Quantitative Claim-Centric Reasoning in Logic-Based Argumentation

Markus Hecher[1] , **Yasir Mahmood**[2] , **Arne Meier**[3] , **Johannes Schmidt**[4]

[1]CSAIL, Massachusetts Institute of Technology, United States
[2]DICE group, Department of Computer Science, Paderborn University, Germany
[3]Institut für Theoretische Informatik, Leibniz Universität Hannover, Germany
[4]Department of Computer Science and Informatics, Jönköping University, Sweden
hecher@mit.edu, yasir.mahmood@uni-paderborn.de,
meier@thi.uni-hannover.de, johannes.schmidt@ju.se

Abstract

Argumentation is a well-established formalism for nonmonotonic reasoning with popular frameworks being Dung's abstract argumentation (AFs) or logic-based argumentation (Besnard-Hunter's framework). Structurally, a set of formulas forms support for a claim if it is consistent, subset-minimal, and implies the claim. Then, an argument comprises a support and a claim. We observe that the computational task (ARG) of asking for support of a claim in a knowledge base is "brave", since many claims with a single support are accepted. As a result, ARG falls short when it comes to the question of confidence in a claim, or claim strength. In this paper, we propose a concept for measuring the (acceptance) strength of claims, based on counting supports for a claim. Further, we settle classical and structural complexity of counting arguments favoring a given claim in propositional knowledge bases (KBs). We introduce quantitative reasoning to measure the strength of claims in a KB and to determine the relevance strength of a formula for a claim.

1 Introduction

Argumentation is a nonmonotonic formalism in artificial intelligence around which an active research community has evolved [Atkinson *et al.*, 2017; Amgoud and Prade, 2009; Rago *et al.*, 2018; Baroni *et al.*, 2018]. The popular frameworks dealing with the theory of argumentation include the *abstract* [Dung, 1995] and the *logic-based* [Besnard and Hunter, 2001; Besnard and Hunter, 2008; Chesnevar *et al.*, 2000; Prakken and Vreeswijk, 2002] approach. The abstract setting, also known as Dung's theory of argumentation, focuses on formalizing the interaction between arguments in a graph-theoretic way. In other words, arguments are nodes in a directed graph, and the 'attack relation' determines the interaction between arguments. Given such a framework, one is interested in computing the sets of arguments (called *extensions*) simultaneously accepted under a given semantics.

In the logic-based setting, one starts with a knowledge base (KB) and searches for inclusion-minimal and consistent sets of formulas Φ (the *support*) that logically entail a claim α. The pair (Φ, α) is then called an *argument* for the claim α. The computational problem of interest is the existence of an argument supporting a claim (ARG), and the relevance problem (ARG-Rel) is to determine whether a formula ψ is *relevant* to the claim α. Formally, ARG asks, given a set of formulas Δ (the *knowledge base*) and a formula α, whether there exists a subset $\Phi \subseteq \Delta$ such that (Φ, α) is an argument in Δ, and ARG-Rel asks, whether there is an argument (Φ, α) in Δ such that $\psi \in \Phi$. The complexity of both reasoning problems is well understood for KBs formulated in propositional logic [Parsons *et al.*, 2003] as well as in fragments thereof [Creignou *et al.*, 2011; Creignou *et al.*, 2014]. In addition, a structural complexity analysis has recently been employed to better understand the source of intractability [Mahmood *et al.*, 2023] using different parts of the input as parameters.

Despite the significant interest in the reasoning problems for logic-based argumentation, no prior research has explored the associated counting complexity. We argue that counting the number of arguments in support of a given claim is a logical choice when investigating the extent to which a claim is favored within the given knowledge base.

Example 1. *We want to evaluate the arguments for and against candidates in an election. The argument A favors candidate X based on her education policy, which leads us to accept that "X is favorable". Another argument (B) opposes this claim targeting the health care policy. Therefore, both claims are equally likely and have an acceptance strength of 50%. Subsequently, an argument C supports X in foreign policy. This increases the probability and thus the strength of the claim that X is indeed a favorable candidate.*

Reasoning problems focusing on claims (instead of arguments) have received growing interest in abstract argumentation [Dvořák and Woltran, 2020; Dvořák *et al.*, 2023; Bernreiter *et al.*, 2023; Baumann *et al.*, 2023]. Specifically, claim-augmented argumentation frameworks (CAF) [Dvořák and Woltran, 2020] extend Dung's AF by employing a conclusion-oriented perspective. Further, with the aim of exploring the justification status of claims, Fichte *et al.* (2023) introduced a quantitative measure of extensions supporting a particular claim. The increasing interest in the abstract setting focused on claim-based reasoning naturally motivates the exploration of claim strength in logic-based argumentation.

This leads to the definition of a quantitative measure of the arguments supporting a given claim in the knowledge base.

Contributions. As our first contribution, we categorize the classical and parameterized complexity of counting arguments to a claim (#ARG), as well as counting arguments where a given formula is also relevant (#ARG-Rel). We prove that both these problems are intractable and $\# \cdot$ **coNP**-complete. Then we turn to parameterized complexity, which considers different parameters. We observe that both problems become **FPT** when parameterized by the number of variables in Δ or the treewidth of the primal graph [Flum and Grohe, 2006] of the input KB. Table 1 gives a summary. For omitted proof details (marked by $\star$), we refer to the appendix.

We introduce a quantitative measure to determine the strength or likelihood of a claim in the given KB. Note that the argument existence problem (ARG) concerns only whether a claim is accepted in a given argument, and does not convey any further information about the given claim, such as how likely it is to be accepted within a given KB. As the following example illustrates, in certain corner cases, the problem ARG in logic-based argumentation can accommodate *disputing* claims that may not be acceptable even under the most lenient (admissible: defended against all attacking arguments) argumentation semantics. Accepting claims based solely on support in a knowledge base leads to the undesirable consequence of being unable to distinguish between weak and unlikely claims and those that are more likely.

Example 2. *Let* $\Delta = \{x\} \cup \{\neg x \wedge \phi_i \mid i \leq n\}$ *be a KB where each* ϕ_i *is a CNF. Then the claim* x *is accepted since* $A = (\{x\}, x)$ *is an argument. Now, for* $y \neq x$, *any argument that supports* y *in* Δ *also supports* $y \wedge \neg x$. *In other words, a single argument* A *over* Δ *claims* x *and every other argument* B *claims* $\neg x$. *Consequently, while both* x *and* $\neg x$ *are acceptable (as per* ARG*),* $\neg x$ *is clearly* more likely *than* x.

The main ingredient for our quantitative reasoning involves counting the arguments for a claim. Then, counting the relevant arguments for a claim with respect to a formula ψ informs how relevant ψ is for the given claim. We establish that one can also determine the claim strength in **FPT**-time when parameterized by the treewidth of an extended primal graph.

Related Work. Probabilistic argumentation [Hunter and Thimm, 2017; Fazzinga *et al.*, 2015; Alfano *et al.*, 2020] focuses on weighting arguments with probabilities. Furthermore, a similar work in logic-based argumentation [Hunter, 2012; Hunter, 2013; Riveret *et al.*, 2007] provides a probability distribution for arguments to model the degree of an argument being true. This whole research area is orthogonal to our setting, since we focus on the (conditional) plausibility of claims, as in the Bayesian setting. Further, our approach does not focus on the *uncertainty* of arguments, since we evaluate the strength of claims corresponding to *certain* arguments.

