Diversity of Extensions in Abstract Argumentation

Johannes K. Fichte[1] , **Markus Hecher**[2,3] , **Yasir Mahmood**[4] , **Zhengjun Wang**[4]

[1]Department of Computer and Information Science (IDA), Linköping University, Sweden
[2]University of Potsdam, Germany
[3]University of Artois, CNRS, UMR8188 (CRIL), France
[4]Data Science Group, Heinz Nixdorf Institute, Paderborn University, Germany
johannes.fichte@liu.se, hecher@cril.fr, yasir.mahmood@uni-paderborn.de, zwang@mail.uni-paderborn.de

Abstract

Argumentation is an important topic of AI for modeling and reasoning about arguments. In abstract argumentation, we consider directed graphs, so-called *argumentation frameworks (AF)*, that express conflicts between arguments. The semantics is defined by the notion of *extensions*, which are sets of arguments that satisfy particular relationship conditions in the AF. Usually, standard reasoning in argumentation do not reveal *how far apart* extensions are.

We introduce a quantitative notion of *diversity of extensions* based on the symmetric-difference and provide a systematic complexity classification. Intuitively, diversity captures whether extensions of a framework (accepted viewpoints) differ only marginally or represent fundamentally incompatible sets of arguments. We study whether an AF admits k-diverse extensions, admits k-diverse extensions covering specific arguments, and to compute the largest k for which an AF admits k-diverse extensions. We outline a prototype and provide an evaluation for computing diversity levels.

1 Introduction

Abstract argumentation [Dung, 1995; Rahwan, 2007; Baroni *et al.*, 2011] is a well-known concept in AI for modeling, relating, and grouping arguments [Amgoud and Prade, 2009; Rago *et al.*, 2018]. To express opposing positions, arguments are related to each other in form of a directed graph, so-called *argumentation framework (AF)*. Sets of arguments (*extensions*), which satisfy certain conditions regarding the relationship among them (semantics), build the solutions to an AF. Usually, an AF admits multiple extensions, in other words, mutually incompatible but internally coherent sets of jointly acceptable arguments, which can be understood as alternative "view-points" on the same argumentative situation.

Research in abstract argumentation has introduced concepts to facilitate insights into diverging extensions, for example, explaining and selecting a set of most representative viewpoints [Bernreiter *et al.*, 2024], deciding between multiple extensions [Dauphin *et al.*, 2018], exploring large solution spaces [Dachselt *et al.*, 2022], deciding the significance of individual arguments [Fichte *et al.*, 2025], conditional probability

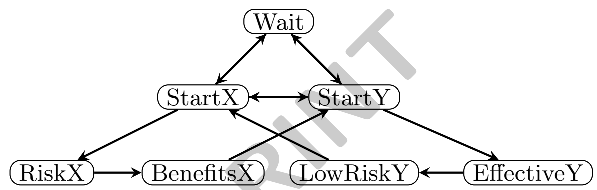

Figure 1: An Example AF modeling the medical decision making.

of one or multiple arguments occurring in an extension [Fichte *et al.*, 2023], counting extensions [Fichte *et al.*, 2024], and incorporating preordering over arguments (extension-ranking semantics) [Skiba *et al.*, 2021]

However, existing concepts do not reveal *how far apart* and *diverse* entire extensions or *how diverse AFs* are. This is particularly interesting in environments where we want to negotiate positions. Different view-points may reasonably endorse different extensions. If we consider only one arbitrarily chosen extension, a representative extension, or parts of extensions, we easily hide genuine trade-offs and impede agreement. To this end, we study diverse extensions and introduce a quantitative notion of extension diversity based on the symmetric-difference. The *symmetric-difference* $A \triangle B$ of two sets A and B takes the elements that are in exactly one of them, but not in both, i.e., $A \triangle B := (A \cup B) \setminus (A \cap B)$. Example 1 illustrates a situation where we are interested in understanding the diversity of extensions.

Example 1. *Take an AF $F = (A, R)$ modeling a medical decision under time pressure, as depicted in Figure 1. The arguments of F have the following intuitive meaning:*

- ***sx** (StartX): Start Treatment X immediately.*
- ***sy** (StartY): Start Treatment Y immediately.*
- ***rx** (RiskX): Given no immediate decision, X's long-term risks must be carefully evaluated first.*
- ***ey** (EffectiveY): Given no immediate decision, Y's effectiveness must be further tested.*
- ***bx** (BenefitsX): X has strong short-term benefits.*
- ***lry** (LowRiskY): Y has fewer side effects.*
- ***w** (Wait): collect more data before deciding.*

Under stable semantics, which intuitively expresses a self-consistent position that defeats all alternatives, F admits exactly three extensions. Precisely, $E_X = \{sx, bx, ey\}, E_Y = \{sy, lry, rx\}$, and $E_W = \{w, rx, ey\}$, each highlighting ei-

Problems/σ	σ_1	σ_2	σ_3
$^{\leq}$k-Div-Ext$_\sigma$	$\in$ **NP**/ Trivial R.13	**NP**-c Th.15	$\mathbf{\Sigma_2^P}$-c Th.19
$^{\geq}$k-Div-Ext$_\sigma$	**NP**-c Th.8	**NP**-c Th.15	$\mathbf{\Sigma_2^P}$-c Th.19
$^{\leq}$k-Div-Arg$_\sigma$	**NP**-c Th.11/$\in$ **P** Th.13	**NP**-c Th.14	$\mathbf{\Sigma_2^P}$-c Th.17
$^{\geq}$k-Div-Arg$_\sigma$	**NP**-c Th.6	**NP**-c Th.14	$\mathbf{\Sigma_2^P}$-c Th.17
Max	**DP**-c Th.22/ Th.23	**DP**-c Th.26	$\mathbf{DP}_2$-c Th.26

Table 1: Overview of our complexity results for the semantics $\sigma_1 \in \{\mathsf{conf}, \mathsf{naiv}\}$, $\sigma_2 \in \{\mathsf{adm}, \mathsf{stab}, \mathsf{comp}\}$, and $\sigma_3 \in \{\mathsf{semiSt}, \mathsf{stag}\}$. The problem "Arg" asks if two arguments are k-diverse and "Ext" asks if an AF has two k-diverse σ extensions. The results for $^{\leq}$k-Div-Arg$_\sigma$ also apply to k-Div-Arg$_\sigma$ for each σ.

ther a type of treatment, or the argument to wait for more data.

