$$\begin{aligned} C &= \{ a \wedge b \wedge c \Rightarrow f, a \wedge b \wedge d \Rightarrow f, a \wedge c \wedge d \Rightarrow f, b \wedge c \wedge d \Rightarrow f \}; \\ \mathcal{W} &= \{ \tau \Rightarrow a, \quad \tau \Rightarrow b, \quad \tau \Rightarrow c, \quad \tau \Rightarrow d \}. \end{aligned}$$

The strong constraints in C (that includes κ of Example 1) filter out from the (16 preferred) extensions of Λ' the solutions where more than two people attend the game, whereas the weak constraints maximize the set (or number) of people attending the game since the “best” extensions are those that satisfy the maximal set (or number) of weak constraints, each asking for the presence of a person. \square

It is worth mentioning that, although in this paper we consider ground constraints, the proposed framework can be easily extended to more general formulae with variables denoting arguments, whose ground version is a propositional formula. For instance, the strong and weak constraints in Example 3 could be written by using only one strong constraint of the form $X \wedge Y \wedge Z \wedge (id(X) \neq id(Y)) \wedge (id(X) \neq id(Z)) \wedge (id(Y) \neq id(Z)) \Rightarrow f$ and only one weak constraint $\tau \Rightarrow X$, where X , Y and Z are variables whose domain is the set of arguments, and $id(X)$ denotes the identifier of the argument associated to X (e.g. the pointer to the object).

1.1. Contributions

In this paper, after introducing CAF and WAF, we investigate the complexity of both credulous and skeptical reasoning in these argumentation frameworks. Credulous and skeptical reasoning are well-known approaches to deal with uncertain information represented by the presence of multiple solutions. In our context, an argument is credulously accepted if there exists a solution (i.e., an extension of the considered framework) containing that argument, whereas an argument is skeptically accepted if it occurs in all solutions.

We provide the complexity results that are summarized in Tables 1 and 2 (reported at the end of Section 5), where CA_S (resp., SA_S) denotes the credulous (resp., skeptical) acceptance problem under one of the following argumentation semantics S : complete (co), stable (st), preferred (pr), and semi-stable (sst). Moreover, since we will consider two alternative 3-valued logics for interpreting the constraints, that is, Kleene logic and Lukasiewicz logic, in the above-mentioned tables we also use the notations CA_S^σ (resp., SA_S^σ) to denote the credulous (resp., skeptical) acceptance problem under semantics S and logic interpretation σ ; herein, $\sigma = K$ (resp., $\sigma = L$; $\sigma = *$) denotes Kleene (resp., Lukasiewicz; either Kleene or Lukasiewicz) logic interpretation of the constraints.

More in detail, we make the following main contributions.

- We propose the CAF framework by relying on a simple yet expressive form of constraints that are interpreted using either Kleene or Lukasiewicz logic, leading to intuitive constraints’ semantics.
- We investigate the complexity of CA_S^σ and SA_S^σ for CAF under four popular semantics, showing that it remains the same as for AF in all cases except the cases of *i*) credulous acceptance under preferred semantics and Lukasiewicz logic, and *ii*) skeptical acceptance under stable semantics (irrespective of the logic considered for interpreting the constraints), where the complexity increases of one level in the polynomial hierarchy.
- We introduce the WAF framework and propose two criteria for interpreting weak constraints, under any argumentation semantics S : *maximal-set* (ms S) and *maximum-cardinality* (mc S) according to which the best/optimal S -extensions are those satisfying a maximal set, or a maximum number, of weak constraints, respectively.
- We investigate the complexity of the credulous and skeptical acceptance problems for WAF, where they are denoted as CA_{msS}^σ and SA_{msS}^σ , and CA_{mcS}^σ and SA_{mcS}^σ , respectively, depending on the considered criterion (*maximal-set* or *maximum-cardinality*) adopted

for interpreting the weak constraints (with $\sigma \in \{K, L\}$). We show that, differently from strong constraints, the introduction of weak constraints typically increases the complexity of the considered problems of one level in the polynomial hierarchy.

- We introduce Stratified WAF (SWAF) and investigate a restriction of SWAF, called Linear WAF (LWAF), where constraints are linearly ordered. It turns out that, in most cases WAF and SWAF have the same complexity under maximal-set semantics, while WAF are in general less expressive than SWAF under maximum-cardinality semantics. Moreover, for LWAF the maximal-set and maximum-cardinality semantics coincide, thus we simply use the notations CA_S^σ and SA_S^σ (with $\sigma \in \{K, L\}$) for denoting the credulous and skeptical acceptance, respectively. For LWAF, the complexity of CA_S^σ and SA_S^σ generally decreases w.r.t. that of WAF under maximal-set semantics, though it is higher than that of CAF and that of WAF under maximum-cardinality semantics.
- Finally, we investigate the case of NCAF and NWAF, that is, CAF and WAF, respectively, where constraints are expressed by negative constraints, i.e., denials constraints whose body is a conjunction of literals (used in several contexts such as databases and logic programming), and show that the complexity of CA_S^* for the preferred semantics decreases (irrespective of the logic considered for interpreting the constraints).

This paper refines and substantially extends the work in [5]. In particular, we have extended the form of constraints considered, which are defined through formulae of one of the two forms $\varphi \Rightarrow v$ and $v \Rightarrow \varphi$, where v is a truth value ($\text{f}, \text{u}, \text{t}$) and φ is a first order formula built over the alphabet of arguments. The formula φ can now also contain the implication \Rightarrow (which is a primitive operator in the Lukasiewicz logic) and equivalence \Leftrightarrow operators, whereas in [5] φ could be built by using the \wedge , \vee and \neg operators only. We investigate CAF, WAF, SWAF, LWAF and NWAF under two alternative 3-valued logics which differ in the interpretation of the implication operator: Kleene logic and Lukasiewicz logic. In contrast, in [5] only Lukasiewicz logic is considered for interpreting the constraints; moreover, the complexity of SWAF is not addressed at all in [5]. We provide tight complexity bounds and close a gap left open in [5] for the complexity of the credulous acceptance problem in CAF (interestingly, although we provide a stronger hardness result, our result holds even for the simpler form of constraints considered in [5]). Overall, we provide a detailed analysis of AF with strong and weak constraints interpreted under Kleene or Lukasiewicz logic, with maximal-set and maximum-cardinality interpretations of weak constraints, under four popular argumentation semantics, by also considering several restrictions on the forms of the constraints: stratified, linearly ordered, and negative constraints. We also show that some preference-based AFs can be encoded in WAF and that CAF (and thus WAF) is more expressive than LabCAF, that is a CAF framework where constraints are defined over the alphabet of labelled arguments [21]. The new material includes all the proofs of the results stated in the core of the paper as well as the proofs of useful auxiliary results stated in the appendix (it is worth mentioning that some of those results are of independent interest, e.g., the mapping from DLPs to logic programs with weak constraints under maximal-set semantics, which entails that the latter are no less expressive than DLPs, see Appendix C.1).

1.2. Organization

The rest of the paper is organized as follows. The abstract argumentation framework and the complexity classes used in the paper are recalled in Section 2. In Section 3 we discuss the syntax and semantics of the forms of CAF presented in the literature, whereas in Section 4 we introduce a simple yet expressive form of constraints that are interpreted under either Kleene or Lukasiewicz logic and lead to the CAF frameworks on which we focus in this paper. This section also analyzes the computational complexity of the credulous and skeptical acceptance problems in CAF. Next, in Section 5, we introduce WAF, which extends CAF through the introduction of weak constraints, and formally define the meaning of WAF under the maximal-set and maximum-cardinality semantics, and investigate the complexity of credulous and skeptical reasoning (Sections 5.1 and 5.2, respectively). In Section 6 we introduce SWAF and investigate the computational complexity of SWAF and of the special case of LWAF, whereas in Section 7 we deal with the credulous and skeptical acceptance for NCAF and NWAF. In Section 8, we discuss the relationship between WAF and preferences in AF, showing that some preference-based AFs can be encoded in WAFs. Related work is discussed in Section 9, while in Section 10 conclusions are drawn and directions for future work are outlined.


To ease readability, in the core of the paper we provide only the proofs regarding CAF (that is, the basic AF extension studied in Section 4) as well as the proof concerning the encoding of PAF into WAF given in Section 8. All the other proofs concerning the remaining results on WAF, SWAF, LWAF and NWAF are given in Appendix A. The paper also contains four further appendixes, organized as follows. Since some proofs given in Appendix A rely on exploiting some results from disjunctive logic programming (DLP) and logic programming with weak constraints (LPWC), to make the paper self-contained, Appendix B contains useful material on DLP, whereas Appendix C introduces LPWC and its relationships with DLP and WAF. Moreover, in Appendix D, we show the relationship between CAF (and thus WAF) and AF with labelled constraints (namely, LabCAF), which is a kind of Epistemic Argumentation Framework [89] with a restricted modal operator. Finally, Appendix E briefly recalls the syntax and the semantics of the Abstract Dialectical Framework (ADF) [36], whose relationship with CAF and WAF is discussed in Section 9.

2. Preliminaries

In this section, we briefly review Dung's framework and some basic notions about computational complexity.

2.1. Argumentation framework

An abstract *Argumentation Framework* (AF) is a pair $\langle \mathcal{A}, \mathcal{R} \rangle$, where \mathcal{A} is a set of *arguments* and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ is a set of *attacks*. If $(a, b) \in \mathcal{R}$ then we say that a attacks b .

Fig. 2. AF Λ of Example 4.

Given an AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ and a set $S \subseteq \mathcal{A}$ of arguments, an argument $a \in \mathcal{A}$ is said to be i) *defeated* w.r.t. S iff $\exists b \in S$ such that $(b, a) \in \mathcal{R}$, and ii) *acceptable* w.r.t. S iff for every argument $b \in \mathcal{A}$ with $(b, a) \in \mathcal{R}$, there is $c \in S$ such that $(c, b) \in \mathcal{R}$. The sets of defeated and acceptable arguments w.r.t. S are defined as follows (where Λ is fixed):

- i) $Def(S) = \{a \in \mathcal{A} \mid \exists b \in S. (b, a) \in \mathcal{R}\}$;
- ii) $Acc(S) = \{a \in \mathcal{A} \mid \forall b \in \mathcal{A}. (b, a) \notin \mathcal{R} \vee b \in Def(S)\}$.

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ of arguments is said to be:

- *conflict-free* iff $S \cap Def(S) = \emptyset$;
- *admissible* iff it is conflict-free and $S \subseteq Acc(S)$.

Different argumentation semantics have been proposed to characterize collectively acceptable sets of arguments, called *extensions* [48,39]. Every extension is an admissible set satisfying additional conditions. Specifically, the complete, preferred, stable, semi-stable, and grounded extensions of an AF are defined as follows.

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$, a set $S \subseteq \mathcal{A}$ is an *extension* called:

- *complete* (co) iff it is an admissible set and $S = Acc(S)$;
- *preferred* (pr) iff it is a maximal (w.r.t. \subseteq) complete extension;
- *stable* (st) iff it is a total preferred extension, i.e., a preferred extension such that $S \cup Def(S) = \mathcal{A}$;
- *semi-stable* (sst) iff it is a preferred extension such that $S \cup Def(S)$ is maximal (w.r.t. \subseteq);
- *grounded* (gr) iff it is the smallest (w.r.t. \subseteq) complete extension.

Arguments occurring in an extension are said to be accepted, whereas arguments attacked by accepted arguments are said to be rejected; remaining arguments are said to be undecided (w.r.t. the considered extension).

The set of complete (resp. preferred, stable, semi-stable, grounded) extensions of an AF Λ will be denoted by $co(\Lambda)$ (resp. $pr(\Lambda)$, $st(\Lambda)$, $ss(\Lambda)$, $gr(\Lambda)$). It is well-known that the set of complete extensions forms a complete semilattice w.r.t. \subseteq , where $gr(\Lambda)$ is the meet element, whereas the greatest elements are the preferred extensions [48]. All the above-mentioned semantics except the stable semantics admit at least one extension. The grounded semantics, that admits exactly one extension, is said to be a *unique status* semantics, while the others are *multiple status* semantics. With a little abuse of notation, in the following we also use $gr(\Lambda)$ to denote the grounded extension. For any AF Λ the following inclusion relations hold: i) $st(\Lambda) \subseteq ss(\Lambda) \subseteq pr(\Lambda) \subseteq co(\Lambda)$, ii) $gr(\Lambda) \in co(\Lambda)$, and iii) $st(\Lambda) \neq \emptyset$ implies that $st(\Lambda) = ss(\Lambda)$.

Example 4. Let $\Lambda = \langle \mathcal{A} = \{a, b, c\}, \mathcal{R} = \{(a, b), (b, a), (b, c), (c, c)\} \rangle$ be the AF shown in Fig. 2. AF Λ has three complete extensions: $E_1 = \emptyset$, $E_2 = \{a\}$, $E_3 = \{b\}$. Moreover, the set of preferred extensions is $\{E_2, E_3\}$, whereas the set of stable (and semi-stable) extensions is $\{E_3\}$, and the grounded extension is E_1 . \square

Credulous and skeptical acceptance Given an AF framework Λ , an argument a , and an argumentation semantics $S \in \{co, pr, st, sst, gr\}$,

- the *credulous acceptance* problem, denoted as CA_S , is the problem of deciding whether argument a is credulously accepted, that is, deciding whether a belongs to at least an S -extension of Λ .
- the *skeptical acceptance* problem, denoted as SA_S , is the problem of deciding whether argument a is skeptically accepted, that is, deciding whether a belongs to every S -extension of Λ .

Clearly, for the grounded semantics, which admits exactly one extension, these problems become identical. The above-defined notions of credulous and skeptical acceptance will be also used in the context of the frameworks extending AF discussed in the paper (e.g., CAF and WAF), that is, by ranging on extensions of those frameworks when checking for the presence of a given argument a in any or all extensions in the process of credulous and skeptical acceptance, respectively.

2.2. Complexity classes

We recall here the main complexity classes used in the paper and, in particular, the definition of the classes Σ_k^p , Π_k^p and Δ_k^p , with $k \geq 0$ (see e.g. [82]):

- $\Sigma_0^p = \Pi_0^p = \Delta_0^p = P$;

- $\Sigma_1^p = NP$ and $\Pi_1^p = coNP$;
- $\Delta_k^p = P^{\Sigma_{k-1}^p}$, $\Sigma_k^p = NP^{\Sigma_{k-1}^p}$, and $\Pi_k^p = co\Sigma_k^p$, $\forall k > 0$.

Thus, P^C (resp., NP^C) denotes the class of problems that can be solved in polynomial time using an oracle in the class C by a deterministic (resp., non-deterministic) Turing machine. Θ_k^p denotes the subclass of Δ_k^p containing the problems that can be solved in polynomial time by a deterministic Turing machine by performing a number of calls bounded by $O(\log n)$ to an oracle in the class Σ_{k-1}^p , that is, $\Theta_k^p = \Delta_k^p[\log n]$. Under standard complexity-theoretic assumptions, we have that:

- $\Sigma_k^p \subset \Theta_{k+1}^p \subset \Delta_{k+1}^p \subset \Sigma_{k+1}^p \subset PSPACE$ and
- $\Pi_k^p \subset \Theta_{k+1}^p \subset \Delta_{k+1}^p \subset \Pi_{k+1}^p \subset PSPACE$.

For AF, the complexity of the credulous and skeptical acceptance problems has been investigated in [48] for the grounded semantics, in [45] for the stable semantics, in [45,49] for the preferred semantics, and in [51,56] for the semi-stable semantics. These results are thoroughly discussed in [53], and summarized in the second column of Tables 1 and 2.

3. Constrained argumentation frameworks

We review the Constrained Argumentation Framework (CAF) introduced in [41] and further investigated in [17].

We assume that, given a set of propositional symbols S , \mathcal{L}_S denotes the propositional language defined in the usual inductive way from S using the built-in constants f , u , and t denoting the truth values *false*, *undef* (*undefined*),¹ and *true*, and the connectives \wedge , \vee , \neg , \Rightarrow and \Leftrightarrow .

Definition 1 (*Constrained Argumentation Framework*). A *Constrained Argumentation Framework* (CAF) is a triple $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ where $\langle \mathcal{A}, \mathcal{R} \rangle$ is an AF and C is a set of propositional formulae built from \mathcal{L}_A .

3.1. CAF semantics

Given an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ and a conflict-free set $S \subseteq \mathcal{A}$, the truth value of an argument $a \in \mathcal{A}$ w.r.t. S is denoted by $\vartheta_S^\sigma(a)$, where σ denotes the underlying 3-valued logic, or simply $\vartheta^\sigma(a)$ whenever S is given, and is defined as follows:

$$\vartheta^\sigma(a) = \begin{cases} \text{true} & \text{if } a \in S \\ \text{false} & \text{if } \exists b \in S \text{ s.t. } (b, a) \in \mathcal{R} \\ \text{undef} & \text{otherwise} \end{cases}$$

Observe that, for a given (complete) extension E , $\vartheta^\sigma(a)$ is true (resp., false, undef) iff $a \in \text{Acc}(E)$ (resp., $a \in \text{Def}(E)$, $a \in \mathcal{A} \setminus (\text{Acc}(E) \cup \text{Def}(E))$).

Assuming that $\neg\text{undef} = \text{undef}$, and the ordering on truth values $\text{false} < \text{undef} < \text{true}$, using a 3-valued logic σ we have that $\vartheta^\sigma(\varphi \wedge \psi) = \min(\vartheta^\sigma(\varphi), \vartheta^\sigma(\psi))$, $\vartheta^\sigma(\varphi \vee \psi) = \max(\vartheta^\sigma(\varphi), \vartheta^\sigma(\psi))$ and $\vartheta^\sigma(\neg a) = \neg\vartheta^\sigma(a)$. It is important to note that, regarding the operator \Rightarrow , there is no consensus on how its semantics should be defined. In the following, we first review the semantics proposed in [41] and [17] (in Sections 3.1.1 and 3.1.2, respectively); the former relying on classical 2-valued semantics, the latter relying on 3-valued semantics. Then, we introduce new three-valued semantics based on Kleene's logic and Lukasiewicz's logic in Section 4.

After that we have defined the semantics of the implication operator \Rightarrow , i.e. have fixed the underlying 3-valued logic, we can also define the semantics of the equivalence operator \Leftrightarrow as follows: $\vartheta^\sigma(\varphi \Leftrightarrow \psi) = \vartheta^\sigma(\varphi \Rightarrow \psi) \wedge \vartheta^\sigma(\psi \Rightarrow \varphi)$. Moreover, we say that a given set S satisfies a set of constraints C (written $S \models C$) if $\vartheta^\sigma(\bigwedge_{\varphi \in C} \varphi) = \text{true}$. We also say that C is satisfiable, under a given logic σ , if there exists a set S such that $S \models C$ according to σ .

In the following, for the different semantics we redefine the evaluation function only for the cases where it differs from the generic function ϑ^σ previously discussed and for formulae using the implication operator.


3.1.1. Semantics of Coste-Marquis et al.

The semantics proposed in [41] is based on a (2-valued) evaluation of the truth value of an argument a w.r.t. a given set S of arguments (denoted as $\vartheta_S^2(a)$), which is defined as follows:

$$\vartheta_S^2(a) = \begin{cases} \text{true} & \text{if } a \in S \\ \text{false} & \text{if } a \notin S \end{cases}$$

Recalling that under 2-valued interpretation $\varphi \Rightarrow \psi \equiv \neg\varphi \vee \psi$, we have that $\vartheta_S^2(\varphi \Rightarrow \psi) = \vartheta_S^2(\neg\varphi \vee \psi)$.

¹ As it will be clearer in the following, undefined (*undef*) is a third value (in addition to the two classical values used in Boolean logic) which intuitively means neither true nor false.

Fig. 3. AF $\langle \mathcal{A}, \mathcal{R} \rangle$ underlying CAF Ω of Example 5.

Definition 2 (*C-admissible set*). Given a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ and a set $S \subseteq \mathcal{A}$, S is a C -admissible set for Ω if and only if S is an admissible set for $\langle \mathcal{A}, \mathcal{R} \rangle$ and $S \models C$.

Example 5. As an example, consider the CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C = \{b \Rightarrow f\} \rangle$, where the AF $\langle \mathcal{A} = \{a, b\}, \mathcal{R} = \{(a, b), (b, a)\} \rangle$ is shown in Fig. 3, and the three sets $S_0 = \emptyset$, $S_1 = \{a\}$, and $S_2 = \{b\}$. We have that (i) $\vartheta_{S_0}^2(a) = \vartheta_{S_0}^2(b) = \text{false}$, (ii) $\vartheta_{S_1}^2(a) = \text{true}$ and $\vartheta_{S_1}^2(b) = \text{false}$, (iii) $\vartheta_{S_2}^2(a) = \text{false}$ and $\vartheta_{S_2}^2(b) = \text{true}$. Therefore, $S_0 \models C$ and $S_1 \models C$, meaning that they are C -satisfiable, whereas S_2 is not C -satisfiable since $S_2 \not\models C$. \square

A constrained argumentation framework $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ is *consistent* when it has a C -admissible set for Ω .

Definition 3 (*Preferred/Stable C-extension*). Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF. A C -admissible set $E \subseteq \mathcal{A}$ for Ω is

- a *preferred C-extension* of Ω if and only if $\nexists E' \subseteq \mathcal{A}$ such that $E \subset E'$ and E' is C -admissible for Ω ;
- a *stable C extension* if and only if it is a total preferred C -extension.

A drawback of the semantics proposed in [41] is that in checking whether an extension E satisfies a set of constraints it does not distinguish between false and undefined arguments. Thus, a constraint of the form $a \wedge \neg a \Rightarrow f$ is always satisfied, even if the truth value of a would be undefined w.r.t. a 3-valued logic.

3.1.2. Arieli's semantics

The semantics proposed in [17] for checking constraints' satisfaction assumes a 3-valued interpretation based on the Slupecki's logic. In particular, it assumes the standard interpretation for the assignment of truth values to atoms (i.e., arguments) and expressions using the \neg , \wedge and \vee operators, whereas for the implication operator \Rightarrow (for which there is no standard interpretation) it assumes the Slupecki's interpretation which is defined as follows:

$$\vartheta^{Sl}(\varphi \Rightarrow \psi) = \begin{cases} \text{true} & \text{if } \vartheta^{Sl}(\varphi) \in \{\text{false}, \text{undef}\} \\ \vartheta^{Sl}(\psi) & \text{otherwise} \end{cases}$$

A natural requirement for constraints applied to argumentation frameworks is that they should have admissible interpretations: the constraints themselves should not be contradictory and every argument that is satisfied by an extension should not be exposed to undefended attacks (w.r.t. the extension).

Definition 4 (*Admissible constraint*). Let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ be an AF. A set C of formulae is called *admissible* (for Λ) if there exists an admissible set $S \subseteq \mathcal{A}$ for $\langle \mathcal{A}, \mathcal{R} \rangle$ such that $S \models C$.

Assuming that constraints are admissible, extensions for a CAF are defined as follows.

Definition 5 (*S-extension of a CAF*). Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF, where C is an admissible set of constraints, and let S be a semantics for $\langle \mathcal{A}, \mathcal{R} \rangle$. Then $E \subseteq \mathcal{A}$ is an S -extension of Ω if it is an S -extension of $\langle \mathcal{A}, \mathcal{R} \rangle$ and $E \models C$.

The main difference between the two CAF semantics briefly reviewed in this section is as follows. In Coste-Marquis et al. (2006) [41] the truth value of arguments is false for every argument not belonging to the considered extension (even for those that are undecided) and satisfiability of constraints is evaluated with respect to two-valued semantics. It follows, e.g., that a constraint of the form $t \Rightarrow a \vee \neg a$ is useless according to [41] (since it is always satisfied). In contrast, in the Arieli's 3-valued semantics this constraint indicates that argument a cannot have a neutral, undefined, status. The use of 3-valued semantics allows us to distinguish between different conditions on arguments. For instance, the constraint $t \Rightarrow \neg a$ means that a should be rejected, while the constraint $a \Rightarrow f$ is a somewhat weaker demand: a should not be accepted, and so its status may be undecided.

