

A Methodology for Incompleteness-Tolerant and Modular Gradual Semantics for Argumentative Statement Graphs

Antonio Rago^{1,2}, Stylianos Loukas Vasileiou^{3,4}, Son Tran³, Francesca Toni² and William Yeoh⁴

¹King's College London

²Imperial College London

³New Mexico State University

⁴Washington University in St. Louis

antonio.rago@kcl.ac.uk, {stelios, stran}@nmsu.edu, ft@imperial.ac.uk, wyeoh@wustl.edu

Abstract

Gradual semantics (GS) have demonstrated great potential in argumentation, in particular for deploying *quantitative bipolar argumentation frameworks* (QBAFs) in a number of real-world settings, from judgmental forecasting to explainable AI. In this paper, we provide a novel methodology for obtaining GS for *statement graphs*, a form of structured argumentation framework, where arguments and relations between them are built from logical statements. Our methodology differs from existing approaches in the literature in two main ways. First, it naturally accommodates incomplete information, so that arguments with partially specified premises can play a meaningful role in the evaluation. Second, it is modularly defined to leverage on any GS for QBAFs. We also define a set of novel properties for our GS and study their suitability alongside a set of existing properties (adapted to our setting) for two instantiations of our GS, demonstrating their advantages over existing approaches.

1 Introduction

Argumentation frameworks have emerged as powerful tools for reasoning about and resolving conflicting information in complex environments (see (Baroni et al. 2018; Atkinson et al. 2017) for overviews). In recent years, *gradual semantics* (GS) have shown great promise in extending these frameworks, particularly in the realm of *quantitative bipolar argumentation frameworks* (QBAFs) (Baroni, Rago, and Toni 2018), where arguments have weights and may be related by *attack* and *support* relations. These approaches thus allow for both negative and positive, respectively, influences between arguments to be accounted for, arguably aligning more closely with human judgment than traditional argumentation frameworks (Polberg and Hunter 2018). GS have found applications in diverse (and often human-centric) areas, from *judgmental forecasting* (Irwin, Rago, and Toni 2022) to *explainable AI* (Vassiliades, Bassiliades, and Patkos 2021; Cyras et al. 2021), by allowing for quantitative, and thus in some cases more nuanced, evaluations of arguments beyond the traditional accept/reject dichotomy in extension-based semantics (Dung 1995; Cayrol and Lagasque-Schiex 2005). These quantitative evaluations, typically referred to as *strengths*, can be particularly useful in cases where the information in argumentation

frameworks is incomplete, since strengths can account for the resulting uncertainty quantitatively.

However, most of the existing work on GS (see (Baroni, Rago, and Toni 2019) for an overview) has focused on formalisms based on abstract arguments (Dung 1995), treated as atomic entities without internal structure. This abstraction, while useful in many contexts, fails to capture the rich logical structure that is often present in real-world arguments. Meanwhile, various forms of *structured argumentation frameworks* (see (Besnard et al. 2014) for an overview) allow to represent the internal composition of arguments and the relationships between their components. They thus offer a more expressive and realistic model in agent-based domains, such as user modeling (Hadoux, Hunter, and Polberg 2023), scientific debates (Cramer and Dauphin 2020), and model reconciliation settings (Vasileiou et al. 2024). However, this greater reliance on the logical (e.g., background or contextual) knowledge can lead to more problematic cases where this information is incomplete, where the *completeness* of a statement (as we formally define in §4) refers to its grounding in facts. Yet, it is only recently that the study of GS for structured argumentation has received attention, e.g., for *deductive argumentation* (Besnard and Hunter 2014) by Heyninck et al. (2023), for *assumption-based argumentation* (ABA) (Cyras et al. 2017) by Skiba et al. (2023), for *ASPIC⁺* (Modgil and Prakken 2014) by Spaans (2021) and Prakken (2024), and for a restricted type of structured argumentation in the form of *statement graphs* (SGs) (Hecham, Bisquert, and Croitoru 2018) by Jedwabny et al. (2020).

We selected SGs as our targeted form of structured argumentation in this paper. We chose this starting point for our analysis, rather than other forms of structured argumentation, because they naturally accommodate attack and support relations. Their GS are thus naturally relatable to those of QBAFs, which are the most widely studied form of argumentation when it comes to GS (Baroni, Rago, and Toni 2019). This allows us to undertake an interesting analysis of properties of our GS in direct comparison with those for QBAFs, given their widespread application.

To illustrate the need for GS in structured argumentation, Figure 1 shows a debate about climate change represented as an SG. SGs use a restricted type of structured arguments in the form of *statements*, each consisting of a




Figure 1: Example SG (with arbitrary weights τ , see §3). Green edges with a plus denote support; red edges with a minus denote attack. The SG includes incomplete information about the statement α_4 , since its premise d is not grounded in facts. However, the given weight for α_4 provides information that our GS can use.

premise (a compound of literals, which may possibly be τ , i.e., true) and a claim (a literal). A typical statement (α_1) may combine evidence about rising CO₂ emissions (a) and increasing global temperatures (b), given that they are both supported by facts (α_2 and α_3), to support the claim that climate change is human-caused (c). Such an SG may be the result of mining arguments from text, e.g., as overviewed in (Lawrence and Reed 2019). Such argument mining techniques coupled with suitable GS for SGs could allow for more expressive representation of arguments while handling incomplete information in the modelling of debates, a rapidly growing research area utilising many forms of argumentation (Visser et al. 2020; Goffredo et al. 2023; Ruiz-Dolz, Heras, and García-Fornes 2023; Budan et al. 2023; Bistarelli, Taticchi, and Meo 2023; Bezou-Vrakatseli, Cocarascu, and Modgil 2024), especially given the groundbreaking potential of large language models (Oluokun et al. 2024).