Here, three most important arguments regarding the final decision are sx, sy and w which either select a type of treatment or argue to wait, respectively. Each of these arguments occur in some extension (credulously accepted). Hence, we cannot tell much about either of the three options. But looking from a global perspective on all the available options (stable extensions), we observe that the extensions for either treatment do not share any argument, whereas they do share at least one argument with the extension that forces us to wait (E_W). Precisely, $E_X \triangle E_Y = \{sx, bx, ey, sy, lry, rx\}$ meaning we consider options "StartX" and "StartY" to be 6-diverse. However, we have $E_X \triangle E_W = \{sx, bx, rx, w\}$, and $E_Y \triangle E_W = \{sy, lry, ey, w\}$. Consequently, we observe that the options "StartX" ("StartY") and "Wait" are only 4-diverse. Intuitively, this means that the justification for either type of treatment uses very "different" line of reasoning, whereas both treatment options share at least some reasoning with the third option to wait before deciding anything. Alternatively, the reasoning for "StartX" is somewhat similar *to the reasoning for "Wait" compared to "StartY".* ◁

Contributions. Our main contributions are as follows.

1. We introduce the notion of symmetric-difference-based diversity to abstract argumentation and study its properties. We show how diversity differs from existing notions and that it provides a more fine-grained notion where concepts based on credulous and skeptical reasoning fail.
2. We study the computational complexity of central problems (above, below, exactly k-diverse, max diversity) and provide a thorough picture of the complexity landscape, which are listed in Table 1.
3. We provide a prototypical implementation and an initial experimental evaluation, where we compute the range of diversity via logic programming (ASP).

Proofs to our statements and details on the code implementation can be found in the technical report [Fichte *et al.*, 2026].

Related Work. Čyras *et al.* [2021] survey argumentation-based explanation under the perspective of XAI. Liao and van der Torre [2020] introduce and study explanation semantics, which associates with each accepted argument a set of such explanation arguments. Niskanen and Järvisalo [2020a] and [Saribatur *et al.*, 2020] study the opposite case where arguments are rejected and analyze reasons for the rejection. Although counting complexity for argumentation has been studied [Fichte *et al.*, 2024], the results do not apply here, as our focus is on the distances between elements in the extension space. There are various practical solvers for decision and reasoning tasks [Egly *et al.*, 2008; Niskanen and Järvisalo, 2020b; Thimm *et al.*, 2021; Alviano, 2018; Alviano, 2021] that compete in the biannual ICCMA competition [Thimm *et al.*, 2024].

Fichte *et al.* [2025] propose facets to abstract argumentation, which enable quantifying significance of decisions and study the computational complexity. Facets do not capture information on extensions containing them, whereas diversity measures the (dis)similarity of accepted argument sets in an AF. Böhl *et al.* [2023] consider diversity and representative sets in answer-set programming. Ulbricht and Wallner [2021] consider strong explanation, which ensure that a target set of arguments is acceptable in each sub-framework containing the explaining set. Besnard *et al.* [2022] investigate concepts of contrastive and non-contrastive questions with respect to explanations. Finding similar and diverse solutions has been considered for constraint programming [Hebrard *et al.*, 2005] and answer-set programming [Eiter *et al.*, 2013]. Rodrigues [2019] compare different alternative representations of extensions for systematic enumeration. Coste-Marquis *et al.* [2015] employed Hamming distance to define enforcement operators as previously also done in the context of minimal change [Baumann, 2012]. Dachselt *et al.* [2022] developed a tool to navigate argumentation frameworks using ASP-facets and more fine-grained reasoning modes have also been studied in ASP [Fichte *et al.*, 2022]. Note that ASP-navigation is based on forbidding or enforcing atoms in programs via integrity constraints.

2 Preliminaries

We assume familiarity with computational complexity [Pippenger, 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 **P** (**NP**) for the class of decision problems solvable in (non-deterministic) polynomial time. Additionally, we let co**NP** be the class of decision problems whose complement is in **NP**, and let **DP** be the class of decision problems representable as the intersection of a problem in **NP** and a problem in co**NP**. On top, we use more classes from the polynomial hierarchy [Stockmeyer and Meyer, 1973; Stockmeyer, 1976; Wrathall, 1976], $\mathbf{\Delta_0^P} := \mathbf{\Pi_0^P} := \mathbf{\Sigma_0^P} := \mathbf{P}$ and $\mathbf{\Delta_i^P} := \mathbf{P}^{\mathbf{\Sigma_{i-1}^P}}$, $\mathbf{\Sigma_i^P} := \mathbf{NP}^{\mathbf{\Sigma_{i-1}^P}}$, and $\mathbf{\Pi_i^P} := \mathrm{co}\mathbf{NP}^{\mathbf{\Sigma_{i-1}^P}}$ for $i > 0$ where C^D is the class C of decision problems augmented by an oracle for some complete problem in class D. The complexity class $\mathbf{DP}_k$ is defined as $\mathbf{DP}_k := \{L_1 \cap L_2 \mid L_1 \in \mathbf{\Sigma_k^P}, L_2 \in \mathbf{\Pi_k^P}\}$, **DP**=$\mathbf{DP}_1$ [Lohrey and Rosowski, 2023].

The Boolean *satisfiability* problem (SAT) for formulas in *conjunctive normal form* (CNF), asks: given $\varphi = \bigwedge_{i=1}^m C_i$ where each C_i is a clause, decide whether φ admits at least one satisfying assignment. SAT is the one of the most famous **NP**-complete problems, whereas its complementary problem to check *unsatisfiability*, is co**NP**-complete. Moreover, the problem to check whether φ is satisfiable and ψ unsatisfiable for a given pair of formulas (φ, ψ) (the SAT-UNSAT problem) is **DP**-complete. For $\mathbf{\Pi_2^P}$ one usually considers 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. Finally, for $\mathbf{\Sigma_2^P}$ the problem is to check whether $\forall X \exists Y.\varphi$ 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. Let $E \subseteq A$ and $a \in A$ be an argument. Then we say that a is *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 $\mathrm{def}_F(E)$ is defined by $\mathrm{def}_F(E) := \{\ a \mid a \in A, a \text{ is defended by } E \text{ in } F\ \}$. In abstract argumentation, one aims at 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 $\mathrm{def}_F(E) = E$; (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 $\sigma \in \{\mathsf{conf}, \mathsf{naiv}, \mathsf{adm}, \mathsf{comp}, \mathsf{stab}, \mathsf{pref}, \mathsf{semiSt}, \mathsf{stag}\}$, 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 $\textsc{Exist}_\sigma$ 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. The complexity of AF reasoning is well-known, e.g., [Dvorák and Dunne, 2017].

3 Diversity in Abstract Argumentation

In this section, we introduce a quantitative notion of *diversity of extensions* based on the symmetric-difference.

Definition 2. *Let $F = (A, R)$ be an AF, σ be a semantics and $k \geq 0$ be an integer. For two sets $S \subseteq A$ and $T \subseteq A$ of arguments, we define by $d(S, T)$ the* distance *between S and T. Formally, we let*

$$d(S, T) := |S \triangle T| = |(S \cup T) \setminus (S \cap T)|$$

*We say that two extensions $S, T \in \sigma(F)$ are k-*diverse *if $d(S, T) = k$. Moreover, for an extension $S \in \sigma(F)$ and a set X, we say that S* covers *X if $X \subseteq S$.*

We continue our example from above.