A drawback of Arieli's semantics, due to the assumption of the Slupecki's logic for interpreting the implication operator, is that it does not distinguish two constraints of the form $\varphi \Rightarrow f$ and $\varphi \Rightarrow u$, though it distinguishes two constraints of the form $t \Rightarrow \varphi$ and $u \Rightarrow \varphi$.

4. Revisiting the CAF semantics

In this section, we investigate two new 3-valued semantics for constraints satisfaction in CAF. The reason for considering 3-valued satisfaction is that all the AF semantics, except the stable one, are 3-valued and constraints satisfaction under 2-valued logic (obtained by interpreting the undefined truth value as either true or false) such as the one discussed in Section 3.1.1 is not satisfactory. We restrict

$\vartheta^K(\varphi \Rightarrow \psi)$	$\vartheta^K(\psi)$		
	f	u	t
$\vartheta^K(\varphi)$	f	t	t
	u	u	t
	t	f	t

Kleene

$\vartheta^L(\varphi \Rightarrow \psi)$	$\vartheta^L(\psi)$		
	f	u	t
$\vartheta^L(\varphi)$	f	t	t
	u	u	t
	t	f	t

Lukasiewicz

$\vartheta^{Sl}(\varphi \Rightarrow \psi)$	$\vartheta^{Sl}(\psi)$		
	f	u	t
$\vartheta^{Sl}(\varphi)$	f	t	t
	u	t	t
	t	f	t

Slupecki

Fig. 4. Semantics of the implication operator $\varphi \Rightarrow \psi$.

our attention to logics which extend the classical 2-valued logic and differ one from the other in the semantics of the implication operator only.² Among these we focus our attention to the most well-known 3-valued logics: Kleene's logic and Lukasiewicz's logic.

The tables in Fig. 4 report three different semantics for the implication operator: Kleene and Lukasiewicz logics are at the basis of the semantics studied in this paper, whereas Slupecki logic is at the basis of the semantics studied in [17]. In the following, $\mathcal{L}_{\mathcal{A}}$ denotes the propositional language defined from a set of arguments \mathcal{A} and truth values (f, u and t) and the standard connectives (\wedge , \vee , \neg , \Rightarrow and \Leftrightarrow).

Moreover, we use the evaluation functions ϑ^K whenever we refer to the Kleene's logic, and ϑ^L whenever we refer to the Lukasiewicz's logic.

Thus, under Kleene's logic we have that $\vartheta^K(\varphi \Rightarrow \psi) = \vartheta^K(\neg\varphi \vee \psi)$, whereas under Lukasiewicz's logic $\vartheta^L(\varphi \Rightarrow \psi) = \vartheta^L(\neg\varphi \vee \psi) \vee (\vartheta^L(\varphi) = \vartheta^L(\psi))$.³ A nice property of Kleene's logic is that it preserves the equivalence $\varphi \Rightarrow \psi \equiv \neg\varphi \vee \psi$. Moreover, concerning the implication operator, differently from other logics (e.g. Lukasiewicz, Slupecki's and Priest's), Kleene's logic preserves equivalence of formulae when elements of the disjunctive head are moved to the body (after negating them), or elements of the conjunctive body are moved to the head (after negating them), analogously to the case of 2-valued semantics. On the other side, Kleene's logic does not preserve the axiom $\varphi \Rightarrow \varphi$, which is instead valid under Lukasiewicz's logic (as well as Slupecki and Godel logics). For formulae defining constraints, Lukasiewicz logic allows to distinguish $\varphi \Rightarrow f$ from $\varphi \Rightarrow u$ and $t \Rightarrow \varphi$ from $u \Rightarrow \varphi$, while Kleene's logic does not. Another reason for investigating CAF under the two different above-mentioned 3-valued logics is the complexity of the expressivity of the two derived frameworks. Indeed, as we will show in the paper, the fact that Lukasiewicz's logic allows us to express finer constraints gives rise, for some semantics (e.g. the preferred one), to a more expressive framework characterized by a higher computational complexity.

Definition 6 ((Strong) constraint). A (strong) constraint is a formula of one of the following forms: (i) $\varphi \Rightarrow v$, or (ii) $v \Rightarrow \varphi$, where φ is a propositional formula in $\mathcal{L}_{\mathcal{A}}$ and $v \in \{f, u, t\}$. A constraint is said *boolean* when $v \in \{f, t\}$. A boolean constraint of the form $\varphi \Rightarrow f$ where φ is a conjunction containing arguments or negated arguments is called *denial* (or *negative*) constraint.

In the following, we refer to both Kleene and Lukasiewicz logics. We assume that the set of constraints is satisfiable, that is, that there is an assignment of truth values to the arguments that makes all constraints true under the given logic (Kleene's or Lukasiewicz's). Checking whether a constraint is satisfied by a set of arguments S under Kleene logic consists in checking whether the body is false or the head is true (w.r.t. S), whereas under the Lukasiewicz logic is equivalent to check whether the body is false or the head is true or the truth values of head and body coincide (w.r.t. S). We also assume that C is a set of (satisfiable) constraints built from $\mathcal{L}_{\mathcal{A}}$ as defined in Definition 6.

Example 6. Under Lukasiewicz logic we have that:

- the constraint $a \wedge b \wedge c \Rightarrow f$ states that at least one of the arguments a , b and c must be false, whereas $a \wedge b \wedge c \Rightarrow u$ states that a , b and c cannot be all true.
- the constraint $(a \Rightarrow b) \Rightarrow f$ states that the implication $a \Rightarrow b$ must be false, that is a must be true and b must be false.
- the constraint $(a \Rightarrow b) \Rightarrow u$ states that the implication $a \Rightarrow b$ cannot be true, that is the truth value of a must be greater than that of b . \square

Clearly, constraints of the forms $f \Rightarrow \varphi$ and $\varphi \Rightarrow t$ are useless because always satisfied. Regarding the stable semantics, which is 2-valued, only the symbols f and t can be used and (all) interpretations of the implication operator coincide with the classical 2-valued interpretation. Thus, a constraint $\varphi \Rightarrow u$ is interpreted as $\varphi \Rightarrow f$, whereas a constraint $u \Rightarrow \psi$ is interpreted as $t \Rightarrow \psi$.

The next definition introduces two new semantics for CAF, where $\sigma = K$ (resp., $\sigma = L$) denotes the Kleene (resp., Lukasiewicz) semantics, i.e. the implication operator \Rightarrow is interpreted according to the Kleene's (resp., Lukasiewicz's) logic.

² Most of the 3-valued semantics differ in the assignment of the truth value to implications of the form $u \Rightarrow f$, which can be either f (e.g. Priest, Godel), or u (Kleene, Lukasiewicz, Bochvar), or t (e.g. Slupecki), or of the form $u \Rightarrow u$, which can be either u (e.g. Kleene, Priest) or t (e.g. Lukasiewicz, Slupecki, Godel, Bochvar). Moreover Bochvar's logic differs from Slupecki's logic as it assigns f to the implication $t \Rightarrow u$, while all other logics considered here assign the value u. We refer the interested reader to [20] for an overview on different 3-valued logics.

³ Recall that the (primitive) propositional connectives of Lukasiewicz's logic are \Rightarrow and the constant f. Additional connectives are defined in terms of these as follows: $\neg A =_{def} A \Rightarrow f$, $A \vee B =_{def} (A \Rightarrow B) \Rightarrow B$, $A \wedge B =_{def} \neg(A \vee B)$ and $A \Leftrightarrow B =_{def} (A \Rightarrow B) \wedge (B \Rightarrow A)$.

Definition 7 ((Revised) CAF semantics). Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF, $S \in \{\text{co}, \text{gr}, \text{pr}, \text{st}, \text{ss}\}$ a semantics, and $\sigma \in \{K, L\}$ the underlying logic (either Kleene's or Lukasiewicz's logic). A set $E \subseteq \mathcal{A}$ is an S^σ -extension for Ω if E is an S -extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ and $E \vDash C$ under the σ logic.

The set of S^σ -extensions for a CAF Ω , where $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}, \text{gr}\}$, will be denoted by $S^\sigma(\Omega)$. Note that, given a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$, if we consider the corresponding AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, then the set of complete extensions of Λ that satisfy C does not always form a complete meet-semilattice. This is an important difference between CAF and AF, and it also holds for the CAF semantics reviewed in the previous section. Roughly speaking, the constraints may break the lattice by marking as unfeasible some extensions. As a consequence, even the grounded extension is not guaranteed to exist, as shown below.

Example 7. Consider the CAF $\Omega = \langle \{a, b, c\}, \{(a, b), (b, a), (b, c), (c, c)\}, \{\tau \Rightarrow a \wedge b\} \rangle$ derived from the AF Λ of Example 4 (shown in Fig. 2) by adding the strong constraint $\tau \Rightarrow a \wedge b$. As shown in Example 4, AF Λ has three complete extensions, $E_1 = \emptyset$, $E_2 = \{a\}$ and $E_3 = \{b\}$, but all extensions do not satisfy the constraint stating that both a and b must belong to them. Thus Ω has no complete extensions, and thus no grounded extension, under both Kleene and Lukasiewicz logics. \square

It is worth noting that for any CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$, we have that $\text{st}^K(\Omega) = \text{st}^L(\Omega)$ and that $S^K(\Omega) = S^L(\Omega)$ for $S \in \{\text{gr}, \text{co}, \text{pr}, \text{ss}\}$ whenever all constraints $\varphi \Rightarrow v$ and $v \Rightarrow \varphi$ in C are such that $i) v \neq u$ and $ii) \varphi$ does not contain the implication and equivalence operators (i.e., \Rightarrow and \Leftrightarrow). The reason is that the stable semantics is 2-valued and does not make use of the undefined truth value; moreover, the evaluation of constraints is the same under both logics if the implication and equivalence operators are not used in the body or in the head of the constraints (condition ii)), and the truth value v differs from u (condition i)). Intuitively, these two conditions exclude the case $u \Rightarrow u$ where the Kleene and Lukasiewicz logics differ (cf. Fig. 4).

4.1. Complexity of credulous and skeptical acceptance

In this section, we investigate the complexity of CAF under Kleene and Lukasiewicz interpretations of the constraints, and in particular of the implication operator whose semantics is different in the two logics. We recall that we use CA_S^K and SA_S^K (resp., CA_S^L and SA_S^L) to denote the credulous and skeptical acceptance problems under Kleene (resp., Lukasiewicz) logic.

We start with the following lemma that intuitively states that Kleene interpretation can be captured by Lukasiewicz logic.

Lemma 1. For every CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ there exists a CAF $\Omega' = \langle \mathcal{A}, \mathcal{R}, C' \rangle$ such that $S^K(\Omega) = S^L(\Omega')$ and C' can be derived from C in linear time.

Proof. It is sufficient to first rewrite every equivalence $a \Leftrightarrow b$ into $(a \Rightarrow b) \wedge (b \Rightarrow a)$ and then replace every implication $a \Rightarrow b$ with $(\neg a \Rightarrow b) \Rightarrow b$. \square

The following lemma states a monotonic property that holds under Kleene logic: if an extension (of an AF underlying a given CAF) satisfies a set of constraints, then the set of constraints continues to be satisfied for larger extensions.

Lemma 2. Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF and $E_1, E_2 \in \text{co}(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $E_1 \subseteq E_2$. Then, under Kleene logic, $E_1 \vDash C$ implies $E_2 \vDash C$.

Proof. First, recall that $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$ and that $E_1 = \text{Acc}(E_1) \subseteq \text{Acc}(E_2) = E_2$. Under Kleene logic every constraint κ can be rewritten in standard form as a disjunction of conjunction of literals, that is, in the form $\kappa : \tau \Rightarrow (\ell_1^1 \wedge \dots \wedge \ell_{n_1}^1) \vee \dots \vee (\ell_1^k \wedge \dots \wedge \ell_{n_k}^k)$. If $E_1 \vDash \kappa$, it means that there must be $i \in [1, k]$ such that $E_1 \vDash (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$. Moreover, as $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$ and $\text{Acc}(E_1) \subseteq \text{Acc}(E_2)$, it holds that $E_2 \vDash (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$ as well. \square

Observe that the previous lemma does not hold under Lukasiewicz interpretation of the implication operator. As an example, consider the CAF $\Omega = \langle \{a, b\}, \{(a, b), (b, a)\}, \{a \vee \neg a \Rightarrow u\} \rangle$. The underlying AF has three complete extensions $E_0 = \emptyset$, $E_1 = \{a\}$ and $E_2 = \{b\}$. We have that $\text{co}^K(\Omega) = \emptyset$, since the constraint is not satisfied by any extension of the underlying AF. In contrast, $\text{co}^L(\Omega) = \{E_0\}$ since the constraint is satisfied by E_0 , but not by E_1 and E_2 even if $E_0 \subseteq E_1$ and $E_0 \subseteq E_2$.

Although the presence of constraints in CAF breaks the meet-semilattice of complete extensions, reasoning under the grounded semantics remains tractable.

Proposition 1. The complexity of checking whether a CAF admits a grounded extension is in PTIME under both Kleene and Lukasiewicz logics.

Proof. Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF. A set of arguments $S \subseteq \mathcal{A}$ is the grounded extension of Ω if S is the grounded extension of $\langle \mathcal{A}, \mathcal{R} \rangle$ and $S \vDash C$. Computing the grounded extension S is in PTIME [48]. Checking whether the grounded extension $E \subseteq \mathcal{A}$ satisfies a given finite set of constraints C is also in PTIME (in the size of Ω), under both Kleene and Lukasiewicz logics, from which the statement follows. \square




Fig. 5. Representation of the AF $\langle \mathcal{A}', \mathcal{R}' \rangle$ in the CAF given in the construction of the hardness proof of Theorem 1 concerning the problem CA_{pr}^L . The dashed ellipse represents the starting AF $\langle \mathcal{A}, \mathcal{R} \rangle$, with $\mathcal{A} = \{a_1, \dots, a_n\}$, from which the construction of the AF $\langle \mathcal{A}', \mathcal{R}' \rangle$ is build.

Therefore, since if a grounded extension for a CAF exists then it is unique, computing the credulous (or, equivalently, the skeptical) acceptance of an argument under the grounded semantics is still polynomial.

However, the credulous and skeptical acceptance of an argument w.r.t a CAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$ may differ from that of the associated AF $\langle \mathcal{A}, \mathcal{R} \rangle$, independently of the semantics adopted, as shown in the following example.

Example 8. Continuing from Example 7, there are no arguments in the CAF that are credulously accepted under the complete semantics. In contrast, for the AF of Example 4, argument a is credulously accepted under the complete and preferred semantics, whereas argument b is credulously accepted under complete, preferred, stable and semi-stable semantics. Moreover, b is skeptically accepted under stable and semi-stable semantics, whereas arguments a and c are not skeptically accepted under any of the semantics considered in the paper. \square

As discussed earlier, the complete semantics in CAF may admit no extensions. This is analogous to what happens in AF for the stable semantics, where the requirement that the extensions must be total may not be satisfied by any set of arguments. Intuitively, the problem of non-existence of complete extensions in CAF may arise because the constraints may contradict the extensions prescribed by the complete semantics. As an example, consider a CAF whose underlying AF consists of a single (unattacked) argument a , and a (strong) constraint prescribing that the acceptance status of a must be false (i.e., $a \Rightarrow \perp$). Clearly, the complete extension of the underlying AF (that is, $\{a\}$) does not satisfy the constraint, and thus the CAF has no complete extensions. However, we can find special cases where it is possible to guarantee the existence of at least one complete extension for CAF. An interesting case is when $gr(\langle \mathcal{A}, \mathcal{R} \rangle) \models C$, i.e., when the constraints in C do not exclude the existence of the grounded extension. Notably, this condition can be checked in polynomial time (cf. Proposition 1) and implies the existence of complete, preferred, and semi-stable extensions, as it holds for AF. Moreover, thanks to the result of Lemma 2, the condition $gr(\langle \mathcal{A}, \mathcal{R} \rangle) \models C$ also guarantees that (i) under the Kleene’s logic the meet-semilattice of complete extensions is preserved, that is $co^K(\langle \mathcal{A}, \mathcal{R}, C \rangle) = co(\langle \mathcal{A}, \mathcal{R} \rangle)$, whereas (ii) under the Lukasiewicz’s logic the meet-semilattice exists, though we have that $co^L(\langle \mathcal{A}, \mathcal{R}, C \rangle) \subseteq co(\langle \mathcal{A}, \mathcal{R} \rangle)$ since extensions of the underlying AF which are larger than the grounded one could be filtered out as they may not satisfy the constraints (as illustrated in the example in the paragraph after the proof of Lemma 2).

In general, the fact that the grounded extension may not exist for CAFs impacts on the complexity of the skeptical acceptance problem under complete semantics (irrespective of the logic considered for interpreting the constraints), which cannot be longer decided by simply looking at the grounded extension as for the case of AFs (where an argument is skeptically accepted under complete semantics if and only if it is in the grounded extension). Similarly, credulous acceptance under preferred semantics for CAFs under Lukasiewicz logic can no longer be decided by checking credulous acceptance under complete semantics (under Kleene logic, the complexity of credulous acceptance under preferred semantics remains the same of that of AF thanks to the property stated in Lemma 2). In fact, it turns out that the complexity of the above-mentioned problems for CAF increases of one level in the polynomial hierarchy w.r.t. that for AF. In all the other cases we can show that the complexity of credulous and skeptical reasoning for CAF and AF coincides, as stated in the following theorem which provides tight complexity results for all problems considered.

Theorem 1. For any CAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$, the problem

- CA_S^σ is: (i) NP-complete for $S \in \{co, st\}$ and $\sigma \in \{K, L\}$,
 (ii) NP-complete for $S = pr$ and $\sigma = K$,
 (iii) Σ_2^p -complete for $S = pr$ and $\sigma = L$,
 (iv) Σ_2^p -complete for $S = sst$ and $\sigma \in \{K, L\}$.
- SA_S^σ is: (i) coNP-complete for $S \in \{co, st\}$ and $\sigma \in \{K, L\}$, and
 (ii) Π_2^p -complete for $S \in \{pr, sst\}$ and $\sigma \in \{K, L\}$.

Proof. (Hardness.) Except that for the coNP-hardness result concerning skeptical acceptance under complete semantics (irrespective of the logic considered for interpreting the constraints), and the Σ_2^p -hardness result concerning the credulous acceptance under preferred semantics and Lukasiewicz logic, which are considered below, the other hardness results derive from the fact that they hold for any CAFs $\langle \mathcal{A}, \Sigma, C \rangle$ where $C = \emptyset$, that is, for AFs. In fact, it has been shown that the complexity of the credulous and skeptical acceptance problems for AFs is [45,49,51,56]:

- under complete semantics, NP-complete and in PTIME, respectively;
- under stable semantics, NP-complete and coNP-complete, respectively;

- under preferred semantics, NP-complete and Π_2^p -complete, respectively;
- under semi-stable semantics, Σ_2^p -complete and Π_2^p -complete, respectively.

The lower bound for skeptical acceptance under complete semantics can be proved by reducing SA_{st} for AF to SA_{co}^* for CAF as follows. Given an AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ we build a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ where $C = \{a \wedge \neg a \Rightarrow \text{f} \mid a \in \mathcal{A}\}$. Observe that the set of stable extensions of Λ coincides with the set of complete extensions of Ω , as the constraints force to select only complete extensions not containing undefined arguments. Therefore, SA_{co}^* is coNP-hard. It is worth noting that the same strategy can be used to provide an alternative proof for the NP-hardness of CA_{co}^* .

The lower bound for credulous acceptance under preferred semantics and Lukasiewicz logic can be proved by reducing from the complement of the Π_2^p -complete problem of checking whether a given AF Λ is *coherent* [49], that is, checking whether $\text{pr}(\Lambda) = \text{st}(\Lambda)$. We use $\overline{\text{CH}}$ to denote the problem of checking whether a given AF Λ is not coherent. Hence, $\overline{\text{CH}}$ is Σ_2^p -complete.

Let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$. Assume w.l.o.g. that $\text{pr}(\Lambda) \neq \text{gr}(\Lambda)$. We show that $\overline{\text{CH}}(\Lambda)$ is true iff $CA_{pr}^L(\langle \mathcal{A}', \mathcal{R}', C \rangle, \phi)$ is true, where:

- $\mathcal{A}' = \mathcal{A} \cup \{\bar{a}, C_a \mid a \in \mathcal{A}\} \cup \{\phi, \bar{\phi}, \psi, \bar{\psi}\}$;
- $\mathcal{R}' = \mathcal{R} \cup \{(a, \bar{a}), (a, C_a), (\bar{a}, C_a), (C_a, \bar{\psi}), (a, \bar{\phi}) \mid a \in \mathcal{A}\} \cup \{(\bar{\psi}, \psi), (\bar{\phi}, \phi)\}$;
- $C = \{u \Rightarrow \psi; \psi \Rightarrow u\}$.

The AF $\langle \mathcal{A}', \mathcal{R}' \rangle$ of the above-defined CAF is shown in Fig. 5.

Given $\Lambda' = \langle \mathcal{A}', \mathcal{R}' \rangle$, there is a one-to-one correspondence between $\text{pr}(\Lambda)$ and $\text{pr}(\Lambda')$. Particularly, it holds that $\text{pr}(\Lambda) = \{E' \cap \mathcal{A} \mid E' \in \text{pr}(\Lambda')\}$ and $\text{pr}(\Lambda') = E \cup \{\bar{a} \mid a \in \text{Def}(E)\} \cup \{\bar{\psi} \mid E \in \text{st}(\Lambda)\} \cup \{\phi \mid E \in \text{pr}(\Lambda)\}$.

(\Rightarrow) $\overline{\text{CH}}(\Lambda)$ is a true instance, that is Λ is not coherent. Thus, there exists at least one preferred extension $E \in \text{pr}(\Lambda)$ s.t. $E \notin \text{st}(\Lambda)$. Thus, by construction there exists $E' \in \text{pr}(\Lambda')$ such that $\psi \in \mathcal{A}' \setminus (E' \cup \text{Def}(E'))$ and $\phi \in E'$. Observe that $E' \models C$ and thus $CA_{pr}^L(\langle \mathcal{A}', \mathcal{R}', C \rangle, \phi)$ is true.

(\Leftarrow) $\overline{\text{CH}}(\Lambda)$ is a false instance, that is Λ is coherent. Thus, all preferred extensions $E \in \text{pr}(\Lambda)$ are s.t. $E \in \text{st}(\Lambda)$. Thus, by construction, all preferred extensions $E' \in \text{pr}(\Lambda')$ contain both ϕ and $\bar{\psi}$. Thus, any preferred extension $E' \in \text{pr}(\Lambda')$ is such that $E' \not\models C$, and thus $\text{pr}(\langle \mathcal{A}', \mathcal{R}', C \rangle) = \emptyset$ and consequently $CA_{pr}^L(\langle \mathcal{A}', \mathcal{R}', C \rangle, \phi)$ is false.

(Membership.) We now provide the membership results for each considered semantics, problem, and logic. Let $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ be a CAF, $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ be an AF, and $a \in \mathcal{A}$ be the argument for which we want to decide either credulous or skeptical acceptance w.r.t. Ω .