The GS proposed by Jedwabny *et al.* (2020) for SGs uses T-norms and T-conorms from fuzzy logic to aggregate the strengths of a statement’s components, but this GS faces challenges when dealing with incomplete information as it relies on complete knowledge and the presence of a statement’s supporting elements. This requirement can limit its applicability in real-world scenarios, where such information is almost always incomplete. For example, in Figure 1, the strength of the statement α_1 is reduced (e.g., in the eyes of a modelled audience) by a statement (α_4) questioning the reliability of CO₂ measurements, without being evidenced (there is no support for d). To date, existing approaches to GS in structured argumentation struggle to handle such incompleteness, potentially leading to oversimplified or incorrect conclusions, and thus poor modelling of debates.

To address this limitation, we introduce a novel methodology for GS in SGs that can naturally accommodate complete and incomplete information. Our approach is *modular* in that it allows for different GS to be instantiated based on the requirements of a given task and context. Specifically, our methodology for GS can be instantiated with any existing GS for QBAFs. This existing GS is used to dialectically evaluate the literals in a premise, based on their attacking and supporting evidence, before these “intermediate” evaluations are aggregated (depending on the statement’s construction) to give a final strength for the statement.

Thus, our methodology combines some of the expressiveness of structured argumentation with the flexibility of GS for QBAFs, while accounting for the practical reality of in-

complete knowledge. Our contributions are as follows:

- We introduce a novel methodology for constructing GS in SGs with incomplete information, presenting two instantiations based on different existing GS for QBAFs.
- We present a collection of novel and existing (by Jedwabny *et al.* (2020)) properties for GS in SGs and discuss their suitability in different settings.
- Using these properties, we conduct a comprehensive theoretical analysis of our two proposed instantiations, two instantiations of the T-norm semantics (Jedwabny, Croitoru, and Bisquert 2020) and two existing GS for QBAFs, discussing the strengths and limitations of each approach.

2 Related Work

GS have received significant attention within abstract argumentation in recent years, given their professed alignment with human reasoning (Vesic, Yun, and Teovanovic 2022; de Tarlé, Bonzon, and Maudet 2022) and applicability to real-world settings (Potyka 2021; Rago et al. 2021). This has given rise to a rich field of research, with a broad repertoire of GS (Gonzalez et al. 2021; Yun and Vesic 2021; Wang and Shen 2024), in-depth analyses of GS’ behaviour (Delobelle and Villata 2019; Oren et al. 2022; Yin, Potyka, and Toni 2023; Kampik et al. 2024), and a number of open-source computational toolsets (Potyka 2022; Alarcón and Kampik 2024). One issue in this area is that many abstract GS, e.g., those we use (Rago et al. 2016; Potyka 2018), do not satisfy the uniqueness property in assigning strengths in cyclic graphs due to their recursive nature (Gabbay and Rodrigues 2015a; Anaissy et al. 2024). This must be considered if our acyclicity restriction is relaxed in future work.

There have been other works related to GS in structured argumentation, differing from our work as follows. Within ASPIC⁺, Spaans (2021) introduced *initial strength functions*, with corresponding properties, though they use weights on rules and do not consider explicit support relations. Further, Prakken (2024) gives a formal model of argument strength across graphs and deploys it. A major difference between this work and our approach is our leveraging of existing bipolar abstract GS for handling incomplete information. We leave to future work the assessment of this approach with respect to our proposed properties. Also for ASPIC⁺, extension-based semantics have been used for cases with incomplete information (Odekerken et al. 2023).

Meanwhile, Heyninck *et al.* (2023) compare properties of GS with those concerning culpability when applied to abstract argumentation frameworks extracted from deductive argumentation, whereas we lift abstract GS to the structured level. Skiba *et al.* (2023) assess a family of GS for ranking arguments in structured argumentation, specifically ABA. Both of these works do not consider the support relation. Finally, Amgoud and Ben-Naim (2015) give a method for ranking logic-based argumentation frameworks, but do not consider an explicit strength or a support relation.

Finally, *probabilistic argumentation* (Kohlas 2003; Dung and Thang 2010; Thimm 2012; Hunter 2013; Gabbay and Rodrigues 2015b; Polberg, Hunter, and Thimm 2017;

Hunter, Polberg, and Thimm 2020; Fazzinga, Flesca, and Furfaro 2018; Thimm, Polberg, and Hunter 2018; Mantadelis and Bistarelli 2020), is a related but distinct subfield where strengths, which may be over (sets of) arguments, relations or both, must adhere to strict probabilistic principles, whereas GS may not.

