

Facets in Argumentation: A Formal Approach to Argument Significance

Johannes K. Fichte¹, Nicolas Fröhlich², Markus Hecher³, Victor Lagerkvist¹, Yasir Mahmood⁴, Arne Meier², Jonathan Persson¹

¹Linköping University, Sweden

²Leibniz Universität Hannover, Germany

³Univ. Artois, CNRS, France

⁴Paderborn University, Germany

{johannes.klaus.fichte, victor.lagerkvist}@liu.se, {nicolas.froehlich, meier}@thi.uni-hannover.de, hecher@mit.edu, yasir.mahmood@uni-paderborn.de, jonpe481@student.liu.se

Abstract

Argumentation is a central subarea of Artificial Intelligence (AI) for modeling and reasoning about arguments. The semantics of abstract argumentation frameworks (AFs) is given by sets of arguments (extensions) and conditions on the relationship between them, such as stable or admissible. Today’s solvers implement tasks such as finding extensions, deciding credulous or skeptical acceptance, counting, or enumerating extensions. While these tasks are well charted, the area between decision, counting/enumeration and fine-grained reasoning requires expensive reasoning so far. We introduce a novel concept (facets) for reasoning between decision and enumeration. Facets are arguments that belong to some extensions (credulous) but not to all extensions (skeptical). They are most natural when a user aims to navigate, filter, or comprehend the significance of specific arguments, according to their needs. We study the complexity and show that tasks involving facets are much easier than counting extensions. Finally, we provide an implementation, and conduct experiments to demonstrate feasibility.

1 Introduction

Abstract argumentation [Dung, 1995; Bench-Capon and Dunne, 2007] is a formalism for modeling and evaluating arguments and its reasoning problems has many applications in artificial intelligence (AI) [Amgoud and Prade, 2009; Rago *et al.*, 2018]. The semantics is based on sets of arguments that satisfy certain conditions regarding the relationship among them, such as being stable or admissible [Dung, 1995]. Such sets of arguments are then called *extensions* of a framework and various practical solvers for decision and reasoning tasks [Egly *et al.*, 2008; Niskanen and Jarvisalo, 2020; Thimm *et al.*, 2021; Alviano, 2021] compete biannually in the ICCMA competition [Thimm *et al.*, 2024].

Qualitative reasoning problems such as finding an extension or deciding credulous or skeptical acceptance are reasonably fast to compute [Dvořák, 2012] but have limitations. Namely, these two reasoning modes represent extremes on the reasoning spectrum, as they provide no insight into preferences




Figure 1: An example argumentation framework.

among arguments for further analysis. As a result, enumeration, counting, and fine-grained quantitative reasoning modes have been studied and computationally classified [Fichte *et al.*, 2023a; Fichte *et al.*, 2024] enabling probabilistic reasoning over arguments. While enumeration is well suited when the total number of extensions is small, some argumentation semantics easily result in a vast number of extensions. However, users might still want to investigate the space of possible extensions in more detail. Possible examples are restricting or diversifying extensions, identifying resilient arguments or sanity checks, evaluating outcomes in argumentation frameworks generated by LLMs, or gaining insights into specific frameworks through explanations.

In all such scenarios evaluating the *significance* for individual arguments in a framework is central. In existing proposals, computing the significance for arguments relied on quantitative measures over extensions containing certain arguments or supporting particular claims. These notions rely on counting all extensions containing a particular argument (or claim), which is computationally expensive [Fichte *et al.*, 2024]. Example 1 illustrates difficulties when comparing significance of certain arguments in the overall world of extensions.

Example 1. Consider the argumentation framework F , depicted in Figure 1, illustrating the choice between a sweet or savory breakfast, that is between maple syrup and burrito. Intuitively, if one prefers a savory flavor, they would not choose maple syrup; likewise, sweet attacks burrito. One does not go to taco bell expecting maple syrup, nor are small portions typical at taco bell. While it is possible to make burritos at home, doing so requires buying expensive ingredients. Making a small portion or going to taco bell avoids this.

The stable extensions of F are: $\{w, m, p\}$, $\{s, b, p\}$ and

$\{s, b, t\}$. Now it is not immediate to compare the significance of accepting/rejecting certain arguments to each other.

In this paper, we propose a combination of credulously and skeptically accepted arguments, which ask whether a given argument belongs to some extension (credulous) but not all extensions (skeptical). We call arguments that are credulously but not skeptically accepted *facets*. Facets quantify the uncertainty of arguments in extensions, providing a measure of their indeterminacy within the framework. They can be utilized to evaluate the significance of specific arguments. Example 2 provides a brief intuition.

Example 2. We return to the argumentation framework F from Example 1. Six of the seven arguments are facets under stable semantics, with only e being a non-facet, indicating a substantial degree of uncertainty.

Assume we aim to compare the relative significance of an argument. Consider the extensions of F rejecting (not containing) the argument “sweet”. There are two such stable extensions $\{s, b, p\}$ and $\{s, b, t\}$. Therefore, rejecting the argument w leaves us with two facets p and t . In contrast, consider the extensions accepting (containing) the argument “sweet”. This results in one stable extension $\{w, m, p\}$ and hence no facets. Accepting the argument w eliminates any uncertainty, whereas rejecting w does not. Consequently, we consider accepting “sweet” to be more significant than rejecting “sweet”.

While the computational complexity of credulous and skeptical reasoning is well studied [Dvořák, 2012], we ask for the concrete complexity of facets and whether counting facets provides a theoretical benefit over projected counting and projected enumerating extensions.

Contributions. In more details, we establish the following.

1. We introduce facets to abstract argumentation as a reasoning tool for significance and filtering extensions in a directed way. By this, we fill a gap in the literature between quantitative and qualitative reasoning.
2. We present a comprehensive complexity analysis for various qualitative and quantitative problems involving facets. Table 1 provides an overview on our results.
3. Finally, we present experiments that demonstrate the feasibility of our framework. We evaluate our implementation on instances of the ICCMA competition.

Related Work. The computational complexity in abstract argumentation is well understood for decision problems [Dunne and Bench-Capon, 2002; Dvořák and Woltran, 2010; Dvořák, 2012], parameterized complexity involving treewidth [Fichte et al., 2021a], for (projected) counting [Fichte et al., 2024] as well as for fine-grained reasoning based on counting [Fichte et al., 2023a]. The decision complexity ranges from \mathbf{P} to $\Sigma_2^{\mathbf{P}}$ for credulous reasoning, and from \mathbf{P} to $\Pi_2^{\mathbf{P}}$ for skeptical reasoning. The complexity of counting extensions ranges between $\#\mathbf{P}$ and $\#\mathbf{coNP}$ depending on the semantics. In theory, we know that $\text{PH} \subseteq \mathbf{P}^{\#\mathbf{P}}$ [Toda, 1991] where $\bigcup_{k \in \mathbb{N}} \Delta_k^{\mathbf{P}} = \text{PH}$ and $\mathbf{NP} \subseteq \Delta_2^{\mathbf{P}} = \mathbf{P}^{\mathbf{NP}}$ [Stockmeyer, 1976]. This renders counting extensions theoretically significantly harder. Approximate counting is in fact easier, i.e., $\text{approx-}\#\mathbf{P} \subseteq \text{BPP}^{\mathbf{NP}} \subseteq \Sigma_3^{\mathbf{P}}$ [Lautemann, 1983;