- (CA_{co}^*). Recall that a complete extension of an AF is an admissible set that contains all the arguments it defends [47]. Hence, a guess-and-check strategy to decide whether a belongs to a complete extension of Ω is as follows. First, guess a set $S \subseteq \mathcal{A}$ of arguments containing a (in *PTIME*). Then, check that (i) S is an admissible set for Λ (in *PTIME*), (ii) S contains all the arguments that it defends w.r.t. Λ (in *PTIME*), and (iii) $S \models C$ (in *PTIME*). Therefore, CA_{co}^* is in NP.
- (SA_{co}^*). Using a strategy similar to that given for CA_{co}^* , it can be shown that the complementary problem of checking whether there exists a complete extension for Ω not containing a is in NP. Therefore SA_{co}^* is in coNP.
- (CA_{st}^*). Recall that a stable extension of an AF is a conflict-free set that attacks every other argument in the AF [47]. We say that a set S attacks an argument b if there is $c \in S$ such that c attacks b . A guess-and-check strategy to decide whether a belongs to a stable extension of Ω is as follows. First, guess a set $S \subseteq \mathcal{A}$ of arguments containing a (in polynomial time). Then, check that (i) S is a conflict-free set for Λ , (ii) S attacks each argument in $\mathcal{A} \setminus S$, and (iii) $S \models C$. Since all these steps can be accomplished in *PTIME*, it follows that CA_{st}^* is in NP.
- (SA_{st}^*). The complementary problem of checking whether there exists a stable extension for Ω not containing a is in NP. Therefore SA_{st}^* is in coNP.
- (CA_{pr}^L). A preferred extension of an AF is a maximal (w.r.t. \subseteq) admissible set for it [47]. A guess-and-check strategy for deciding whether a belongs to a preferred extension of Ω is as follows. First, guess a set $S \subseteq \mathcal{A}$ of arguments containing a (in polynomial time). Then, check that (i) S is an admissible set for Λ (in *PTIME*), (ii) $S \models C$ (in *PTIME*), and (iii) there is no admissible set S' for Λ such that $S' \supset S$ and $S' \models C$. It can be shown that (iii) is in coNP. Indeed, a guess-and-check strategy for the complementary problem of deciding whether there is an admissible set S' for Λ such that $S' \supset S$ and $S' \models C$ is as follows: guess a set $S' \subseteq \mathcal{A}$ such that $S' \supset S$, and check that S' is an admissible set for Λ and $S' \models C$ (in *PTIME*). Therefore, CA_{pr}^L is in Σ_2^p .
- (CA_{pr}^K). Recalling that a preferred extension of an AF is a \subseteq -maximal complete extension [47], a guess-and-check strategy for deciding whether a belongs to a preferred extension of Ω is as follows. First, guess a set $S \subseteq \mathcal{A}$ of arguments containing a . Then, check that (i) S is a complete extension for Λ (in *PTIME*), and (ii) $S \models C$ (in *PTIME*). The fact that it suffices to check that S is a complete (rather than a preferred) extension for Λ follows from the result of Lemma 2, which entails that also \subseteq -maximal complete extensions (w.r.t. the guessed one) satisfy the set constraints. Therefore, CA_{pr}^K is in NP.
- (SA_{pr}^L). The complementary problem of checking whether there exists a preferred extension for Ω not containing a is in Σ_2^p . Therefore SA_{pr}^L is in Π_2^p .
- (CA_{sst}^*). A semi-stable extension of an AF is an admissible set S such that $S \cup \text{Def}(S)$ is maximal (w.r.t. \subseteq) [51]. A guess-and-check strategy for deciding whether a belongs to a semi-stable extension of Ω is as follows. First, guess a set $S \subseteq \mathcal{A}$ of arguments containing a (in polynomial time). Then, check that (i) S is an admissible set for Λ (in *PTIME*), (ii) $S \models C$ (in *PTIME*), and (iii) there is no admissible set S' for Λ such that $S' \cup \text{Def}(S') \supset S \cup \text{Def}(S)$ and $S' \models C$. It can be shown that (iii) is in coNP.



Fig. 6. AF $\langle \mathcal{A}, \mathcal{R} \rangle$ underlying WAF Υ of Example 9 (and the PAF of Example 12).

A guess-and-check strategy for the complementary problem of deciding whether there is an admissible set S' for Λ such that $S' \cup Def(S') \supset S \cup Def(S)$ and $S' \vDash C$ is as follows: guess a set $S' \subseteq \mathcal{A}$ such that $S' \cup Def(S') \supset S \cup Def(S)$, and check that S' is an admissible set for Λ and $S' \vDash C$ (in *PTIME*). Therefore, CA_{sst}^* is in Σ_2^p .

- (SA_{sst}^*) . The complementary problem of checking whether there exists a semi-stable extension for Ω not containing a is in Σ_2^p . Therefore SA_{sst}^* is in Π_2^p . \square

5. Weak constrained AF

In this section, we present a generalization of CAF with *weak* constraints. Differently from the strong constraints previously discussed, weak constraints are propositional formulae that should be satisfied *if possible*. Specifically, weak constraints are logical formulae having the same syntax as strong constraints, but they do not necessarily all have to be satisfied, and we give preference to extensions that better satisfy them (called *best extensions*) according to a given criterion.

Definition 8. (Weak constrained AF) A *Weak constrained Argumentation Framework (WAF)* is a tuple $\Delta = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, where $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a CAF and \mathcal{W} is a set of weak constraints built from $\mathcal{L}_{\mathcal{A}}$.

The semantics of a WAF is defined by considering two possible criteria for selecting the preferable extensions w.r.t. weak constraints—only weak constraints are considered when selecting the preferable extensions since strong constraints must be all satisfied. The two criteria considered for assessing to which extent an extension satisfies a set of weak constraints are: (i) *maximal set* criterion, considering as preferable (or “best”) extensions the ones that satisfy a maximal set of weak constraints, and (ii) *maximum-cardinality* criterion, considering as preferable (or “optimal”) extensions the ones that satisfy a maximal number of weak constraints. Clearly, the selection of preferable extensions makes sense only for semantics admitting multiple extensions, that is, complete, preferred, stable, and semi-stable semantics. Thus, in the following, whenever we consider a generic semantics S , we refer to $S \in \{\text{co, pr, st, sst}\}$.

In the next subsections, after formally defining the meaning of a WAF under the maximal-set and maximum-cardinality semantics, we investigate the complexity of credulous and skeptical reasoning in the new framework.

5.1. Maximal-set semantics

The semantics of a WAF using the maximal-set criterion is defined as follows.

Definition 9 (*Maximal-Set Semantics*). Given a WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, an S^σ -extension E for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ is a maximal-set S^σ -extension (ms- S^σ -extension) for Υ if, let $\mathcal{W}_E \subseteq \mathcal{W}$ be the set of weak constraints that are satisfied by E (that is, $E \vDash \mathcal{W}_E$), there is no S^σ -extension F for $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$ and $\mathcal{W}_F \subseteq \mathcal{W}$ such that $F \vDash \mathcal{W}_F$ and $\mathcal{W}_E \subset \mathcal{W}_F$.

Given a semantics S and a logic $\sigma \in \{K, L\}$ for the interpretation of the constraints, ms- S^σ denotes the maximal-set version of S^σ (e.g., ms-co^K denotes the ms complete semantics under Kleene interpretation).

Example 9. Consider the WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, \mathcal{C}, \mathcal{W} \rangle$, where $\langle \mathcal{A} = \{a, b, c, d\}, \mathcal{R} = \{(a, b), (b, a), (c, d), (d, c)\} \rangle$ is the AF shown in Fig. 6, $\mathcal{C} = \emptyset$, and $\mathcal{W} = \{w_1 = c \Rightarrow \text{f}, w_2 = a \vee \neg a \Rightarrow \text{u}\}$. Under Lukasiewicz logic w_1 and w_2 state that, preferably, c should be false, and a should be undefined, respectively. Under Kleene logic, w_1 states that c should be preferably false, whereas w_2 becomes useless as it is never satisfied (recall that it coincides with $a \vee \neg a \Rightarrow \text{f}$).

Υ has 9 complete extensions: $E_0 = \{\}$, $E_1 = \{a\}$, $E_2 = \{b\}$, $E_3 = \{c\}$, $E_4 = \{d\}$, $E_5 = \{a, c\}$, $E_6 = \{a, d\}$, $E_7 = \{b, c\}$ and $E_8 = \{b, d\}$. In particular, E_0 is the grounded extension, whereas E_5, E_6, E_7, E_8 are preferred, stable, and semi-stable extensions of $\langle \mathcal{A}, \mathcal{R}, \mathcal{C} \rangle$. These are also extensions of AF $\langle \mathcal{A}, \mathcal{R} \rangle$, since $\mathcal{C} = \emptyset$.

Regarding the satisfaction of weak constraints, first observe that argument c is false in E_4, E_6 and E_8 , whereas argument a is undefined in E_0, E_3 and E_4 . Thus, since under Lukasiewicz logic w_1 states that c should be preferably false, and w_2 states a should be preferably undefined, under Lukasiewicz interpretation we have that $E_0 \vDash \{w_2\}$, $E_3 \vDash \{w_2\}$, $E_4 \vDash \{w_1, w_2\}$, $E_6 \vDash \{w_1\}$, and $E_8 \vDash \{w_1\}$, whereas the other complete extensions do not satisfy any constraint. Therefore, the maximal-set preferred (stable, semi-stable) extensions are E_6 and E_8 (i.e. ms-pr^L(Υ) = ms-st^L(Υ) = ms-ss^L(Υ) = $\{E_6, E_8\}$), whereas there is only one maximal-set complete extension, which is E_4 (i.e. ms-co^L(Υ) = $\{E_4\}$).

Considering Kleene interpretation, recalling that c is false in E_4, E_6 and E_8 , and that w_1 states that c should be preferably false (whereas w_2 is useless in this case), we have that ms-pr^K(Υ) = ms-st^K(Υ) = ms-ss^K(Υ) = $\{E_6, E_8\}$, whereas ms-co^L(Υ) = $\{E_4, E_6, E_8\}$. \square

Table 1

Complexity of CA_S^σ under complete (co), stable (st), preferred (pr), and semi-stable (sst) semantics. For any complexity class C , we use C -c to denote C -complete, and C to denote Σ_2^P -hard and in C . All the results except those for AF are new.

S	Framework														
	AF		CAF		NCAF		(S)WAF		WAF		LWAF		SWAF		NWAF
	CA_S	CA_S^K	CA_S^L	CA_S^c	CA_{ms-S}^K	CA_{ms-S}^L	CA_{ms-S}^K	CA_{ms-S}^L	CA_{ms-S}^K	CA_{ms-S}^L	CA_S^K	CA_S^L	CA_{ms-S}^K	CA_{ms-S}^L	CA_{ms-S}^c
Semantics	co	NP-c	NP-c	NP-c	NP-c	Σ_2^P -c	Σ_2^P -c	Θ_2^P -c	Θ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Σ_2^P -c
	st	NP-c	NP-c	NP-c	NP-c	Σ_2^P -c	Σ_2^P -c	Θ_2^P -c	Θ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Δ_2^P -c	Σ_2^P -c
	pr	NP-c	NP-c	Σ_2^P -c	NP-c	Σ_2^P -c	Σ_3^P -c	Θ_2^P -c	Θ_3^P -c	Δ_2^P -c	Δ_3^P -c	Δ_2^P -c	Δ_3^P -c	Δ_3^P -c	Σ_2^P -c
	sst	Σ_2^P -c	Σ_2^P -c	Σ_2^P -c	Σ_2^P -c	Σ_2^P -c	Σ_4^P -c	Θ_3^P -c	Θ_4^P -c	Σ_2^P -c	Δ_3^P -c	Δ_3^P -c	Δ_3^P -c	Δ_3^P -c	Σ_2^P -c

Intuitively, given an S^σ -extension, checking satisfaction of a maximal-set of weak constraints means ensuring that no other S^σ -extension is better according to that criterion. This is an additional source of complexity that makes, in most cases, credulous and skeptical reasoning in WAFs one level higher in the polynomial-time hierarchy than CAFs.

Theorem 2. For any WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- CA_{ms-S}^σ is: (i) Σ_2^P -complete for any semantics $S \in \{co, st, pr, sst\}$ and $\sigma = K$,
(ii) Σ_2^P -complete for $S \in \{co, st\}$ and $\sigma = L$,
(iii) Σ_2^P -hard and in Σ_3^P for $S = pr$ and $\sigma = L$, and
(iv) Σ_3^P -complete for $S = sst$ and $\sigma = L$.
- SA_{ms-S}^σ is: (i) Π_2^P -complete for $S \in \{co, st, pr, sst\}$ and $\sigma = K$,
(ii) Π_2^P -complete for $S \in \{co, st\}$ and $\sigma = L$, and
(iii) Π_3^P -complete for $S \in \{pr, sst\}$ and $\sigma = L$.

5.2. Maximum-cardinality semantics

Maximum-cardinality semantics for WAFs prescribes as preferable extensions those satisfying the highest number of weak constraints. This is similar to the semantics of weak constraints in DLV [9] where, in addition, each constraint has assigned a weight.

Definition 10 (Maximum-Cardinality Semantics). Given a WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, an S^σ -extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ is a maximum-cardinality S^σ -extension (mc S^σ -extension) for Υ if, let $\mathcal{W}_E \subseteq \mathcal{W}$ be the set of weak constraints in \mathcal{W} that are satisfied by E , there is no S^σ -extension F for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and $\mathcal{W}_F \subseteq \mathcal{W}$ such that $F \models \mathcal{W}_F$ and $|\mathcal{W}_E| < |\mathcal{W}_F|$.

The next theorem provides complexity results for credulous and skeptical reasoning in WAF with maximum-cardinality semantics under Kleene and Lukasiewicz interpretation of constraints.

Theorem 3. For any WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- CA_{mc-S}^σ is: (i) Θ_2^P -complete for any semantics $S \in \{co, st, pr\}$ and $\sigma = K$,
(ii) Σ_2^P -hard and in Θ_3^P for semantics $S = sst$ and $\sigma = K$,
(iii) Θ_3^P -complete for any semantics $S \in \{co, st\}$ and $\sigma = L$, and
(iv) Σ_2^P -hard and in Θ_3^P for any semantics $S \in \{pr, sst\}$ and $\sigma = L$.
- SA_{mc-S}^σ is: (i) Θ_2^P -complete for any semantics $S \in \{co, st\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^P -hard and in Θ_3^P for any semantics $S \in \{pr, sst\}$ and $\sigma \in \{K, L\}$.

The complexity results stated so far, as well as those that will be given in the next sections, are summarized in Tables 1 and 2. It turns out that, under standard complexity assumptions, computing credulous and skeptical acceptance in WAFs under maximum-cardinality semantics is easier than using maximal-set semantics. Roughly speaking, this follows from the fact that a binary search strategy can be used for deciding whether the cardinality of a set of constraints satisfied by an S^σ -extension containing a given argument is maximum.

6. Stratified constraints in WAF

In this section, we explore the impact of considering a form of stratification over the set of constraints. We first consider WAF where weak constraints define a partial order, and then focus on linearly ordered sets of constraints. Similarly to other contexts such as logic programming, the concept of stratification allows defining classes of WAF with different complexity and expressivity. In particular, after formally defining the syntax and semantics of stratified WAF, we investigate their complexity under the maximal-set and maximum-cardinality interpretation of the weak constraints, showing that the complexity does not impact on the maximal-set interpretation while it increases with the maximum-cardinality interpretation. In general, for both interpretations, the stratification

Table 2

Complexity of SA_S^σ under complete (co), stable (st), preferred (pr), and semi-stable (sst) semantics. For any complexity class C , we use C -c to denote C -complete, and C to denote Π_2^p -hard and in C . All the results except those for AF are new.

		Framework									
		AF	CAF	NCAF	(S)WAF		WAF	LWAF	SWAF	NWAF	
S		SA_S	SA_S^*	SA_S^*	SA_{msS}^K	SA_{msS}^L	SA_{mcS}^*	SA_S^K	SA_S^L	SA_{msS}^*	SA_{msS}^*
Semantics	co	P	coNP-c	coNP-c	Π_2^p -c	Π_2^p -c	Θ_3^p -c	Δ_3^p -c	Δ_2^p -c	Δ_2^p -c	Π_2^p -c
	st	coNP-c	coNP-c	coNP-c	Π_2^p -c	Π_2^p -c	Θ_3^p -c	Δ_3^p -c	Δ_2^p -c	Δ_2^p -c	Π_2^p -c
	pr	Π_2^p -c	Π_2^p -c	Π_2^p -c	Π_2^p -c	Π_2^p -c	Θ_3^p -c	Π_2^p -c	Δ_3^p -c	Δ_3^p -c	Π_2^p -c
	sst	Π_2^p -c	Π_2^p -c	Π_2^p -c	Π_2^p -c	Π_3^p -c	Θ_3^p -c	Π_2^p -c	Δ_3^p -c	Δ_3^p -c	Π_2^p -c

approach provides a more flexible way to express different strata of constraints. Finally, we focus on a particular form of stratified WAFs where every stratum is a singleton, providing tighter complexity bounds under Kleene logic.

The partial order on the set \mathcal{W} of weak constraints is defined by partitioning \mathcal{W} into strata $\mathcal{W}_1, \dots, \mathcal{W}_n$ (with $n \geq 1$) so that weak constraints in a stratum i have higher priority than those in a stratum $j > i$.

Definition 11. (Stratified WAF) A *Stratified Weak constrained Argumentation Framework (SWAF)* is a tuple $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ where $\langle \mathcal{A}, \mathcal{R}, C \rangle$ is a CAF and \mathcal{W} is a list of sets of weak constraints $(\mathcal{W}_1, \dots, \mathcal{W}_n)$ built from $\mathcal{L}_{\mathcal{A}}$.

Note that whenever $n = 1$ we have a unique stratum and, then, SWAFs coincide with standard WAFs, which in turn implies that SWAFs are at least as hard as WAFs from a computational standpoint.

Regarding the semantics of a SWAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the underlying idea is that weak constraints are applied one stratum at a time. Therefore, given a set S of S^σ -extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$, the best/optimal S^σ -extensions are obtained by first computing the set $S_1 \subseteq S$ which are best/optimal solutions w.r.t. \mathcal{W}_1 , then the set $S_2 \subseteq S_1$ of S^σ -extensions which are best/optimal solutions w.r.t. \mathcal{W}_2 is selected, and so on.

Definition 12 (SWAF Semantics). Let $Y = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ be a SWAF and S^σ a semantics under logic σ . An S^σ -extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ is an ms/mcS^σ -extension for Y if:

- E is an ms/mcS^σ -extension for $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}_1 \rangle$, and
- if $n > 1$ there is no ms/mcS^σ -extension E' for $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}_1 \rangle$ such that E' is a ms/mcS^σ -extension for $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_2, \dots, \mathcal{W}_n) \rangle$ and E is not.

Thus, to determine the set of ms/mcS^σ -extensions S_n for $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, we first compute the set S_0 of S^σ -extensions for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and next, for each stratum $i \in [1, n]$, we compute the set of ms/mcS^σ -extensions S_i from S_{i-1} , by discarding the extensions which do not satisfy a *maximal set/maximal number* of constraints in \mathcal{W}_i .

Example 10. Consider the SWAF derived from the AF Λ of Example 1 by adding the following list of sets of weak constraints $(\{a \wedge b \wedge c \Rightarrow f\}, \{t \Rightarrow a\}, \{t \Rightarrow b\}, \{t \Rightarrow c\})$. Since each stratum contains only one weak constraint, maximal-set and maximum-cardinality semantics give the same result. The weak constraints are applied one (set) at a time to discard extensions. After applying the first constraint the extension containing a, b and c is discarded. At the second step only extensions containing a are selected from the ones selected at the first step. At the third step only the extension containing a and b is selected. Thus, the best/optimal extension is the one containing exactly a and b .

Note that, assuming that weak constraints are not stratified, we would have the three extensions $\{a, b\}, \{a, c\}$ and $\{b, c\}$ under both maximal-set and maximum-cardinality preferred and stable semantics. \square

The next two theorems state the complexity for SWAF under the maximal-set and maximum-cardinality interpretation of weak constraints, respectively.

Theorem 4. For any SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem

- CA_{ms-S}^σ is: (i) Σ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Σ_3^p -hard and in Σ_3^p for $S = \text{pr}$ and $\sigma = L$,
(iii) Σ_2^p -complete for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = K$, and
(iv) Σ_3^p -complete for $S = \text{sst}$ and $\sigma = L$.
- SA_{ms-S}^σ is: (i) Π_2^p -complete for $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma = K$,
(ii) Π_2^p -complete for $S \in \{\text{co}, \text{st}\}$ and $\sigma = L$, and
(iii) Π_3^p -complete for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.

Theorem 5. For any SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem:

- CA_{mcS}^σ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma = K$,
(ii) Σ_2^p -hard and in Δ_3^p for semantics $S = \text{sst}$ and $\sigma = K$,
(iii) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{pr}\}$ and $\sigma = L$, and
(iv) Σ_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.
- SA_{mcS}^σ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

In light of Theorem 4 and Theorem 5, it follows that the introduction of multiple layers of weak constraints under maximal-set semantics does not increase the computational complexity, though it provides a more general and flexible approach to express constraints. In contrast, in the case of maximal-cardinality semantics, introducing multiple layers of weak constraints generally increases the complexity w.r.t. a single layer of constraints.

A particular form of SWAF are the ones used in Example 10, where every stratum is a singleton, meaning that weak constraints define a total order.

Definition 13 (LWAF). A SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ is said to be linearly ordered if every \mathcal{W}_i ($1 \leq i \leq n$) contains only one weak constraint.

Observe that for linearly ordered SWAFs, that we denote by LWAF, $CA_{msS}^\sigma = CA_{mcS}^\sigma$ and $SA_{msS}^\sigma = SA_{mcS}^\sigma$. Thus, for this class of constrained AFs, we simply use the notations CA_S^σ and SA_S^σ to denote the credulous and skeptical acceptance problems, respectively.

Theorem 6. For any LWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem:

- CA_S^σ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma = K$,
(ii) Σ_2^p -complete for semantics $S = \text{sst}$ and $\sigma = K$,
(iii) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{pr}\}$ and $\sigma = L$, and
(iv) Σ_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.
- SA_S^σ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -complete for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = K$,
(iii) Π_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.

Thus, limiting the number of weak constraints in each layer does not result in reducing the complexity bounds under Lukasiewicz's logic, i.e., SWAF and LWAF have the same complexity bounds.

Finally, it is worth noting that extending CAF by stratifying the constraints does not make sense as all constraints must be satisfied.

7. CAF and WAF with denial constraints

In several contexts (e.g., database, logic programming, inconsistent knowledge management) constraints are expressed by denial constraints. In this section, we investigate credulous and skeptical acceptance when constraints are expressed by negative constraints only. In the following, we use the acronym NCAF (resp., NWAF) to denote a CAF (resp., WAF) where all constraints are defined as denials.

Example 11. The WAF of Example 3 is an NWAF, since the constraints in C are denials and those in \mathcal{W} can be equivalently written as the denials $\neg a \Rightarrow f$, $\neg b \Rightarrow f$, $\neg c \Rightarrow f$, and $\neg d \Rightarrow f$. Moreover, if $\mathcal{W} = \emptyset$ then we obtain an NCAF. \square

The following lemma states that for NWAF (and thus NCAF) the semantics of denial constraints under Kleene and Lukasiewicz logics coincide.

Lemma 3. For any NWAF Υ and semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}, \text{gr}\}$, it holds that $S^K(\Upsilon) = S^L(\Upsilon)$.

As a consequence, the result of Lemma 2 also holds for NWAF and NCAF under both Lukasiewicz and Kleene semantics as stated below.

Lemma 4. Let $\Upsilon = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ be a NWAF, $E_1, E_2 \in \text{co}(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $E_1 \subseteq E_2$, and $\mathcal{W}' \subseteq \mathcal{W}$. Then, under both Kleene and Lukasiewicz logics, $E_1 \models C \cup \mathcal{W}'$ implies $E_2 \models C \cup \mathcal{W}'$.

The following theorem shows that, if only denial constraints are used, the complexity of the credulous acceptance problem under preferred semantics does not increase for NCAF w.r.t. that for AF, which is different from what happens for (general) CAF (see Theorem 1). Moreover, the complexity of the skeptical acceptance problem for NCAF does not change w.r.t. that for CAF.