3 Preliminaries

We consider a propositional language $\mathcal{L}$ comprising a finite set of atoms, including a special atom $\top$ to represent truth. A *literal* is an atom ϕ or its negation $\neg\phi$. Note that $\neg(\neg\phi) \equiv \phi$. A *compound* is an expression, limited in this paper to be constructed from literals using the connective $\wedge$.

A *statement* is defined as a pair $\langle \Phi, \Psi \rangle$, where $\Phi \in \mathcal{L}$ is either $\top$ or a logically consistent compound and $\Psi \in \mathcal{L}$ is a literal. We refer to Φ as the *premise* and to Ψ as the *claim* of statement $\alpha = \langle \Phi, \Psi \rangle$. We define $Prem(\alpha) = \{x \mid x \in \Phi\} \setminus \{\top\}$ as the set of literals in the premise of statement α . Note that if the premise of a statement α is $\top$ then $Prem(\alpha) = \emptyset$.

For simplicity and ease of presentation, we restrict attention to statements with conjunctive premises, similarly to (Jedwabny, Croitoru, and Bisquert 2020). In §4.1, we will define our methodology generally, before instantiating it for these restricted statements, pointing towards other instantiations which could be used for more complex compounds, such as those in conjunctive normal form.

A statement *attacks* or *supports* another if the former gives justification against or for the latter's premise:

Definition 1. Let $\alpha_1 = \langle \Phi_1, \Psi_1 \rangle$ and $\alpha_2 = \langle \Phi_2, \Psi_2 \rangle$ be two statements. We say that:

- α_1 *supports* α_2 iff $\exists x \in Prem(\alpha_2)$ s.t. $\Psi_1 \equiv x$;
- α_1 *attacks* α_2 iff $\exists x \in Prem(\alpha_2)$ s.t. $\Psi_1 \equiv \neg x$.

We can then generate an SG structured with respect to the support and attack relations between statements, as in the works by Hecham *et al.* (2018) and Jedwabny *et al.* (2020).¹ However, differently to these works, we include the quantitative weights of each statement within the SG, rather than applying them separately. This reformulation allows for direct comparisons with properties of QBAFs (Baroni, Rago, and Toni 2019), which provided inspiration for some of our properties.

Definition 2. An SG is a quadruple $\langle \mathcal{X}, \mathcal{A}, \mathcal{S}, \tau \rangle$ s.t.:

- $\mathcal{X}$ is a set of statements;
- $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$ is s.t. $\forall \alpha_1, \alpha_2 \in \mathcal{X}, (\alpha_1, \alpha_2) \in \mathcal{A}$ iff α_1 attacks α_2 ;
- $\mathcal{S} \subseteq \mathcal{X} \times \mathcal{X}$ is s.t. $\forall \alpha_1, \alpha_2 \in \mathcal{X}, (\alpha_1, \alpha_2) \in \mathcal{S}$ iff α_1 supports α_2 ;
- $\tau : \mathcal{X} \rightarrow [0, 1]$ is a weight over the statements.

For any $\alpha_1 \in \mathcal{X}$, we call $\mathcal{A}(\alpha_1) = \{\alpha_2 \in \mathcal{X} \mid (\alpha_2, \alpha_1) \in \mathcal{A}\}$ α_1 's attackers and $\mathcal{S}(\alpha_1) = \{\alpha_2 \in \mathcal{X} \mid (\alpha_2, \alpha_1) \in \mathcal{S}\}$ α_1 's

¹Note that we chose to use SGs similarly to Hecham *et al.* (2018) and Jedwabny *et al.* (2020) in order to allow for direct comparisons with the gradual semantics in the latter work. Their focus on attacks and supports on premises, rather than conclusions, also aligns with other works in the literature, e.g., ABA.

supporters. For any $\alpha_x, \alpha_y \in \mathcal{X}$, let a path from α_x to α_y via $\mathcal{R} \subseteq \mathcal{A} \cup \mathcal{S}$ be $(\alpha_0, \alpha_1), \dots, (\alpha_{n-1}, \alpha_n)$ for some $n > 0$, where $\alpha_0 = \alpha_x$, $\alpha_n = \alpha_y$ and, for any $1 \leq i \leq n$, $(\alpha_{i-1}, \alpha_i) \in \mathcal{R}$.

Corollary 1. Given a SG $\langle \mathcal{X}, \mathcal{A}, \mathcal{S}, \tau \rangle$, $\mathcal{A} \cap \mathcal{S} = \emptyset$.²

An SG thus gives a restricted form of structured argumentation framework with conjunctive premises, undercutting attacks, deductive/evidential supports and weights. Note that weights may be fixed, i.e., identical for all arguments which are considered a-priori equal as in classical argumentation frameworks, or variable, representing, e.g., different attributes such as source authority, premise credibility, or goal importance.³ We envisage a usage of weights similar to that in (Alsinet *et al.* 2008) within the realm of *defeasible logic programming* (García and Simari 2004), where weighted certainties on arguments (including facts) are used to introduce *possibilistic argumentation*. Also, we focus on acyclic graphs, following several works on GS (Rago *et al.* 2016; de Tarlé, Bonzon, and Maudet 2022), leaving cyclic graphs to future work.

For the remainder of the paper, we assume as given a generic SG $\mathcal{G} = \langle \mathcal{X}, \mathcal{A}, \mathcal{S}, \tau \rangle$.