Existing research on the justification status of arguments [Wu and Caminada, 2010; Baroni *et al.*, 2016; Baroni and Riveret, 2019] explores *different levels* of acceptance or rejection for arguments. In addition, Baroni *et al.* (2016) also emphasized *statement justification* as a topic of independent interest. Nevertheless, their approach explores the status of an argument (or statement) in different extensions by finding all

Parameter	Problem	Complexity	Ref.
–	#ARG(-Rel)	$\# \cdot$ **coNP**-c.	T. 12, C. 13
$\mid \mathsf{var}(\Delta) \mid$	#ARG(-Rel)	**FFPT**	C. 14
$\mid \mathsf{var}(\alpha) \mid$	#ARG(-Rel)	#**W[1]**-h.	C. 14
Problem	**Runtime UB**	**Runtime LB**	**Ref.**
#ARG(-Rel)	$\exp(2, \mathcal{O}(k))$	$\exp(2, o(k))$	T. 19, C. 20/22

Table 1: Results overview. For the runtime bounds, $k = \mathsf{tw}(\mathcal{G}_\mathcal{I}^e)$ and the polynomial factors in $\mid \mathsf{var}(\Delta) \cup \mathsf{var}(\alpha) \mid$ have been omitted.

the acceptance labels for an argument (statement). Recently, Fichte *et al.* (2023) introduced a fine-grained mode between *credulous* and *skeptical* reasoning (some, respectively, any extension covering a claim) for the acceptability of claims. They propose assigning a probability value in $[0, 1]$ depending on how many extensions in an AF cover the given claim. In this work, our focus lies (as the name suggests) in the quantitative aspect of claim justification in logic-based setting.

There exists research aiming to assess the strength of an argument based on the number of attackers and defenders of the target argument. Besnard & Hunter (2001) proposed aggregation functions to evaluate the relative strength of an argument, considering its (recursive) argument tree (attackers, their attackers, etc.). Although the authors suggest counting arguments for and against a claim, a complexity analysis for counting arguments is still missing. Similarly, Pu et al. (2015) considered the problem of counting the number of attackers and defenders for a topic argument in the abstract setting. This distinguishes their work from our logic-based approach. Gradual semantics [Amgoud *et al.*, 2018; Amgoud and Doder, 2018] are methods for evaluating the overall strength of arguments in a (weighted) argumentation framework, which assumes a set of arguments, the relationship between them, and an initial weight of each argument.

Finally, Amgoud *et al.* (2008) presented an approach to compare different arguments with the goal of evaluating the *status* for each claim (option in their terminology). Their proposed system takes as input a set of arguments together with an attack and a preference relation between arguments, and returns a status for each option. While our approach determines the acceptance status for each claim, we do not consider argumentation semantics from the AFs and focus on the logic-based approach instead. So, we assume a KB and a set of claims as input, rather than the set of arguments and the interaction between them, as required by the former.

2 Preliminaries

We assume familiarity with basic notions in complexity theory (cf. [Sipser, 1997]) and use complexity classes $\mathbf{P}, \mathbf{NP}, \mathbf{coNP}, \mathbf{NP^{NP}} = \Sigma_2^{\mathbf{P}}$. For a set S, we write $|S|$ for its *cardinality*. Abusing notation, we will use $|w|$, for a string w, to denote its *length*. If φ is a formula, $\mathsf{var}(\varphi)$ denotes its set of variables. Symbols $\top$ and $\bot$ denote the constants 1 and 0, respectively. Further, we assume a reasonable encoding computable in polynomial time encoding variables in binary.

Quantified Boolean Formulas. For a Boolean formula F, we write $F(X_1, \ldots, X_l)$ to indicate that

$X_1, \ldots, X_l \subseteq \mathsf{var}(F)$. A *quantified Boolean formula (qBf)* ϕ (in prenex normal form) is of the form $\phi = Q_1 X_1.Q_2 X_2. \cdots Q_\ell X_\ell.F(X_1, \ldots, X_\ell)$, where for $1 \leq i \leq \ell$ (ℓ is the *(quantifier) rank*), we have $Q_i \in \{\forall, \exists\}$ and $Q_i \neq Q_{i+1}$, the X_i are disjoint, non-empty sets of Boolean variables, and F is a Boolean formula. We let $\mathsf{matrix}(\phi) := F$ and say that ϕ is *closed* if $\mathsf{var}(\mathsf{matrix}(F)) = \bigcup_{i \in \ell} X_i$. The semantics are defined as $\exists x.\phi \equiv \phi[x \mapsto 1] \lor \phi[x \mapsto 0]$ and $\forall x.\phi \equiv \phi[x \mapsto 1] \land \phi[x \mapsto 0]$ for a variable x. W.l.o.g., assume that $\mathsf{matrix}(\phi) = \psi_{\mathrm{CNF}} \land \psi_{\mathrm{DNF}}$, where ψ_{CNF} is in CNF (conjunctive normal form) and ψ_{DNF} is in DNF (disjunctive normal form). Note that, depending on Q_ℓ, either ψ_{CNF} or ψ_{CNF} is optional, more precisely, ψ_{CNF} might be $\top$, if $Q_\ell = \forall$, and ψ_{DNF} is allowed to be $\top$, otherwise. The problem ℓ-QBF asks for a closed qBf $\phi = \exists X_1.\phi'$ of rank ℓ, whether $\phi \equiv 1$. The problem #ℓ-QBF asks for a closed qBf $\exists X_1.\phi$ of rank ℓ, to count assignments α to X_1 such that $\phi[\alpha] \equiv 1$. For brevity, we sometimes omit ℓ.

Counting Complexity. A *witness function* w is a mapping $w \colon \{0,1\}^* \to P^{<\omega}(\{0,1\}^*)$, i.e., mapping to a finite subset of $\{0,1\}^*$ (note that this notion can easily be generalized to arbitrary alphabets, not just the binary one). Such functions associate with the counting problem "given $x \in \Sigma^*$, find $|w(x)|$". If C is a decision complexity class then $\# \cdot C$ is the class of all counting problems whose witness function w satisfies (1.) there is a polynomial p such that for all $y \in w(x)$, we have that $|y| \leq p(|x|)$, and (2.) the decision problem "given x and y, is $y \in w(x)$?" is in C. A *parsimonious reduction* between two problems $\#A$, $\#B$ preserves the cardinality between the witness sets and is computable in **P**.

Typical counting complexity classes associated with hard counting problems are $\# \cdot \mathbf{P}$ and $\# \cdot \mathbf{coNP}$. Both classes comprise complete problems that are believed not to be polynomial time solvable. For more background on counting complexity classes we refer the reader to [Durand *et al.*, 2005].

Parameterized Complexity. We give a brief introduction to parameterized complexity [Downey and Fellows, 1999].

A *parameterized problem (PP)* is a subset of $\Sigma^* \times \mathbb{N}$ for some alphabet Σ. For an input (x, k) of a PP, we call k the *parameter* of that input. The parameterized complexity class **FPT** is the class of all parameterized problems that can be solved in time $f(k) \cdot |x|^c$ for some computable function f and constant c. An **FPT**-*reduction* from a parameterized problem A to a parameterized problem B is a **FPT**-time computable function f such that (i) $(x, k) \in A$ if and only if $f(x, k) \in B$, and (ii) if there exists some computable function g such that for all $f(x, k) = (x', k')$ we have $k' \leq g(k)$.

The complexity class **W[1]** is defined as the set of all problems that can be accepted by a tail-nondeterministic and k-bounded Turing machine[1].

Parameterized Counting. The area of parameretized counting complexity has been initiated by Flum and

Grohe [Flum and Grohe, 2004] as well as McCartin [McCartin, 2006]. Similarly, as on the decision side, the complexity class #**W[1]** is the class of all parameterized counting problems reducing to counting k-cliques in a given graph. Two good surveys in this area are [Curticapean, 2019; Haak *et al.*, 2023]. A *parameterized function* is a function $F \colon \{0,1\}^* \times \mathbb{N} \to \mathbb{N}$. If $\mathcal{C}$ is a complexity class and a parameterized function F belongs to $\mathcal{C}$, we say that F is $\mathcal{C}$-computable. Intuitively, we count the number of "solutions" of a parameterized function F and associate them directly with the function value $F(x, k)$. We let **FFPT** be the class of parameterized functions that are computable in **FPT**-time.