Theorem 7. For any NCAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$, the problem

- CA_S^σ is: (i) NP-complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma \in \{K, L\}$,
(ii) Σ_2^p -complete for $S = \text{sst}$ and $\sigma \in \{K, L\}$.
- SA_S^σ is: (i) coNP-complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -complete for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Finally, as stated next, the introduction of weak denial constraints increases the complexity of one level in the polynomial hierarchy w.r.t. that of NCAF, for credulous acceptance under complete, stable and preferred semantics, as well as for skeptical acceptance under complete and stable semantics, independently of the chosen logic.

Theorem 8. For any NAAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- CA_{msS}^σ is Σ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.
- SA_{msS}^σ is Π_2^p -complete for $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Overall, the use of denials makes the two logics equally expressive regardless of the interpretation of constraints as strong or weak constraints. Moreover, under semi-stable semantics, employing only denials (i.e., considering NAAF) results in a lower complexity of credulous acceptance compared to the case of general constraints (WAF) under Lukasiewicz logic. In fact, the complexity of NCAF and NAAF coincide under semi-stable semantics.

8. Encoding preferences through WAF

Several extensions of Dung's framework for handling preferences over arguments have been proposed [10,11,13,14,27,44,90]. In this section, after recalling the syntax of preference-based AF (denoted by PAF), we first propose a new semantics for PAF based on a well-known semantics for Answer-Set Programming (ASP) with preferences, called Answer Set Optimization (ASO) [34], and then show that such PAFs can be encoded in WAF.

Definition 14. A *Preference-based Argumentation Framework (PAF)* is a triple $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$ such that $\langle \mathcal{A}, \mathcal{R} \rangle$ is an AF and \succ is a strict partial order (i.e. an irreflexive, asymmetric and transitive relation) over \mathcal{A} , called preference relation.

For arguments a and b , $a \succ b$ means that a is better than b . As discussed in Section 9, a “best extensions” semantics approach for PAF has been proposed in [14,71], where classical argumentation semantics are used to obtain extensions of the underlying AF $\langle \mathcal{A}, \mathcal{R} \rangle$ and then the preference relation \succ is used to obtain a preference relation \sqsupseteq over such extensions, so that the *best* extensions w.r.t. \sqsupseteq are eventually selected. With the same spirit of selecting the best extensions by following an induced relation from the user-defined preferences, the *Answer Set Optimization (ASO)* approach has been proposed [34], whose semantics is based on the *degree* to which preferences are satisfied. Thus, we propose an intuitive PAF semantics that extends the ASO approach to deal with the fact that argumentation semantics are 3-valued. Then, we show that any PAF under this approach of handling preferences can be equivalently rewritten (in terms of extensions) to a WAF under maximal-set semantics.


For any PAF $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$ under ASO semantics, the set of preferences \succ determine a preference ordering on the set of S -extensions of the underlying AF $\langle \mathcal{A}, \mathcal{R} \rangle$, for any semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}\}$. First, we need to identify the case where the preference is satisfied w.r.t. a particular extension. To this end, given a PAF $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$, an S -extension $E \in S(\langle \mathcal{A}, \mathcal{R} \rangle)$, and a preference $a \succ b$, then $a \succ b$ is said to be *satisfied* w.r.t. E iff $a \in E$ or $b \in \text{Def}(E)$. Thus, it is possible to define the satisfaction degree of $a \succ b$ in E by setting $d_E(a \succ b) = 1$ if $a \in E$ or $b \in \text{Def}(E)$, $d_E(a \succ b) = 0$ otherwise. Let $\langle \mathcal{A}, \mathcal{R}, \{a_1 \succ b_1, \dots, a_n \succ b_n\} \rangle$ be a PAF and $E \in S(\langle \mathcal{A}, \mathcal{R} \rangle)$ be an S -extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ under semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}\}$. We say that E induces a satisfaction vector $D_E = \langle d_E(a_1 \succ b_1), \dots, d_E(a_n \succ b_n) \rangle$, where the position of preferences in the list is determined according to a predefined order (e.g. the lexicographic order or the order according to which they are listed).

We illustrate the ASO semantics for preferences in the following example.

Example 12. Consider the PAF $\langle \mathcal{A} = \{a, b, c, d\}, \mathcal{R} = \{(a, b), (b, a), (c, d), (d, c)\} \{a \succ b, b \succ c\} \rangle$ where the AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ is shown in Fig. 6. Under preferred and stable semantics, for the AF $\langle \mathcal{A}, \mathcal{R} \rangle$ there are 4 extensions: $E_1 = \{a, c\}$, $E_2 = \{a, d\}$, $E_3 = \{b, c\}$, and $E_4 = \{b, d\}$. Comparing the four extensions with respect to the two preferences, we get the following satisfaction vectors:

- $D_{E_1} = \langle 1, 0 \rangle$, as $d_{E_1}(a \succ b) = 1$ (since $a \in E_1$), and $d_{E_1}(b \succ c) = 0$ (since neither $b \in E_1$ nor $c \in \text{Def}(E_1)$);
- $D_{E_2} = \langle 1, 1 \rangle$, as $d_{E_2}(a \succ b) = 1$ (since $a \in E_2$) and $d_{E_2}(b \succ c) = 1$ (since $c \in \text{Def}(E_2)$);
- $D_{E_3} = \langle 0, 1 \rangle$, as $d_{E_3}(a \succ b) = 0$ (since neither $a \in E_3$ nor $b \in \text{Def}(E_3)$) and $d_{E_3}(b \succ c) = 1$ (since $b \in E_3$); and
- $D_{E_4} = \langle 0, 1 \rangle$, as $d_{E_4}(a \succ b) = 0$ (since neither $a \in E_4$ nor $b \in \text{Def}(E_4)$) and $d_{E_4}(b \succ c) = 1$ (since $b \in E_4$). \square

We extend the preorder on satisfaction degrees to preorders on satisfaction vectors and extensions as follows.

Fig. 7. AF $\langle \mathcal{A}, \mathcal{R} \rangle$ underlying PAF Δ of Example 14.

Definition 15 (PAF Semantics). Let $\Delta = \langle \mathcal{A}, \mathcal{R}, \{a_1 > b_1, \dots, a_n > b_n\} \rangle$ be a PAF, and $E, F \in S(\langle \mathcal{A}, \mathcal{R} \rangle)$ be two S -extensions for $\langle \mathcal{A}, \mathcal{R} \rangle$ under semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}\}$. We say that $E \sqsupseteq F$ iff $d_E(a_i > b_i) \geq d_F(a_i > b_i)$ for each $i \in [1, n]$, and write $E \sqsupset F$ iff $E \sqsupseteq F$ and $F \not\sqsupseteq E$. Moreover, the best S -extensions of Δ (denoted as $S(\Delta)$) are the extensions $E \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ such that there is no $F \in \sigma(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $F \sqsupset E$.

Example 13. Continuing with Example 12, comparing the D_{E_i} -vectors associated with extensions E_i , with $i \in [1..4]$, we have the E_2 is the best one since $E_2 \sqsupset E_1$, $E_2 \sqsupset E_3$, and $E_2 \sqsupset E_4$. \square

Example 14. Consider the PAF $\Delta = \langle \mathcal{A} = \{a, b, c\}, \mathcal{R} = \{(x, y) \mid x, y \in \mathcal{A} \wedge x \neq y\}, \{(a > b), (b > c), (a > c)\} \rangle$. The preferred extensions for the underlying AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$ (shown in Fig. 7), obtained from Δ by ignoring the preferences, are $\text{pr}(\Lambda) = \{E_1 = \{a\}, E_2 = \{b\}, E_3 = \{c\}\}$. The satisfaction vector of E_1, E_2 and E_3 is $D_{E_1} = \langle 1, 1, 1 \rangle$, $D_{E_2} = \langle 0, 1, 1 \rangle$, and $D_{E_3} = \langle 1, 0, 0 \rangle$, respectively. As $E_1 \sqsupset E_2$ and $E_1 \sqsupset E_3$, we have that E_1 is the only best preferred extension of Δ . \square

Theorem 9. For any PAF $\Delta = \langle \mathcal{A}, \mathcal{R}, > \rangle$ and semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{ss}\}$, there exists a WAF Υ_Δ (derivable from Δ in linear time) s.t. $S(\Delta) = \text{ms-}S^*(\Upsilon_\Delta)$.

Proof. Consider the case of Kleene logic. Given a PAF $\Delta = \langle \mathcal{A}, \mathcal{R}, > \rangle$, we denote by $\Upsilon_\Delta = \langle \mathcal{A}, \mathcal{R}, \emptyset, \mathcal{W}_\Delta \rangle$ the WAF derived from Δ by replacing every preference $a > b$ with the constraint $\omega_{a,b} : \tau \Rightarrow (b \Rightarrow a)$. Thus, observe that, for any S -extension $E \in S(\langle \mathcal{A}, \mathcal{R} \rangle)$ we have that $E \vDash \omega_{a,b}$ iff $d_E(a > b) = 1$.

(\Rightarrow) Reasoning by contradiction, assume that there exists $E \in S(\Lambda)$, with $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, such that $E \in S(\Delta) \setminus \text{ms-}S(\Upsilon_\Delta)$. So there must exist $F \in S(\Lambda) \cap \text{ms-}S(\Upsilon_\Delta)$ s.t. $F \vDash \mathcal{W}_F \subseteq \mathcal{W}_\Delta$, $E \vDash \mathcal{W}_E \subseteq \mathcal{W}_\Delta$, and $\mathcal{W}_F \supset \mathcal{W}_E$. Thus, the satisfaction vectors for E and F , respectively $D_E = \langle d_E^1, \dots, d_E^n \rangle$ and $D_F = \langle d_F^1, \dots, d_F^n \rangle$, are such that:

- $\forall w_i \in \mathcal{W}_E, d_E^i = d_F^i = 1$;
- $\forall w_i \in \mathcal{W} \setminus \mathcal{W}_E, d_E^i = 0$;
- $\forall w_j \in \mathcal{W} \setminus \mathcal{W}_F, d_E^j = d_F^j = 0$;
- $\forall w_j \in \mathcal{W}_F \setminus \mathcal{W}_E, d_E^j = 0$ and $d_F^j = 1$.

Thus $F \sqsupset E$, that is an absurd.

(\Leftarrow) Reasoning by contradiction, assume that there exists $E \in S(\Lambda)$ s.t. $E \in \text{ms-}S(\Upsilon_\Delta) \setminus S(\Delta)$. So there must exist $F \in S(\Lambda) \cap S(\Delta)$ s.t. $F \sqsupset E$. This implies that, $F \vDash \mathcal{W}_F = \{\omega_{a,b} \in \mathcal{W}_\Delta \mid d_F(a > b) = 1\}$ and $E \vDash \mathcal{W}_E = \{\omega_{a,b} \in \mathcal{W}_\Delta \mid d_E(a > b) = 1\}$. Thus, $\mathcal{W}_E \subset \mathcal{W}_F$ and thus $F \sqsupset E$, that is an absurd.

As by Lemma 1 any constraint under Kleene logic can be equivalently rewritten into a new one under Lukasiewicz logic, the result also follows under Lukasiewicz logic. \square

Example 15. Continuing with the previous example, the WAF derived from Δ is $\Upsilon_\Delta = \langle \mathcal{A}, \mathcal{R}, \emptyset, \mathcal{W}_\Delta = \{w_1, w_2, w_3\} \rangle$ where:

$$w_1 : \tau \Rightarrow (b \Rightarrow a), \quad w_2 : \tau \Rightarrow (c \Rightarrow b), \quad \text{and} \quad w_3 : \tau \Rightarrow (c \Rightarrow a).$$

We have that $E_1 \vDash w_1, w_2, w_3$, $E_2 \vDash w_2, w_3$, while $E_3 \vDash w_1$. Thus, E_1 is the only best preferred extension of Δ and the only ms-pr extension of Υ_Δ . \square

9. Related work

We start our discussion of related work by observing that an important difference between the semantics of CAF introduced in (Coste-Marquis et al. (2006)) [41] and our work concerns the meaning of constraints. Indeed, constraints in [41] (under Kleene logic) are imposed on the admissibility of sets of arguments (i.e., over admissible sets) that are at the basis of S -extensions, with $S \in \{\text{gr}, \text{co}, \text{pr}, \text{st}, \text{ss}\}$. As a consequence, a drawback of this approach is that S -extensions of a given CAF are no longer guaranteed to be S -extensions of the underlying AF, that is, we may have $E \in S(\langle \mathcal{A}, \mathcal{R}, C \rangle)$, but $E \notin S(\langle \mathcal{A}, \mathcal{R} \rangle)$. Differently, the semantics proposed in our work under Kleene (as well as under Lukasiewicz) logic prescribes S -extensions that are S -extensions of underlying AF.

Besides being related to the proposals for CAF in [41,17], our work is also related to the approach presented in [29] that provides a method for generating non-empty conflict-free extensions for CAFs. Constraints have been also used in the context of dynamic AFs to refer to the enforcement of some change [46]. In this context, extension enforcement aims at modifying an AF to ensure that a given set of arguments becomes (part of) an extension for the chosen semantics [23,42,93,81]. This is different from our approach where integrity constrains are used to discard unfeasible solutions (extensions), without enforcing that a new set of arguments becomes an extension.

Weak constraints allow for selecting “best” or “optimal” extensions satisfying some conditions on arguments, if possible. This can be viewed as expressing a kind of preference over the set of extensions. Dung’s framework has been extended in several ways for allowing preferences over arguments [11,70,78]. Two main approaches have been proposed to handle preferences in argumentation, where a Preference-based Argumentation Framework (PAF) is defined as a triple $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$ (cf. Definition 14). The first approach considers AF-based semantics and consists in first defining a defeat relation \mathcal{R}_i that combines attacks in \mathcal{R} and preferences in \succ , and then applying the usual semantics on the AF $\langle \mathcal{A}, \mathcal{R}_i \rangle$. Here \mathcal{R}_i (with $i \in [1,4]$) denotes one of the four mappings proposed in the literature [11,14,71]. However, in some cases these semantics fail to capture the expected meaning as, in some cases, extensions of the resulting PAF could not be conflict-free w.r.t. the original AF. We point out that for these PAF semantics the complexity of acceptance problems does not increase as the mapping to AF (i.e., building \mathcal{R}_i) is polynomial time. Anyway, it is clear that any PAF $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$ under this approach can be encoded into a WAF $\langle \mathcal{A}, \mathcal{R}_i, \emptyset, \emptyset \rangle$ under any semantics.

The second approach to handle preferences relies on a “best extensions” semantics for PAF [14,71]. In particular, given a PAF $\langle \mathcal{A}, \mathcal{R}, \succ \rangle$, classical argumentation semantics are used to obtain extensions of the underlying AF $\langle \mathcal{A}, \mathcal{R} \rangle$, and then the preference relation \succ is used to obtain a preference relation \sqsupseteq over such extensions, so that the *best* extensions w.r.t. \sqsupseteq are eventually selected. Clearly, \sqsupseteq is not trivial for multiple-status semantics only (for the grounded semantics, its extension is trivially the best one). Even if not excluded from the complexity standpoint, as the semantics of this approach is to compare pairs of extensions to filter-out those that are not best, it is not straightforward to encode a PAF into a WAF sharing the same underlying AF and having the same set of extensions.

Preferences can be also expressed in value-based AFs [25,50], where each argument is associated with a numeric value, and a set of possible orders (preferences) among the values is defined. In [52] weights are associated with attacks, and new semantics extending the classical ones on the basis of a given numerical threshold are proposed. [43] extends [52] by considering other aggregation functions over weights apart from sum. Except for weighted solutions under grounded semantics (that prescribes more than one weighted solution), the complexity bounds of credulous and skeptical reasoning in the above-considered frameworks are lower or equal than those of WAFs, which suggests that WAFs are more expressive and can be used to model those frameworks. In this regard, we have proposed a novel (3-valued) PAF semantics based on the (2-valued) ASO semantics for answer set programs [34]. Differently from [34], w.l.o.g. we have not considered preferences that are conditioned by a boolean conjunctive formula in the body, e.g. $a > b \leftarrow c \wedge \neg d$. In fact, these preferences can be encoded by imposing an additional condition concerning the body, that is $c \in E \wedge d \in Def(E)$ in the case of the preference in the previous example. ASO preferences can be also generalized to express meta-preferences specifying a sequence of pairwise disjoint sets of preferences. In the same spirit, a SWAF can be used to encode those preferences as done in the specific case of a single set of preferences. Finally, preferences in ASO can be also modeled as preferences of the form $C_1 > C_2 > \dots > C_k$ (e.g. $a > b > c$) where C_i s are boolean formulas built using arguments and standard connectives \wedge, \vee , and \neg . However, to simplify the presentation and the translation to WAF, we assumed w.l.o.g. that preferences $C_1 > C_2 > \dots > C_k$ are rewritten as $C_i > C_j$ such that $i < j$ and $i, j \in [1, k]$.

An interesting extension of Dung’s framework with *epistemic constraints* called *Epistemic Argumentation Framework* (EAF) has been recently proposed [89]. An epistemic constraint is a propositional formula over labeled arguments ($\{\text{in}(a), \text{out}(a), \text{undef}(a) \mid a \text{ is an argument}\}$)⁴ extended with the modal operators **K** and **M**. Intuitively, **K** ϕ (resp. **M** ϕ) states that the considered agent believes that ϕ is always (resp. possibly) true. The semantics of an EAF is given by the set of so-called *S-epistemic labelling* sets. Intuitively, an *S-epistemic labelling* set is a collection of *S*-labellings that reflects the belief of an agent. More in detail, every *S-epistemic labelling* set consists of *S*-labellings of the underlying AF and it is a maximal set of *S*-labellings that satisfy the epistemic constraint. Epistemic constraints without modal operators play the same role of strong constraints, that is, they play the same role of strong constraints in CAF. In Appendix D, we formally show that any EAF without modal operators (or with restricted modal operators), that is an AF with constraints defined over the alphabet of arguments’ labels [21], can be rewritten into an equivalent CAF with the semantics defined in this paper.

Constraints having the form of explicit acceptance conditions over arguments have been firstly explored in the *Abstract Dialectical Framework* (ADF) [36], whose semantics can be captured by the (monotonic three-valued) possibilistic logic [67]. In particular, the semantics of an ADF D relies on a characteristic operator, namely Γ_D , which takes as an input a three-valued interpretation ν and returns an interpretation by considering all possible two-valued completions of ν .⁵ To explain the connection between ADF and CAF, let us illustrate how the CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ introduced in Section 1, where $\langle \mathcal{A}, \mathcal{R} \rangle$ is the AF shown in Fig. 1(a) and $C = \{\kappa = a \wedge b \wedge c \Rightarrow \text{f}\}$ is the considered strong constraint, can be modeled through an ADF of the form $D = \langle S = \{a, b, c\}, L = \{(x, y) \mid x, y \in S \wedge x \neq y\}, C = \{C_a = a \wedge \neg(b \wedge c), C_b = b \wedge \neg(a \wedge c), C_c = c \wedge \neg(a \wedge b)\}$ whose set of complete interpretations coincides with the set of complete extensions of Ω . The fact that CAF can be modeled by ADF, and in particular that credulous and skeptical reasoning in CAF can be reduced to ADF is backed by the computational complexity of the two frameworks [92]. In fact, ADF is at least as expressive as CAF under complete, preferred and stable semantics, and strictly more expressive (one level higher in the polynomial hierarchy) for credulous acceptance under complete semantics (which is Σ_2^p -complete for ADF) and skeptical acceptance under preferred semantics (Π_3^p -complete for ADF). As for the semi-stable semantics, to the best of our knowledge, the complexity of semi-stable semantics for ADF has not been studied yet. Although a reduction from CAF to ADF is not ruled out by our complexity analysis, providing such a mapping means translating CAF constraints which work at the global level of a set of extensions to (local) acceptance conditions that are specifically defined for arguments—this deserves more investigation, and is left for future work.

⁴ **in**, **out**, and **undef** are synonyms of **true**, **false** and **undef**, respectively.

⁵ The interested reader can find an overview of ADF’s syntax and semantics in Appendix E.

However, it is worth mentioning that under both complete and stable semantics, ADF is not as expressive as WAF from a complexity standpoint. In fact, skeptical acceptance under complete semantics is coNP-complete in ADF [33] while it is Π_2^P -complete for WAF (cf. Theorem 2). Moreover, under stable semantics, both credulous and skeptical acceptance problems in ADF are one level lower in the polynomial hierarchy w.r.t. WAF (the same happens for CAF). This suggests that WAF cannot be encoded through ADF under complete and stable semantics. Also, a reduction from ADF to WAF is not ruled out by our complexity analysis, though in this case an approach to translate (local) acceptance condition over arguments to global constraints needs to be devised. Finally, analogously to what is done in this paper concerning the exploration of less expressive subclasses within WAF and CAF (e.g. denial constraints), similar analyses have been conducted within ADF. For instance, the subclass called bipolar ADFs (BADFs) has been shown to exhibit complexity comparable to that of classical AFs, as is it possible to avoid considering all the possible two-valued completions through the application of Kleene logic [24]. Exploring the connection between subclasses of WAF/CAF and ADF is another possible direction for future work.

As mentioned earlier, with the aim of allowing for a more straightforward and compact encoding of knowledge w.r.t. AF, several frameworks extending AF have been proposed, such as the *argumentation framework with collective attacks* (SETAF) [80,62,54,55]. SETAF generalizes AF by allowing for collective attacks, i.e., attacks from non-empty sets of arguments to a single argument. Intuitively, a collective attack from set $a = \{a, b\}$ to argument c means that neither a nor b is strong enough to defeat c by themselves. To illustrate the semantics of SETAF, let us consider again the situation of Example 1, which can be modeled by the CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$, where $\langle \mathcal{A}, \mathcal{R} \rangle$ is the AF shown in Fig. 1(a) and $C = \{\kappa = a \wedge b \wedge c \Rightarrow \mathfrak{f}\}$ is a strong constraint. This situation can be also modeled by the SETAF $\langle \{a, b, c\}, \{(\{a, b\}, c), (\{a, c\}, b), (\{b, c\}, a)\} \rangle$, whose set of preferred extensions is $\{\{a, b\}, \{a, c\}, \{b, c\}\}$. It is worth noting that, the notion of collective attacks can also be represented in ADF by using proper acceptance conditions [75,85,54]. As for ADF, SETAF is not as expressive as WAF from the complexity viewpoint, suggesting that WAF cannot be encoded through SETAF.

Finally, concerning the hierarchy of constraints in SWAF, we point out that there is a connection with the preferred subtheory approach introduced in [32], where a default theory is a tuple (T_1, \dots, T_n) such that each T_i is a set of classical first order formulas with the property that information in T_i is more reliable than that in T_j if $i < j$. This is analogous to the way strata in SWAFs are defined, as weak constraints in a stratum i have higher priority than those in a stratum $j > i$.

We believe that the findings of this study may also apply to other well-known related AI fields, such as logic programming with 3-valued semantics [87,40,4], computing repairs and consistent answers over inconsistent data [15,61,76,77] (see e.g. [31,57,65] for the relationship between logic programming and consistent query answering), and integration of different AI formalisms with (strong and weak) constraints and preferences [35,38,83].

10. Conclusions and future work

We have introduced a general argumentation framework where both strong and weak constraints can be easily expressed. Our complexity analysis shows how the several forms of constrains (including restricted forms, e.g., denials) impact on the complexity of credulous and skeptical reasoning. It turns out that constraints, especially weak ones, generally increase the expressivity of AFs. In fact, WAFs allow us to model optimization problems such as, for instance, Min Coloring and Maximum Satisfiability, where some kind of preferences (e.g., use the minimum number of colors) are expressed on solutions. This is not possible for AFs/CAFs whose expressivity is lower than that of WAFs (AFs/CAFs can capture simpler problems such as k -coloring and SAT).

We envisage implementations of the proposed WAF semantics by exploiting ASP-based systems and analogies with logic programs with weak constraints [37,66] (the relationship between the semantics of some frameworks extending AF and that of logic programs has been investigated in [4]). For WAFs, DLV system [9] could be used for computing maximum-cardinality stable semantics.