Problems/ σ	σ_1	σ_2	σ_3
ISFACET $_{\sigma}$	$\mathbf{P}^{\mathbf{R3}}$	$\mathbf{NP}^{\mathbf{T5}}$	$\Sigma_2^{\mathbf{P}^{\mathbf{T5/7}}}$
FACETS $_{\sigma}^{\geq k}$	$\mathbf{P}^{\mathbf{T8}}$	$\mathbf{NP}^{\mathbf{T11}}$	$\Sigma_2^{\mathbf{P}^{\mathbf{T12}}}$
FACETS $_{\sigma}^{\leq k}$	$\mathbf{P}^{\mathbf{T8}}$	$\mathbf{coNP}^{\mathbf{T15}}$	$\Pi_2^{\mathbf{P}^{\mathbf{T15}}}$
FACETS $_{\sigma}^{=k}$	$\mathbf{P}^{\mathbf{T16}}$	$\mathbf{DP}^{\mathbf{T17}}$	$\in \mathbf{DP}_2^{\mathbf{T18}}$
$\mathbb{S}_{\sigma}[F, \ell]$	$\text{“}\in \mathbf{P}^{\text{”}}^{\mathbf{T16}}$	$\text{“}\in \Delta_2^{\mathbf{P}^{\text{”}}^{\mathbf{T21}}}$	$\text{“}\in \Delta_3^{\mathbf{P}^{\text{”}}^{\mathbf{T21}}}$

Table 1: Overview of our complexity results for the semantics $\sigma_1 \in \{\text{conf, naiv}\}$, $\sigma_2 \in \{\text{adm, stab, comp}\}$, and $\sigma_3 \in \{\text{pref, semiSt, stag}\}$. All results depict completeness except for \mathbf{P} -cases or when stated o/w. ISFACET $_{\sigma}$ asks whether a given argument is a facet; FACETS $_{\sigma}^{\geq k}$ asks whether there are at least ($\geq k$), at most ($\leq k$), or exactly ($= k$) facets. $\mathbb{S}_{\sigma}[F, \ell]$ asks what the significance of approving ($\ell = a$) or disapproving ($\ell = \bar{a}$) of a σ -facet a in AF F is. Superscripts behind the complexity classes refer to Remark (R) and Theorem (T), with the proof. “ $\in \Delta_i^{\mathbf{P}}$ ” slightly abuses notation meaning that it can be computed by a deterministic polynomial-time Turing machine with access to a $\Sigma_{i-1}^{\mathbf{P}}$ oracle.

Sipser, 1983; Stockmeyer, 1983], but turns out to be still harder than counting facets. Facets were initially proposed for answer-set programming (ASP) by Alrabbaa et al. [2018] as a tool to navigate large solution spaces. Their computational complexity has been systematically classified [Rusovac et al., 2024]. ASP is a popular problem solving paradigm to model and solve hard combinatorial problems in form of a logic program that expresses constraints [Gebser et al., 2012]. Plausibility reasoning has been developed for ASP based on full counting [Fichte et al., 2022a]. Dachsel et al. [2022] developed a tool to navigate argumentation frameworks using ASP-facets. Note that ASP-navigation is based on forbidding or enforcing atoms in programs via integrity constraints. In contrast, argumentation facets enable approving or disapproving arguments, while not necessarily removing the extensions entirely leading to a natural notion of significance of an argument (see Section 4). Finally, our complexity analysis (Table 1) indicates a computational gain for reasoning with facets compared to separately asking credulous/skeptical reasoning in each case (see e.g., [Dvořák and Dunne, 2017]). Facets have recently also been applied to planning [Speck et al., 2025].

Proofs of statements marked with (\star) can be found in [Fichte et al., 2025].

2 Preliminaries

We assume familiarity with computational complexity [Pipenger, 1997], graph theory [Bondy and Murty, 2008], and Boolean logic [Biere et al., 2021].

Complexity Classes. We use standard notation for basic complexity classes and for example write \mathbf{P} (\mathbf{NP}) for the class of decision problems solvable in (non-deterministic) polynomial time. Additionally, we let \mathbf{coNP} be the class of decision problems whose complement is in \mathbf{NP} , and let \mathbf{DP} be the class of decision problems representable as the intersection of a problem in \mathbf{NP} and a problem in \mathbf{coNP} . On top, we use more classes from the polynomial hierarchy [Stockmeyer and Meyer, 1973; Stockmeyer, 1976; Wrathall, 1976], $\Delta_0^{\mathbf{P}} := \Pi_0^{\mathbf{P}} := \Sigma_0^{\mathbf{P}} := \mathbf{P}$ and $\Delta_i^{\mathbf{P}} := \mathbf{P}^{\Sigma_{i-1}^{\mathbf{P}}}$,

$\Sigma_i^P := \mathbf{NP}^{\Sigma_{i-1}^P}$, and $\Pi_i^P := \mathbf{coNP}^{\Sigma_{i-1}^P}$ for $i > 0$ where C^D is the class C of decision problems augmented by an oracle for some complete problem in class D . Recall that $\mathbf{PH} := \bigcup_{i \in \mathbb{N}} \Delta_i^P$ [Stockmeyer, 1976]. The canonical NP-complete problem is the Boolean *satisfiability* problem for formulas in *conjunctive normal form* (CNF), i.e., given $\varphi := \bigwedge_{i=1}^m C_i$ where each C_i is a clause, decide whether φ admits at least one satisfying assignment. For \mathbf{coNP} the corresponding problem is simply to check *unsatisfiability*, and for \mathbf{DP} to check whether φ is satisfiable and ψ unsatisfiable for a given pair of formulas (φ, ψ) (the SAT-UNSAT problem). The complexity class \mathbf{DP}_k is defined as $\mathbf{DP}_k := \{L_1 \cap L_2 \mid L_1 \in \Sigma_k^P, L_2 \in \Pi_k^P\}$, $\mathbf{DP} = \mathbf{DP}_1$ [Lohrey and Rosowski, 2023]. For Π_2^P we may e.g. consider the evaluation problem for a *quantified Boolean formula* of the form $\forall X \exists Y. \varphi$ where X and Y are two disjoint sets of variables and φ a formula in CNF over X and Y . For Σ_2^P the problem is instead to check that $\forall X \exists Y. \varphi$ is false.