Example 3. *Consider the AF from Example 1, its extensions, and their symmetric difference. We can see that one extension supports treatment X while another supports Y under stable semantics. These extensions are not only* different*, but are very* diverse *due to their large symmetric difference.*

The arguments "StartX" and "StartY" are both credulously accepted under the stable semantics. However, asking whether both arguments are k-diverse for some $k \in \mathbb{N}$ reveals how incompatible these acceptable choices actually are, or how radically different two justified medical decisions in this setting can be. In this context, smaller values of k highlight that the two treatments differ only marginally, i.e., same medical reasoning, but different preference. In contrast, larger values of k depict two entirely different diagnostic strategies. Thus, in decision-making, higher diversity indicates deeper disagreement between choices. ◁

This consideration leads us to the following problem.

> $\mathsf{k\text{-}Div\text{-}Arg}_\sigma(F, K, a, b)$
> Input: an AF F, $k \in \mathbb{N}$, and arguments a, b.
> Question: does F admit two k-diverse σ-extensions S and T such that $a \in S, b \in T$?

Note that we can extend $\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ from individual arguments to the sets of arguments, which further allows us to compute the diversity when "absence" (or even complements) of arguments are considered. Since obtained complexity results are identical, we did not explicitly mention these results.

We can also ask the diversity-related question at the level of an AF. That is, does an AF have two k-diverse σ-extensions for a given semantics σ. Here, diversity measures how deep the disagreement encoded in the AF is, beyond mere multiplicity of extensions. Our notion of diversity applies to the extension space (all acceptable sets of arguments) under some semantics. Therefore, disagreement is interpreted in terms of extensions (or accepted viewpoints). High diversity implies that the given AF supports radically different and internally coherent positions, thus indicating higher level of disagreement and low robustness of conclusions. On the other hand, low diversity implies that extensions differ only marginally with most reasons being shared. Consequently, they indicate near-consensus and robust conclusions with shallower disagreements.

> $\mathsf{k\text{-}Div\text{-}Ext}_\sigma(F, K)$
> Input: an AF F and $k \in \mathbb{N}$
> Question: does F admit two **non-empty** k-diverse σ-extensions?

The $\mathsf{k\text{-}Div\text{-}Ext}_\sigma(F, K)$ problems immediately give rise to whether there are σ-extension that are at least/most k-diverse, which we abbreviate by $^{\geq}\mathsf{k\text{-}Div\text{-}Ext}$ and $^{\leq}\mathsf{k\text{-}Div\text{-}Ext}$, respectively. Similarly, $\mathsf{k\text{-}Div\text{-}Arg}_\sigma(F, K, a, b)$ gives rise to $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}$ and $^{\leq}\mathsf{k\text{-}Div\text{-}Arg}$, respectively.

Finally, we define the problem to obtain the maximum k.

> $\mathsf{max\text{-}k\text{-}Div\text{-}Ext}_\sigma(F, A, B)$
> Input: an AF F and sets A,B of arguments
> Task: Output the maximum k for which A and B are covered by k-diverse extensions.

Similarly, we can define the corresponding minimization task $\mathsf{min\text{-}k\text{-}Div\text{-}Ext}_\sigma(F, A, B)$.

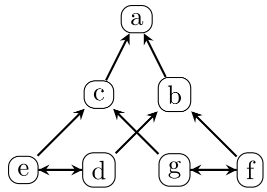

(a) Stable 4 and 6-diversity and admissible 1–3 and 5-diversity.

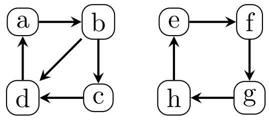

(b) Neither 3-diverse nor 5-diverse adm-extensions.

Figure 2: Example AFs showing different levels of diversity.

Levels of Diversity. It could be that some arguments are k-diverse and k'-diverse for some $k' < k \in \mathbb{N}$ without being ℓ-diverse for any $k' \leq \ell \leq k$. We illustrate how an AF might have different *levels* of diversity among its extensions.

Example 4. *Consider the AF $F_1 = (A, R)$ as depicted in Figure 2a. F_1 admits the following stable sets $E_1 = \{a, e, f\}, E_2 = \{b, e, g\}, E_3 = \{a, d, g\}$, and $E_4 = \{c, d, f\}$. The pair-wise symmetric difference of these sets yield diversity values in $\{4, 6\}$. In particular F_1 does not admit k-diverse stable extensions for any $k \in \{1, 2, 3, 5\}$. When considering admissible semantics, F_1 additionally admits the following singleton $\{e\}$, $\{d\}$, $\{g\}$, $\{f\}$, and two-element extensions $\{e, g\}$, $\{e, f\}$, $\{d, g\}$, and $\{d, f\}$. Here, it can be seen that under admissible semantics, F_1 also admits diversity values of 1 (due to extensions $\{e\}$, $\{e, f\}$), 2 (due to $\{e\}, \{d\}$), and 3 (due to $\{e\}$, $\{d, f\}$). Finally, by considering sets $\{a, e, f\}$ and $\{d, g\}$, we also obtain two 5-diverse* adm*-extension.*

Next, consider the AF $F_2 = (A, R)$ as depicted in Figure 2b. The AF F_2 admits the admissible sets $E_1 = \{a, c\}$, $E_2 = \{e, g\}$, $E_3 = \{f, h\}$, $E_4 = \{a, c, f, h\}$, and $E_5 = \{a, c, e, g\}$. The pair-wise symmetric difference of these sets yield diversity values in $\{2, 4, 6\}$. Notable, there are neither 3-diverse nor 5-diverse adm*-extensions.* ◁

4 Complexity of Reasoning about Diversity

We first consider the complexity of deciding whether an AF admits two k-diverse extensions ($\mathsf{k\text{-}Div\text{-}Ext}_\sigma$) for a semantics σ, or whether additionally contains two arguments in the question ($\mathsf{k\text{-}Div\text{-}Arg}_\sigma$). Moreover, we also consider the complexity of deciding whether an AF admits some diversity level above, or below k, i.e., ${}^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ and ${}^{\leq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$, respectively. Finally, we consider analogous problems (${}^{\geq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ and ${}^{\leq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$) for extensions.

We find it convenient to present our hardness reductions on top of the *standard translation* of SAT/QBF instances to AFs. The following reduction is relevant for some hardness results.