Future work will be also devoted to considering more general forms of constraints, not only using variables ranging on the sets of arguments, but also constraints allowing to express conditions on aggregates [8] (e.g., at least n arguments from a given set S should be accepted/rejected). We believe that the basic idea of adding weak constraints could be also applicable for structured argumentation formalisms [28,64], which is another direction for future research.

Finally, given the inherent nature of argumentation and the typical high computational complexity of most of the reasoning tasks, there have been several efforts toward the investigation of incremental techniques that use AF solutions (e.g., extensions, skeptical acceptance) at time t to recompute updated solutions at time $t + 1$ after that an update (e.g., adding/ removing an attack) is performed [2,46]. These approaches have been extended to argumentation frameworks more general than AFs [3,1]. Following this line of research, we plan to investigate incremental techniques for recomputing CAF and WAF semantics after performing updates consisting of changes to the AF component or to the sets of strong and weak constraints.

CRediT authorship contribution statement

Gianvincenzo Alfano: Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Sergio Greco:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Domenico Mandaglio:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Francesco Parisi:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization. **Irina Trubitsyna:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

The authors wish to thank the anonymous referees for providing detailed comments and suggestions that helped to substantially improve the paper. Moreover, the authors acknowledge the support from *i*) PNRR project FAIR - Future AI Research (PE00000013), Spoke 9 - Green-aware AI, under the NRRP MUR program funded by the Next Generation EU; *ii*) PNRR project SERICS (PE00000014), under the NRRP MUR program funded by the Next Generation EU; *iii*) PNRR project Tech4You (ECS00000009) - Spoke 6, under the NRRP MUR program funded by the Next Generation EU; and *iv*) MUR project PRIN 2022 EPICA (H53D23003660006), Enhancing Public Interest Communication with Argumentation, funded by the European Union - Next Generation EU.

Appendix A. Proofs

In this appendix we provide the proofs not already given in the core of the paper.

To ease readability, we restate the results and organize them in sections by following the order used in the core of the paper.

Since some proofs use results from disjunctive logic programs (DLPs) and logic programs with weak constraints, in Appendix B we recall DLPs and needed results, whereas in Appendix C we introduce logic programs with weak constraints, and show their relationship with WAFs and DLPs. In Appendix D, we show that any EAF without modal operators can be rewritten into an equivalent CAF with the semantics defined in the paper. Finally, in Appendix E we recall the Abstract Dialectical Framework [36], which is discussed in the related work.

For a better understanding of some concepts used in the appendix, we introduce some lemmas and additional examples.

A.1. WAF with maximal-set semantics

We start by introducing three technical lemmas whose results will be used in the following.

Lemma 5. *Let $Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ be a WAF, and let E be a set of arguments. Deciding whether $E \in \text{ms-co}^*(Y)$ (or $E \in \text{ms-st}^*(Y)$) is in *coNP*.*

Proof. Consider the complementary problem: decide whether E is not a maximal-set complete (resp., stable) extension for Y (under Lukasiewicz or Kleene logic). A guess-and-check strategy to decide this problem is as follows. Guess a set $S \subseteq \mathcal{A}$ and check that (i) S is a complete (resp., stable) extension for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and (ii) the set $\mathcal{W}' = \{w \in \mathcal{W} \mid E \vDash w\}$ is such that $\mathcal{W}' \subset \mathcal{W}'' = \{w \in \mathcal{W} \mid S \vDash w\}$ (that is, \mathcal{W}' is not maximal). The complexity of checking (i) is polynomial for both complete and stable semantics, since both checking whether E is a complete (resp., stable) extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ and checking whether $E \vDash C$ can be accomplished in polynomial time. Checking (ii) is in *PTIME* too. Therefore the complementary problem is in *NP*, from which the statement follows. \square

The following lemma states a result analogous to that of Lemma 2 but for the case of WAF (instead of CAF).

Lemma 6. *Let $Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ be a WAF and $E_1, E_2 \in \text{co}^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$ with $E_1 \subseteq E_2$. Then, for any $\omega \in \mathcal{W}$ under Kleene logic, $E_1 \vDash \omega$ implies $E_2 \vDash \omega$.*

Proof. Firstly recall that, given two complete extensions E_1 and E_2 for $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$, $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$. Under Kleene logic every weak constraint ω can be rewritten in standard form as a disjunction of conjunction of literals of the form $\omega : \tau \Rightarrow (\ell_1^1 \wedge \dots \wedge \ell_{n_1}^1) \vee \dots \vee (\ell_1^k \wedge \dots \wedge \ell_{n_k}^k)$. If $E_1 \vDash \omega$, it means that there must be a value $i \in [1, k]$ such that $E_1 \vDash (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$. Moreover, as $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$, it holds that $E_2 \vDash (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$ as well. \square

Lemma 7. *Let $Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ be a WAF, and let E be a set of arguments. Deciding whether $E \in \text{ms-pr}^K(Y)$ is in *coNP*.*

Proof. Consider the complementary problem, that is, deciding whether E is not a maximal-set preferred extension for Y under Kleene logic. This problem can be decided as follows. Guess a set $S \subseteq \mathcal{A}$ and check that (i) S is a complete extension for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and (ii) $\mathcal{W}' = \{w \in \mathcal{W} \mid E \vDash w\} \subset \mathcal{W}'' = \{w \in \mathcal{W} \mid S \vDash w\}$. The fact that it suffices to check that S is a complete (rather than preferred) extension for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ comes from the results of Lemma 2 and Lemma 6. Both (i) and (ii) can be decided in *PTIME*, as checking

whether a set of arguments is a complete extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ is in P and this holds even if we additionally check satisfaction of the sets of constraints. Thus, this problem is in NP . Consequently, deciding whether E is a maximal-set preferred extension for Υ under Kleene logic is in $coNP$. \square

Theorem 2. For any WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- CA_{ms-S}^σ is: (i) Σ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma = K$,
(ii) Σ_2^p -complete for $S \in \{\text{co}, \text{st}\}$ and $\sigma = L$,
(iii) Σ_2^p -hard and in Σ_3^p for $S = \text{pr}$ and $\sigma = L$, and
(iv) Σ_2^p -complete for $S = \text{sst}$ and $\sigma = L$.
- SA_{ms-S}^σ is: (i) Π_2^p -complete for $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma = K$, and
(ii) Π_2^p -complete for $S \in \{\text{co}, \text{st}\}$ and $\sigma = L$,
(iii) Π_3^p -complete for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.

Proof. (Membership.) The membership results come from the general case of SWAF (cf. Theorem 4) as any WAF is a SWAF with a single stratum of constraints.

(Hardness.) All the lower bound results, except those for CA_{ms-sst}^L , SA_{ms-pr}^L , and SA_{ms-sst}^L derive from the fact that they hold for WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ where \mathcal{W} and C are sets of denial constraints, that is for NWAFF (cf. Theorem 8).

The hardness results for CA_{ms-sst}^L , SA_{ms-pr}^L , and SA_{ms-sst}^L are obtained by mapping disjunctive datalog programs under partial stable model semantics [59] to WAFs under complete semantics and Lukasiewicz logic.

It is important to recall that for disjunctive datalog programs under partial stable model semantics the complexity of credulous and skeptical acceptance are as follows [59]:

- CA_S is: (i) Σ_2^p -complete for any semantics $S \in \{\text{ps}, \text{st}, \text{ms}\}$, and (iii) Σ_3^p -complete for $S = \text{ls}$.
- SA_S is: (i) Π_2^p -complete for $S \in \{\text{ps}, \text{st}\}$, and (ii) Π_3^p -complete for $S \in \{\text{ms}, \text{ls}\}$.

where $\text{ps}, \text{st}, \text{ms}$ and ls denote the semantics partial stable, (total) stable, maximal stable, and least-undefined stable, respectively;

Since every disjunctive datalog programs under partial (resp., total, maximal, least-undefined) stable model semantics can be mapped into an equivalent WAF under complete (resp. stable, preferred, semi-stable) semantics and Lukasiewicz logic, the hardness results follow. We show how disjunctive logic programs can be mapped to normal logic programs with constraints and then to WAFs in Appendix C. The partial stable semantics for normal and disjunctive logic programs are recalled in Appendix B. \square

A.2. WAF with maximum-cardinality semantic

We first introduce a technical lemma which will be used in the proof of Theorem 3.

Lemma 8. Given a WAF $\Upsilon = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, and a natural number $k \leq |\mathcal{W}|$, deciding whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k constraints of \mathcal{W} is in

- NP (resp., NP , Σ_2^p , Σ_2^p) under Lukasiewicz logic.
- NP (resp., NP , NP , Σ_2^p) under Kleene logic.

The result still holds if it is required that the extension E contains a given argument $a \in \mathcal{A}$, that is for the problem of deciding whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k constraints of \mathcal{W} .

Proof. First, consider Lukasiewicz logic. Guess a set $E \subseteq \mathcal{A}$ and check that (i) E is a complete (resp., stable, preferred, semi-stable) extension in PTIME (resp., PTIME, NP , NP) [45,49,51], and (ii) E satisfies at least k constraints of \mathcal{W} in PTIME. Thus, the considered problem is in NP (resp., NP , Σ_2^p , Σ_2^p) under complete (resp., stable, preferred, semi-stable) semantics.

As for the Kleene logic, the difference with respect to the above-described procedure is that in the case of the preferred semantics the result of Lemma 6 can be used. Thus, at step i), it suffices to check that E is a complete extension rather than preferred, since if there is a complete extension satisfying at least k constraints, then there is a preferred extension satisfying at least k constraints.

Finally, consider the case where we additionally require that the extension E contains an argument $a \in \mathcal{A}$. The results continue to hold in such a case, since what is said earlier continues to hold if we start by guessing a set $E \subseteq \mathcal{A}$ containing a . \square

Theorem 3. For any WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- CA_{mcS}^σ is: (i) Θ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma = K$,
(ii) Σ_2^p -hard and in Θ_3^p for semantics $S = \text{sst}$ and $\sigma = K$,
(iii) Θ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma = L$, and
(iv) Σ_2^p -hard and in Θ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.

- SA_{mcS}^σ is: (i) Θ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -hard and in Θ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Proof. (Membership.) In the following, we use n to denote the number of constraints, i.e. $n = |\mathcal{W}|$.

- $(CA_{mc-co}^*, CA_{mc-st}^*, SA_{mc-co}^*$ and $SA_{mc-st}^*)$.
We prove the results only for CA_{mcS}^* , as for each semantics $S \in \{\text{co}, \text{st}\}$ the complexity of the complementary problem of SA_{mcS}^* , that is checking whether there exists a maximum-cardinality S -extension not containing a , can be shown by reasoning analogously to the case of CA_{mcS} .
We first call an NP oracle to check that $\langle \mathcal{A}, \mathcal{R}, C \rangle$ admits a complete (resp., stable) extension containing a (this corresponds to checking credulous acceptance for a CAF, which is in NP for complete and stable semantics). Then, we perform a binary search in $[0, n]$ to find the maximum number k_{max} of constraints that are satisfied by the extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$. In particular, in the binary search we use an NP oracle to decide whether there exist a complete (resp., stable) extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k constraints (Lemma 8), where k is the middle value in the search interval. The number of calls to the oracle is bounded by $O(\log n)$, as at the first step the search space is $[0, n]$ and we call the oracle with $k_1 = n/2$, at the second step the search space is one half of the previous step (either $[0, k_1 - 1]$ or $[k_1, n]$, with $k_2 = (k_1 - 1)/2$ or $k_2 = \lfloor (n - k_1)/2 \rfloor$), and so on. Finally, given the maximum number k_{max} of constraints that are satisfied by the extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$, we use another call to an NP oracle to decide whether there exists a complete (resp., stable) extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k_{max} constraints (as stated in Lemma 8, checking this is still in NP). Therefore, deciding whether there is a maximum-cardinality complete (resp. stable) extension containing a is in Θ_2^p .
- (CA_{mc-pr}^K) . Using the result of Lemma 8, the Θ_2^p algorithm for CA_{mc-co}^K also applies to the case of preferred semantics under Kleene logic.
- (SA_{mc-pr}^K) .
Consider the complementary problem of SA_{mc-pr}^K , that is checking whether there exists a maximum-cardinality pr-extension not containing a . We first show that this problem is in Θ_3^p .
First, call an NP oracle to check that $\langle \mathcal{A}, \mathcal{R}, C \rangle$ admits a preferred extension (this has the same complexity of checking credulous acceptance for a CAF, which is in NP for preferred semantics under Kleene logic). Then, we perform a binary search in $[0, n]$ to find the maximum number k_{max} of constraints that are satisfied by the preferred extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$. In particular, in the binary search we use an NP oracle to decide whether there exist a preferred extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k constraints under Kleene logic (Lemma 8), where k is the middle value in the search interval. Finally, given the maximum number k_{max} of constraints that are satisfied by the extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$, we use a call to a Σ_2^p oracle to decide whether there exist a preferred extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ not containing a and satisfying at least k_{max} constraints. Observe that deciding whether there exist a preferred extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ not containing a and satisfying at least k constraints is in Σ_2^p . In fact, differently from the case of Lemma 8, where we require that argument a is contained in the extension, here the result of Lemma 6 cannot be used. To show the membership is in Σ_2^p , it suffices to consider the following procedure: guess a set $E \subseteq \mathcal{A}$ not containing a and check that (i) E is a preferred extension (in NP), and (ii) E satisfies at least k constraints of \mathcal{W} (in PTIME). Therefore, the considered problem is Σ_2^p under preferred semantics and, thus, deciding whether there is a maximum-cardinality preferred extension not containing a is in Θ_3^p . Since Θ_3^p is closed under complement, the result follows.
- $(CA_{mc-pr}^L$ and $SA_{mc-pr}^L)$.
We prove the results only for CA_{mc-pr}^L , as the complexity of the complementary problem of SA_{mc-pr}^L , that is checking whether there exists a maximum-cardinality pr-extension not containing a , can be shown reasoning analogously to the case of CA_{mc-pr} .
We first call a Σ_2^p oracle to check that $\langle \mathcal{A}, \mathcal{R}, C \rangle$ admits a preferred extension containing a (this corresponds to checking credulous acceptance for a CAF, which is in Σ_2^p for preferred semantics under Lukasiewicz logic). Then, we perform a binary search in $[0, n]$ to find the maximum number k_{max} of constraints that are satisfied by the preferred extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$. In the binary search, we can use a Σ_2^p oracle to decide whether there exist a preferred extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k constraints (Lemma 8). The number of calls to the oracle is bounded by $O(\log n)$. Finally, we use another call to Σ_2^p oracle to decide whether there exist a preferred extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k_{max} constraints (as stated in Lemma 8, checking this is still in Σ_2^p). Therefore, deciding whether there is a maximum-cardinality preferred extension containing a is in Θ_3^p .
- $(CA_{mc-sst}^*$ and $SA_{mc-sst}^*)$. We prove the results only for CA_{mc-sst}^* , as the complexity of the complementary problem of SA_{mc-sst}^* , that is checking whether there exists a maximum-cardinality S -extension not containing a , can be shown reasoning analogously to CA_{mc-sst}^* .
The proof for CA_{mc-sst}^* is analogous to that for CA_{mc-pr}^L . We first call a Σ_2^p oracle to check that $\langle \mathcal{A}, \mathcal{R}, C \rangle$ admits a semi-stable extension containing a (this corresponds to checking credulous acceptance for a CAF, which is in Σ_2^p for semi-stable semantics under Lukasiewicz logic). Then, we perform a binary search in $[0, n]$ to find the maximum number k_{max} of constraints that are satisfied by the semi-stable extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$. In the binary search, we can use a Σ_2^p oracle to decide whether there exist a semi-stable extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k constraints (Lemma 8). The number of calls to the oracle is bounded by $O(\log n)$. Finally, we use another call to Σ_2^p oracle to decide whether there exist a semi-stable extension of $\langle \mathcal{A}, \mathcal{R}, C \rangle$

containing a and satisfying at least k_{max} constraints (Lemma 8). Therefore, deciding whether there is a maximum-cardinality semi-stable extension containing a is in Θ_3^p .

(Hardness.) We now prove the hardness results.

- $(CA_{mc-st}^*$ and SA_{mc-st}^*). The hardness results derive from analogous results for logic programs with weighted weak constraints, stating that the complexity credulous and skeptical reasoning in logic programs with weighted weak constraints is Θ_2^p -complete [37] (Theorem 21). In fact, since every logic program LP under total stable model semantics can be translated into an AF Λ , where the set of stable models of LP (restricted to positive literals) coincide with the set of stable extensions of Λ [40], a logic program with strong and weak constraints (LP, C, \mathcal{W}) can be translated into an equivalent WAF $\langle \mathcal{A}, \mathcal{R}, C', \mathcal{W}' \rangle$, from which the result follows.
- $(CA_{mc-co}^*$ and SA_{mc-co}^*). The hardness result follows from the Θ_2^p -complete problem CA_{mc-st}^* for WAF. Given a WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, $E \in mc-st^*(Y)$ iff $E \in mc-co^*(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C' = C \cup C'', \mathcal{W} \rangle$ and $C'' = \{t \Rightarrow a \vee \neg a \mid a \in \mathcal{A}\}$. Thus, $CA_{mc-st}^*(Y, g)$ is true iff $CA_{mc-co}^*(Y', g)$ is true and $SA_{st}^*(Y, g)$ is true iff $SA_{co}^*(Y', g)$ is true.
- (CA_{mc-pr}^K) . The hardness result follows from the Θ_2^p -complete problem CA_{mc-st}^K for WAF. Given a WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, $E \in mc-st^K(Y)$ iff $E \in mc-pr^K(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C' = C \cup C'', \mathcal{W} \rangle$ and $C'' = \{t \Rightarrow a \vee \neg a \mid a \in \mathcal{A}\}$. Thus, $CA_{mc-st}^K(Y, g)$ is true iff $CA_{mc-pr}^K(Y', g)$ is true.
- $(CA_{mc-pr}^L, CA_{mc-sst}^*, SA_{mc-pr}^*$ and $SA_{mc-sst}^*)$. The hardness result derives from the fact that they hold for any WAF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ where $\mathcal{W} = \emptyset$, that is for CAF (cf. Theorem 1). \square

A.3. WAF with stratified weak constraints

In this section, we provide the proofs of Theorem 4 and Theorem 5.

Theorem 4. For any SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem

- CA_{ms-S}^σ is: (i) Σ_2^p -complete for any semantics $S \in \{co, st\}$ and $\sigma \in \{K, L\}$,
(ii) Σ_2^p -hard and in Σ_3^p for $S = pr$ and $\sigma = L$,
(iii) Σ_2^p -complete for $S \in \{pr, sst\}$ and $\sigma = K$, and
(iv) Σ_3^p -complete for $S = sst$ and $\sigma = L$.
- SA_{ms-S}^σ is: (i) Π_2^p -complete for $S \in \{co, st, pr, sst\}$ and $\sigma = K$,
(ii) Π_2^p -complete for $S \in \{co, st\}$ and $\sigma = L$, and
(iii) Π_3^p -complete for $S \in \{pr, sst\}$ and $\sigma = L$.

Proof. (Hardness.) The lower bound results derive from the fact that they hold for any SWAFs $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ where $n = 1$, that is, for WAFs (cf Theorem 2).

(Membership.) We now provide the membership results for each considered semantics and problem. Let $Y = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ be a SWAF and $a \in \mathcal{A}$ the argument for which we want to decide either credulous or skeptical acceptance w.r.t. Y .

- $(CA_{ms-co}^*$ and $CA_{ms-st}^*)$. We first prove that deciding whether E is a maximal-set complete (resp., stable) extension for Y (under Lukasiewicz or Kleene logic) is in coNP. A guess-and-check strategy to decide the complementary problem is as follows. Guess a set $S \subseteq \mathcal{A}$ and an index $j \in [1, n]$ and check that (i) S is a complete (resp., stable) extension for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and (ii) for each $i \in [1, j]$ the sets $\mathcal{W}'_i = \{w \in \mathcal{W}_i \mid E \vDash w\}$ and $\mathcal{W}''_i = \{w \in \mathcal{W}_i \mid S \vDash w\}$ are such that $\mathcal{W}'_k = \mathcal{W}''_k$ with $k \in [1, j-1]$ and $\mathcal{W}'_j \subset \mathcal{W}''_j$ (that is, \mathcal{W}'_j is not maximal w.r.t the j -th stratum). The complexity of checking (i) is polynomial for both complete and stable semantics, since checking whether E is a complete (resp., stable) extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ and checking whether $E \vDash C$ can be accomplished in polynomial time. Checking (ii) is in PTIME too. Therefore, the problem of deciding whether E is not a maximal-set complete (resp., stable) extension is in NP, and thus the complementary problem is in coNP. Given this, to prove that CA_{ms-co}^* (resp., CA_{ms-st}^*) is in Σ_2^p , it suffices to consider the following guess-and-check strategy: guess a set $E \subseteq \mathcal{A}$ of arguments containing a and check that E is a maximal-set complete (resp. stable) extension for Y by using the above-shown coNP oracle. Thus CA_{ms-co}^* (resp. CA_{ms-st}^*) is in Σ_2^p .
- $(CA_{ms-pr}^K$ and $SA_{ms-pr}^K)$. We show that, given a SWAF $Y = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, it is the case that $E \in ms-pr^K(Y)$ iff $E \in ms-co^K(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}'_1, \mathcal{W}_2, \dots, \mathcal{W}_n) \rangle$ with $\mathcal{W}'_1 = \{t \Rightarrow x \mid x \in \mathcal{A}\}$. First recall that, by Lemma 2 we have that $pr^K(\langle \mathcal{A}, \mathcal{R}, C \rangle) \subseteq co^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Then, since the constraints in \mathcal{W}'_1 select from the extensions in $co^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$ those that are maximal w.r.t. \subseteq (i.e., \mathcal{W}'_1 has the effect filtering out the preferred extensions), we have that $ms-co^K(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}'_1 \rangle) = pr^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Then the result of applying $(\mathcal{W}_1, \dots, \mathcal{W}_n)$ over $ms-co^K(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}'_1 \rangle)$ is the same as that of applying $(\mathcal{W}_1, \dots, \mathcal{W}_n)$ to $pr^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$, from which the result follows.
- $(CA_{ms-sst}^K$ and $SA_{ms-sst}^K)$. The strategy of the proof is similar to that of the previous item, except that a different set \mathcal{W}'_1 of weak constraints is used to simulate the semi-stable semantics. In particular, we show that, given a SWAF $Y = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, it holds that $E \in ms-sst^K(Y)$ iff $E \in ms-co^K(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}'_1, \mathcal{W}_2, \dots, \mathcal{W}_n) \rangle$ with $\mathcal{W}'_1 = \{t \Rightarrow x \vee \neg x \mid x \in \mathcal{A}\}$. Recall that, by Lemma 2, we have that $sst^K(\langle \mathcal{A}, \mathcal{R}, C \rangle) \subseteq co^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Then, since \mathcal{W}'_1 selects from $co^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$

the extensions that are maximal w.r.t. the presence of arguments whose status is either true or false (or equivalently that are minimal w.r.t. undecided arguments) we have that $\text{ms-co}^K(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}'_1 \rangle) = \text{sst}^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Therefore, applying $(\mathcal{W}_1, \dots, \mathcal{W}_n)$ over $\text{ms-co}^K(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W}'_1 \rangle)$ is the same as applying $(\mathcal{W}_1, \dots, \mathcal{W}_n)$ to $\text{sst}^K(\langle \mathcal{A}, \mathcal{R}, C \rangle)$, from which the result follows.