Abstract Argumentation. We use Dung’s argumentation framework (1995) and consider only non-empty and finite sets of arguments A . An (*argumentation*) *framework* (AF) is a directed graph $F = (A, R)$, where A is a set of arguments and $R \subseteq A \times A$, consisting of pairs of arguments representing direct attacks between them. An argument $a \in E$, is called *defended* by E in F if for every $(a', a) \in R$, there exists $a'' \in E$ such that $(a'', a') \in R$. The family $\text{def}_F(E)$ is defined by $\text{def}_F(E) := \{a \mid a \in A, a \text{ is defended by } E \text{ in } F\}$. In abstract argumentation, one strives for computing so-called *extensions*, which are subsets $E \subseteq A$ of the arguments that have certain properties. The set E of arguments is called *conflict-free* in E if $(E \times E) \cap R = \emptyset$; E is *admissible* in F if (1) E is *conflict-free* in F , and (2) every $a \in E$ is *defended* by E in F . Let $E_R^+ := E \cup \{a \mid (b, a) \in R, b \in E\}$ and E be conflict-free. Then, E is (1) *naive* in F if no $E' \supset E$ exists that is conflict-free in F , and (2) *stage* in F if there is no conflict-free set $E' \subseteq A$ in F with $E_R^+ \subsetneq (E')_R^+$. An admissible set E is (1) *complete* in F if $\text{def}_F(E) = E$; (2) *preferred* in F , if no $E' \supset E$ exists that is *admissible* in F ; (3) *semi-stable* in F if no admissible set $E' \subseteq A$ in F with $E_R^+ \subsetneq (E')_R^+$ exists; and (4) *stable* in F if every $a \in A \setminus E$ is *attacked* by some $a' \in E$. For a semantics $\sigma \in \{\text{conf}, \text{naiv}, \text{adm}, \text{comp}, \text{stab}, \text{pref}, \text{semiSt}, \text{stag}\}$, we write $\sigma(F)$ for the set of all *extensions* of semantics σ in F . Let $F = (A, R)$ be an AF. Then, the problem EXIST_σ asks if $\sigma(F) \neq \emptyset$. The problems c_σ and s_σ question for $a \in A$, whether a is in some $E \in \sigma(F)$ (“*credulously accepted*”) or every $E \in \sigma(F)$ (“*skeptically accepted*”), respectively. We let C_σ (resp., S_σ) denote the set of all credulously (skeptically) accepted arguments under semantics σ .

3 Facet Reasoning

Central reasoning problems in argumentation include deciding whether an argument is credulously (or skeptically) accepted. In the following, we will see how the problem of deciding facets can be seen as a generalization of these two modes.

We begin by defining reasoning problems pertaining to *facets* in argumentation. Intuitively, a σ -facet is an argument which is accepted in some, but not all σ -extensions for the considered semantics σ . Formally, given a semantics σ , then

an argument a is a σ -*facet* if $a \in C_\sigma \setminus S_\sigma$. Given an AF F and semantics σ , then $\mathcal{F}_\sigma(F)$ denotes the set of all σ -facets in F . In this work, we consider the following reasoning problems parameterized by a semantics σ .

- The problem ISFACET_σ asks, given an AF $F = (A, R)$ and an argument $a \in A$, is a a σ -facet in F ?
- The problems $\text{FACETS}_\sigma^{\geq k}$, $\text{FACETS}_\sigma^{\leq k}$ and $\text{FACETS}_\sigma^{=k}$ has an integer k as an additional input, and ask whether an input $F = (A, R)$ has exactly, at least, or at most k σ -facets, respectively.

We continue by analyzing the complexity of these problems, beginning with ISFACET_σ in Section 3.1, $\text{FACETS}_\sigma^{\geq k}$ and $\text{FACETS}_\sigma^{\leq k}$ in Section 3.2, and complete the study with $\text{FACETS}_\sigma^{=k}$ in Section 3.3.

3.1 Complexity of Deciding Facets

Regarding the complexity classification, for conflict-free and naive semantics, the problem ISFACET is rather straightforward to classify. To see this, for conf , each argument not attacking itself is a facet (provided there are at least two such arguments in the AF), and for naiv , one additionally has to remove each argument not in conflict with any other argument since it can not be a facet.

Remark 3. ISFACET_σ is in \mathbf{P} for $\sigma \in \{\text{conf}, \text{naiv}\}$.

For the remaining semantics, we get hardness by observing that ISFACET_σ is as hard as the credulous reasoning (c_σ) for each considered semantics σ .

Lemma 4. Let σ be any semantics. Then $c_\sigma \leq_m^P \text{ISFACET}_\sigma$.

Proof. We provide a polynomial time many-one reduction from c_σ to ISFACET_σ for each semantics σ as follows. Let $F = (A, R)$ be an AF, and $a \in A$ be an argument. Our reduction yields an AF F' where we duplicate the argument a which has all the incoming and outgoing attacks similar to $a \in A$. Precisely, $F' = (A', R')$ is as follows:

- $A' := A \cup \{a'\}$ for a fresh $a' \notin A$,
- $R' := R \cup \{(a, a'), (a', a)\} \cup \{(a', x) \mid (a, x) \in R\} \cup \{(x, a') \mid (x, a) \in R\}$.

Then, for any semantics σ , the argument a is credulously accepted under σ in F iff a is a facet under σ in F' . Indeed, let a be credulously accepted, then there is a σ -extension $E \subseteq A$ such that $a \in E$. Since a defends itself against a' in F' , we have that a is also credulously accepted in F' as the other attacks remain the same. Finally, a can not be in all σ -extensions E in F' , as $E \setminus \{a\} \cup \{a'\}$ is also a σ -extension. Therefore the claim follows. \square

We specifically obtain the following characterization.

Theorem 5. ISFACET_σ is \mathbf{NP} -complete for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$ and Σ_2^P -complete for $\sigma \in \{\text{semiSt}, \text{stag}\}$.

Proof. The hardness in each case follows due to Lemma 4 and the complexity for credulous reasoning under corresponding semantics.

The membership follows, since one can guess two σ -extensions for an input F , one containing the argument

in the question, and another without it. The verification (of σ -extensions) requires (1) polynomial time for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$, and (2) coNP-oracle for $\sigma \in \{\text{semiSt}, \text{stag}\}$. This establishes the membership results. \square

The following observation is necessary to establish the hardness proof in Theorem 7. Let $\Phi = \forall X \exists Y. \varphi$ be a QBF instance, where $\varphi := \bigwedge_{i=1}^n C_i$ is a CNF. If φ is not satisfiable, the formula Φ can not be true. Whereas, the problem to check whether φ is satisfiable, is NP-complete. Given Φ , one can construct a new formula $\Phi' = \forall X' \exists Y. \varphi'$ such that φ' is satisfiable and Φ is true iff Φ' is true. To this aim, we let $X' = X \cup \{z\}$ for a fresh variable $z \notin X \cup Y$. Then $\varphi' := \bigwedge_{C \in \varphi} (\neg z \vee C)$. Notice that $\Phi[z \mapsto 0]$ is trivially true whereas $\Phi[z \mapsto 1]$ is true iff Φ is true. Consequently, we have the following observation.

Remark 6. Given a QBF instance $\Phi = \forall X \exists Y. \varphi$, where φ is a CNF. One can assume w.l.o.g. that φ is satisfiable.

Next, we establish that $\text{ISFACET}_{\text{pref}}$ is Σ_2^{P} -complete. Notice that this case is not covered by Theorem 5 as the credulous reasoning for preferred semantics (c_{pref}) is only NP-complete.

Theorem 7. $\text{ISFACET}_{\text{pref}}$ is Σ_2^{P} -complete.

Proof. The membership follows since one can guess two preferred extensions for an input F , one containing the argument in the question, and another without it. The verification of pref-extensions requires an NP-oracle. This establishes the mentioned membership results.