Definition 3. A GS $\sigma : \mathcal{X} \rightarrow [0, 1]$ assigns each statement $\alpha \in \mathcal{X}$ a strength $\sigma^{\mathcal{G}}(\alpha) \in [0, 1]$.⁴

In essence, GS offer a principled approach to *evaluating* statements quantitatively. We refer informally throughout the paper to any GS that takes into account the premises of any statement as a *structured* GS, and any that does not as an *abstract* GS. It should be noted that while abstract GS can be applied to SGs, structured GS cannot be applied to frameworks using abstract argumentation such as QBAFs, since the statements therein are abstract entities.

We refer to any statement $\alpha \in \mathcal{X}$ with $\sigma^{\mathcal{G}}(\alpha) = 1$ as having/being *top-strength*, and $\sigma^{\mathcal{G}}(\alpha) = 0$ as having/being *bottom-strength*. We note here that other notions of acceptance and rejection, e.g., with thresholds over the strength range as in (de Tarlé, Bonzon, and Maudet 2022; Rago, Li, and Toni 2023), would be possible in our setting, but we leave the study of this to future work.

Next, we recall two notions from (Jedwabny, Croitoru, and Bisquert 2020).⁵ The first gives the notion indicating when an argument's premises is grounded in facts, while the second defines operators used in the GS.

Definition 4. A complete support tree (CST) for some $\alpha_1 \in \mathcal{X}$ is a set of statements $T \subset \mathcal{X}$ s.t.:

- $\alpha_1 \in T$;
- $\forall \alpha_2 \in T, \forall x \in Prem(\alpha_2), \exists \alpha_3 = \langle \Phi_3, \Psi_3 \rangle \in T$ s.t. $\Psi_3 \equiv x$;

²The proofs for all theoretical results can be found in Appendix B here: <https://arxiv.org/abs/2410.22209>.

³While methods for learning these weights represent an important research direction, e.g., as in (Rago *et al.* 2025) within the setting of review aggregation, they fall outside the scope of this work.

⁴With an abuse of notation, when referring to GS generally, we drop the $\mathcal{G}$ superscript for clarity.

⁵We omit the acyclicity condition since we use acyclic graphs.

- $\nexists \alpha_2, \alpha_3 \in T$ s.t. $(\alpha_2, \alpha_3) \in \mathcal{A}$;
- $\nexists T' \subset T$ s.t. T' is a CST for α_1 (minimal wrt set inclusion).

We denote with $\mathcal{T}^G(\alpha_1)$ the set of CSTs for α_1 wrt $\mathcal{G}$. Given two CSTs $T \in \mathcal{T}^G(\alpha_1)$ and $T' \in \mathcal{T}^G(\alpha_2)$ for $\alpha_1, \alpha_2 \in \mathcal{X}$, respectively, T attacks T' iff $\exists \alpha_3 \in T'$ s.t. $(\alpha_1, \alpha_3) \in \mathcal{A}$.

Definition 5. A triple $(\otimes, \oplus, \neg)$, where $\otimes$ is a T-norm, $\oplus$ is a T-conorm, and $\neg$ is a negation, is called a De Morgan triple iff $\forall v_1, v_2 \in [0, 1]$, $\otimes(v_1, v_2) = \neg(\oplus(\neg(v_1), \neg(v_2)))$, and $\oplus(v_1, v_2) = \neg(\otimes(\neg(v_1), \neg(v_2)))$.

Then, Jedwabny *et al.* (2020) introduce the following GS.

Definition 6. Given a De Morgan triple $(\otimes, \oplus, \neg)$, the T-norm semantics σ_T is a structured GS s.t. for any $\alpha \in \mathcal{X}$, where $\sigma_T^G(\alpha) = \bigoplus_{T \in \mathcal{T}^G(\alpha)} \mathcal{O}(T)$, $\mathcal{O}(T) = \mathcal{I}(T) \otimes \neg \bigoplus_{T' \in \mathcal{T}^G(\alpha' \in \mathcal{X}), T' \text{ attacks } T} \mathcal{I}(T')$ and $\mathcal{I}(T) = \bigotimes_{\alpha \in T} \tau(\alpha)$.

While the T-norm semantics is modular wrt the De Morgan triple, the methodology is based on a statement's CSTs, and thus it can be seen that it inherently relies on complete information, i.e., a statement's grounding in facts.

In this paper, we instantiate two T-norm semantics. For the first, which we name the *T-norm-p* semantics, denoted by σ_{T_p} , we let, as in the illustrations by Jedwabny *et al.* (2020): $\otimes$ be the product, i.e., $\otimes(v_1, v_2) = v_1 \times v_2$; $\oplus$ be the probabilistic sum, i.e., $\oplus(v_1, v_2) = v_1 + v_2 - (v_1 \times v_2)$; and $\neg$ be standard negation, i.e., $\neg(v_1) = 1 - v_1$. For the second, which we name the *T-norm-m* semantics, denoted by σ_{T_m} , we let: $\otimes$ be the minimum, i.e., $\otimes(v_1, v_2) = \min(v_1, v_2)$; $\oplus$ be the maximum, i.e., $\oplus(v_1, v_2) = \max(v_1, v_2)$; and $\neg$ be standard negation as for the T-norm-p semantics.