- $(CA_{\text{ms-pr}}^L$ and $CA_{\text{ms-sst}}^L$). We first prove that deciding whether E is a maximal-set preferred (resp., semi-stable) extension for Y under Lukasiewicz logic is in Π_2^p . A guess-and-check strategy to decide the complementary problem is as follows. Guess a set $S \subseteq \mathcal{A}$ and an index $j \in [1, n]$ and check that (i) S is a preferred (resp., semi-stable) extension for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ and (ii) for each $i \in [1, j]$ the sets $\mathcal{W}'_i = \{w \in \mathcal{W}_i \mid E \vDash w\}$ and $\mathcal{W}''_i = \{w \in \mathcal{W}_i \mid S \vDash w\}$ are such that $\mathcal{W}'_k = \mathcal{W}''_k$ with $k \in [1, j-1]$ and $\mathcal{W}'_j \subset \mathcal{W}''_j$ (that is, \mathcal{W}'_j is not maximal w.r.t the j -th stratum). The complexity of checking (i) is coNP for both preferred and semi-stable semantics, since checking whether E is a preferred (resp., semi-stable) extension for $\langle \mathcal{A}, \mathcal{R} \rangle$ is coNP and checking whether $E \vDash C$ can be accomplished in polynomial time. Checking (ii) is in P TIME too. Therefore, the complement of the above-stated problem is in Σ_2^p . Thus, $CA_{\text{ms-pr}}^L$ (resp., $CA_{\text{ms-sst}}^L$) is in Σ_3^p since it suffices to guess a set $E \subseteq \mathcal{A}$ of arguments containing a and check that E is a maximal-set preferred (resp. semi-stable) extension for Y by using the above-described Π_2^p oracle. Hence, it follows that $CA_{\text{ms-pr}}^L$ (resp. $CA_{\text{ms-sst}}^L$) is in Σ_3^p .
- Skeptical acceptance $(SA_{\text{ms-co}}^\sigma, SA_{\text{ms-st}}^\sigma, SA_{\text{ms-pr}}^\sigma, SA_{\text{ms-sst}}^\sigma)$. For each semantics $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$ consider the complementary problem of checking whether there exists a maximal-set S^σ -extension for Y not containing a . Reasoning as in the cases of the credulous acceptance considered earlier, it can be shown that this problem is in Σ_2^p for (i) $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$ and (ii) $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = K$, and in Σ_3^p for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$. Therefore, $SA_{\text{ms-S}}^L$ is in Π_2^p for $S \in \{\text{co}, \text{st}\}$ and in Π_3^p for $S \in \{\text{pr}, \text{sst}\}$ while $SA_{\text{ms-S}}^K$ is in Π_2^p for $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$. \square

The following lemma generalizes the result of Lemma 8 to the case of SWAF.

Lemma 9. Given a SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, and n natural numbers $k_1 \leq |\mathcal{W}_1|, \dots, k_n \leq |\mathcal{W}_n|$, deciding whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k_i constraints of \mathcal{W}_i , for each $i \in [1..n]$ is in:

- NP (resp., NP, Σ_2^p , Σ_2^p) under Lukasiewicz logic, and
- NP (resp., NP, NP, Σ_2^p) under Kleene logic.

The result still holds if it is required that the extension E contains a given argument $a \in \mathcal{A}$, that is for the problem of deciding whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k_i constraints of \mathcal{W}_i , for each $i \in [1..n]$.

Proof. The result can be proved by reasoning analogously to the proof of Lemma 8. \square

Theorem 5. For any SWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem:

- $CA_{\text{mc-S}}^\sigma$ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma = K$,
(ii) Σ_2^p -hard and in Δ_3^p for semantics $S = \text{sst}$ and $\sigma = K$,
(iii) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{pr}\}$ and $\sigma = L$, and
(iv) Σ_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma = L$.
- $SA_{\text{mc-S}}^\sigma$ is: (i) Δ_2^p -complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -hard and in Δ_3^p for any semantics $S \in \{\text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Proof. (Hardness.) The lower bound results derive from the fact that they hold for any SWAFs $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ where $|\mathcal{W}_i| = 1$ for each $i \in [1, n]$, that is, for LWAFs.

(Membership.) We now provide the membership results for each considered semantics and problem. Let $Y = \langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ be a SWAF and $a \in \mathcal{A}$ the argument for which we want to decide either credulous or skeptical acceptance w.r.t. Y .

- $(CA_{\text{mc-S}}^*$ with $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$). We consider the complete (resp., stable, preferred, semi-stable) semantics. We first call an NP (resp., NP, Σ_2^p , Σ_2^p) oracle to check that $\langle \mathcal{A}, \mathcal{R}, C \rangle$ admits a complete (resp., stable, preferred, semi-stable) extension containing a . This corresponds to checking credulous acceptance for a CAF; note that this problem is in Σ_2^p for preferred semantics under Lukasiewicz logic, but it is in NP under Kleene logic— we will reconsider this at the end of this proof to provide a better upper bound for $CA_{\text{mc-pr}}^K$. Let $m_i = |\mathcal{W}_i|$ be the number of constraints in the i -th stratum, with $i \in [1..n]$. We perform n consecutive binary search in the intervals $[0, m_i]$ to find the maximum number of constraints that are satisfied by complete (resp., stable, preferred, semi-stable) extensions at each stratum, given the maximum number of constraints that are satisfied at previous strata, as follows. In the first execution of the binary search in $[0, m_1]$, we use an NP (resp., NP, Σ_2^p , Σ_2^p) oracle to decide whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k constraints of \mathcal{W}_1 (cf. Lemma 9). Let k_{max}^1 be the maximum number of constraints in \mathcal{W}_1 that are satisfied by the extensions of $\langle \mathcal{A}, \mathcal{R}, C \rangle$. During the

i -th execution of the binary search in $[0, m_i]$, with $i \in [2..n]$, we use an NP (resp., NP, Σ_2^p , Σ_2^p) oracle to decide whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ satisfying at least k_{max}^j constraints of \mathcal{W}_j , for each $j \in [1..i-1]$, and at least k constraints of \mathcal{W}_i (cf. Lemma 9). Hence, at the end of the i -th binary search in $[0, m_i]$, we find the maximum number k_{max}^i of constraints in \mathcal{W}_i , that are satisfied by any extension E , given that k_{max}^j is the maximum number of constraints in \mathcal{W}_j that are satisfied by E , with $j \in [1..i-1]$. Finally, given the numbers $k_{max}^1, \dots, k_{max}^n$, we use another call to an NP (resp., NP, Σ_2^p , Σ_2^p) oracle to decide whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ containing a and satisfying at least k_{max}^i constraints of \mathcal{W}_i , for each $i \in [1..n]$; as stated in Lemma 9, checking this is still in NP (resp., NP, Σ_2^p , Σ_2^p).

We now discuss the complexity of the above-described procedure. For the i -th execution of the binary search in $[0, m_i]$, the number of calls to the oracle is bounded by $O(\log m_i)$, with $i \in [1..n]$. By observing that the overall number of calls to the oracle is bounded by $O(n \log m)$, where $m = \max\{m_i \mid i \in [1..n]\}$, we obtain that the complexity of CA_{mc-co}^* (resp. CA_{mc-st}^* , CA_{mc-pr}^* , CA_{mc-sst}^*), is in the class Δ_2^p (resp., Δ_2^p , Δ_3^p , Δ_3^p). Furthermore, since for CA_{mc-pr}^K it is sufficient to use NP oracles only (cf. Lemma 9), a better upper bound can be found, that is, CA_{mc-pr}^K can be decided in Δ_2^p .

- (SA_{mc-S}^* with $S \in \{co, st, pr, sst\}$). We consider the complete (resp. stable, preferred, semi-stable) semantics. Consider the complementary problem of SA_{mc-co}^* (resp. SA_{mc-st}^* , SA_{mc-pr}^* , SA_{mc-sst}^*), that is checking whether there exists a maximum-cardinality complete (resp. stable, preferred, semi-stable) extension of $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ not containing a . We can show that this problem is in Δ_2^p (resp., Δ_2^p , Δ_3^p , Δ_3^p) by reasoning as in the proof of CA_{mc-S}^* , with $S \in \{co, st, pr, sst\}$, in the previous item. The only difference is that in the last call to the oracle we have to decide whether there exists a complete (resp., stable, preferred, semi-stable) extension E for $\langle \mathcal{A}, \mathcal{R}, C \rangle$ not containing a and satisfying at least k_{max}^i constraints of \mathcal{W}_i , for each $i \in [1..n]$, where each k_{max}^i is determined by a binary search as described above. Since the complexity of this problem is in NP (resp., NP, Σ_2^p , Σ_2^p), we obtain that complexity of the complementary problem of SA_{mc-co}^* (resp. SA_{mc-st}^* , SA_{mc-pr}^* , SA_{mc-sst}^*), is in the class Δ_2^p (resp., Δ_2^p , Δ_3^p , Δ_3^p). Finally, since Δ_i^p is closed under complement, the statement follows. \square

A.4. WAF with linearly ordered weak constraints

We now provide the proof of Theorem 6, whose statement is recalled below.

Theorem 6. For any LWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, the problem:

- CA_S^σ is: (i) Δ_2^p -complete for any semantics $S \in \{co, st, pr\}$ and $\sigma = K$,
(ii) Σ_2^p -complete for semantics $S = sst$ and $\sigma = K$,
(iii) Δ_2^p -complete for any semantics $S \in \{co, pr\}$ and $\sigma = L$, and
(iv) Σ_2^p -hard and in Δ_3^p for any semantics $S \in \{pr, sst\}$ and $\sigma = L$.
- SA_S^σ is: (i) Δ_2^p -complete for any semantics $S \in \{co, st\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^p -complete for any semantics $S \in \{pr, sst\}$ and $\sigma = K$,
(iii) Π_2^p -hard and in Δ_3^p for any semantics $S \in \{pr, sst\}$ and $\sigma = L$.

Proof. (Membership.) Recall that for LWAF the maximal-set and maximum-cardinality semantics coincide, that is, $CA_{msS}^\sigma = CA_{mcS}^\sigma$ and $SA_{msS}^\sigma = SA_{mcS}^\sigma$. Therefore, the membership results for CA_{sst}^K , SA_{sst}^K , and SA_{pr}^K follows from the fact that a LWAF is SWAF under maximal-set semantics, and thus these results follow from Theorem 4. Moreover, all the other results follow from Theorem 5, since a LWAF is SWAF under maximum-cardinality semantics.

(Hardness.) We now prove the hardness results.

- (CA_{st}^* and SA_{st}^*). The hardness result derives from analogous results for logic programs with weighted weak constraints with priorities, for which credulous and skeptical reasoning is Δ_2^p -complete [37]. Thus, since every logic program LP under total stable model semantics can be translated into an AF Λ , where the set of stable models of LP (restricted to positive literals) coincide with the set of stable extensions of Λ [40], a logic program with weighted weak constraints with priorities $(LP, \{w_1\}, \dots, \{w_k\})$, can be translated into a LWAF Ω where weak constraints are linearly ordered. The result presented in [37] holds even if the weight of every weak constraint w_i is 2^{i-1} , with $i \in [1, k]$, meaning that a linear order is imposed.
- (CA_{co}^* and SA_{co}^*). The hardness result follows from the Δ_2^p -complete problem CA_{st}^* for LWAF. Given a LWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, $E \in st^*(Y)$ iff $E \in co^*(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C' = C \cup C'', (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ and $C'' = \{t \Rightarrow a \vee \neg a \mid a \in \mathcal{A}\}$. Thus, $CA_{st}^*(Y, g)$ is true iff $CA_{co}^*(Y', g)$ is true and $SA_{st}^*(Y, g)$ is true iff $SA_{co}^*(Y', g)$ is true.
- (CA_{pr}^*). The hardness result follows from the Δ_2^p -complete problem CA_{st}^* for LWAF. Given a LWAF $\langle \mathcal{A}, \mathcal{R}, C, (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$, $E \in st^K(Y)$ iff $E \in pr^K(Y')$ where $Y' = \langle \mathcal{A}, \mathcal{R}, C' = C \cup C'', (\mathcal{W}_1, \dots, \mathcal{W}_n) \rangle$ and $C'' = \{t \Rightarrow a \vee \neg a \mid a \in \mathcal{A}\}$. Thus, $CA_{st}^*(Y, g)$ is true iff $CA_{pr}^*(Y', g)$ is true and $SA_{st}^*(Y, g)$ is true iff $SA_{pr}^*(Y', g)$ is true.
- (CA_{sst}^* , CA_{pr}^L , SA_{pr}^L , and SA_{sst}^*). The hardness results follow from Theorem 1 since any CAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$ is an LWAF $\langle \mathcal{A}, \mathcal{R}, C, \emptyset \rangle$. \square

A.5. WAF with denial constraints

We start providing the proofs of Lemma 3 and Lemma 4 whose statements are recalled below.

Lemma 3. For any NWF Υ and semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}, \text{gr}\}$, it holds that $S^K(\Upsilon) = S^L(\Upsilon)$.

Proof. Let $Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, and assume that $S \in \text{co}(\langle \mathcal{A}, \mathcal{R} \rangle)$ is a complete extension of $\langle \mathcal{A}, \mathcal{R} \rangle$. Let $\kappa : \ell_1 \wedge \dots \wedge \ell_n \Rightarrow f \in C \cup \mathcal{W}$ be a denial constraint. Then, $S \models \kappa$ under Lukasiewicz (and Kleene) logic iff there exists at least one positive (resp. negative) argument ℓ_i s.t. $\ell_i \in \text{Def}(S)$ (resp. $\ell_i \in S$). Thus $T \models \kappa$ under Lukasiewicz logic iff $T \models \kappa$ under Kleene logic. \square

Lemma 4. Let $Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ be a NWF, $E_1, E_2 \in \text{co}(\langle \mathcal{A}, \mathcal{R} \rangle)$ with $E_1 \subseteq E_2$, and $\mathcal{W}' \subseteq \mathcal{W}$. Then, under both Kleene and Lukasiewicz logic, $E_1 \models C \cup \mathcal{W}'$ implies $E_2 \models C \cup \mathcal{W}'$.

Proof. First, recall that $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$ and that $E_1 = \text{Acc}(E_1) \subseteq \text{Acc}(E_2) = E_2$. Considering Kleene's logic, every constraint $\kappa \in C \cup \mathcal{W}$ can be rewritten in standard form as a disjunction of conjunction of literals, that is, in the form $\kappa : \tau \Rightarrow (\ell_1^1 \wedge \dots \wedge \ell_{n_1}^1) \vee \dots \vee (\ell_1^k \wedge \dots \wedge \ell_{n_k}^k)$. If $E_1 \models \kappa$, it means that there must be $i \in [1, k]$ such that $E_1 \models (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$. Moreover, as $E_1 \subseteq E_2$ implies that $\text{Def}(E_1) \subseteq \text{Def}(E_2)$ and $\text{Acc}(E_1) \subseteq \text{Acc}(E_2)$, it holds that $E_2 \models (\ell_1^i \wedge \dots \wedge \ell_{n_i}^i)$ as well. As for Lemma 3 Kleene's logic and Lukasiewicz's logic coincide, the results hold also under Lukasiewicz's logic. \square

We now provide the proofs of Theorem 7 and Theorem 8, whose statements are recalled below.

Theorem 7. For any NCAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$, the problem

- CA_S^σ is: (i) NP-complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}\}$ and $\sigma \in \{K, L\}$,
(ii) Σ_2^P -complete for $S = \text{sst}$ and $\sigma \in \{K, L\}$.
- SA_S^σ is: (i) coNP-complete for any semantics $S \in \{\text{co}, \text{st}\}$ and $\sigma \in \{K, L\}$,
(ii) Π_2^P -complete for $S \in \{\text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Proof. (Hardness.) The lower bound results for CA_S^* with $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and for SA_S^* with $S \in \{\text{st}, \text{pr}, \text{sst}\}$ derive from the fact that they hold for any CAF $\langle \mathcal{A}, \mathcal{R}, C \rangle$ where $C = \emptyset$, that is, for AF. As regard the coNP-hardness for SA_{co}^* , it suffices to observe that, in the proof of Theorem 1, to show the coNP-hardness of SA_{co}^* , a CAF where C consists only of denial constraints is used. That is, the proof provided in proof of Theorem 1 still holds also for NCAFs.

(Membership.) All the membership results except that for CA_{pr}^* derive from the analogous ones of Theorem 1. Regarding CA_{pr}^* , it is in NP since after guessing a set $S \subseteq \mathcal{A}$ of arguments containing a , we only need to check that (i) S is a complete extension for $\langle \mathcal{A}, \mathcal{R} \rangle$, i.e., S is admissible and contains all the arguments it defends (in PTIME) and (ii) $S \models C$ (in PTIME). In fact, since $C \cup \mathcal{W}$ consists of denial constraints only, by Lemma 4, we do not have to check that S is maximal, as it is sufficient to check that S is a complete extension satisfying C . Indeed, by Lemma 4, if there is a complete extension containing a and satisfying C , then there is also a preferred extension containing a and satisfying C . \square

Theorem 8. For any NWF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$, the problem

- $CA_{\text{ms}S}^\sigma$ is Σ_2^P -complete for any semantics $S \in \{\text{co}, \text{st}, \text{pr}, \text{ss}\}$ and $\sigma \in \{K, L\}$.
- $SA_{\text{ms}S}^\sigma$ is Π_2^P -complete for $S \in \{\text{co}, \text{st}, \text{pr}, \text{sst}\}$ and $\sigma \in \{K, L\}$.

Proof. (Membership.) The membership results under Kleene logic follow from the general case of WAF (cf. Theorem 2) as any NWF is also a WAF. Moreover, since the Lukasiewicz and Kleene logics coincide for denial constraints (cf. Lemma 3), the results also hold under Lukasiewicz logic.

(Hardness.)

- ($SA_{\text{ms-pr}}^*$, $CA_{\text{ms-sst}}^*$ and $SA_{\text{ms-sst}}^*$) The lower bound results derive from the fact that they hold for any NWF $\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle$ where $C = \mathcal{W} = \emptyset$, that is, for AF.
- ($CA_{\text{ms-co}}^*$ and $SA_{\text{ms-co}}^*$) The lower bound for credulous acceptance under complete semantics can be proved by showing that, let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, $\text{sst}(\Lambda) = \text{ms-co}(Y = \langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle)$ where $C = \emptyset$, and $\mathcal{W} = \{x \wedge \neg x \Rightarrow f \mid x \in \mathcal{A}\}$. First, recall that $\text{sst}(\Lambda) \subseteq \text{co}^*(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Then, since \mathcal{W} selects from the extensions in $\text{co}^*(\langle \mathcal{A}, \mathcal{R}, C \rangle)$ those that are maximal w.r.t. the presence of arguments whose status is either true or false (or equivalently that are minimal w.r.t. undecided arguments), we have that $\text{ms-co}^*(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle) = \text{sst}(\langle \mathcal{A}, \mathcal{R} \rangle)$, from which the result follows.
- ($CA_{\text{ms-pr}}^*$) The lower bound for credulous acceptance under preferred semantics can be proved by showing that, let $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, $\text{sst}(\Lambda) = \text{ms-pr}(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle)$ where $C = \emptyset$, and $\mathcal{W} = \{x \wedge \neg x \Rightarrow f \mid x \in \mathcal{A}\}$. Recall that $\text{sst}(\Lambda) \subseteq \text{pr}^*(\langle \mathcal{A}, \mathcal{R}, C \rangle)$. Again, since \mathcal{W} selects from the extensions in $\text{pr}^*(\langle \mathcal{A}, \mathcal{R}, C \rangle)$ those that are maximal w.r.t. true or false arguments (or equivalently that are minimal w.r.t. undecided arguments), we have that $\text{ms-pr}^*(\langle \mathcal{A}, \mathcal{R}, C, \mathcal{W} \rangle) = \text{sst}(\langle \mathcal{A}, \mathcal{R} \rangle)$, from which the result follows.

- $(CA_{ms-st}^*$ and SA_{ms-st}^*). The hardness results derive from the fact that in reducing disjunctive programs to logic programs with strong and weak constraints (see Appendix C), which in turn can be rewritten into a WAF, all constraints used are denials. Recall that under stable semantics, a constraint $t \Rightarrow a \vee b$ can be rewritten into the equivalent constraint $\neg a \wedge \neg b \Rightarrow f$. \square

Appendix B. Normal and disjunctive logic programs

We briefly review (normal) logic programs and disjunctive logic programs, recalling how partial stable models can be computed by reducing to total stable models, as well as the relationship between logic programs and abstract argumentation frameworks.

B.1. Normal logic programs

The semantics of a logic program is given by the set of its partial stable models (PSMs) (corresponding to complete extensions of AFs [40]). We summarize the basic concepts which underly the notion of PSMs [88].

A (normal) logic program (LP) is a set of rules of the form $A \leftarrow B_1 \wedge \dots \wedge B_n$, with $n \geq 0$, where A is an atom, called head, and $B_1 \wedge \dots \wedge B_n$ is a conjunction of literals, called body. We consider programs without function symbols. Given a program P , $ground(P)$ denotes the set of all ground instances of the rules in P . The Herbrand Base of a program P , i.e., the set of all ground atoms which can be constructed using predicate and constant symbols occurring in P , is denoted by B_P , whereas $\neg B_P$ denotes the set $\{\neg A \mid A \in B_P\}$. Analogously, for any set $S \subseteq B_P \cup \neg B_P$, $\neg S$ denotes the set $\{\neg A \mid A \in S\}$, where $\neg \neg A = A$. Given $I \subseteq B_P \cup \neg B_P$, $pos(I)$ (resp., $neg(I)$) stands for $I \cap B_P$ (resp., $\neg I \cap B_P$). I is consistent if $pos(I) \cap neg(I) = \emptyset$, otherwise I is inconsistent.

Given a program P , $I \subseteq B_P \cup \neg B_P$ is an interpretation of P if I is consistent. Also, I is total if $pos(I) \cup neg(I) = B_P$, partial otherwise. A partial interpretation M of a program P is a partial model of P if for each $\neg A \in M$ every rule in $ground(P)$ having as head A contains at least one body literal B such that $\neg B \in M$. Given a program P and a partial model M , the positive instantiation of P w.r.t. M , denoted by P^M , is obtained from $ground(P)$ by deleting: (a) each rule containing a negative literal $\neg A$ such that $A \in pos(M)$; (b) each rule containing a literal B such that neither B nor $\neg B$ is in M ; (c) all the negative literals in the remaining rules. M is a partial stable model of P iff M is the minimal model of P^M . Alternatively, P_M could be built by replacing every negated body literal in $ground(P)$ by its truth value.

The set of partial stable models of a logic program P , denoted by $PS(P)$, define a meet semi-lattice. The well-founded model (denoted by $WF(P)$) and the maximal-stable models $MS(P)$,⁶ are defined by considering \subseteq -minimal and \subseteq -maximal elements. The set of (total) stable models (denoted by $ST(P)$) is obtained by considering the maximal-stable models which are total, whereas the least-undefined models (denoted by $LS(P)$) are obtained by considering the maximal-stable models with a \subseteq -minimal set of undefined atoms (i.e., atoms which are neither true or false). The max-deterministic model (denoted by $MD(P)$) is the \subseteq -maximal PSM contained in every maximal-stable model [87]. To denote a specific semantics, we use the acronyms ps, st, ms, ls, wf and md for the semantics partial stable, (total) stable, maximal stable, least-undefined stable, well-founded, and max-deterministic, respectively.

The semantics of a logic program is given by the set of its partial stable models or by one of the restricted sets above recalled.

B.2. Disjunctive logic programs

The partial stable model semantics has been extended to disjunctive logic programs (DLPs), that is, programs whose rules allow disjunctive heads.

Positive disjunctive programs may have more than one minimal model. A set of literals M is a partial stable model of P iff M is a minimal model of P^M , the positive disjunctive program derived through the same steps defined earlier for normal programs.

Example 11. Consider the disjunctive program P :

$$\begin{array}{ll} a \leftarrow \neg b & c \vee d \leftarrow a \\ b \leftarrow \neg a & d \leftarrow c \end{array}$$

There are three partial stable models for P : $M_1 = \{\neg c\}$, $M_2 = \{\neg a, b, \neg c, \neg d\}$ and $M_3 = \{a, \neg b, \neg c, d\}$. M_2 and M_3 are maximal stable models, as well as total stable and least undefined stable models. \square

B.2.1. Computing partial stable models

A technique for computing partial stable models using Answer Set Programming (ASP) solvers (solvers computing total stable models) has been proposed in [68]. For the sake of presentation, here we consider ground programs.