For hardness, we utilize the following reduction proving Π_2^{P} -hardness of skeptical acceptance with preferred semantics [Dvorák and Dunne, 2017, Reduction 3.7]. Given a QBF $\Phi = \forall Y \exists Z. \varphi$ where $\varphi := \bigwedge_{i=1}^m C_i$ is a CNF-formula with clauses C_i over variables $X = Y \cup Z$. We construct an AF $F_\Phi = (A, R)$, where $A = \{\varphi, \bar{\varphi}\} \cup \{C_1, \dots, C_m\} \cup X \cup \bar{X}$. The relation R includes the following attacks:

$$\begin{aligned} & \{(C_i, \varphi) \mid 1 \leq i \leq m\} \cup \\ & \{(x, C_i) \mid x \in C_i\} \cup \{(\bar{x}, C_i) \mid \bar{x} \in C_i\} \cup \\ & \{(x, \bar{x}), (\bar{x}, x) \mid x \in \text{var}(\varphi)\} \cup \\ & \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi)\} \cup \{(\bar{\varphi}, z), (\bar{\varphi}, \bar{z}) \mid z \in Z\}. \end{aligned}$$

Then, it holds that there is a preferred extension in F_Φ not containing the argument φ iff the formula Φ is false. Furthermore, we have that φ is satisfiable. Therefore, there exists a preferred extension S containing φ . Namely, S corresponds to a satisfying assignment θ for φ as $S = \{\varphi, \ell \mid \ell \in X \cup \bar{X}, \theta(\ell) = 1\}$. As a result, φ is a pref-facet iff there is a preferred extension S' with $\varphi \notin S'$ iff the formula Φ is false. This results in Σ_2^{P} -hardness. \square

3.2 Atleast/Atmost k Facets Complexity

We begin by proving that for conflict-free and naive semantics, one can count all the facets in polynomial time.

Theorem 8. $\text{FACETS}_{\sigma}^{\geq k}, \text{FACETS}_{\sigma}^{\leq k} \in \mathbf{P}$, $\sigma \in \{\text{conf}, \text{naiv}\}$.

Proof. For conf, the AF F has at least k facets if F contains at least k non self-conflicting arguments (without self-attacks). For naiv, one additionally has to remove arguments (N) not in conflict with any other argument since those can not be a

facet. Thus, F has at least k naiv-facets if $F \setminus N$ contains at least k non self-conflicting arguments, where N includes those arguments not participating in any attack. \square

The following reduction is essential for proving Lemma 10 which we utilize later in achieving certain lower bounds.

Definition 9 ([Dvorák and Dunne, 2017]). Let $\varphi := \bigwedge_{i=1}^m C_i$, be a CNF-formula where each C_i is a clause. Consider the AF $F_\varphi = (A, R)$ constructed as follows:

$$\begin{aligned} A &:= \{\varphi, C_1, \dots, C_m\} \cup \{x, \bar{x} \mid x \in \text{var}(\varphi)\} \\ R &:= \{(C_i, \varphi) \mid i \leq m\} \cup \{(x, \bar{x}), (\bar{x}, x) \mid x \in \text{var}(\varphi)\} \\ &\quad \cup \{(x, C_i) \mid x \in C_i\} \cup \{(\bar{x}, C_i) \mid \bar{x} \in C_i\}. \end{aligned}$$

We call F_φ the argumentation framework of φ generated via the standard translation.

It is known that φ is satisfiable iff the argument φ is credulously accepted in F_φ under semantics $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$. We next prove the following intermediate lemma. Essentially, the standard translation allows us to characterize exactly the number of facets in F_φ based on whether the formula φ is satisfiable or not.

Lemma 10 (*). Let φ be a CNF-formula involving m clauses and n variables. Moreover, let F_φ be the AF of φ as depicted in Definition 9 and let $k = 2n + m + 1$. Then the following statements are true for every $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$.

1. φ is satisfiable iff F_φ admits exactly k σ -facets.
2. φ is not satisfiable iff F_φ admits exactly $k - 1$ σ -facets.

We continue with $\text{FACETS}_{\sigma}^{\geq k}$ and prove that it is either NP-complete or Σ_2^{P} -complete depending on the choice of σ .

Theorem 11. $\text{FACETS}_{\sigma}^{\geq k}$ is NP-c., $\sigma \in \{\text{adm}, \text{stab}, \text{comp}\}$.

Proof. For membership, we guess k distinct arguments $\{a_1, \dots, a_k\}$, and simultaneously $2k$ σ -extensions $P_1, \dots, P_k, N_1, \dots, N_k$ such that: $a_i \in P_i$ and $a_i \notin N_i$. The verification that each $S \in \{P_i, N_i \mid i \leq k\}$ is a σ -extension can be done in polynomial time. Then, F has at least k σ -facets iff each argument in $\{a_i \mid i \leq k\}$ is a σ -facet.

For hardness, we utilize Lemma 10. Indeed, φ is satisfiable iff the AF F_φ has at least k facets where $k = 2n + m + 1$ for the formula φ with n variables and m clauses. \square

Theorem 12. $\text{FACETS}_{\sigma}^{\geq k}$ is Σ_2^{P} -c., $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.

Proof. For membership, we guess k distinct arguments $\{a_1, \dots, a_k\}$, and simultaneously $2k$ σ -extensions $P_1, \dots, P_k, N_1, \dots, N_k$ such that: $a_i \in P_i$ and $a_i \notin N_i$. The verification that each $S \in \{P_i, N_i \mid i \leq k\}$ is a σ -extension can be done via an NP-oracle. Then, F has at least k σ -facets iff each argument in $\{a_i \mid i \leq k\}$ is a σ -facet. This yields membership in NP^{NP} (equivalently, Σ_2^{P}).

For hardness, we reduce from ISFACET_{σ} for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$. To this aim, let $F = (A, R)$ be an AF and $a \in A$ be an argument in the question. Assume that $|A| = n$. We let $n - 1$ additional copies of a and consider the set $C_a = \{a_1, \dots, a_n\}$ of arguments where $a_1 = a$

and $a_i \notin A$ are fresh arguments for $i \geq 2$. Then, we construct the AF $F' = (A', R')$ where $A' = A \cup C_a$. The relation R' consists of R and additionally the following attacks: $\{(a_i, x) \mid (a, x) \in R, i \leq n\} \cup \{(x, a_i) \mid (x, a) \in R, i \leq n\}$. That is, F' simply copies the argument a together with all its incoming and outgoing attacks for each of the $n - 1$ fresh arguments.

We first prove that for a conflict-free (resp., admissible) set S in F , adding arguments from C_a to S does not change its conflict-freeness (admissibility) in F' as long as S contains a .

Claim 13. *A set $S \subseteq A$ containing a is conflict-free (resp., admissible) in F iff $S \cup C_a$ is conflict-free (admissible) in F' .*

Proof of Claim 13. We prove the case for conflict-freeness, the case for admissible semantics follows analogously. If S is not conflict-free in F then S is also not conflict-free in F' since $R \subseteq R'$.