Tree Decompositions (TD) and Treewidth. We use rooted (directed) trees $T = (N, A)$ with a *root* $\mathrm{root}(T)$ and a node $t \in N$. Here, we denote by $\mathsf{children}(t)$ the set of all nodes t', which have an edge $(t, t') \in A$. Let $G = (V, E)$ be a graph. A *tree decomposition (TD)* of a graph G is a pair $\mathcal{T} = (T, \chi)$, where T is a rooted tree, and χ is a mapping assigning to each node t of T a set $\chi(t) \subseteq V$, called *bag*, s.t.:

1. $V = \bigcup_{t \text{ of } T} \chi(t)$ and $E \subseteq \bigcup_{t \text{ of } T} \{\{u, v\} \mid u, v \in \chi(t)\}$
2. for each s lying on any r-t-path: $\chi(r) \cap \chi(t) \subseteq \chi(s)$

Then, define $\mathsf{width}(\mathcal{T}) := \max_{t \text{ of } T} |\chi(t)| - 1$. The *treewidth* $\mathsf{tw}(G)$ of G is the minimum $\mathsf{width}(\mathcal{T})$ over all tree decompositions $\mathcal{T}$ of G. Notice that for all $v \in V$, we have a unique node t^* with $v \in \chi(t^*)$ such that either $t^* = \mathrm{root}(T)$ or there is a node t of T with $\mathsf{children}(t) = \{t^*\}$ and $v \notin \chi(t)$. Denote the node t^* by $\mathsf{last}(v)$. Finally, notice that for any fixed $w \geq 1$, one can decide in linear time if a graph has treewidth at most w and, if so, to compute a TD of width w [Bodlaender, 1996]. Here, we assume only TDs (T, χ), where for every node t of T, we have that $|\mathsf{children}(t)| \leq 2$. Such a TD, while maintaining the same width, can be computed from any TD in linear time [Bodlaender and Koster, 2008].

Treewidth and qBfs. For a qBf ϕ with $\mathsf{matrix}(\phi) = \psi_{\mathrm{CNF}} \land \psi_{\mathrm{DNF}}$, define the *primal graph* $\mathcal{G}_\phi := \mathcal{G}_{\mathsf{matrix}(\phi)}$, whose vertices are $\mathsf{var}(\mathsf{matrix}(\phi))$. Two vertices of $\mathcal{G}_\phi$ are adjoined by an edge, whenever the corresponding variables share a clause/term of $\psi_{\mathrm{CNF}}/\psi_{\mathrm{DNF}}$. Let $\exp(i, p)$ be $\exp(i - 1, 2^p)$ if $i > 0$ and p otherwise. For the following result, assume that $\mathsf{poly}(n)$ is any polynomial for given positive integer n.

Proposition 3 (Chen, 2004). *For any arbitrary qBf ϕ of quantifier rank $\ell > 0$, the problem ℓ-QBF can be solved in time $\exp(\ell, \mathcal{O}(\mathsf{tw}(\mathcal{G}_\phi))) \cdot \mathsf{poly}(|\mathsf{var}(\phi)|)$.*

Note that, under the *exponential time hypothesis (ETH)* [Impagliazzo *et al.*, 2001], one cannot significantly improve this runtime: ETH implies that SAT $=$ 1-QBF can not be decided in time better than $2^{o(|\mathsf{var}(\varphi)|)}$ for a formula φ.

Proposition 4 (Fichte *et al.*, 2020). *Under ETH, for any arbitrary qBf φ of quantifier rank $\ell > 0$, problem ℓ-QBF cannot be solved in time $\exp(\ell, o(\mathsf{tw}(\mathcal{G}_\varphi))) \cdot \mathsf{poly}(|\mathsf{var}(\varphi)|)$.*

Logic-based Argumentation (LBA). Formulas are propositional; we follow the notion of Creignou et al. [2014].

Definition 5 (Besnard and Hunter [2001]). *Given a set of formulas Φ and a formula α, one says that (Φ, α) is an* argument *(for α) if (1) Φ is consistent, (2) $\Phi \models \alpha$, and (3) Φ is subset-minimal w.r.t. (2). In case of $\Phi \subseteq \Delta$, (Φ, α) is an* argument

---

[1]A tail-nondeterministic TM reads all nondet. bits at the end of the computation while for a k-bounded one, the number of nondet. bits is limited to $f(k) \cdot \log |x|$ many for all inputs (x, k).

in Δ. *We call α the* claim, *Φ the* support *of the argument, and Δ the* knowledge-base.

We consider two problems from the area of logic-based argumentation, namely ARG, and ARG-Rel. The problem ARG asks, given a set of formulas Δ and a formula α, is there a set $\Phi \subseteq \Delta$ such that (Φ, α) is an argument in Δ? The problem ARG-Rel asks, given a set of formulas Δ, and formulas $\psi \in \Delta$ and α, is there a set $\Phi \subseteq \Delta$ with $\psi \in \Phi$ such that (Φ, α) is an argument in Δ? Given a KB Δ and formula α, then $\mathsf{Args}(\alpha, \Delta)$ denotes the set of all arguments for α in Δ. Further, given a formula ψ, then $\mathsf{Rel\text{-}Args}(\psi, \alpha, \Delta)$ denotes the set of (*relevant*) arguments for α in Δ containing ψ. We define the counting problem #ARG of computing $|\mathsf{Args}(\alpha, \Delta)|$ for input (Δ, α) and the counting variant of #ARG-Rel to compute $|\mathsf{Rel\text{-}Args}(\psi, \alpha, \Delta)|$.

TDs for LBA. If φ is a formula in CNF, then the *primal graph* $\mathcal{G}_\varphi$ is the undirected graph $G = (V, E)$ such that the vertex set V is $\mathsf{var}(\varphi)$ and $\{u, v\} \in E$ if and only if the variables u, v share a clause in φ. Similarly, if φ is in DNF, then the edges are drawn w.r.t. shared term variables instead. For a graph $G = (V', E')$, we also write $V(G)$ for V' and $E(G)$ for E'. For an LBA-instance (Δ, α) the *primal graph* $\mathcal{G}_{(\Delta, \alpha)}$, is defined as the union of the primal graphs of each formula from $\Delta \cup \{\alpha\}$. That is, the vertex set $V(\mathcal{G}_{(\Delta, \alpha)}) = \bigcup_{\varphi \in \Delta \cup \{\alpha\}} V(\mathcal{G}_\varphi)$ and $E(\mathcal{G}_{(\Delta, \alpha)}) = \bigcup_{\varphi \in \Delta \cup \{\alpha\}} E(\mathcal{G}_\phi)$.

Decomposition-Guided Reductions. Inspired by work of Fichte et al. (2021), a *decomposition-guided (DG) reduction* $\mathcal{R}$ is a function that takes both a problem instance $\mathcal{I}$ and a TD $\mathcal{T} = (T, \chi)$ of $\mathcal{G}_\mathcal{I}$, and returns a qBf φ with $\mathsf{matrix}(\phi) := F$. A DG reduction has to yield a TD $\mathcal{T}' = (T, \chi')$ of $\mathcal{G}_F$. Such a reduction has to construct φ from a TD node's point of view. In that point of view, for any node t of T, the constructed bag $\chi'(t)$ functionally depends on the original bag $\chi(t)$, but also on the constructed bags $\chi'(t_1), \ldots, \chi'(t_o)$ of its child nodes $\{t_1, \ldots, t_o\} = \mathsf{children}(t)$. This results in a function f with an input of a TD node t, its bag $\chi(t)$, and a set $\chi'(\mathsf{children}(t)) = \{ \chi'(t_i) \mid t_i \in \mathsf{children}(t) \}$ of constructed bags for the child nodes of t. It follows that, as the $\mathsf{width}(\mathcal{T})$ is bounded by $\mathcal{O}(\max_{t \text{ of } T}(|\chi(t)|))$, also the treewidth of the resulting qBf is at most $\mathcal{O}(\max_{t \text{ of } T}\{|f(t, \chi(t), \chi'(\mathsf{children}(t))|\})$.