For each atom a in P we consider a dummy atom a^* whose meaning is *atom a is potentially true*. The program P^* is defined as follows:

$$\begin{aligned} P^* = & \{a \leftarrow b_1, \dots, b_m, \neg c_1^*, \dots, c_n^* \mid a \leftarrow b_1, \dots, b_m, \neg c_1, \dots, c_n \in P\} \cup \\ & \{a^* \leftarrow b_1^*, \dots, b_m^*, \neg c_1^*, \dots, c_n^* \mid a \leftarrow b_1, \dots, b_m, \neg c_1, \dots, c_n \in P\} \cup \\ & \{a^* \leftarrow a \mid a \text{ occurs in } P\} \end{aligned}$$

⁶ Corresponding to Dung's preferred extensions [47].

It has been shown that for each partial stable model M of P there is a total stable N of P^* and vice versa. M is obtained from N as follows $M = \{a \mid a^* \in N \wedge a \in N\} \cup \{\neg a \mid \neg a^* \in N \wedge \neg a \in N\}$. On the other side N can be derived from M as follows $N = \{\ell, \ell^* \mid \ell \in M\} \cup \{\ell^*, \neg \ell \mid \ell, \neg \ell \notin M\}$.

Example 12. Consider the program P of Example 11. The program P^* is defined as follows:

$$\begin{array}{lll} a \leftarrow \neg b^* & a^* \leftarrow \neg b & a^* \leftarrow a \\ b \leftarrow \neg a^* & b^* \leftarrow \neg a & b^* \leftarrow b \\ c \vee d \leftarrow a & c^* \vee d^* \leftarrow a^* & c^* \leftarrow c \\ d \leftarrow c & d^* \leftarrow c^* & d^* \leftarrow d \end{array}$$

P^* has three total stable models:

- $N_1 = \{a^*, \neg a, b^*, \neg b, \neg c^*, \neg c, d^*, \neg d\}$,
- $N_2 = \{\neg a^*, \neg a, b^*, b, \neg c^*, \neg c, \neg d^*, \neg d\}$,
- $N_3 = \{a^*, a, \neg b^*, \neg b, \neg c^*, \neg c, d^*, d\}$,

corresponding to PSMs M_1, M_2 and M_3 of Example 11, respectively.

B.3. Logic programs and argumentation frameworks

It is well-knowns that there is a tight relationship between AFs and LPs under partial stable model semantics. In particular, for each AF Λ there is a normal logic program P_Λ (derived from Λ) such that the set complete (resp., grounded, stable, preferred, semi-stable) extensions of Λ is equivalent to the set of partial (resp., well-founded, total, maximal, least-undefined) stable models of P_Λ [40]. It has been shown also the reverse result for all semantics, except for the least-undefined stable model semantics.

Appendix C. Weak constrained logic programs

Logic programs with weak constraints have been proposed in [37] and implemented in the well-known DLV system [9]. Here, we consider weak constraints with a syntax similar to that defined in the core of the paper for AFs and a maximal-set semantics.⁷

Definition 16. A (ground) logic program with weak constraints (WLP) is a triple $\langle LP, C, \mathcal{W} \rangle$, where LP is a (ground) normal logic program, C is a set of (ground, strong) constraints and \mathcal{W} is a set of (ground) weak constraints.

The semantics of a weak constrained logic program is given by the partial (resp., maximal, total, least-undefined) stable models that satisfy all strong constraints in C and a maximal set of weak constraints in \mathcal{W} .

The set of maximal-set partial (resp., total, maximal, least-undefined) stable models of a WLP P is denoted by $MS\text{-}\mathcal{PM}(P)$ (resp., $MS\text{-}\mathcal{ST}(P)$, $MS\text{-}\mathcal{MS}(P)$, $MS\text{-}\mathcal{LS}(P)$).

Example 13. Consider the weak constrained program P derived from the WAF of Example 8 (example in the core of the paper):

- $LP = \{a \leftarrow \neg b; \quad b \leftarrow \neg a; \quad c \leftarrow \neg d; \quad d \leftarrow \neg c\}$;
- $C = \emptyset$;
- $\mathcal{W} = \{w_1 = c \Rightarrow f, w_2 = a \vee \neg a \Rightarrow u\}$.

It is easy to check that P has 9 partial stable models: $M_0 = \{\}$, $M_1 = \{a, \neg b\}$, $M_2 = \{\neg a, b\}$, $M_3 = \{c, \neg d\}$, $M_4 = \{\neg c, d\}$, $M_5 = \{a, \neg b, c, \neg d\}$, $M_6 = \{a, \neg b, \neg c, d\}$, $M_7 = \{\neg a, b, c, \neg d\}$ and $M_8 = \{\neg a, b, \neg c, d\}$. In particular, M_0 is the well-founded model, whereas E_5, E_6, E_7, E_8 are total, maximal and least-undefined stable models of P . These models correspond to the complete extensions of the AF in Example 8.

Regarding the satisfaction of weak constraints, we have that $M_0 \models \{w_2\}$, $M_4 \models \{w_1, w_2\}$, $M_6 \models \{w_1\}$, and $M_8 \models \{w_1\}$, whereas the other partial stable models do not satisfy any constraint. Therefore, the maximal-set maximal (total, least-undefined) stable models are M_6 and M_8 , whereas there is only one maximal-set partial stable model, which is M_4 . These models correspond to the maximal-set extensions of the WAF in Example 8. \square

C.1. Mapping disjunctive programs to weak constrained logic programs

It has been shown that the introduction of disjunctive heads increases the expressivity of logic programs of one level in the polynomial hierarchy [58,59]. Restricted cases of weak constraints under (total) stable model semantics have been studied in [66], where it is shown that the expressivity of LPs grows of one level (CA_{st} is Σ_2^P -complete and SA_{st} is Π_2^P -complete) and in [37], where it is shown that the expressivity of DLP grows to Δ_2^P (the paper assumes a maximum-cardinality based semantics). We now show that (normal) WLPs, under maximal-set semantics, are no less expressive than DLPs. This is shown by mapping DLPs to WLPs.

⁷ Weak constraints implemented in DLV have a similar syntax, but a maximum-cardinality based semantics.

Definition 17. For any disjunctive program P , $T(P) = \langle LP, C, \mathcal{W} \rangle$ is the program with constraints (and without disjunctive rules) derived as follows:

- LP is derived from P by replacing every disjunctive rule $a_1 \vee \dots \vee a_n \leftarrow \varphi$ with $2 \times n$ normal rules of the form: $a_i \leftarrow \varphi, \neg \bar{a}_i$ and $\bar{a}_i \leftarrow \varphi, \neg a_i$
- $C = \{ \varphi \Rightarrow a_1 \vee \dots \vee a_n \mid a_1 \vee \dots \vee a_n \leftarrow \varphi \in P, n > 1 \}$ ⁸;
- $\mathcal{W} = \{ a \Rightarrow f \mid a \text{ occurs in } \text{ground}(P) \}$.

Given a disjunctive logic program P and an interpretation I of $T(P)$, then $I[P]$ denotes the subset of I whose atoms occur in P .

Theorem 10. For any disjunctive logic program P , it is the case that $ST(P) = \{ M[P] \mid M \in \text{MS-ST}(T(P)) \}$.

Proof. Let $T(P) = \langle LP, C, \mathcal{W} \rangle$, we first prove that for every total stable model M of P there is M' such that $M \cup M'$ is a total stable model of LP and that such model is a best model of $T(P)$. Given M , for each overlined atom \bar{a} occurring in $T(P)$, M' contains either \bar{a} or $\neg \bar{a}$. More specifically, for every \bar{a} occurring in $T(P)$, if there are in $T(P)$ two rules $a \leftarrow \varphi, \neg \bar{a}$ and $\bar{a} \leftarrow \varphi, \neg a$ such that $M \models \varphi$, then $\neg a \in M$ implies $\bar{a} \in M'$, otherwise $\neg \bar{a} \in M'$. Clearly, if M is a minimal model of P^M , $M \cup M'$ is the minimal model of $LP^{M \cup M'}$. Indeed, for each rule having an atom $a \in M$ occurring in the head of a rule r of P such that $M \models \text{body}(r)$, there is a rule r' in $T(P)$ (derived from r) such that $M \cup M' \models \text{body}(r')$. Moreover, $M \cup M' \models C$ as M is a model of P and is a best model since M is a minimal model that satisfies a maximal set of constraints in \mathcal{W} .

We now show that for each $M \in \text{MS-ST}(T(P))$, $M[P]$ is a stable model of P . The set $ST(LP)$ could contain stable models which do not satisfy rules of P with disjunctive heads. However, as for each disjunctive rule $a_1 \vee \dots \vee a_n \leftarrow \varphi$ we have a strong constraint $\varphi \Rightarrow a_1 \vee \dots \vee a_n$, these models are not feasible. The maximization of weak constraints guarantees that our best models are minimal model of P . \square

Observe that for normal logic programs P , the corresponding WLP is $T(P) = \langle LP, C, \mathcal{W} \rangle$, where we have that $LP = P$ and $C = \emptyset$. Notably, the set of weak constraints \mathcal{W} is useless as minimality is implicit in the total stable model semantics.

Example 14. Consider the program of Example 11. The corresponding WLP is $\langle LP, C, \mathcal{W} \rangle$, where LP consists of the following rules:

- $a \leftarrow \neg b$
- $b \leftarrow \neg a$
- $c \leftarrow a, \neg \bar{c}$
- $\bar{c} \leftarrow a, \neg c$
- $d \leftarrow a, \neg \bar{d}$
- $\bar{d} \leftarrow a, \neg d$
- $d \leftarrow c$

whereas $C = \{ \zeta = a \Rightarrow c \vee d \}$ and $\mathcal{W} = \{ w_1 = a \Rightarrow f; w_2 = b \Rightarrow f; w_3 = c \Rightarrow f; w_4 = d \Rightarrow f \}$.

The total stable models of LP are:

- $T_1 = \{ \neg a, b, \neg c, \neg d, \neg \bar{c}, \neg \bar{d} \}$,
- $T_2 = \{ a, \neg b, \neg c, \neg d, \bar{c}, \bar{d} \}$,
- $T_3 = \{ a, \neg b, \neg c, d, \bar{c}, \neg \bar{d} \}$,
- $T_4 = \{ a, \neg b, c, d, \neg \bar{c}, \neg \bar{d} \}$.

As $T_2 \not\models \zeta$, the feasible models are T_1, T_3, T_4 . Considering the satisfaction of weak constraints, we have that:

- $T_1 \models \{ w_1, w_3, w_4 \}$,
- $T_3 \models \{ w_2, w_3 \}$,
- $T_4 \models \{ w_2 \}$.

Therefore, the best models are T_1 and T_3 . Comparing the best models of $T(P)$ with the stable model of P (see Example 11) we have $T_1[P] = M_2$ and $T_3[P] = M_3$. \square

The next final example shows how a disjunctive logic program P is first translated into a normal logic programs P^* so that $C\mathcal{O}(P) \equiv ST(P^*)$ (using the approach in B.2.1), and then how P^* is mapped into a logic program with weak constraints.

⁸ Under stable semantics or Kleene logic, $\kappa : \varphi \Rightarrow a_1 \vee \dots \vee a_n$ can be equivalently rewritten as a denial constraint $\kappa' : \varphi \wedge \neg a_1 \wedge \dots \wedge \neg a_n \Rightarrow \bar{f}$.

Example 15. Consider again the program P of Example 11, and the corresponding program P^* of Example 12. The corresponding WLP in this case is $T(P^*) = \langle LP, C, \mathcal{W} \rangle$, where LP is the set of rules:

$$\begin{array}{lll} a \leftarrow \neg b^* & a^* \leftarrow \neg b & a^* \leftarrow a \\ b \leftarrow \neg a^* & b^* \leftarrow \neg a & b^* \leftarrow b \\ c \leftarrow a, \neg \bar{c} & c^* \leftarrow a^*, \neg \bar{c}^* & c^* \leftarrow c \\ \bar{c} \leftarrow a, \neg c & \bar{c}^* \leftarrow a^*, \neg c^* & d^* \leftarrow d \\ d \leftarrow a, \neg \bar{d} & d^* \leftarrow a^*, \neg \bar{d}^* & \\ \bar{d} \leftarrow a, \neg d & \bar{d}^* \leftarrow a^*, \neg d^* & \\ d \leftarrow c & d^* \leftarrow c^* & \end{array}$$

Moreover, $C = \{\zeta = a \Rightarrow c \vee d; \quad \zeta^* = a^* \Rightarrow c^* \vee d^*\}$ and $\mathcal{W} = \{w_1 = a \Rightarrow f; w_2 = b \Rightarrow f; w_3 = c \Rightarrow f; w_4 = d \Rightarrow f; w_1^* = a^* \Rightarrow f; w_2^* = b^* \Rightarrow f; w_3^* = c^* \Rightarrow f; w_4^* = d^* \Rightarrow f\}$.

The total stable models of LP are:

- $N_0 = \{\neg a, a^*, \neg b, b^*, \neg c, \neg c^*, \neg d, \neg d^*, \neg \bar{c}, \neg \bar{d}, \bar{c}^*, \bar{d}^*\}$,
- $N_1 = \{\neg a, a^*, \neg b, b^*, \neg c, \neg c^*, \neg d, d^*, \neg \bar{c}, \neg \bar{d}, \bar{c}^*, \neg \bar{d}^*\}$,
- $N_2 = \{\neg a, a^*, \neg b, b^*, \neg c, c^*, \neg d, d^*, \neg \bar{c}, \neg \bar{d}, \neg \bar{c}^*, \neg \bar{d}^*\}$,
- $N_3 = \{\neg a, \neg a^*, b, b^*, \neg c, \neg c^*, \neg d, \neg d^*, \neg \bar{c}, \neg \bar{d}, \neg \bar{c}^*, \neg \bar{d}^*\}$,
- $N_4 = \{a, a^*, \neg b, \neg b^*, \neg c, \neg c^*, \neg d, \neg d^*, \bar{c}, \bar{d}, \bar{c}^*, \bar{d}^*\}$,
- $N_5 = \{a, a^*, \neg b, \neg b^*, \neg c, \neg c^*, \neg d, d^*, \bar{c}, \bar{c}^*, \bar{d}, \neg \bar{d}^*\}$,
- $N_6 = \{a, a^*, \neg b, \neg b^*, \neg c, \neg c^*, d, d^*, \bar{c}, \bar{c}^*, \neg \bar{d}, \neg \bar{d}^*\}$,
- $N_7 = \{a, a^*, \neg b, \neg b^*, \neg c, c^*, \neg d, d^*, \bar{c}, \neg \bar{c}^*, \bar{d}, \neg \bar{d}^*\}$,
- $N_8 = \{a, a^*, \neg b, \neg b^*, \neg c, c^*, d, d^*, \bar{c}, \neg \bar{c}^*, \neg \bar{d}, \neg \bar{d}^*\}$,
- $N_9 = \{a, a^*, \neg b, \neg b^*, c, c^*, d, d^*, \neg \bar{c}, \neg \bar{c}^*, \neg \bar{d}, \neg \bar{d}^*\}$.

Moreover, as $N_0 \not\models \zeta^*$ and $N_4, N_5, N_7 \not\models \zeta$, the feasible models are $N_1, N_2, N_3, N_6, N_8, N_9$. Considering the satisfaction of weak constraints, we have that:

- $N_1 \models \{w_1, w_2, w_3, w_4, w_3^*\}$,
- $N_2 \models \{w_1, w_2, w_3, w_4\}$,
- $N_3 \models \{w_1, w_3, w_4, w_1^*, w_3^*, w_4^*\}$,
- $N_6 \models \{w_2, w_3, w_2^*, w_3^*\}$,
- $N_8 \models \{w_2, w_3, w_2^*\}$,
- $N_9 \models \{w_2, w_2^*\}$.

Therefore, the best models are N_1, N_3 and N_6 . Comparing the best models of $T(P^*)$ with the stable model of P (see Example 11) we have $N_1[P] = M_1, N_3[P] = M_2$ and $N_6[P] = M_3$. \square

Appendix D. Abstract argumentation framework with epistemic constraints

We now review the *Epistemic Argumentation Framework* [89], which extends Dung's framework with epistemic constraints, and then show a relationship with the framework proposed in this paper.

D.1. Labelling

Argumentation semantics can be also defined in terms of *labelling* [21]. A labelling for an AF $\langle \mathcal{A}, \mathcal{R} \rangle$ is a total function $Lab : \mathcal{A} \rightarrow \{\mathbf{in}, \mathbf{out}, \mathbf{undec}\}$ assigning to each argument a label: $Lab(a) = \mathbf{in}$ means that a is accepted, $Lab(a) = \mathbf{out}$ means that a is rejected, and $Lab(a) = \mathbf{undec}$ means that a is undecided.

Let $\mathbf{in}(Lab) = \{a \mid a \in \mathcal{A} \wedge Lab(a) = \mathbf{in}\}$, $\mathbf{out}(Lab) = \{a \mid a \in \mathcal{A} \wedge Lab(a) = \mathbf{out}\}$, and $\mathbf{undec}(Lab) = \{a \mid a \in \mathcal{A} \wedge Lab(a) = \mathbf{undec}\}$, a labelling Lab can be represented by means of a triple $\langle \mathbf{in}(Lab), \mathbf{out}(Lab), \mathbf{undec}(Lab) \rangle$.

Given an AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, a labelling Lab for \mathcal{A} is said to be *conflict-free* if there are no two arguments $a, b \in \mathbf{in}(Lab)$ such that $(a, b) \in \mathcal{R}$, and *admissible* (or *legal*) if $\forall a \in \mathbf{in}(Lab) \cup \mathbf{out}(Lab)$ it holds that:

- (i) $Lab(a) = \mathbf{out}$ iff $\exists (b, a) \in \mathcal{R}$ such that $Lab(b) = \mathbf{in}$; and
- (ii) $Lab(a) = \mathbf{in}$ iff $\forall (b, a) \in \mathcal{R}, Lab(b) = \mathbf{out}$ holds.

Moreover, Lab is a *complete* labelling iff conditions (i) and (ii) hold for all arguments $a \in \mathcal{A}$.

Between complete extensions and complete labellings there is a bijective mapping defined as follows: for each extension E there is a unique labelling $Lab(E) = \langle E, Def(E), \mathcal{A} \setminus (E \cup Def(E)) \rangle$ and for each labelling Lab there is a unique extension, that is $\mathbf{in}(Lab)$. We say that $Lab(E)$ is the labelling *corresponding* to E . Moreover, we say that $Lab(E)$ is an S -labelling for a given AF Λ and semantics $S \in \{\text{co}, \text{pr}, \text{st}, \text{ss}, \text{gr}\}$ iff E is an S -extension of Λ .

We say that the *status of an argument* a w.r.t. a labelling Lab (or its corresponding extension $\mathbf{in}(Lab)$) is **in** (resp. **out**, **undec**) iff $Lab(a) = \mathbf{in}$ (resp. $Lab(a) = \mathbf{out}$, $Lab(a) = \mathbf{undec}$). We will avoid to mention explicitly the labelling (or the extension) whenever it is understood.

D.2. AF with epistemic constraints

Given an AF $\Lambda = \langle \mathcal{A}, \mathcal{R} \rangle$, an epistemic atom over Λ is of the form $\mathbf{K}\varphi$ or $\mathbf{M}\varphi$, where φ is a propositional formula built from $\lambda_{\mathcal{A}} = \{\mathbf{in}(a), \mathbf{out}(a), \mathbf{undec}(a) \mid a \in \mathcal{A}\}$ by using the connectives \neg , \vee , and \wedge . An epistemic literal is an epistemic atom or its negation. An *epistemic formula* (over $\lambda_{\mathcal{A}}$) is a propositional formula constructed over epistemic literals and connectives \wedge and \vee . Intuitively, $\mathbf{K}\varphi$ (resp. $\mathbf{M}\varphi$) means that the considered agent believes that φ is always true (resp. φ is possibly true).

A labelling Lab satisfies a formula φ (denoted as $Lab \models \varphi$) if the formula obtained from φ by replacing every atom occurring in Lab with \mathbf{t} (true), and every atom not occurring in Lab with \mathbf{f} (false), evaluates to true.

A set SL of labellings satisfies an epistemic formula φ , denoted as $SL \models \varphi$, if one of the following conditions holds:

- $\varphi = \mathbf{t}$,
- $\varphi = \mathbf{K}\psi$ and $Lab \models \psi$ for every $Lab \in SL$,
- $\varphi = \mathbf{M}\psi$ and $Lab \models \psi$ for some $Lab \in SL$,
- $\varphi = \neg\psi$ and $SL \not\models \psi$,
- $\varphi = \varphi_1 \wedge \varphi_2$ and ($SL \models \varphi_1$ and $SL \models \varphi_2$),
- $\varphi = \varphi_1 \vee \varphi_2$ and ($SL \models \varphi_1$ or $SL \models \varphi_2$).

An epistemic formula φ is consistent if there exists a (non-empty) set SL of labellings such that $SL \models \varphi$; otherwise, φ is inconsistent. The following basic properties hold:

- $SL \models \neg\mathbf{M}\varphi$ iff $SL \models \mathbf{K}\neg\varphi$,
- $SL \models \neg\mathbf{K}\varphi$ iff $SL \models \mathbf{M}\neg\varphi$,
- $SL \models \mathbf{M}(\varphi_1 \vee \varphi_2)$ iff $SL \models \mathbf{M}\varphi_1 \vee SL \models \mathbf{M}\varphi_2$,
- $SL \models \mathbf{K}(\varphi_1 \wedge \varphi_2)$ iff $SL \models \mathbf{K}\varphi_1 \wedge SL \models \mathbf{K}\varphi_2$.

Definition 18 (EAF Syntax). An *Epistemic AF (EAF)* is a triple $\langle \mathcal{A}, \mathcal{R}, \varphi \rangle$ where $\langle \mathcal{A}, \mathcal{R} \rangle$ is an AF and φ is an epistemic formula to be satisfied, also called epistemic constraint.

The semantics of EAF relies on the concept of S -epistemic labelling, that is a maximal set of labellings of the underlying AF satisfying the epistemic constraint.

Definition 19 (EAF Semantics). Let $\mathcal{E} = \langle \mathcal{A}, \mathcal{R}, \varphi \rangle$ be an EAF and $S \in \{\text{gr}, \text{co}, \text{pr}, \text{st}, \text{sst}\}$ a semantics. A set SL of labellings is an *S -epistemic labelling set* of \mathcal{E} if (i) each $Lab \in SL$ is an S -labelling of $\langle \mathcal{A}, \mathcal{R} \rangle$, and (ii) SL is a \subseteq -maximal set of S -labellings of $\langle \mathcal{A}, \mathcal{R} \rangle$ that satisfies φ .

An EAF may have multiple S -epistemic labelling sets. In fact, an S -epistemic labelling set is a collection of S -labellings that represent the belief of an agent. In particular, EAF $\mathcal{E} = \langle \mathcal{A}, \mathcal{R}, \mathbf{K}\mathbf{t} \rangle$ has a unique S -epistemic labelling set that coincides with the set of S -labellings of the underlying AF. By definition, an EAF always has a (possibly empty) S -epistemic labelling set.

D.3. AF with labelled constraints

An EAF is called *Labelled CAF (LabCAF)* if it is of the form $\langle \mathcal{A}, \mathcal{R}, \mathbf{K}\varphi \rangle$, where φ is a propositional formula built from $\Lambda_{\mathcal{A}}$ and using the operators \wedge , \vee and \neg . Any LabCAF $\langle \mathcal{A}, \mathcal{R}, \mathbf{K}\varphi \rangle$ has a unique labelling set which consists of the set of S -extensions for $\langle \mathcal{A}, \mathcal{R} \rangle$ satisfying φ .