Conversely, suppose S is conflict-free in F . Suppose to the contrary, there exists $x, y \in S \cup C_a$ such that $(x, y) \in R'$. Recall that C_a is conflict-free in F' by definition and S is conflict-free in $F' \setminus C_a$. Then, it must be the case that $x \in S$ and $y \in C_a$ (or vice versa). But this leads to a contradiction to the conflict-freeness of S since $a \in S$ and $(x, a_i) \in R'$ iff $(x, a) \in R$. Analogous case holds if $x \in C_a$ and $y \in S$. Thus $S \cup C_a$ is conflict-free in F' . \square

Claim 14. *The argument a is a σ -facet in F iff each argument $a_i \in C_a$ is a σ -facet in F' for each $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.*

Proof of Claim 14. Claim Proof. Suppose a is a σ -facet in F . Then, there are σ -extensions S_1, S_2 in F such that $a \in S_1$ and $a \notin S_2$. Then, we prove that each argument in C_a belongs to some, but not all σ -extensions of F' . Notice first that S_1 is not a σ -extension in F' for any $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$. This holds due to the arguments in C_a . Indeed, if S_1 is a σ -extension in F , then $S_1 \cup \{x \mid x \in C_a\}$ is a counter-example to S_1 being σ -set in F' due to Claim 13.

SomeE: $S_1 \cup C_a$ is a σ -extension containing each $a_i \in C_a$ (again, due to Claim 13).

NotAllE: We prove that S_2 is a σ -extension in F' and $a_i \notin S_2$ for each $a_i \in C_a$. We prove the claim for preferred semantics, other cases can be proven analogously. Since S_2 is a subset maximal admissible set in F and $a \notin S_2$, either $S_2 \cup \{a\}$ is not conflict-free, or not admissible in F . Consequently, either $S_2 \cup \{a_i\}$ is not conflict-free, or not admissible (by Claim 13). Since $A' \setminus A = C_a$, this proves that S_2 is a preferred set in F' . Similar arguments (with Claim 13) yield results for the remaining two semantics. As a result, each $a_i \in C_a$ is a σ -facet since $a_i \notin S_2$ for each $a_i \in C_a$.

Conversely, suppose a is not a σ -facet in F . If a is not contained in any σ -extension of F , then no argument from C_a can be in any σ -extension of F' (using the same argument as in **SomeE**). Hence, no argument $a_i \in C_a$ is a σ -facet in F' . Similarly, if a belongs to every σ -extension of F , once again we have that every $a_i \in C_a$ is contained in every σ -extension in F' (due to **NotAllE**). This results once again in no argument $a_i \in C_a$ being a σ -facet in F' . \square

We next observe that the argument a is a σ -facet in F iff the AF F' admits at least n σ -facets for each $\sigma \in$

$\{\text{pref}, \text{semiSt}, \text{stag}\}$. Indeed, a is a σ -facet in F iff each $a_i \in C_a$ is also a facet in F' due to Claim 14. Thus resulting in at least n facets in F' . In contrast, if a is not a facet in F then no argument in $a_i \in C_a$ is a facet. Hence F' has at most $n - 1$ facets. \square

It is worth remarking that the reduction from ISFACET_σ to $\text{FACETS}_\sigma^{\geq k}$ presented in the proof of Theorem 12 does not work for admissible semantics. This holds since the converse direction of Claim 14 ‘if a is not a facet in F then no argument $a_i \in C_a$ is a facet in F' ’ is no longer true. Suppose a is not a facet. Assume further that a belongs to all admissible sets and there is at least one such set. Now, take any admissible set S in F . Clearly, $a \in S$, however, $S \setminus \{a\} \cup \{a_i\}$ is admissible in F' and does not contain a . This results in every $a_i \in C_a$ being adm-facet in F' . Consequently, a is not an admissible-facet in F although each of its copy in F' is an admissible-facet. Moreover, F' also admits at least $n - 1$ adm-facets in this case, thus violating the proof from Theorem 12.

Theorem 15 (\star). $\text{FACETS}_\sigma^{\leq k}$ is coNP-c., $\sigma \in \{\text{adm}, \text{stab}, \text{comp}\}$, whereas Π_2^{P} -c., $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.

3.3 Exact k Facets

Perhaps unsurprisingly, $\text{FACETS}_\sigma^{\leq k}$ also turns out to be easy for conflict-free and naive semantics.

Theorem 16 (\star). $\text{FACETS}_\sigma^{\leq k}$ is in P for $\sigma \in \{\text{conf}, \text{naiv}\}$.

For $\sigma \in \{\text{adm}, \text{stab}, \text{comp}\}$ the problem $\text{FACETS}_\sigma^{\leq k}$ turns out to be more interesting since it is complete for the comparably esoteric class DP.

Theorem 17. $\text{FACETS}_\sigma^{\leq k}$ is DP-c., $\sigma \in \{\text{adm}, \text{stab}, \text{comp}\}$.

Proof. The membership follows directly from $\text{FACETS}_\sigma^{\geq k}$ (Thm. 11) and $\text{FACETS}_\sigma^{\leq k}$ (Thm. 15).

For hardness, we reduce from SAT-UNSAT. To this aim, we utilize Lemma 10 for an instance (φ, ψ) of SAT-UNSAT. We assume w.l.o.g. that φ and ψ do not share variables. Then, φ is satisfiable and ψ is not satisfiable iff the AF F_φ has k_1 facets and F_ψ has $k_2 - 1$ facets, where F_i is the corresponding AF for $i \in \{\varphi, \psi\}$ with k_i arguments. However, there is a small technical issue as the AF $F_\varphi \cup F_\psi$ can not distinguish the failure of the satisfaction of φ from that of ψ . Therefore, we can not simply take the union $F_\varphi \cup F_\psi$ and let the number of facets be $k_1 + k_2 - 1$. Nevertheless, we duplicate the argument φ in F_φ to yield the pair φ, φ' . Then F_φ includes the additional attacks $(\varphi, \varphi'), (\varphi', \varphi)$ as well as (C, φ') for each $C \in \varphi$. The resulting AF F_φ has $k_1 + 1$ arguments and φ is satisfiable iff F_φ has $k_1 + 1$ facets. Note that the newly added argument φ' is a σ -facet iff φ is σ -facet for each $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$. Then, the theorem follows since (φ, ψ) is a positive instance of SAT-UNSAT iff φ is satisfiable and ψ is not satisfiable iff $F_\varphi \cup F_\psi$ has exactly $(k_1 + k_2)$ σ -facets for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$. \square

We conclude our complexity analysis by stating the non-tight bounds for $\text{FACETS}_\sigma^{\leq k}$ for the remaining semantics.

Theorem 18 (\star). $\text{FACETS}_\sigma^{\leq k}$ is in DP₂ for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.