2.1 Quantitative Claim-Centric Reasoning

We propose a mode of reasoning that captures the quantitative (probabilistic) aspects of claim acceptance in logic-based argumentation. Accordingly, rather than asking whether a claim is accepted or not, a natural reasoning problem is to inquire how strongly (or how likely) the claim is supported in the KB. This mode of reasoning in argumentation allows to draw reasonable conclusions about claims, rather than just deciding whether a claim is accepted or not. Furthermore, this quantitative reasoning also allows us to directly compute a justification score or probability for claims at the level of KBs. The question we pose is: given a claim α (for simplicity, a literal), how strongly (or likely) is α acceptable in Δ. We observe that counting supports (Φ) for a claim (α) has not been considered (up to our knowledge). We first characterize its computational complexity.

Example 6 (Counting Arguments)**.** *Let C consist of all the candidates in an election, $\alpha := X$ be our favorite candidate, and the KB Δ contain poll results appropriately encoded in propositional logic. For example, a proposition x_s encodes that candidate X promotes policy s, along with p_s (participant p desires s) and d_s (s is desirable), where $s \in \{education, health, taxes, welfare\}$. Then for each participant p and candidate X, Δ contains formulas of the form $(x_s \wedge d_s \to X)$ and $p_s \to d_s$. Clearly, $\{x_{ed}, p_{ed}, p_{ed} \to d_{ed}, (x_{ed} \wedge d_{ed} \to X)\}$ yields a support for X if Δ contains both x_{ed} and p_{ed}. Then $|\mathsf{Args}(X, \Delta)|$ counts the number of arguments from all participants in favor of X and $|\mathsf{Rel\text{-}Args}(d_{ed}, X, \Delta)|$ gives the number of arguments in favor of X relevant to their education policy.*

We want to point out some interesting properties concerning the counting of (relevant) arguments for claims in propositional KBs. Notice that the support for a claim is a subset of the KB Δ, and there are at most exponentially many of them (precisely, 2^n many where n is the number of formulas in Δ). It is easy to construct examples where the number of arguments for a claim is indeed exponential.

Example 7. *Let $\Delta = \{x_i \to (y_i \wedge z_i), y_i \to x_{i+1}, z_i \to x_{i+1} \mid 0 \le i \le n\}$ and $\alpha = (x_0 \to x_n)$. Then, a support for α contains each $x_i \to (y_i \wedge z_i)$ and at least one of $y_i \to x_{i+1}$ or $z_i \to x_{i+1}$ for $i \le n$. Clearly, there are exponentially many (minimal) supports for α in total.*

Also note that counting relevant arguments is indeed an interesting problem, since certain formulas in a KB may be more *essential* than others when constructing support for a claim. This phenomenon is further highlighted next.

Example 8. *Let $\Delta = \{x \to y_i, (y_1 \vee \ldots \vee y_n) \to z, \mid i \le n\}$ and $\alpha = (x \to z)$. Then, a support for α contains the formula $(x \to y_i)$ for some $i \le n$ and the formula $(\bigvee_i y_i) \to z$. Clearly, the formula $(\bigvee_i y_i) \to z$ is more relevant than any other formula in Δ for α.*

While the number of supports is an interesting measure for claims, it falls short when it comes to assessing and comparing the likelihood of different claims. Consequently, one would like to compute the strength or likelihood of a claim compared to other claims in a given knowledge base.

Claim Strength and Likelihood. We assume a set C of possible claims is given as input. This is indeed a reasonable assumption, since in the abstract setting for AFs and claim-augmented AFs [Dvořák and Woltran, 2020], a set of arguments and their claims is always given. The set C can be all the claims for which arguments can be constructed in Δ, or only the claims we want to compare against the given α. Furthermore, if we do not restrict the set of claims, one can construct infinitely many arguments from a finite KB, since an argument for x also yields an argument for $x \wedge x$, and so on. Indeed, it seems natural to consider all claims and compare the strength of justification relative to each other.

Given a KB Δ, claim α and a set C of claims, we denote $\beta_C = \bigvee_{c \in C} c$. Then, the following formula allows us to compute the *strength* of α relative to all the claims in C:

$$\mathsf{str}(\alpha, \Delta, C) = |\mathsf{Args}(\alpha, \Delta)| / |\mathsf{Args}(\beta_C, \Delta)|.$$

Moreover, we define the *likelihood* of α relative to claims in C via the following formula:

$$\mathsf{lhd}(\alpha, \Delta, C) = |\mathsf{Args}(\alpha, \Delta)| / \Sigma_{c \in C} |\mathsf{Args}(c, \Delta)|.$$

The term in the denominator is a notable factor that distinguishes the strength of a claim from its likelihood. Interestingly, the presence of the disjunction (β_C) in $\mathsf{Args}(\beta_C, \Delta)$ results in an undercount of arguments for claims due to the minimality of the supports. In contrast, summing over all arguments for claims in set C, expressed as $\Sigma_{c \in C} |\mathsf{Args}(c, \Delta)|$, leads to overcounting. This is further emphasized below.

Example 9. *Let $\Delta = \{x, x \rightarrow c, c \rightarrow z_i, z_i \rightarrow d \mid i \leq 3\}$ and $C = \{c, d\}$. Then $\mathsf{Args}(c) = \{(\{x, x \rightarrow c\}, c)\}$ and $\mathsf{Args}(d)$ contains $(\{x, x \rightarrow c, c \rightarrow z_i, z_i \rightarrow d\})$ for $i \leq 3$. Now, $\mathsf{Args}(c \vee d) = \{(\{x, x \rightarrow c\}, c)\}$ since arguments corresponding to d are not subset-minimal. As a result, $\mathsf{str}(d, \Delta, C) = 3 \cdot \mathsf{str}(c, \Delta, C)$. Intuitively, arguments supporting claim d inherently serve as arguments for $d \wedge c$, establishing d as a stronger claim than c.*

Example 10. *Let $\Delta = \{c \wedge d\}$ and $C = \{c, d\}$. Then, $\mathsf{Args}(c) = \{(\{c \wedge d\}, c)\}$ and $\mathsf{Args}(d) = \{(\{c \wedge d\}, d)\}$. Now, $\mathsf{lhd}(c, \Delta, C) = 0.5 = \mathsf{lhd}(d, \Delta, C)$. Both claims are equally likely and equally strong as $\mathsf{str}(c, \Delta, C) = \mathsf{str}(d, \Delta, C)$. Intuitively, both claims are acceptable with a likelihood of 50%.*

Indeed, both measures are valuable for comparing the justification of claims within logic-based argumentation. In addition, the probability of a claim is further motivated from a fact-checking perspective: the truth value of a fact is calculated by considering the number of arguments in favor of the fact and then determining the ratio of all possible arguments over the given knowledge base [Syed *et al.*, 2019].

The *relevance strength* of a formula ψ for a claim α is based on counting relevant arguments ($|\mathsf{Rel\text{-}Args}(\psi, \alpha, \Delta)|$):

$$\mathsf{str}(\psi, \alpha, \Delta) = |\mathsf{Rel\text{-}Args}(\psi, \alpha, \Delta)| / |\mathsf{Args}(\alpha, \Delta)|.$$

Example 11 (Claim Strength & Likelihood). *We count all the arguments in Δ favoring X and take the ratio with the count of arguments favoring any candidate in C. This indicates how much X is favored in our KB compared to other candidates in C. Similarly, the ratio of $|\mathsf{Args}(X, \Delta)|$ and $\Sigma_{y \in C} |\mathsf{Args}(y, \Delta)|$ gives the probability that X is the favorite among C. Interestingly, the two measures for each candidate X coincide in Δ. Finally, $\mathsf{str}(d_{ed}, X, \Delta)$ informs the impact of the policy on education (d_{ed}) behind X being the favorite.*

3 Complexity of Counting Arguments

We now establish the complexity of counting arguments for knowledge bases formulated in propositional logic.

3.1 Classical Complexity

We characterize the classical complexity of counting (resp., relevant) arguments for a claim (with respect to a formula).