Definition 5 ([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:*

$$A := \{\varphi, C_1, \ldots, C_m\} \cup \{\, x, \bar{x} \mid x \in \mathsf{var}(\varphi) \}$$

$$R := \{\, (C_i, \varphi) \mid i \leq m \,\} \cup \{\, (x, \bar{x}), (\bar{x}, x) \mid x \in \mathsf{var}(\varphi) \,\} \cup \{\, (x, C_i) \mid x \in C_i \,\} \cup \{\, (\bar{x}, C_i) \mid \bar{x} \in C_i \,\}.$$

We call F_φ the argumentation framework of φ generated via the standard translation *for admissible semantics.*

Then, φ is satisfiable iff the argument φ is credulously accepted in F_φ under semantics $\sigma \in \{\mathsf{adm}, \mathsf{comp}, \mathsf{stab}\}$.

Above, Below, and Exact k Diversity

The problem $\mathsf{k\text{-}Div\text{-}Arg}_\sigma(F, k, a, b)$ asks: given two arguments a and b in an AF $F = (A, R)$ and a number $k \in \mathbb{N}$, are there two σ-extensions S_1, S_2 with $a \in S_1, b \in S_2$, and $d(S_1, S_2) = k$? We prove that this problem (i) shows higher complexity than credulous reasoning for conflict-free, naive, and preferred semantics, whereas (ii) is of same complexity as credulous reasoning under all other considered semantics.

We also consider the problem $\mathsf{k\text{-}Div\text{-}Ext}_\sigma(F, k, a, b)$ asking: given an AF $F = (A, R)$ and a number $k \in \mathbb{N}$, are there two σ-extensions S_1, S_2 with $d(S_1, S_2) = k$?

4.1 Conflict-free and Naive Semantics

We first consider the *lightest* semantics. Surprisingly, here we get hardness already for the conflict-free semantics. This case is interesting as almost all other reasoning problems (extension existence, credulous and skeptical acceptance, etc.) are decidable in polynomial time under conflict-free semantics. Nevertheless, the hardness is more involved in some cases (e.g., for below k diversity) and in fact both semantics incur different complexity in those cases (under reasonable complexity assumptions).

Theorem 6 ($\star$). *The problems* $\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ *and* ${}^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ *are* **NP***-complete for $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$.*

Proof. For membership, one guesses two witness sets $S_a, S_b \subseteq A$ such that (i) S_a and S_b are σ-extension, (ii) $a \in S_a$, $b \in S_b$, and (iii) $d(S_a, S_b) = k$. The verification requires polynomial time for each $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$. For naive, one checks conflict-freeness and additionally that adding any argument not in the set breaks the conflict-freeness. For ${}^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$, we substitute (iii) by $d(S_a, S_b) \geq k$.

For hardness, we reduce from the canonical **NP**-complete problem 3SAT. Here, the reduction in Definition 5 can not be directly used as it is based on the notion of admissibility, which both semantics $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$ lack. Our reduction builds on the idea of the standard reduction for proving the hardness of independent sets existence (IS) problem implicitly. However, one can not directly reduce from IS as we do not have control over the size of solution sets.

Let $\varphi = \bigwedge_{i \leq m} C_i$ where each $C_i = (\ell_{i1} \vee \ell_{i2} \vee \ell_{i3})$ is a clause with three literals over the set X of variables. In our AF, we only require symmetric edges. Hence we find it convenient to write $\{u, v\}$ for two attacks (u, v) and (v, u). We construct an AF $F = (A, R)$ where:

$$A := \{\ell_{ij} \mid 1 \leq i \leq m, 1 \leq j \leq 3\} \cup \{a, b\},$$

$$R := \{\{\ell_{i1}, \ell_{i2}\}, \{\ell_{i2}, \ell_{i3}\}, \{\ell_{i3}, \ell_{i1}\} \mid 1 \leq i \leq m\} \cup \{\{\ell, \ell'\} \mid \ell = x, \ell' = \bar{x} \text{ for some } x \in X\} \cup \{\{a, b\}\} \cup \{\{b, x\} \mid x \in A \setminus \{b\}\},$$

for fresh arguments $a, b \notin X$. Moreover, we let $k = m + 2$. Here, we observe that $\{b\}$ is a naive and hence also a conflict-free set in F. Intuitively, we let each position of a literal in a clause as an argument thus resulting in three arguments for each clause of φ. For attacks, we add a "triangle" of symmetric attacks between all three arguments for a clause, and further attacks between two literals of opposite parity. For correctness, we prove the following claim.

Claim 7 ($\star$). *Arguments a and b are $(m+2)$-diverse in F iff formula φ is satisfiable.*

We conclude by observing that the presented reduction also applies to the problem that asks above k diverse extensions covering a and b, that is, to $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$. $\square$

Next, we consider the problem of (general) k-diverse extensions in an AF (hence, to determine whether an AF admits two at least k-diverse extensions). The reduction in the proof above does not apply directly as F may have two σ extensions only coming from the arguments in triangles. This has the consequence that φ might not be satisfiable but still F has two k-diverse extensions. Luckily, we can still utilize the ideas and expand our reduction to obtain the following.

Theorem 8. *The problems* $\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ *and* $^{\geq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ *are* **NP***-complete for* $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$.

Proof. For membership, guess two witness sets as before.

For hardness, we reuse the reduction from the proof of Theorem 6. The intuition why the same reduction may fail is the following. An AF F might have two small conflict-free/naive sets such that they give jointly a diversity of k. However, we are looking for one satisfying assignment for all the m clause in φ. To overcome this issue, we duplicate each argument in F and ask for a larger diversity value instead.

Let $F = (A, R)$ be the AF constructed in the proof of Theorem 6 for a formula φ with m clauses. Here, we do not need arguments a and b, hence our arguments set A and attacks in R does not include these auxiliary arguments. As before, we only consider symmetric edges in our AF and write $\{u, v\}$ for attacks. We consider an AF $F' = (A', R')$ where

$$\begin{aligned} A' :=& A \cup \{s_* \mid s \in A\}, \\ R' :=& R \cup \{\{s, s_*\} \mid s \in A\} \\ & \cup \{\{s, t_*\}, \{s_*, t_*\}, \{s_*, t\} \mid \{s, t\} \in R\}. \end{aligned}$$

Intuitively, each argument s now has a *shadow* vertex s_* which participates in all the edges as s. Additionally, s_* also have edges to the shadow vertices of each neighbor t of s in F. This encodes that whenever a set $S \subseteq A$ is conflict-free in F', then its shadow set $S_* = \{s_* \mid s \in S\}$ is also conflict-free whereas neither $S \cup \{t_*\}$ nor $S_* \cup \{t\}$ is conflict-free for any $t \in S$ and its shadow $t_* \in S_*$, respectively. Moreover, in this case, the set $S \setminus \{t\} \cup \{t_*\}$ and $S_* \setminus \{t_*\} \cup \{t\}$ are both still conflict-free in F'. We first prove this claim in the following. For brevity, we denote $\ell_* = a_*$ if $\ell = a$ and $\ell_* = a$ for $\ell = a_*$ (intuitively, the shadow of a shadow argument is the argument).