4 Significance

Our notion of significance adopts a decision-driven perspective. We define significance of arguments in terms of the influence of a decision to eliminate the degree of freedom (on choices of remaining arguments). While counting approaches assess the plausibility of arguments in terms of their likelihood of being accepted, we measure how much the acceptance of an argument decreases freedom (or increases the significance of the decision). Intuitively, a higher significance score indicates that a specific decision does have a huge influence on the remaining facets. Furthermore, the number of facets directly measures the *amount of uncertainty* in extensions. Consider an argument a (as an *opinion* or a *view point*) and denote by \bar{a} the complement/negation of a . E.g., an argument is approved (a) versus not approved (\bar{a}). Then, a facet $\ell \in \{a, \bar{a}\}$ can be seen as the *uncertainty* regarding a , since a can either be included in, or be excluded from certain extensions. We next introduce the notion of *approving* and *disapproving* an argument. Let $F = (A, R)$ be an AF, and σ be a semantics. Recall that $\sigma(F)$ denotes the collection of σ -extensions in F . Moreover, C_σ (resp., S_σ) denotes the collection of credulously (skeptically) accepted arguments in F under semantics σ . For an argument a , we let $\sigma^a(F)$ denote the σ -extensions in F *approving* the argument a . Precisely, we define $\sigma^a(F) = \{E \in \sigma(F) \mid a \in E\}$. Moreover, $\sigma^{\bar{a}}(F) = \{E \in \sigma(F) \mid a \notin E\}$ represents the σ -extensions in F *disapproving* a . Now, let C_σ^a (resp., S_σ^a) be the arguments in some (all) $E \in \sigma^a(F)$. Finally, $\mathcal{F}_\sigma^a(F)$ denotes the σ -facets by considering only extensions in $\sigma^a(F)$ (i.e., $C_\sigma^a \setminus S_\sigma^a$).

For an argument $a \in A$ and $\ell \in \{a, \bar{a}\}$, we denote $\ell = \bar{a}$ if $\ell = a$ and $\ell = a$ for $\ell = \bar{a}$. We say that ℓ is approved iff ℓ is disapproved. Approving a facet $\ell \in \{a, \bar{a}\}$ reduces the uncertainty regarding the remaining arguments in A by restricting the extensions space to sets (not) containing a . Further, approving ℓ can render a facet argument $b \in A$ non-facet. This holds since, either (C1) $b \in E$ for each $E \in \sigma^\ell(F)$ but $b \notin E$ for each $E \in \sigma(F)$, or (C2) $b \notin E$ for any $E \in \sigma^\ell(F)$ but $b \in E$ for some $E \in \sigma(F)$. Intuitively, we say that the uncertainty of such an argument b has been *resolved* by approving ℓ . Further, we say that approving ℓ results in the approval of b in the case of (C1), and disapproval of b if (C2) is the case.

Notice that approving (or disapproving) an argument results in fewer facets for every semantics σ . That is, the (dis)approval of any argument can not generate new facets. Intuitively, we have less uncertainty than before after we (dis)approve certain arguments. Precisely, we have the following lemma.

Lemma 19. *For any argument $a \in A$ and semantics σ , $\mathcal{F}_\sigma^a(F) \subseteq \mathcal{F}_\sigma(F)$.*

Let σ be a semantics, $a \in A$ be a σ -facet and $\ell \in \{a, \bar{a}\}$. The observation that “ ℓ reduces the uncertainty among remaining arguments” leads to the notion of *significance* of ℓ under semantics σ . For an AF F , we define:

$$\mathbb{S}_\sigma[F, \ell] := \frac{|\mathcal{F}_\sigma(F)| - |\mathcal{F}_\sigma^\ell(F)|}{|\mathcal{F}_\sigma(F)|}. \quad (1)$$

Intuitively, approving an argument a is less significant if many uncertain arguments (facets) remain in $\mathcal{F}_\sigma^a(F)$. Similarly,

$\ell \in$	$\{w, m, t, \bar{s}, \bar{b}, \bar{p}\}$	$\{s, b, \bar{w}, \bar{m}\}$	$\{p, t\}$
$ \mathcal{F}_{\text{stab}}^\ell(F) $	0	2	4
$\mathbb{S}_{\text{stab}}[F, \ell]$	1	$\frac{2}{3}$	$\frac{1}{3}$

Table 2: Argument significance for the AF from Example 20.

disapproving a (and thus approving \bar{a}) is less significant if many facets remain in $\mathcal{F}_\sigma^{\bar{a}}(F)$.

Example 20 (Arguments Significance). *Reconsider the AF F from Example 1 with stable extensions $\text{stab}(F) = \{\{w, m, p\}, \{s, b, p\}, \{s, b, t\}\}$. While Example 2 gave an intuition of significance, Table 2 presents precise values for each argument. As outlined, the argument w has score 1, and is thus more significant than \bar{w} (score $2/3$). The argument e not being a stab-facet is excluded for significance reasoning.*

4.1 Computing Significance for Arguments

Let $F = (A, R)$ be an AF, σ be a semantics and $a \in A$ be an argument. Observe that the computation of $\mathbb{S}_\sigma[F, \ell]$ (Equation 1) requires counting facets in $\mathcal{F}_\sigma(F)$ and $\mathcal{F}_\sigma^\ell(F)$. We argue that one can count the number of σ -facets in a framework $F = (A, R)$ without having to enumerate or count all σ -extensions explicitly. In fact, $\mathcal{F}_\sigma(F)$ can be computed by asking ISFACET_σ for each argument $a \in A$, which requires $|A|$ -many queries. Moreover, one can also count the exact facets in $\mathcal{F}_\sigma^\ell(F)$ without having to explicitly identify all σ^ℓ -extensions. Observe that $\mathcal{F}_\sigma^\ell(F)$ corresponds to the result of remaining facets after approving ℓ .

Theorem 21. *Let σ be a semantics and $F = (A, R)$ be an AF. For every $a \in A$ and $\ell \in \{a, \bar{a}\}$ the sets $\mathcal{F}_\sigma(F)$ and $\mathcal{F}_\sigma^\ell(F)$ can be computed by a deterministic polynomial-time Turing machine with access to*

- an **NP** oracle for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$.
- a Σ_2^P oracle, for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.

Proof. Given an AF $F = (A, R)$, semantics σ and argument $a \in A$. To compute $\mathcal{F}_\sigma(F)$, we consider the following procedure. For each $b \in A$:

1. Guess two sets $E_1, E_2 \subseteq A$,
2. Check that $b \in E_1, b \notin E_2$,
3. Check that $E_i \in \sigma(F)$,
4. Answer “Yes” if each check is passed in Step 2 – 3.

This procedure is repeated polynomially-many times (precisely $|A|$ -many times). As a post-processing, count the number of arguments $b \in A$, for which Step-4 answers “Yes”. The 2nd step requires non-deterministic guesses, whereas the 3rd step needs (1) **P** for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$ and (2) **coNP**, for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$. Step-4 is again a final post-processing. As a result, the procedure overall runs in the mentioned runtime for corresponding semantics. That is, **P**-time with an **NP** oracle for $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$ and Σ_2^P -oracle, for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$.

To compute $\mathcal{F}_\sigma^\ell(F)$, we additionally include the following check to the Step-2 in the above procedure.

- 2a. Check that $a \in E_1$, $a \in E_2$ if $\ell = a$, and check $a \notin E_1$, $a \notin E_2$ if $\ell = \bar{a}$.

This does not increase runtime, and completes the proof. \square

Observe that one can not expect to lower the runtime in Theorem 21 (e.g., to **NP** in the case of $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$). Intuitively, although each step (Step 1 – 4 in the proof of Theorem 21) requires **NP**-time, the final post-processing needs counting the number of arguments for which each check is passed. In fact, this would contradict Theorem 17 (unless **NP** = **DP**) since one can count all facets in an AF and determine whether this number equals k .