Theorem 12 ($\star$)**.** $\#\mathrm{ARG}$ *is* $\# \cdot \mathbf{coNP}$-*complete.*

Note that the relevance problem (ARG-Rel) has higher complexity than the existence problem (ARG) due to the reduction $(\Delta, \alpha) \mapsto (\Delta \cup \{x\}, \alpha \wedge x, x)$, where x is a new

proposition that does not appear in $\Delta \cup \{\alpha\}$. It turns out that this is in fact a parsimonious reduction, preserving the number of supports Φ for α in Δ and relevant supports for $\alpha \wedge x$ in $\Delta \cup \{x\}$. Furthermore, the membership in $\# \cdot \mathbf{coNP}$ can also be established in a similar way as for $\#\mathrm{ARG}$, resulting in the same complexity of counting the relevant supports for the claim given a formula. This leads to the following corollary.

Corollary 13. *The problem* $\#\mathrm{ARG}$-Rel *of counting relevant arguments is* $\# \cdot \mathbf{coNP}$-*complete.*

3.2 Parameterized and Structural Complexity

We derive various parameterized complexity results.

Parameters: Size of Claim and KB

Mahmood *et al.* (2023) explored the parameterized complexity of the decision problems ARG and ARG-Rel. For parameter $|\mathsf{var}(\Delta)|$ both problems are **FPT** (Thm 5.13) and the algorithm can be easily extended to count the number of solutions. For $|\mathsf{var}(\alpha)|$ both problems are **W[1]**-hard (Lem 5.8, Thm 5.10). Close inspection reveals that the employed reductions are parsimonious, yielding the following corollary.

Corollary 14 ($\star$)**.** $\#\mathrm{ARG}$ *and* $\#\mathrm{ARG}$-Rel *parameterized by* $|\mathsf{var}(\Delta)|$ *are in* **FFPT**, *and by* $|\mathsf{var}(\alpha)|$ *are* $\#\mathbf{W[1]}$-*hard.*

Parameter: Treewidth

We derive treewidth-aware reductions between $\#\mathrm{ARG}$ and $\#2$-$\mathrm{Q_{BF}}$ to count the extensions as well as compute the claim strength and likelihood in **FPT**-time when parameterized by treewidth of a suitable graph representation of the input. To achieve this, we reduce an instance (Δ, α) of ARG to an instance of 2-$\mathrm{Q_{BF}}$ using a similar idea as employed by Fichte *et al.* (2021). Note that the known-reductions between ARG and 2-$\mathrm{Q_{BF}}$ (given below) are restricted to the case when Δ includes clauses and do not preserve the number of solutions for the two instances.

DG-reductions for Clausal KBs. Let $\Delta = \{C_i, \mid 1 \leq i \leq n\}$ be a collection of clauses and α be a Boolean formula in DNF. Towards the reductions, a variable e_i for each $C_i \in \Delta$ encodes whether C_i is contained in the support, resulting in the set $E := \{e_i \mid 1 \leq i \leq n\}$ of *support variables*. Then, let M denote the set of variables over $\mathsf{var}(\Delta) = \bigcup_{i \leq n} \mathsf{var}(C_i)$. Moreover, let $N := \mathsf{var}(\Delta \cup \{\alpha\})$ and let $\tilde{N}$ denote the renaming of variables in N. That is $\tilde{N} := \{\tilde{x}_i \mid x_i \in N\}$ and each $\tilde{x}_i$ is a fresh variable. Finally, $\tilde{\beta}$ denotes each $\beta \in \{C, \alpha, \Delta\}$ over renamed variables $\tilde{N}$. Then, let $\mathsf{cons}_\Delta(E, M) := \bigwedge_{C_i \in \Delta}(\neg e_i \vee C_i)$ and $\mathsf{ent}_{\Delta, \alpha}(E, \tilde{N}) := \bigvee_{C_i \in \Delta}(e_i \wedge \neg \tilde{C}_i) \vee \tilde{\alpha}$. To preserve the treewidth linearly, both subformulas are split into a conjunction from each node's perspective in the tree-decomposition. To achieve this, let $\mathcal{T} = (T, \chi)$ be a TD of $\mathcal{G}_{(\Delta, \alpha)}$. Then, define a *labeled TD (LTD)* $\mathcal{T}' = (T, \chi, \delta)$ of $\mathcal{T}$, where *labeling* $\delta : T \rightarrow \Delta$ is such that $\delta(t) \in \Delta_t$ and $\Delta = \bigcup_{t \text{ of } T} \{\delta(t)\}$. An LTD can be easily obtained from any TD without changing the width by copying nodes accordingly. The labeling function guarantees that $\delta(t)$ contains C_i if, and only if, $C_i \in t$. Next, assume such an LTD $\mathcal{T}' = (T, \chi, \delta)$ of $\mathcal{T}$ and construct the following formula $\mathcal{R}_{\mathrm{ARG} \rightarrow 2\text{-}\mathrm{Q_{BF}}}(\mathcal{I}, \mathcal{T}') :=$
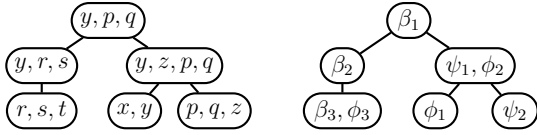
Figure 1: A TD for Δ from Example 15 (left) together with the clause-labeling for each bag (right)

$\exists E, M. \forall \tilde{N}. \varphi_{\text{ARG}}(E, M, \tilde{N})$, where φ_{ARG} is built for every node t of T by $\text{cons}_{\{\delta(t)\}}(E, M) \wedge \text{ent}_{\{\delta(t)\}}(E, \tilde{N})$. We extend the DG-reductions from above in three folds. First, the KB Δ includes formulas in CNF, not just clauses. Second, the presented reduction is a bijection and hence the number of solutions to the resulting 2-QBF-instance is same as the number of arguments for the given claim in Δ. Third, the reduction incorporates the minimality of supports and therefore only counts "valid" arguments as presented in Definition 5.

Generalization to CNFs. The first extension poses a non-trivial challenge, since the clauses of a formula can be *spread* over multiple bags. Recall that (by definition) in the clausal KBs, for each clause $C \in \Delta$ and a TD T for Δ, there is a node $t \in T$ such that the bag χ_t covers C. Then, the labeling (δ) of T relates each bag in T to the clauses it contains. However, if Δ contains formulas ϕ in CNF, then it is not guaranteed that all clauses of ϕ are contained in a bag. This challenge can alternatively be attributed to the graph representation ($\mathcal{G}_\Delta$) of the KB Δ. If $\mathcal{G}_\Delta$ is considered as the union of primal graphs $\mathcal{G}_\phi$ for each formula $\phi \in \Delta$, the graph encodes only the connectivity of variables in clauses. However, it does not represent the participation of variables or clauses in a formula. The following example further illustrates this issue.

Example 15. *Let $\Delta := \{\phi, \psi, \beta\}$ where each $\gamma \in \Delta$ is a CNF with clauses named as γ_i for $i \leq 3$. Precisely, let $\phi_1 = (x \vee \neg y), \phi_2 = (y \vee z \vee \neg p), \phi_3 = (r \vee \neg s \vee \neg t), \psi_1 = (y \vee z \vee \neg p \vee q), \psi_2 = (p \vee q \vee \neg z), \beta_1 = (\neg y \vee \neg p \vee q), \beta_2 = (y \vee r \vee \neg s)$ and $\beta_3 = (\neg r \vee s \vee \neg t)$. Now, consider the TD for $\mathcal{G}_\Delta$ depicted in Figure 1. For clarity, we label bags with corresponding clauses it covers (2nd TD on right). Then, the bag covering ϕ_3 (variables r, s, t) is* disconnected *from the bags covering the remaining clauses (ϕ_1 and ϕ_2) of ϕ.*

As Example 15 shows, clauses within the same formula can be covered by unrelated bags. Now, following a similar strategy as for the clausal KBs, one would use a support variable e_i for $i \leq 3$. Then additionally the bags covering clauses β_1 and β_2 in the TD must contain e_1. But this is undesirable due to the resulting blow-up in the treewidth of the primal graph if Δ contains many formulas stretched in this way.