We now recap two abstract GS from the literature, as defined by Rago *et al.* (2016) and Potyka (2018), respectively.

Definition 7. The DF-QuAD semantics σ_D is an abstract GS s.t. for any $\alpha \in \mathcal{X}$, $\sigma_D^G(\alpha) = c(\tau(\alpha), \Sigma(\sigma_D^G(\mathcal{A}(\alpha))), \Sigma(\sigma_D^G(\mathcal{S}(\alpha))))$ where, for any $S \subseteq \mathcal{X}$, $\sigma_D^G(S) = (\sigma_D^G(\alpha_1), \dots, \sigma_D^G(\alpha_k))$ for $(\alpha_1, \dots, \alpha_k)$, an arbitrary permutation of S , and:

- Σ is s.t. $\Sigma(()) = 0$, where $()$ is an empty sequence, and, for $v_1, \dots, v_n \in [0, 1]$ ($n \geq 1$), if $n = 1$, then $\Sigma((v_1)) = v_1$; if $n = 2$, then $\Sigma((v_1, v_2)) = v_1 + v_2 - v_1 \cdot v_2$; and if $n > 2$, then $\Sigma((v_1, \dots, v_n)) = \Sigma(\Sigma((v_1, \dots, v_{n-1})), v_n)$;
- c is s.t., for $v^0, v^-, v^+ \in [0, 1]$, if $v^- \geq v^+$, then $c(v^0, v^-, v^+) = v^0 - v^0 \cdot |v^+ - v^-|$ and if $v^- < v^+$, then $c(v^0, v^-, v^+) = v^0 + (1 - v^0) \cdot |v^+ - v^-|$.

Definition 8. The QEM semantics⁶ σ_Q is an abstract GS s.t. for any $\alpha \in \mathcal{X}$, $\sigma_Q^G(\alpha) = \tau(\alpha) + (1 - \tau(\alpha)) \cdot h(E_\alpha) - \tau(\alpha) \cdot h(-E_\alpha)$ where $E_\alpha = \sum_{\alpha_1 \in \mathcal{S}(\alpha)} \sigma_Q^G(\alpha_1) - \sum_{\alpha_2 \in \mathcal{A}(\alpha)} \sigma_Q^G(\alpha_2)$ and for all $v \in \mathbb{R}$, $h(v) = \frac{\max\{v, 0\}^2}{1 + \max\{v, 0\}^2}$.

While these abstract GS do not take into account the structure of statements, they do consider a weight on each statement, which gives an intrinsic strength which may be used as a starting point for calculating statement strengths, e.g., in the case of incomplete information.

⁶We define a simplified GS for the case of acyclic graphs.




Figure 2: SG from Example 1 with initial (added) statements shown in solid (striped, respectively) turquoise, no attacks, initial (added) supports as solid (dashed, respectively) green arrows and some arbitrary τ (not shown here). The statement $\alpha'_1 \in \mathcal{X}$ is incomplete when $\mathcal{X} = \{\alpha'_1, \alpha'_2, \alpha'_3\}$, but is partially-complete when α'_4 (or α'_5) is added and complete when both α'_4 and α'_5 are added.

4 Gradual Semantics for Statement Graphs

We now introduce a novel methodology for defining structured GS, leveraging abstract GS, to evaluate the strength of statements in SGs (§4.1). We then redefine and discuss the suitability of existing (redefined for this setting) and novel properties for GS in SGs (§4.2), before we evaluate our approach with the properties, comparing against existing structured and abstract GS (§4.3). These contributions will be discussed with a focus of the following notion.

Definition 9. For any $\alpha_1 \in \mathcal{X}$, we say α_1 is:

- complete iff $\forall \alpha_2 \in \mathcal{X}$, where $\alpha_2 = \alpha_1$ or there exists a path from α_2 to α_1 via $\mathcal{S}$, $\mathcal{T}^G(\alpha_2) \neq \emptyset$;
- partially-complete iff it is not complete and $\mathcal{T}^G(\alpha_1) \neq \emptyset$;
- incomplete otherwise, i.e., iff $\mathcal{T}^G(\alpha_1) = \emptyset$.

Note that partially-complete and incomplete statements are in the same spirit as *potential arguments* in ABA (Toni 2013), but the latter are used for computational purposes, as steps towards complete statements, rather than as representations in their own right, as we do. Further investigation, e.g. on the equivalent in ASPIC⁺, is left to future work.

Example 1. Consider the SG in Figure 2, in which $\mathcal{X} = \{\alpha'_1, \alpha'_2, \alpha'_3\}$, $\alpha'_1 = \langle b, a \rangle$, $\alpha'_2 = \langle c, b \rangle$ and $\alpha'_3 = \langle d, b \rangle$, giving $\mathcal{S} = \{(\alpha'_2, \alpha'_1), (\alpha'_3, \alpha'_1)\}$ and $\mathcal{A} = \emptyset$ (the τ values are irrelevant here). In this first case, α'_1 incomplete as it has no CSTs. If we add the statement $\alpha'_4 = \langle \tau, c \rangle$, leading to $\mathcal{S} = \{(\alpha'_2, \alpha'_1), (\alpha'_3, \alpha'_1), (\alpha'_4, \alpha'_2)\}$, α'_1 becomes partially-complete as it has one CST, but α'_3 , which has a path to α'_1 , does not have a CST. Finally, if we add the statement $\alpha'_5 = \langle \tau, d \rangle$, leading to $\mathcal{S} = \{(\alpha'_2, \alpha'_1), (\alpha'_3, \alpha'_1), (\alpha'_4, \alpha'_2), (\alpha'_5, \alpha'_3)\}$, α'_1 becomes complete as all statements with a path to it have CSTs. Crucial here is the fact that in the incomplete case the T-norm semantics ignores the (potentially useful) information from both α'_2 and α'_3 , regardless of τ , while in the partially-complete case it ignores that of α'_3 .