Claim 9 ($\star$). *(I) S is conflict-free (naive) in F' iff S_* is conflict-free (naive) in F' for any $S \subseteq A$. (II) For any conflict-free set $L \subseteq A'$ and $\ell \in L$: the sets $(L \setminus \{\ell\}) \cup \{\ell_*\}$ is also conflict-free.*

We next prove the correctness via the following claim.

Claim 10 ($\star$). *φ is satisfiable if and only if F admits two σ-extensions which are $2m$-diverse for $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$.*

We conclude by observing that the reduction also applies to $^{\geq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ for both semantics $\sigma \in \{cnf, \mathsf{naiv}\}$. $\square$

We next turn to the problem of asking whether two arguments are *below* k diverse. Here, the hardness holds only for naive semantics as we establish next. Observe that the reduction from before does not apply here as in Claim 3: a "small" naive set in F may yield two above-k-diverse extensions for a and b while the corresponding literals still do not satisfy all the clauses (e.g., when a witness naive set for a has size $< m$). In fact, this also explains why the hardness does not apply to the conflict-free semantics.

Theorem 11. *The problem* $^{\leq}\mathsf{k\text{-}Div\text{-}Arg}_{\mathsf{naiv}}$ *is* **NP***-complete.*

Proof. For membership, guess two witness sets as before.

For hardness, we reduce from the dominating set problem. Given a graph $G = (V, E)$ and an integer $k \in \mathbb{N}$: the question asks whether there are k vertices $S \subseteq V$ such that for every vertex $v \in V$, either $v \in S$ or there exists $u \in S$ with $\{u, v\} \in E$. We consider an AF $F = (A, R)$ with symmetric attacks where, for fresh $a, b \notin V$: $\quad A := V \cup \{a, b\}$,

$$R := \{\{x, y\} \mid \{x, y\} \in E\} \cup \{\{a, b\}, \{b, x\} \mid x \in V\}.$$

Then, we prove the following claim:

Claim 12 ($\star$). *G admits a dominating set of size k iff arguments a, b are below $(k+2)$-diverse in F (naive semantics).*

$\square$

The reduction does not work for conflict-free semantics, but turns out to be trivial as any two non-self-attacking arguments are trivially below k-diverse for any $2 \leq k \in \mathbb{N}$.

Remark 13. *Let $F = (A, R)$ be an AF and $a, b \in A$ be two arguments such that none of these arguments have a self-attack. Then, considering conflict-free semantics a and b are at least ($\geq$) 2-diverse iff either $(a, b) \in R$ or $(b, a) \in R$. Thus, any two arguments in an AF (without self-attacks) are trivially (below) k-diverse for any $k \neq 1$. The answer for $k = 1$ depends on whether two arguments are conflicting. In fact, no two conflicting arguments can be 1-diverse in any AF.*

4.2 Admissible, Complete, and Stable Semantics

We now consider semantics $\sigma \in \{\mathsf{adm}, \mathsf{comp}, \mathsf{stab}\}$ and prove that the complexities are similar to credulous reasoning. That is, reasoning on diverse arguments and extensions does not incur additional complexity compared to argument acceptance.

Theorem 14 ($\star$). *Problems* $\mathsf{k\text{-}Div\text{-}Arg}_\sigma$, $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$*, and* $^{\leq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ *are* **NP***-complete for* $\sigma \in \{\mathsf{adm}, \mathsf{comp}, \mathsf{stab}\}$.

Details are given in the appendix, as we mostly adapt standard reductions. We next establish that the problem with diverse extensions admits the same complexity.

Theorem 15 ($\star$). *Problems* $\mathsf{k\text{-}Div\text{-}Ext}_\sigma$, $^{\geq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$*, and* $^{\leq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ *are* **NP***-complete for* $\sigma \in \{\mathsf{adm}, \mathsf{comp}, \mathsf{stab}\}$.

NP-hardness for preferred semantics follows from admissible semantics. However, we do not have membership, as preferred extensions cannot be verified in polynomial time.

4.3 Semi-stable, Stage, and Preferred Semantics

The following lemma is essential for proving claims regarding semi-stable and stage semantics. Let $F = (A, R)$ and $a, b \in A$ be two arguments. We say that a and b are *R-equivalent* in F if they have identical attacks in F, i.e., if $(a, c) \in R \iff (b, c) \in R$ and $(c, a) \in R \iff (c, b) \in R$ for any $c \in A$.

Lemma 16 ($\star$). *Let $F = (A, R)$ be an AF and $a, b \in A$ be two R-equivalent arguments in F; for any $S \subseteq A$ with $a \in S$:*

1. *S is conflict-free (admissible) in F iff $(S \setminus \{a\}) \cup \{b\}$ is conflict-free (admissible) in F.*
2. *if (a, b) and $(b, a) \in R$, then $S^+_R = S'^+_R$ for $S' = (S \setminus \{a\}) \cup \{b\}$.*

Theorem 17. *The problems* $\mathsf{k\text{-}Div\text{-}Arg}_\sigma$*,* $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ *and* $^{\leq}\mathsf{k\text{-}Div\text{-}Arg}_\sigma$ *are $\boldsymbol{\Sigma}_2^{\mathbf{P}}$-complete for $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$.*

Proof. For membership, observe that one can guess two sets of arguments $S_a, S_b \subseteq A$ and verify that both are σ-extensions, k-diverse, and contain a and b, respectively. The verification for σ-extensions requires an **NP** oracle to decide range-maximality. Moreover, the same also applies to extensions with diversity $\leq k$ and $\geq k$.

We next turn to prove hardness. To this aim, we utilize an existing reduction by [Dvořák and Woltran, 2010] proving credulous acceptance under both considered semantics. 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 let $Y_* = \{y_* \mid y \in Y\}$ and $\bar{Y}_* = \{\bar{y}_* \mid y \in Y\}$ as auxiliary sets of arguments and consider the AF $F_\Phi = (A, R)$ constructed as follows:

$$
\begin{aligned}
A :=& \{\varphi, \bar{\varphi}, b\} \cup \{C_1, \ldots, C_m\} \cup X \cup \bar{X} \cup Y_* \cup \bar{Y}_* \\
R :=& \{\, (C_i, \varphi) \mid i \leq m \,\} \cup \{\, (x, \bar{x}), (\bar{x}, x) \mid x \in X \,\} \\
& \cup \{\, (x, C_i) \mid x \in C_i \,\} \cup \{\, (\bar{x}, C_i) \mid \bar{x} \in C_i \,\} \\
& \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi), (\varphi, b), (b, b)\} \\
& \cup \{(y, y_*), (\bar{y}, \bar{y}_*), (y_*, y_*), (\bar{y}_*, \bar{y}_*) \mid y \in Y\}.
\end{aligned}
$$

It is known that Φ is true iff φ is contained in each semi-stable (stage) extension of F_Φ iff $\bar{\varphi}$ is not contained in any semi-stable (stage) extension of F_Φ. To complete the reduction, we construct an AF $F' = (A', R')$ by taking a fresh argument $\bar{\varphi}_c \notin A$, such that $A' := A \cup \{\bar{\varphi}_c\}$,

$$R' := R \cup \{(\varphi, \bar{\varphi}_c), (\bar{\varphi}_c, \varphi), (\bar{\varphi}, \bar{\varphi}_c), (\bar{\varphi}_c, \bar{\varphi})\}.$$

Intuitively, the argument $\bar{\varphi}_c$ serves as a copy of $\bar{\varphi}$. Whereby, it is accepted in some σ-extension iff $\bar{\varphi}$ is accepted in some σ-extension for $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$ due to Lemma 16. We prove correctness of our reduction in the following claim.