Theorem 11. Let $\mathcal{E} = \langle \mathcal{A}, \mathcal{R}, \mathbf{K}\varphi \rangle$ be an LabCAF, $S \in \{\text{co}, \text{pr}, \text{st}, \text{sst}, \text{gr}\}$ a semantics and SL the unique set of S -labellings for \mathcal{E} . Then, there exist a CAF $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ such that C can be derived from φ in linear time and $SL = S(\Omega)$.

Proof. Let $\mathcal{E} = \langle \mathcal{A}, \mathcal{R}, \mathbf{K}\varphi \rangle$ be an LabCAF, and $\Omega = \langle \mathcal{A}, \mathcal{R}, C \rangle$ where $C = \varphi'$ is obtained as follows. Any atom $\mathbf{in}(x) \in \varphi$ is replaced with $(\mathbf{t} \Rightarrow x)$, $\mathbf{out}(x) \in \varphi$ is replaced with $(x \Rightarrow \mathbf{f})$, and any atom $\mathbf{undec}(x) \in \varphi$ is replaced with $(\mathbf{u} \Rightarrow x \wedge \neg x)$. Observe that $C = \varphi'$ is built in linear time w.r.t. φ .

We now prove that $E \in SL$ iff $E \in S(\Omega)$. For any atom $\mathbf{in}(x) \in \varphi$, we have that $\mathbf{in}(x)$ is true w.r.t. E iff $x \in E$, implying that $\mathbf{t} \Rightarrow x$ is true iff $x \in E$. Analogously, for any atom $\mathbf{out}(x) \in \varphi$, we have that $\mathbf{out}(x)$ is true w.r.t. E iff $x \in Def(E)$, implying that $x \Rightarrow \mathbf{f}$ is true iff $x \in Def(E)$. Assume now that we have an atom $\mathbf{undec}(x) \in \varphi$. We have that $\mathbf{undec}(x)$ is true w.r.t. E iff $x \in \mathcal{A} \setminus (E \cup Def(E))$. This implies that $\mathbf{u} \Rightarrow x \wedge \neg x$ is true iff $x \in \mathcal{A} \setminus (E \cup Def(E))$ or the consequent $(x \wedge \neg x)$ is true. As the latter is a contradiction, we have that $\mathbf{undec}(x)$ is true w.r.t. E implies that $\mathbf{u} \Rightarrow x \wedge \neg x$ is true w.r.t. E . As the inverse direction holds by reasoning analogously, we showed that $E \in SL$ iff $E \in S(\Omega)$. \square

Therefore, CAF is at least expressive as LabCAF.

Appendix E. Background on ADF

In this appendix, we review the syntax and the semantics of the ADF framework. An *Abstract Dialectical Framework (ADF)* [36] is a triple $D = \langle S, L, C \rangle$ where:

- S is a set of *statements* (also called arguments or nodes);
- $L \subseteq S \times S$ is a set of *links*;
- $C = \{C_s\}_{s \in S}$ is a set of total functions $C_s : 2^{par(s)} \rightarrow \{t, f\}$, one for each statement, where $par(s) = \{x \mid x \in S \wedge (x, s) \in L\}$. C_s is called *acceptance condition* of s .

Here, links in L represent dependencies: the status of a node s depends only on the status of its parents, $par(s)$, the nodes with a direct link to s . Each node s is associated with an acceptance condition C_s that specifies the conditions under which s is acceptable. Acceptance conditions are represented by a collection $\{C_s\}_{s \in S}$ of propositional formulae, using atoms from $par(s)$ and the logical connectives \wedge, \vee, \neg . Observe that we could assume that L is understood and simply denote an ADF by a pair $\langle S, C \rangle$ (instead of a triple).

Semantics assign to ADFs a collection of (3-valued) interpretations mapping each statement to truth values $\{t, f, u\}$, denoting *true*, *false*, and *undefined*, respectively. Truth values are partially ordered by \leq_i according to their information content: $u <_i t$ and $u <_i f$ and no other pair is in \leq_i . The information ordering \leq_i extends in a straightforward way to interpretations v_1, v_2 over S in that $v_1 \leq_i v_2$ iff $v_1(s) \leq_i v_2(s)$ for all $s \in S$. An interpretation v is 2-valued if all statements are mapped to t or f . For interpretations v and ω , we say that ω *extends* v iff $v \leq_i \omega$. We denote by $[v]_2$ the set of all *completions* of v , that is, 2-valued interpretations that extend v . For an ADF $D = \langle S, C \rangle$, $s \in S$, and an interpretation v , the characteristic function is $\Gamma_D(v) = v'$, where

$$v'(s) = \begin{cases} t & \text{if } \omega(C_s) = t \text{ for all } \omega \in [v]_2 \\ f & \text{if } \omega(C_s) = f \text{ for all } \omega \in [v]_2 \\ u & \text{otherwise} \end{cases}$$

That is, operator Γ_D returns an interpretation mapping a statement s to t (resp., f) iff all 2-valued interpretations extending v evaluate C_s to t (resp., f). Intuitively, Γ_D checks if truth values can be justified based on the information in v and the acceptance conditions. Note that Γ_D is defined on 3-valued interpretations, while acceptance conditions are evaluated under their 2-valued completions. Given an ADF $D = \langle S, C \rangle$, an interpretation v is (w.r.t. D):

- *admissible*, if $v \leq_i \Gamma_D(v)$;
- *complete*, if $v = \Gamma_D(v)$;
- *preferred*, if v is \subseteq -maximal admissible w.r.t. \leq_i ;
- *grounded*, if v is complete and there is no other complete interpretation v' such that $v' \leq_i v$.

A 2-valued interpretation v is a *model* of D if $v(s) = v(C_s)$ for every $s \in S$. The definition of the *stable* semantics for ADFs is inspired by the stable semantics for logic programs: its purpose is to disallow cyclic supports within a model. In particular, (i) to be a stable model of D , v needs to be a model of D , and (ii) $S^v = \{s \in S \mid v(s) = t\}$ must equal the statements set to *true* in the grounded interpretation of the *reduced* ADF $D^v = \langle S^v, \{C_s^v\}_{s \in S^v} \rangle$, where for $s \in S^v$ we set $C_s^v = C_s[b/f \mid v(b) = f]$. If $v|_{S^v}$ is the interpretation v projected on S^v , that is, $v|_{S^v}(s) = v(s)$ for $s \in S^v$ and undefined otherwise, then the latter amounts to the fact that $v|_{S^v}$ be the grounded interpretation of D^v .

As shown in [36], these semantics generalize the corresponding ones defined for AF.

Example 16. [33] For the ADF $D = \langle \{a, b, c\}, \{C_a = b \vee \neg b, C_b = b, C_c = \neg c \vee b\} \rangle$,⁹ the complete interpretations are $M_0 = \{a\}$ (b and c are undefined), $M_1 = \{a, \neg b\}$ (c is undefined) and $M_2 = \{a, b, c\}$. M_0 is the grounded interpretation, while M_1 and M_2 are preferred. Only M_2 is a model. There is no stable model. \square

References

- [1] G. Alfano, A. Cohen, S. Gottifredi, S. Greco, F. Parisi, G.R. Simari, *Credulous acceptance in high-order argumentation frameworks with necessities: an incremental approach*, *Artif. Intell.* 333 (2024) 104159.
- [2] G. Alfano, S. Greco, F. Parisi, *Efficient computation of extensions for dynamic abstract argumentation frameworks: an incremental approach*, in: *Proc. of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*, 2017, pp. 49–55.
- [3] G. Alfano, S. Greco, F. Parisi, G.I. Simari, G.R. Simari, *Incremental computation for structured argumentation over dynamic DeLP knowledge bases*, *Artif. Intell.* 300 (2021) 103553.
- [4] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, *On the semantics of abstract argumentation frameworks: a logic programming approach*, *Theory Pract. Log. Program.* 20 (2020) 703–718.
- [5] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, *Argumentation frameworks with strong and weak constraints: semantics and complexity*, in: *Proc. of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021, pp. 6175–6184.
- [6] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, *On preferences and priority rules in abstract argumentation*, in: *Proc. of International Joint Conference on Artificial Intelligence (IJCAI)*, 2022, pp. 2517–2524.
- [7] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, *Abstract argumentation framework with conditional preferences*, in: *Proc. of AAAI Conference on Artificial Intelligence (AAAI)*, 2023, pp. 6218–6227.
- [8] G. Alfano, S. Greco, F. Parisi, I. Trubitsyna, *On acceptance conditions in abstract argumentation frameworks*, *Inf. Sci.* 625 (2023) 757–779.
- [9] M. Alviano, F. Calimeri, C. Dodaro, D. Fuscà, N. Leone, S. Perri, F. Ricca, P. Veltri, J. Zangari, *The ASP system DLV2*, in: *Proc. of 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, 2017, pp. 215–221.

⁹ Condition C_c in the original paper is $c \rightarrow b$. According to the Kleene 3-valued semantics, it has been rewritten as $\neg c \vee b$.

- [10] L. Amgoud, C. Cayrol, On the acceptability of arguments in preference-based argumentation, in: Proc. of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI), 1998, pp. 1–7.
- [11] L. Amgoud, C. Cayrol, A reasoning model based on the production of acceptable arguments, *Ann. Math. Artif. Intell.* 34 (2002) 197–215.
- [12] L. Amgoud, H. Prade, Using arguments for making and explaining decisions, *Artif. Intell.* 173 (2009) 413–436.
- [13] L. Amgoud, S. Vesic, A new approach for preference-based argumentation frameworks, *Ann. Math. Artif. Intell.* 63 (2011) 149–183.
- [14] L. Amgoud, S. Vesic, Rich preference-based argumentation frameworks, *Int. J. Approx. Reason.* 55 (2014) 585–606.
- [15] M. Arenas, L.E. Bertossi, J. Chomicki, Consistent query answers in inconsistent databases, in: Proceedings of the Eighteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, 1999, pp. 68–79.
- [16] O. Arieli, Towards constraints handling by conflict tolerance in abstract argumentation frameworks, in: Proc. of the Twenty-Sixth International Florida Artificial Intelligence Research Society Conference (FLAIRS), 2013.
- [17] O. Arieli, Conflict-free and conflict-tolerant semantics for constrained argumentation frameworks, *J. Appl. Log.* 13 (2015) 582–604.
- [18] O. Arieli, On the acceptance of loops in argumentation frameworks, *J. Log. Comput.* 26 (2016) 1203–1234.
- [19] K. Atkinson, T.J.M. Bench-Capon, Argumentation schemes in AI and law, *Argument & Computation* 12 (2021) 417–434.
- [20] A. Avron, Natural 3-valued logics—characterization and proof theory, *J. Symb. Log.* 56 (1991) 276–294.
- [21] P. Baroni, M. Caminada, M. Giacomin, An introduction to argumentation semantics, *Knowl. Eng. Rev.* 26 (2011) 365–410.
- [22] P. Baroni, D. Gabbay, M. Giacomin, L. Van der Torre, Handbook of Formal Argumentation. Volume 1, College Public, 2018.
- [23] R. Baumann, What does it take to enforce an argument? Minimal change in abstract argumentation, in: Proc. of the 20th European Conference on Artificial Intelligence (ECAI), 2012, pp. 127–132.
- [24] R. Baumann, M. Heinrich, Bipolar Abstract Dialectical Frameworks Are Covered by Kleene’s Three-Valued Logic, *IJCAI*, 2023.
- [25] T.J.M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, *J. Log. Comput.* 13 (2003) 429–448.
- [26] T.J.M. Bench-Capon, K. Atkinson, A.Z. Wyner, Using argumentation to structure e-participation in policy making, *Transactions on Large-Scale Data and Knowledge-Centered Systems* 18 (2015) 1–29.
- [27] M. Bernreiter, W. Dvorák, S. Woltran, Abstract argumentation with conditional preferences, in: Computational Models of Argument - Proceedings of COMMA, 2022, pp. 92–103.
- [28] A. Bondarenko, P.M. Dung, R.A. Kowalski, F. Toni, An abstract, argumentation-theoretic approach to default reasoning, *Artif. Intell.* 93 (1997) 63–101.
- [29] R. Booth, S. Kaci, T. Rienstra, L.W.N. van der Torre, A logical theory about dynamics in abstract argumentation, in: Proc. of International Conference Scalable Uncertainty Management (SUM), 2013, pp. 148–161.
- [30] M.E.B. Brarda, L.H. Tamargo, A.J. García, Using argumentation to obtain and explain results in a decision support system, *IEEE Intell. Syst.* 36 (2021) 36–42.
- [31] L. Bravo, L.E. Bertossi, Logic programs for consistently querying data integration systems, in: G. Gottlob, T. Walsh (Eds.), Proc. of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI), 2003, pp. 10–15.
- [32] G. Brewka, Preferred subtheories: an extended logic framework for default reasoning, in: *IJCAI*, 1989, pp. 1043–1048.
- [33] G. Brewka, M. Diller, G. Heissenberger, T. Linsbichler, S. Woltran, Solving advanced argumentation problems with answer set programming, *Theory Pract. Log. Program.* 20 (2020) 391–431.
- [34] G. Brewka, I. Niemelä, M. Truszczynski, Answer set optimization, in: Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI), 2003, pp. 867–872.
- [35] G. Brewka, I. Niemelä, M. Truszczynski, Preferences and nonmonotonic reasoning, *AI Mag.* 29 (2008) 69–78.
- [36] G. Brewka, H. Strass, S. Ellmauthaler, J.P. Wallner, S. Woltran, Abstract dialectical frameworks revisited, in: *IJCAI*, 2013, pp. 803–809.
- [37] F. Buccafurri, N. Leone, P. Rullo, Enhancing disjunctive datalog by constraints, *IEEE Trans. Knowl. Data Eng.* 12 (2000) 845–860.
- [38] M. Calautti, S. Greco, C. Molinaro, I. Trubitsyna, Preference-based inconsistency-tolerant query answering under existential rules, *Artif. Intell.* 312 (2022) 103772.
- [39] M. Caminada, Semi-stable semantics, in: Proc. of Computational Models of Argument (COMMA), 2006, pp. 121–130.
- [40] M. Caminada, S. Sá, J.F.L. Alcántara, W. Dvorák, On the equivalence between logic programming semantics and argumentation semantics, *Int. J. Approx. Reason.* 58 (2015) 87–111.
- [41] S. Coste-Marquis, C. Devred, P. Marquis, Constrained argumentation frameworks, in: Proc. of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR), 2006, pp. 112–122.
- [42] S. Coste-Marquis, S. Konieczny, J. Maily, P. Marquis, Extension enforcement in abstract argumentation as an optimization problem, in: Proc. of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI), 2015, pp. 2876–2882.
- [43] S. Coste-Marquis, S. Konieczny, P. Marquis, M.A. Ouali, Weighted attacks in argumentation frameworks, in: Proc. of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning (KR), 2012.
- [44] K. Cyras, F. Toni, ABA+: assumption-based argumentation with preferences, in: Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference (KR), AAAI Press, 2016, pp. 553–556.
- [45] Y. Dimopoulos, A. Torres, Graph theoretical structures in logic programs and default theories, *Theor. Comput. Sci.* 170 (1996) 209–244.
- [46] S. Doutre, J. Maily, Constraints and changes: a survey of abstract argumentation dynamics, *Argument & Computation* 9 (2018) 223–248.
- [47] P.M. Dung, Negations as hypotheses: an abductive foundation for logic programming, in: Logic Programming, in: Proc. of the Eighth International Conference (ICLP), 1991, pp. 3–17.
- [48] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *AI* 77 (1995) 321–358.
- [49] P.E. Dunne, T.J.M. Bench-Capon, Coherence in finite argument systems, *Artif. Intell.* 141 (2002) 187–203.
- [50] P.E. Dunne, T.J.M. Bench-Capon, Complexity in value-based argument systems, in: Proc. of 9th European Conference on Logics in Artificial Intelligence (JELIA), 2004, pp. 360–371.
- [51] P.E. Dunne, M. Caminada, Computational complexity of semi-stable semantics in abstract argumentation frameworks, in: Proc. of 11th European Conference on Logics in Artificial Intelligence (JELIA), 2008, pp. 153–165.
- [52] P.E. Dunne, A. Hunter, P. McBurney, S. Parsons, M.J. Wooldridge, Weighted argument systems: basic definitions, algorithms, and complexity results, *Artif. Intell.* 175 (2011) 457–486.
- [53] W. Dvorák, P.E. Dunne, Computational problems in formal argumentation and their complexity, *FLAP* 4 (2017).
- [54] W. Dvorák, A. Keshavarzi Zafarhandi, S. Woltran, Expressiveness of setafs and support-free adfs under 3-valued semantics, *J. Appl. Non-Class. Log.* 33 (2023) 298–327.
- [55] W. Dvorák, M. König, M. Ulbricht, S. Woltran, Principles and their computational consequences for argumentation frameworks with collective attacks, *J. Artif. Intell. Res.* 79 (2024) 69–136.
- [56] W. Dvorák, S. Woltran, Complexity of semi-stable and stage semantics in argumentation frameworks, *Inf. Process. Lett.* 110 (2010) 425–430.
- [57] T. Eiter, M. Fink, G. Greco, D. Lembo, Efficient evaluation of logic programs for querying data integration systems, in: C. Palamidessi (Ed.), Proc. of 19th International Conference on Logic Programming (ICLP), 2003, pp. 163–177.
- [58] T. Eiter, G. Gottlob, On the computational cost of disjunctive logic programming: propositional case, *Ann. Math. Artif. Intell.* 15 (1995) 289–323.
- [59] T. Eiter, N. Leone, D. Sacca, Expressive power and complexity of partial models for disjunctive deductive databases, *Theor. Comput. Sci.* 206 (1998) 181–218.
- [60] W. Faber, M. Vallati, F. Cerutti, M. Giacomin, Solving set optimization problems by cardinality optimization with an application to argumentation, in: Proc. of 22nd European Conference on Artificial Intelligence (ECAI), 2016, pp. 966–973.

- [61] S. Flesca, F. Furfaro, F. Parisi, Preferred database repairs under aggregate constraints, in: Proc. of International Conference on Scalable Uncertainty Management (SUM), 2007, pp. 215–229.
- [62] G. Flouris, A. Bikakis, A comprehensive study of argumentation frameworks with sets of attacking arguments, *Int. J. Approx. Reason.* 109 (2019) 55–86.
- [63] D. Gabbay, M. Giacomini, G.R. Simari, M. Thimm (Eds.), *Handbook of Formal Argumentation*, vol. 2, College Public, 2021.
- [64] A.J. Garcia, H. Prakken, G.R. Simari, A comparative study of some central notions of ASPIC+ and delp, *Theory Pract. Log. Program.* 20 (2020) 358–390.
- [65] G. Greco, S. Greco, E. Zupanò, A logical framework for querying and repairing inconsistent databases, *IEEE Trans. Knowl. Data Eng.* 15 (2003) 1389–1408.
- [66] S. Greco, Non-determinism and weak constraints in datalog, *New Gener. Comput.* 16 (1998) 373–396.
- [67] J. Heyninck, G. Kern-Isberner, T. Rienstra, K. Skiba, M. Thimm, Possibilistic logic underlies abstract dialectical frameworks, in: *IJCAI*, 2022, pp. 2655–2661.
- [68] T. Janhunen, I. Niemelä, D. Seipel, P. Simons, J.H. You, Unfolding partiality and disjunctions in stable model semantics, *ACM Trans. Comput. Log.* 7 (2006).
- [69] S. Kaci, L. van Der Torre, S. Vesic, S. Villata, Preference in abstract argumentation, in: *Handbook of Formal Argumentation*, Volume 2, 2021, pp. 199–236, Chapter 3.
- [70] S. Kaci, L.W.N. van der Torre, Preference-based argumentation: arguments supporting multiple values, *Int. J. Approx. Reason.* 48 (2008) 730–751.
- [71] S. Kaci, L.W.N. van der Torre, S. Villata, Preference in abstract argumentation, in: Proc. of International Conference on Computational Models of Argument (COMMA), 2018, pp. 405–412.
- [72] N. Kökciyan, I. Sassoon, E. Sklar, S. Modgil, S. Parsons, Applying metalevel argumentation frameworks to support medical decision making, *IEEE Intell. Syst.* 36 (2021) 64–71.
- [73] N. Kökciyan, N. Yaglikci, P. Yolum, An argumentation approach for resolving privacy disputes in online social networks, *ACM Trans. Internet Technol.* 17 (2017) 27:1–27:22.
- [74] M.W. Krentel, The complexity of optimization problems, in: *STOC*, 1986, pp. 69–76.
- [75] T. Linsbichler, J. Pührer, H. Strass, A uniform account of realizability in abstract argumentation, in: Proc. of ECAI, 2016, pp. 252–260.
- [76] T. Lukasiewicz, E. Malizia, C. Molinaro, Complexity of inconsistency-tolerant query answering in datalog+/- under preferred repairs, in: *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 2023, pp. 472–481.
- [77] M.V. Martinez, F. Parisi, A. Pugliese, G.I. Simari, V.S. Subrahmanian, Policy-based inconsistency management in relational databases, *Int. J. Approx. Reason.* 55 (2014) 501–528.
- [78] S. Modgil, Reasoning about preferences in argumentation frameworks, *Artif. Intell.* 173 (2009) 901–934.
- [79] S. Modgil, H. Prakken, A general account of argumentation with preferences, *Artif. Intell.* 195 (2013) 361–397.
- [80] S.H. Nielsen, S. Parsons, A generalization of Dung's abstract framework for argumentation: arguing with sets of attacking arguments, in: Proc. of Third International Workshop on Argumentation in Multi-Agent Systems (ArgMAS), 2006, pp. 54–73.
- [81] A. Niskanen, J.P. Wallner, M. Järvisalo, Extension enforcement under grounded semantics in abstract argumentation, in: Proc. of the Sixteenth International Conference on Principles of Knowledge Representation and Reasoning (KR), 2018, pp. 178–183.
- [82] C.H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.
- [83] F. Parisi, J. Grant, On measuring inconsistency in definite and indefinite databases with denial constraints, *Artif. Intell.* 318 (2023) 103884.
- [84] A. Paziienza, D. Grossi, F. Grasso, R. Palmieri, M. Zito, S. Ferilli, An abstract argumentation approach for the prediction of analysts' recommendations following earnings conference calls, *Intell. Artif.* 13 (2019) 173–188.
- [85] S. Polberg, Understanding the abstract dialectical framework, in: Proc. of JELIA, 2016, pp. 430–446.
- [86] A. Ramos, Two new weak constraint qualifications for mathematical programs with equilibrium constraints and applications, *J. Optim. Theory Appl.* 183 (2019) 566–591.
- [87] D. Saccà, The expressive powers of stable models for bound and unbound DATALOG queries, *J. Comput. Syst. Sci.* 54 (1997) 441–464.
- [88] D. Saccà, C. Zaniolo, Stable models and non-determinism in logic programs with negation, in: Proc. of PODS, 1990, pp. 205–217.
- [89] C. Sakama, T.C. Son, Epistemic argumentation framework: theory and computation, *J. Artif. Intell. Res.* 69 (2020) 1103–1126.
- [90] R. Silva, S. Sá, J.F.L. Alcântara, Semantics hierarchy in preference-based argumentation frameworks, in: *COMMA*, 2020, pp. 339–346.
- [91] M. Snaith, R.Ø. Nielsen, S.R. Kotnis, A. Pease, Ethical challenges in argumentation and dialogue in a healthcare context, *Argument & Computation* 12 (2021) 249–264.
- [92] H. Strass, J.P. Wallner, Analyzing the computational complexity of abstract dialectical frameworks via approximation fixpoint theory, *Artif. Intell.* 226 (2015) 34–74.
- [93] J.P. Wallner, A. Niskanen, M. Järvisalo, Complexity results and algorithms for extension enforcement in abstract argumentation, *J. Artif. Intell. Res.* 60 (2017) 1–40.