To solve this issue, we target a methodology which behaves well in all cases, i.e., considering all available information.




Figure 3: The conventional way (left) in which abstract GS (σ_*) are used for SGs (centre), with a strength calculated for α_1 using the strengths of its supporters (α_2 and α_4) and attacker (α_3), and our approach (right), modular structured GS (σ_M), which calculate the strength, using σ_* , of each literal in α_1 's premises (b and c) within the premise graph of α_1 , before σ_M aggregates these strengths in a premise aggregation function to calculate the strength of α_1 .

4.1 Methodology

First, we introduce the following notion.

Definition 10. A premise graph for some $\alpha_1 \in \mathcal{X}$ is an SG $\mathcal{G}^1 = \langle \mathcal{X}^1, \mathcal{A}^1, \mathcal{S}^1, \tau^1 \rangle$ s.t. $\mathcal{X}^1 = \text{Prem}(\alpha_1) \cup \mathcal{A}(\alpha_1) \cup \mathcal{S}(\alpha_1)$ and:

- $\mathcal{A}^1 = \{ \langle \langle \Phi_2, \Psi_2 \rangle, x \rangle \in \mathcal{A}(\alpha_1) \times \text{Prem}(\alpha_1) \mid \neg x \equiv \Psi_2 \} ;$
- $\mathcal{S}^1 = \{ \langle \langle \Phi_2, \Psi_2 \rangle, x \rangle \in \mathcal{S}(\alpha_1) \times \text{Prem}(\alpha_1) \mid x \equiv \Psi_2 \} .$

A statement's premise graph contains a new set of statements representing the statement's premises, attackers and supporters from the original SG, as shown in the dashed blue box on the right of Figure 3. Note that we do not specify τ^1 at this point; we will do so in the specific instantiation of our methodology. In introducing a premise graph, we allow for the literals in statements' premises to be evaluated by abstract gradual semantics, a core tenet of our approach.

Our methodology, illustrated in Figure 3, is as follows.⁷

Definition 11. Given an abstract GS σ_* , a modular structured GS supported by σ_* is a structured GS σ_M s.t. for any $\alpha_1 \in \mathcal{X}$, where $\mathcal{G}^1$ is the premise graph of $\mathcal{G}$ for α_1 :

$$\sigma_M^{\mathcal{G}}(\alpha_1) = \odot(\{ \sigma_*^{\mathcal{G}^1}(x) \mid x \in \text{Prem}(\alpha_1) \})$$

where $\odot : [0, 1]^* \rightarrow [0, 1]$ is a premise aggregation function.

The modular structured GS thus dialectically evaluates the evidence supporting or attacking each literal in a statement's premise via an abstract GS using the premise graph, rather than based on the completeness of its support as in (Jedwabny, Croitoru, and Bisquert 2020). These dialectical evaluations are then used to evaluate the statement with the premise aggregation function. This permits the incorporation of incomplete information, since the evidence for or against premises is combined dialectically starting from a statement's weight, giving an approximation of the strength of a statement based on the *available* evidence, giving desirable behaviour in all three cases in Definition 9. This modular methodology permits instantiations tailored to individual applications based on choices in Definition 11. We now define one such instantiation, tailored to our restricted

SGs of Definition 2, where premises are conjunctions, with its supporting abstract GS left unspecified.

Definition 12. Given an abstract GS σ_* , the dialectical conjunction (DC) semantics supported by σ_* , denoted by $\sigma_{\wedge*}$, is a modular structured GS s.t., for any $\alpha_1 \in \mathcal{X}$, if $\text{Prem}(\alpha_1) = \emptyset$, then $\sigma_{\wedge*}^{\mathcal{G}}(\alpha_1) = \tau(\alpha_1)$; otherwise:

$$\sigma_{\wedge*}^{\mathcal{G}}(\alpha_1) = \bigotimes_{x \in \text{Prem}(\alpha_1)} \sigma_*^{\mathcal{G}^1}(x)$$

where $\tau^1(x) = \sqrt[n]{\tau(\alpha_1)}$, $n = |\text{Prem}(\alpha_1)|$ and $\tau^1(\alpha_2) = \sigma_{\wedge*}^{\mathcal{G}}(\alpha_2) \forall \alpha_2 \in \mathcal{X}^1 \setminus \text{Prem}(\alpha_1)$.

Corollary 2. For any abstract GS σ_* that satisfies existence and uniqueness, $\sigma_{\wedge*}$ also satisfies existence and uniqueness.