Claim 18 ($\star$). *Φ is false iff argument $\bar{\varphi}$ is contained in some σ-extension of F_Φ for $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$ iff arguments $\bar{\varphi}$ and $\bar{\varphi}_c$ are 2-diverse in F' under σ.*

Note that the reduction in the proof above also applies to the case of $\leq k$- and $\geq k$-diverse arguments for $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$. This holds as one can still take $k = 2$. $\square$

For deciding diverse extensions, the reduction in the proof of Theorem 17 does not directly apply for $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$. This holds due to the following issue: when the CNF φ in Φ is satisfiable, F_Φ already admits a σ-extension of size $|X| + 1$ where $X = Y \cup Z$ and hence $|X| > |Y|$. It might happen that Φ admits two satisfying assignments which only differ at one literal. Thus, their corresponding σ-extensions also differ at only one argument and hence are 2-diverse. Nevertheless, we can expand our reduction to consider sufficiently many copies of the argument $\bar{\varphi}$ in F to avoid this case. The proof of the following theorem formalizes this intuition. Interestingly, this trick with multiple copies does not work in the case of $^{\leq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$. This holds since smaller than k diversity is already achievable via any satisfying assignments for φ.

Theorem 19 ($\star$). *The problems* $\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ *and* $^{\geq}\mathsf{k\text{-}Div\text{-}Ext}_\sigma$ *are $\boldsymbol{\Sigma}_2^{\mathbf{P}}$-complete for semantics $\sigma \in \{\mathsf{semiSt}, \mathsf{stag}\}$.*

We next turn towards preferred semantics which possesses the same complexity, modulo a different reduction.

Theorem 20 ($\star$). *The problem* $\mathsf{k\text{-}Div\text{-}Arg}_{\mathsf{pref}}$ *and* $^{\leq}\mathsf{k\text{-}Div\text{-}Arg}_{\mathsf{pref}}$ *are both $\boldsymbol{\Sigma}_2^{\mathbf{P}}$-complete.*

However, the precise complexity of $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_{\mathsf{pref}}$ is open. The issue why the above reduction does not work for $^{\geq}\mathsf{k\text{-}Div\text{-}Arg}_{\mathsf{pref}}$ is the following. If φ is satisfiable, one can construct an extension containing $\bar{b}$ from a satisfying assignment, of size $|X| + 2$, which together with $\{b\}$ gives two $(|X| + 3)$-diverse extensions which is above k anyway.

We also extablish the complexity for k-diverse extensions.

Theorem 21 ($\star$). *The problem* $\mathsf{k\text{-}Div\text{-}Ext}_{\mathsf{pref}}$ *is $\boldsymbol{\Sigma}_2^{\mathbf{P}}$-complete.*

4.4 Maximal Diversity

Finally, we compute the threshold for maximum diversity. Precisely, we ask whether an AF has k-diverse σ-extensions (for a pair of arguments) but not k'-diverse for any $k' > k$.

We reduce SAT-UNSAT (for QBF) to the considered semantics by encoding two formulas into one AF and choosing a diversity threshold K that separates them. The construction ensures that the AF has two k-diverse extensions iff the first formula is satisfiable (true), and no k'-diverse extensions for $k' > k$ iff the second formula is unsatisfiable (false).

Theorem 22 ($\star$). *The problem* $\mathsf{max\text{-}k\text{-}Div\text{-}Ext}_\sigma$ *is* **DP**-*complete for each $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$.*

We next prove that the argument version of the diversity problem is also **DP**-complete for $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$. Our construction works similar as in the proof of Theorem 22.

Theorem 23 ($\star$). *The problem* $\mathsf{max\text{-}k\text{-}Div\text{-}Arg}_\sigma$ *is* **DP**-*complete for each $\sigma \in \{\mathsf{conf}, \mathsf{naiv}\}$.*

Next, we consider the connection between extensions.

Lemma 24 ($\star$). *Let $F_1 = (A_1, R_1)$, $F_2 = (A_2, R_2)$ be two mutually disjoint AFs, and let $F = (A_1 \cup A_2, R_1 \cup R_2)$. Then, for each semantics σ and sets $S_i \subseteq A_i$: $S_1 \cup S_2$ is a σ-extension in F iff S_i is a σ-extension in F_i for $i = 1, 2$.*

By Lemma 24, when considering disjoint AFs, their extensions are sub-extensions in the individual sub-AFs.

We next establish that the argument version of the diversity problem is **DP**-complete for $\sigma \in \{\mathsf{adm}, \mathsf{comp}, \mathsf{stab}\}$.

Semantics	Maximum diversity ($k_{\max}$) avg(k)	med(k)	avg(t) (s)	Minimum diversity ($k_{\min}$) avg(k)	med(k)	avg(t) (s)
adm	102.352	71.0	1.591	1.0	1.0	0.91
comp	88.683	77.0	1.359	6.894	1.0	0.907
stab	101.5	93.0	0.911	12.478	2.0	0.587

Table 2: Summary of experimental results for AFs with at most 1000 arguments. For each semantics, we report the average and median diversity values and the average runtime (in seconds) for both maximal ($k_{\max}$) and minimal ($k_{\min}$) variants.

Theorem 25 ($\star$)**.** *The problem* max-k-Div-Arg$_\sigma$ *is* **DP**-*complete for each* $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$.

Also the extension-version of the problem is **DP** and **DP**$_2$-c.

Theorem 26 ($\star$)**.** *The problem* max-k-Div-Ext$_\sigma$ *is*

1. **DP**-*complete for each* $\sigma \in \{\text{adm}, \text{comp}, \text{stab}\}$ *and*
2. **DP**$_2$-*complete for semantics* $\sigma \in \{\text{semiSt}, \text{stag}\}$.

A Note On Minimal Diversity. Since diversity is neither monotonic, nor monotonic, the problem of determining the smallest k for which an AF admits two k-diverse σ-extensions is not complementary to its maximum counterpart. Moreover, our hardness-reductions in Section 4.4 does not immediately extend to the case of minimal diversity.

5 Experiments

To investigate practical feasibility of diversity, we implemented our approach and conducted experiments.