To overcome this limitation, we extend the graph representation for Δ by considering the *extended* primal graph denoted as $\mathcal{G}_\Delta^e$. More precisely, each formula $\phi_i \in \Delta$ contains a new proposition f_i (called *formula name*), so that the variables in ϕ_i are also connected to f_i. This extension of the primal graph encodes the participation of variables in a formula and the relationship between variables that do not share a clause but a formula. This yields a set $F = \{f_i \mid 1 \leq i \leq n\}$ of variables encoding the *formula names* for a KB Δ. Additionally, the label (δ) in a LTD can be extended to guarantee

that $\delta(t)$ contains part of the formula ϕ_i if and only if $f_i \in t$. This way, we can still access a formula ϕ_i from a TD even if there is no single bag that covers entire ϕ_i. Note also that connecting f_i to each $\text{var}(\phi_i)$ increases the treewidth for ϕ_i by at most one over the treewidth for ϕ (simply add f_i to each bag for the TD of ϕ). To incorporate the addition of these variables into our reduction, we use f_i for $\phi_i \in \Delta$ to also encode whether ϕ_i is contained in the support, and use the set $F := \{f_i \mid 1 \leq i \leq n\}$ to serve the *support variables*. As before, let M denote the set of variables over $\text{var}(\Delta) = \bigcup_{\phi \in \Delta} \text{var}(\phi)$. Moreover, let $N := \text{var}(\Delta \cup \{\alpha\})$ and let $\tilde{N} := \{\tilde{x}_i \mid x_i \in N\}$ denote the renaming of variables in N. Finally, by $\tilde{\beta}$, we denote each $\beta \in \{\phi, C, \alpha, \Delta\}$ over renamed variables. Then, we let $\text{cons}_\Delta(F, M) := \bigwedge_{\phi_i \in \Delta}(\neg f_i \vee \phi_i)$ and $\text{ent}_{\Delta, \alpha}(F, \tilde{N}) := \bigvee_{\phi_i \in \Delta}(f_i \wedge \neg \tilde{\phi}_i) \vee \tilde{\alpha}$. Equivalently, we have $\text{cons}_\Delta(F, M) := \bigwedge_{\phi_i \in \Delta} \bigwedge_{C_{ij} \in \phi_i}(\neg f_i \vee C_{ij})$, and $\text{ent}_{\Delta, \alpha}(F, \tilde{N}) := \bigvee_{\phi_i \in \Delta} \bigvee_{C_{ij} \in \phi_i}(f_i \wedge \neg \tilde{C}_{ij}) \vee \tilde{\alpha}$. Then, $\psi'_{\text{ARG}} = \exists F, M. \forall \tilde{N}.(\text{cons}_\Delta(F, M) \wedge \text{ent}_{\Delta, \alpha}(F, \tilde{N}))$ yields the desired (intermediate) 2-QBF.

If we consider the extended primal graph for (Δ, α), ψ'_{ARG} already preserves the treewidth linearly. This is due to the additional edges between f_i and $\text{var}(\phi_i)$ of each $\phi_i \in \Delta$. However, this linear preservation of treewidth will break down when imposing the minimality condition over arguments. Thus, we split the two formulas.

Let $\mathcal{T} = (T, \chi)$ be a TD of $\mathcal{G}_{(\Delta, \alpha)}^e$. Moreover, let ϕ_i^t (resp., α_t) denote the clauses of ϕ_i (terms of α) contained in bag t and Δ_t denote the collection of all such ϕ_i^t. We define a *labeled TD (LTD)* $\mathcal{T}' = (T, \chi, \delta)$ of $\mathcal{T}$, where *labeling* $\delta \colon T \to \Delta$ is such that $\delta(t) \in \Delta_t \cup \{\alpha_t\}$.

Finally, we let $\delta(t_1) \sqcap \delta(t_2)$ denote the set $\{\phi_i^{t_1} \wedge \phi_i^{t_2} \mid \phi_i \in \Delta\}$, then $\Delta = \sqcap_{t \in T}\{\delta(t)\}$. An LTD can be easily obtained from any TD without changing the width by copying nodes accordingly. Furthermore, our labeling function also guarantees that $\delta(t)$ contains ϕ_i^t if, and only if, $f_i \in t$. Next, we assume such an LTD $\mathcal{T}' = (T, \chi, \delta)$ of $\mathcal{T}$, and let $E = F \cup \{e_i^t \mid i \leq n, t \in T\}, L = \{\ell^t \mid t \in T\}, S = \{s_i^t \mid i \leq n, t \in T\}$ sets of variables. Then, we construct the formula

$$\psi_{\text{ARG}} := \exists E(\exists M. \varphi_{\text{cons}}(E, M) \wedge \forall L\, S\, \tilde{N}. \varphi_{\text{ent}}(E, L, S, \tilde{N})),$$

where $\varphi_{\text{cons}}(E, M)$ is built for every node t of T as the conjunction of Formulas (1)–(2), which can be converted into CNF. Formula $\varphi_{\text{ent}}(E, L, S, \tilde{N}))$ is built for every t of T as the disjunction of (3)–(6), which can be converted into DNF. To obtain CNF and DNF as claimed, one can use Tseitin [Tseitin, 1983] or assume each bag considers at most one clause in ϕ_i^t (easily obtainable by, e.g., just duplicating bags). This does not increase the treewidth and only increases TD size linearly.

$$e_i^t \leftrightarrow \left(\phi_i^t \wedge \bigwedge_{t' \in \text{children}(t)} e_i^{t'}\right) \qquad \text{for } \phi_i^t \in \Delta_t, t \in T \quad (1)$$

$$f_i \leftrightarrow e_i^{\text{last}(i)} \qquad \text{for } \phi_i \in \Delta \quad (2)$$

$$\neg(s_i^t \leftrightarrow (e_i^t \wedge \neg\tilde{\phi}_i^t) \vee \bigvee_{t' \in \text{children}(t)} s_i^{t'}) \qquad \text{for } \phi_i^t \in \Delta_t, t \in T \quad (3)$$

$$\neg\left(\ell^t \leftrightarrow \left(\bigwedge_{D \in \alpha_t} \tilde{D} \wedge \bigwedge_{t' \in \mathsf{children}(t)} \ell^{t'}\right)\right) \qquad \text{for } t \in T \text{ (4)}$$

$$s_i^{\mathsf{last}(i)} \qquad \text{for } \phi_i \in \Delta \text{ (5)}$$

$$\ell^{\mathsf{last}(\alpha)} \qquad (6)$$

The purpose of Eqs. (1)–(2) is to guide along the TD whether a formula ϕ_i is taken in the support or not. Then Eq. (3) guides a contradiction along the TD, which together with Eqs. (4)–(6) guarantee that each term of α is indeed explained.

Lemma 16 (Correctness, $\star$). *Let (Δ, α) be an instance of* ARG. *Then, there is a support $\Phi \subseteq \Delta$ such that (Φ, α) is an argument in Δ if and only if ψ_{ARG} is true.*

Since ψ_{ARG} is constructed for each node of T, it is easy to see that the DG reduction preserves the treewidth linearly.

Lemma 17 (TW-Awareness, $\star$). *Let $\mathcal{I} = (\Delta, \alpha)$ be an instance of* ARG *and $\mathcal{T}' = (T, \chi, \delta)$ be an LTD of a TD $\mathcal{T}$ of $\mathcal{G}^e_{(\Delta,\alpha)}$ of width k. Then, the reduction $\mathcal{R}_{\mathrm{ARG}\to2\text{-}\mathrm{QBF}}(\mathcal{I}, \mathcal{T}')$ constructs a qBf ψ_{ARG} that linearly preserves the width. More precisely, we have that $\mathsf{tw}(\mathcal{G}_{\mathsf{matrix}(\psi_{\mathrm{ARG}})}) \in \mathcal{O}(k)$.*

Counting Arguments. Our reductions work for the decision variant, but when it comes to runtime upper bounds for counting, double exponential bounds are not guaranteed since only E includes free variables in ψ_{ARG}. Nevertheless, observe that if we do not have full quantifier alternations, the validity of quantified Boolean formulas is easier than in the general case. This still holds for treewidth (under ETH) and counting, which we will utilize afterwards.