Note that conjunctive premises represent a challenge for handling incomplete information, as the failure of any single premise invalidates the entire statement. For that matter, the DC semantics has been designed such that certain desirable theoretical properties will be satisfied, as we will see in §4.2. Importantly, the DC semantics represents a natural starting point for our methodology because conjunctive premises are typically abundant in real-world argumentative reasoning. For instance, when people construct arguments, they typically combine multiple pieces of evidence conjunctively to support their claims. This pattern is evident in scientific argumentation, legal reasoning and everyday discourse, where people naturally aggregate supporting facts with “and” connectives (Mercier and Sperber 2011). In future work, modular structured GS could be instantiated for different structures of SGs, e.g., for SGs with disjunctions as premises, a probabilistic sum could be used for the premise aggregation function. This would align with the intuition that multiple independent supporting lines of reasoning may strengthen a statement's overall strength. For other logical structures such as statements in conjunctive normal form, a hierarchical aggregation approach could be employed, i.e., first combining literals within clauses using disjunctive operators, then aggregating clauses conjunctively.

Example 2. Consider the SG from Figure 1, where $\mathcal{S} = \{(\alpha_2, \alpha_1), (\alpha_3, \alpha_1)\}$ and $\mathcal{A} = \{(\alpha_4, \alpha_1)\}$. Given the DC

⁷Given the modular nature of our approach, we leave a study of the complexities of the resulting GS to future work.

semantics supported by DF-QuAD, the strength of statement α_1 is calculated as follows: $\sigma_{\wedge_D}^G(\alpha_1) = \sigma_D^G(a) \times \sigma_D^G(b) = 0.916 \times 0.957 = 0.877$, where $\tau^1(a) = \tau^1(b) = \sqrt[2]{0.8} = 0.894$. In contrast, the T-norm-p semantics produces a strength of: $\sigma_{T_p}^G(\alpha_1) = \tau(\alpha_1) \times \tau(\alpha_2) \times \tau(\alpha_3) = 0.8 \times 0.9 \times 0.6 = 0.432$. Here, we see how the DC and T-norm-p semantics differ. The DC semantics considers the individual strengths of the premises of α_1 , while the T-norm-p semantics takes into account the weight of α_1 and its direct supporters.

4.2 Properties

We now discuss existing (from (Jedwabny, Croitoru, and Bisquert 2020)⁸) and novel properties for GS in SGs. Note that the properties apply to either all or exclusively structured GS as specified.

The first property (Jedwabny, Croitoru, and Bisquert 2020) is a fundamental requirement for GS, stating that a statement's strength only depends on statements that are connected to it, and are thus relevant to it.

Property 1. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X} \cup \{\alpha_1\}, \mathcal{A}', \mathcal{S}', \tau' \rangle$, where $\alpha_2 = \alpha_1$ or $\alpha_3 = \alpha_1 \forall (\alpha_2, \alpha_3) \in (\mathcal{A}' \cup \mathcal{S}') \setminus (\mathcal{A} \cup \mathcal{S})$, a GS σ satisfies directionality iff for any $\alpha_2 \in \mathcal{X}$, where there exists no path from α_1 to α_2 via $\mathcal{A} \cup \mathcal{S}$ and $\tau'(\alpha_3) = \tau(\alpha_3) \forall \alpha_3 \in \mathcal{X}$, it holds that $\sigma^{\mathcal{G}'}(\alpha_2) = \sigma^{\mathcal{G}}(\alpha_2)$.

As in previous works (Amgoud and Ben-Naim 2018; Amgoud, Doder, and Vesic 2022), we would advocate that this property is satisfied in the vast majority of settings since we would not expect information that is unrelated to a statement to have an effect on its strength.

Next, Jedwabny et al. (2020) introduced a property which requires that the length of a proof in the logic, i.e., a chain of reasoning, should not affect the strength of a statement.

Property 2. A structured GS σ satisfies rewriting iff for any $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{X}$, if:

- $\alpha_1 = \langle \Phi_1, \Psi_1 \rangle$, where $\Phi_1 \equiv x_1 \wedge \dots \wedge x_n \wedge x_{n+1}$ and $\Psi_1 \equiv x$;
- $\alpha_2 = \langle \Phi_2, \Psi_2 \rangle$, where $\Phi_2 \equiv x_1 \wedge \dots \wedge x_n$ and $\Psi_2 \equiv x'$;
- $\alpha_3 = \langle \Phi_3, \Psi_3 \rangle$, where $\Phi_3 \equiv x' \wedge x_{n+1}$ and $\Psi_3 \equiv x$;
- $\exists \langle \Phi_4, \Psi_4 \rangle$ where $x' \in \text{Prem}(\langle \Phi_4, \Psi_4 \rangle)$ or $x' \equiv \Psi_4$;
- $\tau(\alpha_1) = \tau(\alpha_3)$ and $\tau(\alpha_2) = 1$;

then $\sigma^{\mathcal{G}}(\alpha_1) = \sigma^{\mathcal{G}}(\alpha_3)$.

This property is indeed intuitive in a number of settings where we would prioritise the logical consistency over approximate evaluations of statements. However, in some settings, e.g., when modelling human discourse, it may be the case that the same logical conclusions, when reached over the course of multiple statements, are weaker, or even stronger, than those stated succinctly (e.g., see the study of priority in (Yin, Potyka, and Toni 2024)). For example,

given the statements instantiated in rewriting, consider the effect of adding a statement $\alpha_4 = \langle \Phi_4, \Psi_4 \rangle$, where $\Psi \equiv \neg x_1$. In some settings it may be desirable that the weakening effect of α_4 on α_3 is lesser than that on α_1 due to the distance. We leave a formalisation of this (non-trivial) property to future work. We would, however, argue that rewriting may not be universally applicable, as illustrated below.