Implementation. We base our implementation on an ASP-encoding and the argumentation solver ASPARTIX [Egly *et al.*, 2008; Egly *et al.*, 2010; Dvorák *et al.*, 2021] for rapid prototyping. Our encoding computes, given an AF F, two extensions that are the minimally or maximally k-diverse in F. The core computational component of our approach is an ASP encoding implementing the proposed diversity computation, which we solve using clingo (5.8.0). We place a simple wrapper to orchestrate the experiments, and collect runtime and optimization statistics. For each instance and semantics, we record the runtime required to solve the corresponding optimization problem, as well as the resulting diversity value k. We consider both the maximal and minimal diversity variants.

Instances and Setup. As benchmark data, we use AFs from the ICCMA 2025 benchmark suite. For each benchmark instance, we evaluate the proposed diversity computation under the admissible, complete, and stable semantics. Each run was subject to a fixed timeout of 300 seconds. Our analysis focuses on instances for which an optimal solution was found within the given timeout. To avoid the influence of a small number of extremely large AFs, we further restrict the evaluation to AFs with at most 1000 arguments, leaving us with 270 instances. We conduct the experiments on a desktop machine equipped with an Intel i9-11900K CPU running Windows 10, which is sufficient as we are primarily interested in obtaining the maximum k and minimum k for the considered frameworks.

Result. In total, we evaluated 270 instances, each under three semantics. Among these, 145 instances reached optimality within the timeout under at least one semantics, while the remaining instances either where unsatisfiable or exceeded the time limit under all semantics. Table 2 provides a summary of the results aggregated over all instances in this range, reporting the average and median diversity values, as well as the average runtime per semantics, for both maximal and minimal variants. To analyze how the minimum and maximum diversity distributes, we provide Figure 3 that plots per instance (X-axis) the min (orange), max (blue), and median (dashed green) k-diversity value (Y-axis).

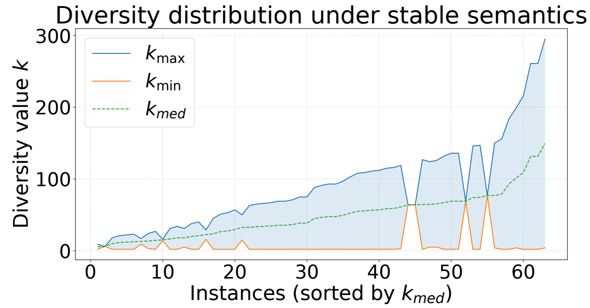

Figure 3: Distribution of min, med, and max diversity for instances under stable semantics. The instances are sorted by median diversity.

Discussion. The experimental results reveal characteristic patterns in how diversity values are distributed under the considered semantics. Figure 3 shows that, diversity values vary substantially across instances, indicating that the structure of extensions can differ markedly even within the same AF. In particular, the instance-wise distribution reveals that for a large portion of instances, extensions remain relatively close to each other, while a smaller set of instances admits extensions that are structurally far apart. This suggests that high diversity is not uniformly present, but rather concentrated in specific frameworks that allow for fundamentally different solutions. Distributions for further semantics are in the appendix.

6 Conclusion

We introduce a quantitative notion of *extension diversity* for abstract argumentation based on symmetric difference, which makes explicit how strongly alternative, internally coherent viewpoints in an AF diverge. We provide a *comprehensive complexity classification* for reasoning about existence of k-diverse extensions, diversity covering specific arguments, and optimization variants. We implement a k-diversity computation and conduct an initial experimental study, which indicates practical feasibility for medium-sized frameworks.

For future work, we are interested in completing the complexity landscape for open cases (e.g., for minimal k-diversity). To tame the high complexity, one could consider parameterized complexity or structure-aware encoding into Boolean satisfiability [Fichte *et al.*, 2021; Mahmood *et al.*, 2026]. Moreover, we propose diversity as a useful measure to consistent query answering in data- or knowledge-bases, following the natural connections to argumentation [Mahmood *et al.*, 2024; Mahmood *et al.*, 2025]. Finally, extending the diversity notion to logic-based argumentation [Mahmood *et al.*, 2023; Hecher *et al.*, 2024] is another venue for future work. Finally we are interested in applications and use-cases of diversity to check disagreement and robustness of real-world instances.

Acknowledgments

Authors are stated in alphabetical order. This research was funded in whole or in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), grant TRR 318/3 2026 – 438445824 and the Ministry of Culture and Science of North Rhine-Westphalia (MKW NRW) within projects WHALE (LFN 1-04) funded under the Lamarr Fellow Network programme and project SAIL, grant NW21-059D, the Austrian Science Fund (FWF) grant J4656, the French Agence nationale de la recherche (ANR) grant ANR-25-CE23-7647; and Excellence Center at Linköping – Lund in Information Technology (ELLIIT) funded by the Swedish government and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

[Alviano, 2018] Mario Alviano. The Pyglaf Argumentation Reasoner. In *ICLP*, volume 58, pages 2:1–2:3. Dagstuhl Publishing, 2018.

[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.

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

[Baumann, 2012] Ringo Baumann. What does it take to enforce an argument? minimal change in abstract argumentation. In *ECAI*, pages 127–132. IOS Press, 2012.

[Bernreiter *et al.*, 2024] Michael Bernreiter, Jan Maly, Oliviero Nardi, and Stefan Woltran. Combining voting and abstract argumentation to understand online discussions. In *AAMAS*, pages 170–179, 2024.

[Besnard *et al.*, 2022] Philippe Besnard, Sylvie Doutre, Théo Duchatelle, and Marie-Christine Lagasquie-Schiex. Explaining semantics and extension membership in abstract argumentation. *Intelligent Systems with Applications*, 16:200118, 2022.

[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.

[Böhl *et al.*, 2023] Elisa Böhl, Sarah Alice Gaggl, and Dominik Rusovac. Representative answer sets: Collecting something of everything. In *ECAI*, pages 271–278. IOS Press, 2023.

[Bondy and Murty, 2008] J. Adrian Bondy and Uppaluri S. R. Murty. *Graph Theory*. Graduate Texts in Mathematics. Springer, 2008.

[Coste-Marquis *et al.*, 2015] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. Extension enforcement in abstract argumentation as an optimization problem. In *IJCAI*, pages 2876–2882. AAAI Press, 2015.

[Čyras *et al.*, 2021] Kristijonas Čyras, Antonio Rago, Emanuele Albini, Pietro Baroni, and Francesca Toni. Argumentative XAI: A survey. In *IJCAI*, pages 4392–4399. ijcai.org, 8 2021.

[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*, pages 116–127. IOS Press, 2022.

[Dauphin *et al.*, 2018] Jérémie Dauphin, Marcos Cramer, and Leon Van der Torre. Abstract and concrete decision graphs for choosing extensions of argumentation frameworks. In *COMMA*, pages 437–444. IOS Press, 2018.