Proposition 18 ($\star$). *Let Q be a QBF over free variables V and with ℓ alternating quantifier blocks such that the innermost formula is of the form $(\exists M.C) \wedge (\forall N.D)$ with C being in CNF and D being in DNF. There is an algorithm for counting the number of assignments θ over V such that $Q[\theta]$ evaluates to true, in time $\exp(\ell+2, \mathcal{O}(k)) \cdot \mathsf{poly}(|\mathsf{var}(Q)|)$, where k is the treewidth of the primal graph $\mathcal{G}_{C \wedge D}$.*

Clearly, in ψ_{ARG}, there is no quantifier alternation between variables M of φ_{cons} and $L, S, \tilde{N}$ of φ_{ent} (except those over E). Consequently, we apply Proposition 18 and reduction $\mathcal{R}_{\mathrm{ARG}\to2\text{-}\mathrm{QBF}}$ above, where E is not existentially quantified, but the set of free variables, to obtain an upper bound $\exp(2, \mathcal{O}(k)) \cdot \mathsf{poly}(|\mathsf{var}(Q)|)$ for counting solutions.

Theorem 19 (Runtime UB, $\star$). *Let $\mathcal{I} = (\Delta, \alpha)$ be an instance of* ARG. *Then, #ARG can be solved in time $\exp(2, \mathcal{O}(k)) \cdot \mathsf{poly}(|\mathsf{var}(\Delta) \cup \mathsf{var}(\alpha)|)$, where $k = \mathsf{tw}(\mathcal{G}^e_{\mathcal{I}})$.*

Unluckily, this can probably not be significantly improved. This follows from the runtime lower bounds for the decision problem ARG [Fichte *et al.*, 2021, Thm. 12].

Corollary 20 (Runtime LB, $\star$). *Let $\mathcal{I} = (\Delta, \alpha)$ be an instance of* ARG. *Then, under ETH, the problem #ARG cannot be solved in time $\exp(2, o(\mathsf{tw}(\mathcal{G}^e_{\mathcal{I}})) \cdot \mathsf{poly}(|\mathsf{var}(\Delta) \cup \mathsf{var}(\alpha)|)$.*

One cannot expect to significantly improve the DG reduction.

Corollary 21 (TW-LB, $\star$). *Let $\mathcal{I} = (\Delta, \alpha)$ be an instance of* ARG. *Under ETH, there is no reduction $\mathcal{R}'$ from* ARG *to 2-QBF yielding a qBf ψ_{ARG} in time $\exp(2, o(\mathsf{tw}(\mathcal{G}^e_{\mathcal{I}}))) \cdot \mathsf{poly}(|\mathsf{var}(\Delta) \cup \mathsf{var}(\alpha)|)$ with $\mathsf{tw}(\mathcal{G}_{\psi_{\mathrm{ARG}}}) \in o(\mathsf{tw}(\mathcal{G}^e_{\mathcal{I}}))$.*

Adding Minimality. Next, we modify our reductions such that only subset-minimal supports are selected and counted. The idea is to use an additional formula ensuring that, whenever E provides a support, then for all other sets $\hat{E}$: either $\hat{E}$ is not a support for α, or $\hat{E} \not\subset E$. To achieve this, we use variables $\hat{E} = \{\hat{e}_i^t \mid i \leq n, t \in T\}$ and $\hat{M}$. Moreover, we use $Q = \{q^t, q_i \mid t \in T, \phi_i \in \Delta\}$, called *equality variables* to guide along the TD whether E and $\hat{E}$ are (un)equal. Let $\psi_m(E)$ denote the formula $\forall \hat{E}, \hat{M}, Q, L, S, \tilde{N}(\neg\varphi_{\mathsf{cons}}(\hat{E}, \hat{M}) \vee \neg\varphi_{\mathsf{ent}}(\hat{E}, L, S, \tilde{N}) \vee \varphi_{\hat{E} \not\subset E}(E, \hat{E}, Q))$, where formulas φ_{cons} and φ_{ent} are built for every node t of T as before, and $\varphi_{\hat{E} \not\subset E}(E, \hat{E}, Q)$ a collection of DNF formulas:

$$\hat{f}_i \wedge \neg f_i \qquad \text{for } \phi_i \in \Delta \text{ (7)}$$

$$\neg(q_i \leftrightarrow f_i \wedge \neg\hat{f}_i) \qquad \text{for } \phi_i \in \Delta \text{ (8)}$$

$$\neg\left(q^t \leftrightarrow \bigvee_{\phi_i^t \in \delta(t)} q_i \vee \bigvee_{t' \in \mathsf{children}(t)} q^{t'}\right) \qquad \text{for } t \in T \text{ (9)}$$

$$\neg q^{\mathsf{root}(T)} \qquad (10)$$

Eq. (7) allows a support given by $\hat{E}$ to contain a formula not in E, then Eq. (8) defines inequality and Eq. (9) guides the inequality along TD. Finally, Eq. (10), allows equal supports over E and $\hat{E}$. Then our reduction yields the formula $\exists E(\psi_{\mathrm{ARG}}(E) \wedge \psi_m(E))$. Again, since the additional formulas are constructed for each node of T, it is easy to see that the DG reduction is correct and preserves the treewidth linearly.

Counting Relevant Arguments. In order to compute #ARG-Rel(ψ, α, Δ) for a formula ψ and claim α in Δ, we force a support to contain ψ by simply plugging-in the formula name (say f_0) for ψ in $\mathcal{R}_{\mathrm{ARG}\to2\text{-}\mathrm{QBF}}(\mathcal{I}, \mathcal{T}')$. This achieves that only those supports are counted that contain ψ.

Corollary 22 (Tight runtime bounds for #ARG-Rel, $\star$). *Let $\mathcal{I} = (\Delta, \psi, \alpha)$ be an instance of* ARG-Rel. *Then #ARG-Rel can be solved in time $\exp(2, \mathcal{O}(k)) \cdot \mathsf{poly}(|\mathsf{var}(\Delta) \cup \mathsf{var}(\alpha) \cup \mathsf{var}(\psi)|)$, where $k = \mathsf{tw}(\mathcal{G}^e_{\mathcal{I}})$. Under ETH, the problem #ARG-Rel cannot be solved in time $\exp(2, o(\mathsf{tw}(\mathcal{G}^e_{\mathcal{I}})) \cdot \mathsf{poly}(|\mathsf{var}(\Delta) \cup \mathsf{var}(\alpha) \cup \mathsf{var}(\psi)|)$.*

4 Conclusion

We explored the classical and structural complexity for counting arguments in propositional knowledge bases. To the best of our knowledge, this is the first paper to address the complexity of counting in logic-based argumentation. Regarding the structural complexity, we established tight runtime bounds for #ARG and #ARG-Rel under the parameter treewidth. The upper bounds were determined using a technique described in Proposition 18, which is interesting and useful in itself for solving QBFs. Finally, we proposed and motivated quantitative reasoning for claims to determine their relative strength and likelihood in a given knowledge base.

Acknowledgments

References

[Alfano *et al.*, 2020] Gianvincenzo Alfano, Marco Calautti, Sergio Greco, Francesco Parisi, and Irina Trubitsyna. Explainable acceptance in probabilistic abstract argumentation: Complexity and approximation. In *KR'20*, pages 33–43. IJCAI Organization, 2020.

[Amgoud and Doder, 2018] Leila Amgoud and Dragan Doder. Gradual semantics for weighted graphs: An unifying approach. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*, 2018.

[Amgoud and Prade, 2009] Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artificial Intelligence*, 173(3-4):413–436, 2009.

[Amgoud *et al.*, 2008] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. Making decisions through preference-based argumentation. *KR*, 8:963–970, 2008.

[Amgoud *et al.*, 2018] Leila Amgoud, Elise Bonzon, Jérôme Delobelle, Dragan Doder, Sébastien Konieczny, and Nicolas Maudet. Gradual semantics accounting for similarity between arguments. In *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*, 2018.