5 Implementation and Experiments

Implementation We implemented counting of extensions and facets for various semantics into our tool called `frame` (Facets for Reasoning and Analyzing Meaningful Extensions). We build on the `Aspartix` system, an ASP-based argumentation system for Dung style abstract argumentation and extensions thereof [Egly *et al.*, 2008]. We employ `Aspartix`’s ASP encoding and take the ASP solver `clingo` version 5.7.1 [Gebser *et al.*, 2012] to compute credulous and skeptical consequences and take set differences. We leverage ASP as the solvers have native support for enumerating consequences without exhaustive enumeration of all answer-sets [Alviano *et al.*, 2023; Gebser *et al.*, 2009].

Design and Expectations. We ran our experiments on a Ubuntu 11.4.0 Linux 5.15 computer with an eight core Intel i7-14700 CPU 1.5 GHz machine with 64GB of RAM. Each run is executed exclusively on the system. To illustrate that we can count facets on practical instances and obtain insights over counting, we take the admissible, stable semantics, and semi-stable semantics over instances from the 3rd International Competition on Computational Models of Argumentation (IC-CMA’19) [Bistarelli *et al.*, 2019]. This gives us 326 different argumentation frameworks of varying sizes in the range of 0 to 10.000 arguments with an average of 800.3 arguments and a median of 160.0 arguments. We take admissible, stable, and semi-stable semantics as representative from each different level of hardness (see Table 1). We take the 2019 competition as instances are of reasonable size, runtime, number for the scope of this experiment, and we can use input instances without further modification. We limit the runtime on each instance to 60 seconds for sustainability reasons, as differences become visible already with the limitation, and as a user might not want to wait long when investigating search spaces. We collect the number of extensions and facets and measure the solver runtime. We have the following expectations: (i) computing facets is faster than enumerating extensions, (ii) facets are still accessible when the number of extensions is very high and enumeration takes longer runtime, (iii) even when there are many extensions, there are reasonably small number of facets.

Observations and Summary. Table 3 presents a survey of our results. Details on the evaluation are available in the supplemental data. We see that the number of admissible extensions can be larger than 10^6 and `clingo` fails to enumerate all extensions (*) if the number of extensions is very high. For

σ	n	I	$\#_e$	$\#_f$	t_e	t_f
adm*	128	$[10^6, \cdot)$	$\dagger 189.0 \cdot 10^6$	165.1	61.0	0.5
adm	11	$[10^6, \cdot)$	$27.1 \cdot 10^6$	40.6	6.2	1.3
adm	187	$(0, 10^6)$	$29.4 \cdot 10^3$	266.7	0.8	1.4
stab*	1	$(0, 10^6)$	$\dagger 3.0 \cdot 10^0$	—	61.0	61.0
stab	$\dagger 310$	$(0, 10^6)$	$36.0 \cdot 10^0$	43.1	0.7	0.8
stab	14	$[0, 0]$	$0.0 \cdot 10^0$	0.0	1.4	1.2
semiSt*	9	$(0, 10^6)$	$\dagger 43.8 \cdot 10^0$	79.9	61.0	49.4
semiSt	177	$(0, 10^6)$	$50.9 \cdot 10^0$	33.9	3.7	2.1
semiSt*	140	$[0, 0]$	$\dagger 0.0 \cdot 10^0$	—	61.0	—

Table 3: Overview of results on enumerating extensions and computing facets for semantics $\sigma \in \{\text{adm}, \text{semiSt}, \text{stab}\}$ where * marks timeouts (in order to distinguish the cases with/without timeout on enumerating extensions) and the columns contain the total number n of argumentation frameworks with the interval $I = [a, b)$ referring to $a \leq \#_e < b$; the average number $\#_e$ and $\#_f$ of extensions/facets; the average runtime t_e and t_f for enumerating extensions (e) / computing facets (f). The symbol \dagger illustrates that there is only a lower bound as computation did not finish. We excluded one instance (\dagger) due to timeout when computing credulous/skeptical extensions.

the admissible extensions, we observe that even if the number of extensions is quite high, the number of facets remains reasonably small making it interesting to diversify extensions, or investigate more details about those arguments, which still allow flexibility. The observations confirm our expectations for the admissible semantics. However, we gain only limited insights for the semantics `stab` and `semiSt`. Here the number of extensions is fairly low and the solver either manages to enumerate all extensions, or already fails to solve one.

6 Conclusion

We defined a new perspective on exploring significance of arguments in extensions of an abstract argumentation framework. We present a comprehensive complexity analysis. We establish that the complexity of deciding whether the number of facets is exactly k ranges between **P** and **DP**₂, including tight lower bounds for most cases (see Table 1). While our primary focus lies on establishing a comprehensive complexity picture, our implementation allows computing the number of facets practically building on top of existing solvers.

For future work, we plan to investigate techniques whether significance originating from facets can be extended to arguments depending on each other and notions of fairness in argumentation frameworks. We also believe that the missing case for $\sigma \in \{\text{pref}, \text{semiSt}, \text{stag}\}$ remains interesting to study. We expect that the problem is also hard, as the decision problems for at least and at most are Σ_2^P - and Π_2^P -complete. From a practical perspective, we believe that it would be interesting to integrate facet-based reasoning and significance computation into modern SAT-based argumentation solvers. Moreover, investigating facets for other formalisms such as abductive reasoning [Mahmood *et al.*, 2020] or default logic [Fichte *et al.*, 2022b] seems interesting as well as closing the gap to the topic of inconsistencies [Fichte *et al.*, 2021b; Fichte *et al.*, 2023b].

Acknowledgments

The work has received funding from the Austrian Science Fund (FWF), grants J 4656 and P 32830, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), grants TRR 318/1 2021 – 438445824 and ME 4279/3-1 (511769688), the European Union’s Horizon Europe research and innovation programme within project ENEXA (101070305), the Society for Research Funding in Lower Austria (GFF, Gesellschaft für Forschungsförderung NÖ), grant ExzF-0004, the Swedish research council under grant VR-2022-03214, as well as the Vienna Science and Technology Fund (WWTF), grant ICT19-065, and by the Ministry of Culture and Science of North Rhine-Westphalia (MKW NRW) within project WHALE (LFN 1-04) funded under the Lamarr Fellow Network programme, and ELLIIT funded by the Swedish government.

Part of the research was carried out while Hecher was visiting the Simons Institute for the Theory of Computing at UC Berkeley.

Johannes and Yasir appreciate valuable discussions with Lydia Blümel and Matthias Thimm (Fern Universität in Hagen) on an alternative proof to Theorem 21.