[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.

[Dvorák and Dunne, 2017] Wolfgang Dvorák and Paul E. Dunne. Computational problems in formal argumentation and their complexity. *J. Logics and their Applications*, 4(8), 2017.

[Dvořák and Woltran, 2010] Wolfgang Dvořák and Stefan Woltran. Complexity of semi-stable and stage semantics in argumentation frameworks. *Information Processing Letters*, 110(11):425–430, 2010.

[Dvorák *et al.*, 2021] Wolfgang Dvorák, Matthias König, Johannes Peter Wallner, and Stefan Woltran. Aspartix-V21. *CoRR*, abs/2109.03166, 2021.

[Egly *et al.*, 2008] Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. ASPARTIX: Implementing argumentation frameworks using answer-set programming. In *ICLP*, pages 734–738. Springer, 2008.

[Egly *et al.*, 2010] Uwe Egly, Sarah Alice Gaggl, and Stefan Woltran. Answer-set programming encodings for argumentation frameworks. *Argument Comput.*, 1(2):147–177, 2010.

[Eiter *et al.*, 2013] Thomas Eiter, Esra Erdem, Halit Erdogan, and Michael Fink. Finding similar/diverse solutions in answer set programming. *Theory Pract. Log. Program.*, 13(3):303–359, 2013.

[Fichte *et al.*, 2021] Johannes Klaus Fichte, Markus Hecher, Yasir Mahmood, and Arne Meier. Decomposition-guided reductions for argumentation and treewidth. In Zhi-Hua Zhou, editor, *IJCAI*, pages 1880–1886. ijcai.org, 2021.

[Fichte *et al.*, 2022] Johannes K. Fichte, Markus Hecher, and Mohamed A. Nadeem. Plausibility reasoning via projected answer set counting - a hybrid approach. In *IJCAI*, pages 2620–2626. ijcai.org, 2022.

[Fichte *et al.*, 2023] Johannes Klaus Fichte, Markus Hecher, Yasir Mahmood, and Arne Meier. Quantitative reasoning and structural complexity for claim-centric argumentation. In *IJCAI*, pages 3212–3220. ijcai.org, 2023.

[Fichte *et al.*, 2024] Johannes K. Fichte, Markus Hecher, and Arne Meier. Counting complexity for reasoning in abstract argumentation. *Journal of Artificial Intelligence Research*, 80, 2024.

[Fichte *et al.*, 2025] Johannes K. Fichte, Nicolas Fröhlich, Markus Hecher, Victor Lagerkvist, Yasir Mahmood, Arne Meier, and Jonathan Persson. Facets in argumentation: A formal approach to argument significance. In *IJCAI*, pages 4491–4499, 2025.

[Fichte *et al.*, 2026] Johannes K. Fichte, Markus Hecher, Yasir Mahmood, and Zhengjun Wang. Diversity of extensions in abstract argumentation, 2026.

[Hebrard *et al.*, 2005] Emmanuel Hebrard, Brahim Hnich, Barry O'Sullivan, and Toby Walsh. Finding diverse and similar solutions in constraint programming. In *AAAI*, pages 372–377. AAAI Press, 2005.

[Hecher *et al.*, 2024] Markus Hecher, Yasir Mahmood, Arne Meier, and Johannes Schmidt. Quantitative claim-centric reasoning in logic-based argumentation. In *IJCAI*, pages 3404–3412. ijcai.org, 2024.

[Liao and van der Torre, 2020] Beishui Liao and Leendert van der Torre. Explanation semantics for abstract argumentation. In *COMMA*, pages 271–282. IOS Press, 2020.

[Lohrey and Rosowski, 2023] Markus Lohrey and Andreas Rosowski. On the Complexity of Diameter and Related Problems in Permutation Groups. In *ICALP*, pages 134:1–134:18. Dagstuhl Publishing, 2023.

[Mahmood *et al.*, 2023] Yasir Mahmood, Arne Meier, and Johannes Schmidt. Parameterized complexity of logic-based argumentation in schaefer's framework. *ACM Trans. Comput. Log.*, 24(3):26:1–26:25, 2023.

[Mahmood *et al.*, 2024] Yasir Mahmood, Jonni Virtema, Timon Barlag, and Axel-Cyrille Ngonga Ngomo. Computing repairs under functional and inclusion dependencies via argumentation. In *FoIKS*, pages 23–42. Springer, 2024.

[Mahmood *et al.*, 2025] Yasir Mahmood, Markus Hecher, and Axel-Cyrille Ngonga Ngomo. Dung's argumentation framework: Unveiling the expressive power with inconsistent databases. In *AAAI*, pages 15058–15066. AAAI Press, 2025.

[Mahmood *et al.*, 2026] Yasir Mahmood, Markus Hecher, Johanna Groven, and Johannes Klaus Fichte. Structure-aware encodings of argumentation properties for clique-width. In *AAAI*, pages 19294–19302. AAAI Press, 2026.

[Niskanen and Järvisalo, 2020a] Andreas Niskanen and Matti Järvisalo. Smallest Explanations and Diagnoses of Rejection in Abstract Argumentation. In *KR*, pages 667–671, 9 2020.

[Niskanen and Järvisalo, 2020b] Andreas Niskanen and Matti Järvisalo. μ-toksia: An Efficient Abstract Argumentation Reasoner. In *KR*, pages 800–804, 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*, pages 1949–1955. IJCAI Organization, 2018.

[Rahwan, 2007] Iyad Rahwan. Argumentation in artificial intelligence. *Artificial Intelligence*, 171(10-15):619–641, 2007.

[Rodrigues, 2019] Odinaldo Rodrigues. Representing and comparing large sets of extensions of abstract argumentation frameworks. In *SAC*, pages 1154–1161. ACM, 2019.

[Saribatur *et al.*, 2020] Zeynep G. Saribatur, Johannes P. Wallner, and Stefan Woltran. Explaining non-acceptability in abstract argumentation. In *ECAI*, pages 881–888. IOS Press, 2020.

[Skiba *et al.*, 2021] Kenneth Skiba, Tjitze Rienstra, Matthias Thimm, Jesse Heyninck, and Gabriele Kern-Isberner. Ranking extensions in abstract argumentation. In Zhi-Hua Zhou, editor, *IJCAI*, pages 2047–2053, 2021.

[Stockmeyer and Meyer, 1973] Larry J. Stockmeyer and Albert R. Meyer. Word problems requiring exponential time. In *STOC*, pages 1–9. ACM, 1973.

[Stockmeyer, 1976] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):1–22, 1976.

[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.

[Ulbricht and Wallner, 2021] Markus Ulbricht and Johannes P. Wallner. Strong explanations in abstract argumentation. *AAAI*, 35(7):6496–6504, May 2021.

[Wrathall, 1976] Celia Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.