[Atkinson *et al.*, 2017] Katie Atkinson, Pietro Baroni, Massimiliano Giacomin, Anthony Hunter, Henry Prakken, Chris Reed, Guillermo Simari, Matthias Thimm, and Serena Villata. Towards artificial argumentation. *AI magazine*, 38(3):25–36, 2017.

[Baroni and Riveret, 2019] Pietro Baroni and Régis Riveret. Enhancing statement evaluation in argumentation via multi-labelling systems. *J. Artif. Intell. Res.*, 66:793–860, 2019.

[Baroni *et al.*, 2016] Pietro Baroni, Guido Governatori, Ho-Pun Lam, and Régis Riveret. On the justification of statements in argumentation-based reasoning. In *KR'16*, pages 521–524. The AAAI Press, 2016.

[Baroni *et al.*, 2018] Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors. *Handbook of Formal Argumentation*. College Publications, 2018.

[Baumann *et al.*, 2023] Ringo Baumann, Anna Rapberger, and Markus Ulbricht. Equivalence in argumentation frameworks with a claim-centric view: Classical results with novel ingredients. *Journal of Artificial Intelligence Research*, 77:891–948, 2023.

[Bernreiter *et al.*, 2023] Michael Bernreiter, Wolfgang Dvorák, Anna Rapberger, and Stefan Woltran. The effect of preferences in abstract argumentation under a claim-centric view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6253–6261, 2023.

[Besnard and Hunter, 2001] Philippe Besnard and Anthony Hunter. A logic-based theory of deductive arguments. *Artificial Intelligence*, 128(1-2):203–235, 2001.

[Besnard and Hunter, 2008] Philippe Besnard and Anthony Hunter. *Elements of argumentation*, volume 47. MIT press Cambridge, 2008.

[Bodlaender and Koster, 2008] Hans L. Bodlaender and Arie M. C. A. Koster. Combinatorial optimization on graphs of bounded treewidth. *Comput. J.*, 51(3):255–269, 2008.

[Bodlaender, 1996] Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.

[Chen, 2004] Hubie Chen. Quantified constraint satisfaction and bounded treewidth. In *ECAI'04*, pages 161–165. IOS Press, 2004.

[Chesnevar *et al.*, 2000] Carlos Iván Chesnevar, Ana Gabriela Maguitman, and Ronald Prescott Loui. Logical models of argument. *ACM Computing Surveys (CSUR)*, 32(4):337–383, 2000.

[Creignou *et al.*, 2011] Nadia Creignou, Johannes Schmidt, Michael Thomas, and Stefan Woltran. Complexity of logic-based argumentation in post's framework. *Argument & Computation*, 2(2-3):107–129, 2011.

[Creignou *et al.*, 2014] Nadia Creignou, Uwe Egly, and Johannes Schmidt. Complexity classifications for logic-based argumentation. *ACM Trans. Comput. Log.*, 15(3):19:1–19:20, 2014.

[Curticapean, 2019] Radu Curticapean. Counting Problems in Parameterized Complexity. In Christophe Paul and Michal Pilipczuk, editors, *13th International Symposium on Parameterized and Exact Computation (IPEC 2018)*, volume 115 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:18, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[Downey and Fellows, 1999] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer Verlag, 1999.

[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.

[Durand *et al.*, 2005] Arnaud Durand, Miki Hermann, and Phokion G. Kolaitis. Subtractive reductions and complete

problems for counting complexity classes. *Theor. Comput. Sci.*, 340(3):496–513, 2005.

[Dvořák and Woltran, 2020] Wolfgang Dvořák and Stefan Woltran. Complexity of abstract argumentation under a claim-centric view. *Artificial Intelligence*, 285:1–22, 2020.

[Dvořák *et al.*, 2023] Wolfgang Dvořák, Alexander Greßler, Anna Rapberger, and Stefan Woltran. The complexity landscape of claim-augmented argumentation frameworks. *Artificial Intelligence*, 317:103873, 2023.

[Fazzinga *et al.*, 2015] Bettina Fazzinga, Sergio Flesca, and Francesco Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Trans. Comput. Log.*, 16(3):22:1–22:39, 2015.

[Fichte *et al.*, 2020] Johannes Klaus Fichte, Markus Hecher, and Andreas Pfandler. Lower Bounds for QBFs of Bounded Treewidth. In *LICS*, pages 410–424. ACM, 2020.

[Fichte *et al.*, 2021] Johannes Klaus 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.*, 2023] Johannes K Fichte, Markus Hecher, Yasir Mahmood, and Arne Meier. Quantitative reasoning and structural complexity for claim-centric argumentation. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 3212–3220, 2023.

[Flum and Grohe, 2004] Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM J. Comput.*, 33(4):892–922, 2004.

[Flum and Grohe, 2006] Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Springer Verlag, 2006.

[Haak *et al.*, 2023] Anselm Haak, Arne Meier, Om Prakash, and B. V. Raghavendra Rao. Parameterised counting in logspace. *Algorithmica*, 85(10):2923–2961, 2023.

[Hunter and Thimm, 2017] Anthony Hunter and Matthias Thimm. Probabilistic reasoning with abstract argumentation frameworks. *J. Artif. Intell. Res.*, 59:565–611, 2017.

[Hunter, 2012] Anthony Hunter. Some foundations for probabilistic abstract argumentation. *Comma*, 2012:117–128, 2012.

[Hunter, 2013] Anthony Hunter. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning*, 54(1):47–81, 2013.

[Impagliazzo *et al.*, 2001] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.

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

[McCartin, 2006] Catherine McCartin. Parameterized counting problems. *Annals of Pure and Applied Logic*, 138(1-3):147–182, 2006.

[Parsons *et al.*, 2003] Simon Parsons, Michael Wooldridge, and Leila Amgoud. Properties and complexity of some formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.

[Prakken and Vreeswijk, 2002] Henry Prakken and Gerard Vreeswijk. *Logics for Defeasible Argumentation*, pages 219–318. Springer Verlag, 2002.

[Pu *et al.*, 2015] Fuan Pu, Jian Luo, Yulai Zhang, and Guiming Luo. Attacker and defender counting approach for abstract argumentation. In David C. Noelle, Rick Dale, Anne S. Warlaumont, Jeff Yoshimi, Teenie Matlock, Carolyn D. Jennings, and Paul P. Maglio, editors, *Proceedings of the 37th Annual Meeting of the Cognitive Science Society, CogSci 2015, Pasadena, California, USA, July 22-25, 2015*. cognitivesciencesociety.org, 2015.

[Rago *et al.*, 2018] Antonio Rago, Oana Cocarascu, and Francesca Toni. Argumentation-based recommendations: Fantastic explanations and how to find them. 2018.

[Riveret *et al.*, 2007] Regis Riveret, Antonino Rotolo, Giovanni Sartor, Roth Bram, and Henry Prakken. Success chances in argument games: a probabilistic approach to legal disputes. *Available at SSRN 1100672*, 2007.

[Sipser, 1997] Michael Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.

[Syed *et al.*, 2019] Zafar Habeeb Syed, Michael Röder, and Axel-Cyrille Ngonga Ngomo. Unsupervised discovery of corroborative paths for fact validation. In Chiara Ghidini, Olaf Hartig, Maria Maleshkova, Vojtěch Svátek, Isabel Cruz, Aidan Hogan, Jie Song, Maxime Lefrançois, and Fabien Gandon, editors, *The Semantic Web – ISWC 2019*, pages 630–646, Cham, 2019. Springer International Publishing.

[Tseitin, 1983] G. S. Tseitin. On the complexity of derivation in propositional calculus. In *Automation of Reasoning: 2: Classical Papers on Computational Logic 1967–1970*, pages 466–483. Springer Berlin Heidelberg, 1983.

[Wu and Caminada, 2010] Yining Wu and Martin Caminada. A labelling based justification status of arguments. *Studies in Logic*, 3:12–29, 01 2010.