References

- [Alrabbaa *et al.*, 2018] Christian Alrabbaa, Sebastian Rudolph, and Lukas Schweizer. Faceted answer-set navigation. In *RuleML+RR’18*, pages 211–225. Springer Verlag, 2018.
- [Alviano *et al.*, 2023] Mario Alviano, Carmine Dodaro, Salvatore Fiorentino, Alessandro Previti, and Francesco Ricca. ASP and subset minimality: Enumeration, cautious reasoning and MUSes. *Artificial Intelligence*, 320:103931:1–25, 2023.
- [Alviano, 2021] Mario Alviano. The pyglaf argumentation reasoner (ICCMA2021). *CoRR*, abs/2109.03162, 2021.
- [Amgoud and Prade, 2009] Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173(3-4):413–436, 2009.
- [Bench-Capon and Dunne, 2007] Trevor J.M. Bench-Capon and Paul E. Dunne. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10–15):619–641, 2007.
- [Biere *et al.*, 2021] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability - Second Edition*, volume 336 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2021.
- [Bistarelli *et al.*, 2019] Stefano Bistarelli, Lars Kotthoff, Theofrastos Mantadelis, Francesco Santini, and Carlo Taticchi. The third international competition on computational models of argumentation (ICCMA’19). 2019.
- [Bondy and Murty, 2008] J. Adrian Bondy and Uppaluri S. R. Murty. *Graph Theory*. Graduate Texts in Mathematics. Springer, 2008.
- [Dachselt *et al.*, 2022] Raimund Dachselt, Sarah Alice Gaggl, Markus Krötzsch, Julián Méndez, Dominik Rusovac, and Mei Yang. NEXAS: A visual tool for navigating and exploring argumentation solution spaces. In *COMMA’22*, 2022.
- [Dung, 1995] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n -person games. *Artificial Intelligence*, 77(2):321–357, 1995.
- [Dunne and Bench-Capon, 2002] Paul E. Dunne and Trevor J. M. Bench-Capon. Coherence in finite argument systems. *Artificial Intelligence*, 141(1/2):187–203, 2002.
- [Dvořák and Dunne, 2017] Wolfgang Dvořák and Paul E. Dunne. Computational problems in formal argumentation and their complexity. *FLAP*, 4(8), 2017.
- [Dvořák, 2012] Wolfgang Dvořák. *Computational aspects of abstract argumentation*. PhD thesis, TU Wien, 2012.
- [Dvořák and Woltran, 2010] Wolfram Dvořák and S. Woltran. Complexity of semi-stable and stage semantics in argumentation frameworks. *Information Processing Letters*, 110(11):425–430, 2010.
- [Egly *et al.*, 2008] Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. Aspartix: Implementing argumentation frameworks using answer-set programming. In *ICLP’08*, pages 734–738, 2008.
- [Fichte *et al.*, 2021a] Johannes K. Fichte, Markus Hecher, Yasir Mahmood, and Arne Meier. Decomposition-guided reductions for argumentation and treewidth. In *IJCAI’21*, pages 1880–1886. IJCAI Organization, 2021.
- [Fichte *et al.*, 2021b] Johannes K. Fichte, Markus Hecher, and Arne Meier. Knowledge-base degrees of inconsistency: Complexity and counting. In *AAAI’21*, pages 6349–6357, 2021.
- [Fichte *et al.*, 2022a] Johannes K Fichte, Markus Hecher, and Mohamed A. Nadeem. Plausibility reasoning via projected answer set counting—a hybrid approach. In *IJCAI’22*, pages 2620–2626. IJCAI Organization, 2022.
- [Fichte *et al.*, 2022b] Johannes K. Fichte, Markus Hecher, and Irina Schindler. Default logic and bounded treewidth. *Information and Computation*, 283, 2022.
- [Fichte *et al.*, 2023a] Johannes K. Fichte, Markus Hecher, Yasir Mahmood, and Arne Meier. Quantitative reasoning and structural complexity for claim-centric argumentation. In *IJCAI’23*, pages 3212–3220. ijcai.org, 2023.
- [Fichte *et al.*, 2023b] Johannes K. Fichte, Markus Hecher, and Stefan Szeider. Inconsistent cores for ASP: The perks and perils of non-monotonicity. In *AAAI’23*, pages 6363–6371, 2023.
- [Fichte *et al.*, 2024] Johannes K. Fichte, Markus Hecher, and Arne Meier. Counting complexity for reasoning in abstract argumentation. *JAIR*, 80, 2024.
- [Fichte *et al.*, 2025] Johannes Fichte, Nicolas Fröhlich, Markus Hecher, Victor Lagerkvist, Yasir Mahmood, Arne Meier, and Jonathan Persson. Facets in argumentation: A formal approach to argument significance. *CoRR*, abs/2505.10982, 2025.
- [Gebser *et al.*, 2009] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Solution enumeration for projected boolean search problems. In *CPAIOR’09*, volume 5547, pages 71–86, 2009.

- [Gebser *et al.*, 2012] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187-188:52–89, 2012.
- [Lautemann, 1983] Clemens Lautemann. Bpp and the polynomial hierarchy. *Information Processing Letters*, 17(4):215–217, 1983.
- [Lohrey and Rosowski, 2023] Markus Lohrey and Andreas Rosowski. On the Complexity of Diameter and Related Problems in Permutation Groups. In *ICALP’23*, volume 261, pages 134:1–134:18. Dagstuhl Publishing, 2023.
- [Mahmood *et al.*, 2020] Yasir Mahmood, Arne Meier, and Johannes Schmidt. Parameterized complexity of abduction in schaefer’s framework. *Journal of Logic and Computation*, 31(1):266–296, 12 2020.
- [Niskanen and Järvisalo, 2020] Andreas Niskanen and Matti Järvisalo. μ -toksia: An Efficient Abstract Argumentation Reasoner. In *KR’20*, pages 800–804, 9 2020.
- [Pippenger, 1997] Nicholas Pippenger. *Theories of computability*. Cambridge University Press, 1997.
- [Rago *et al.*, 2018] Antonio Rago, Oana Cocarascu, and Francesca Toni. Argumentation-based recommendations: Fantastic explanations and how to find them. In *IJCAI’18*, pages 1949–1955. IJCAI Organization, 2018.
- [Rusovac *et al.*, 2024] Dominik Rusovac, Markus Hecher, Martin Gebser, Sarah Alice Gaggl, and Johannes K. Fichte. Navigating and querying answer sets: How hard is it really and why? In *KR’24*, 2024.
- [Sipser, 1983] Michael Sipser. A complexity theoretic approach to randomness. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing (STOC’83)*, pages 330–335, Boston, Massachusetts, USA, 1983.
- [Speck *et al.*, 2025] David Speck, Markus Hecher, Daniel Gnad, Johannes K. Fichte, and Augusto B. Corrêa. Counting and reasoning with plans. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(25):26688–26696, Apr. 2025.
- [Stockmeyer and Meyer, 1973] Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time. In *Proc. of the 5th Annual ACM Symposium on Theory of Computing (STOC’73)*, pages 1–9. ACM, 1973.
- [Stockmeyer, 1976] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.
- [Stockmeyer, 1983] Larry Stockmeyer. The complexity of approximate counting. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing (STOC’83)*, pages 118–126, New York, NY, USA, 1983. Association for Computing Machinery.
- [Thimm *et al.*, 2021] Matthias Thimm, Federico Cerutti, and Mauro Vallati. Fudge: A light-weight solver for abstract argumentation based on SAT reductions. *CoRR*, abs/2109.03106, 2021.
- [Thimm *et al.*, 2024] Matthias Thimm, Jean-Guy Mailly, and Johannes P. Wallner. International competitions on computational models of argumentation. <https://argumentationcompetition.org/2025/index.html>, 2024.
- [Toda, 1991] Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- [Wrathall, 1976] Celia Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.