Example 3. Continuing from Example 2, let us introduce two additional statements in the SG of Figure 1: $\alpha_5 = \langle a, e \rangle$ (“CO₂ emissions are rising, therefore extreme weather events are increasing”), and $\alpha_6 = \langle e \wedge b, c \rangle$ (“Extreme weather events and global temperatures are increasing, therefore climate change is human-caused”). Let $\tau(\alpha_5) = 1$ and $\tau(\alpha_6) = 0.8$. $\mathcal{A} = \{(\alpha_4, \alpha_1), (\alpha_4, \alpha_5)\}$. Under the DC semantics, we get $\sigma_{\wedge_D}^G(\alpha_1) = 0.877$ and $\sigma_{\wedge_D}^G(\alpha_6) = \sigma_D^G(e) \times \sigma_D^G(b) = 1 \times 0.957 = 0.957$. In contrast, under T-norm-p semantics, we get $\sigma_{T_p}^G(\alpha_1) = \sigma_{T_p}^G(\alpha_6) = 0.432$. We can see that the DC semantics violate the rewriting property, as α_1 and α_6 have different strengths despite making the same claim through different reasoning paths. The T-norm-p semantics, however, is agnostic to the different lengths of the reasoning chain.

Next, Jedwabny et al. (2020) state that a statement without at least one supporter for each of its premises, i.e., with no evidence for any of them, is bottom-strength.

Property 3. A structured GS σ satisfies provability iff for any $\alpha_1 \in \mathcal{X}$, if $\exists x \in \text{Prem}(\alpha_1)$ and $\nexists \langle \Phi_2, \Psi_2 \rangle \in \mathcal{S}(\alpha_1)$ s.t. $x \equiv \Psi_2$, then $\sigma^{\mathcal{G}}(\alpha_1) = 0$.

This property is clearly useful in cases under complete information, i.e., first case in Definition 9, but in the other two cases, we posit that this property may be too demanding.

Example 4. In our running example, consider $\alpha_4 = \langle d, \neg a \rangle$. Note that, while α_4 is an attacker of α_1 and α_5 , it is incomplete because its premise d does not have a support. Under the T-norm-p semantics, the strength of α_4 is $\sigma_{T_p}^G(\alpha_4) = 0$. Consequently, it has no effect on the strengths of α_1 and α_5 . In contrast, under the DC semantics, the strength of α_4 defaults to its weight, i.e., $\sigma_{\wedge_D}^G(\alpha_4) = \tau(\alpha_4) = 0.7$. This non-zero strength is then used when computing the strengths of α_1 and α_5 , thus reducing their overall strength.

As an optional alternative, we propose a weaker property (trivially implied by provability), which makes the same demands but only when the statement's weight is zero.

Property 4. A structured GS σ satisfies weak provability iff for any $\alpha_1 \in \mathcal{X}$ such that $\tau(\alpha_1) = 0$, if $\exists x \in \text{Prem}(\alpha_1)$ and $\nexists \langle \Phi_2, \Psi_2 \rangle \in \mathcal{S}(\alpha_1)$ s.t. $x \equiv \Psi_2$, then $\sigma^{\mathcal{G}}(\alpha_1) = 0$.

Corollary 3. A structured GS σ which satisfies provability necessarily satisfies weak provability.

In Example 3, weak provability would not require that α_4 's strength is zero, allowing a form of “trust” in the statement's weight in the absence of complete information providing the logical reasoning to support its premises. This means the information is not lost in the partially-complete or incomplete cases, unless it is assigned the minimal weight by the GS, which essentially indicates it cannot be trusted.

⁸Jedwabny et al. (2020) also defined properties that are specific to CSTs, which we reformulate and assess in Appendix A here: <https://arxiv.org/abs/2410.22209>.

We now propose more novel properties that offer alternatives to provability, firstly considering the “initial state” of statements, i.e., that the strength of a statement with no attackers or supporters should be equal to its weight, as in properties defined for abstract GS, namely *stability* (Amgoud and Ben-Naim 2018) and *balance* (Baroni, Rago, and Toni 2019). It is also related to *logical void precedence* (Heyninck, Raddaoui, and Straßer 2023).

Property 5. A GS σ satisfies stability iff for any $\alpha \in \mathcal{X}$, if $\mathcal{A}(\alpha) = \mathcal{S}(\alpha) = \emptyset$, then $\sigma^G(\alpha) = \tau(\alpha)$.

The desirability of this property is clear: as demonstrated below, it gives an intuitive starting point for the strength of a statement, regardless of its completeness from Definition 9.

Example 5. Using statement $\alpha_4 = \langle d, \neg a \rangle$ again, note that $\mathcal{S}(\alpha_4) = \mathcal{A}(\alpha_4) = \emptyset$, meaning that α_4 has no supporters or attackers. Under the DC semantics, $\sigma_{\wedge D}^G(\alpha_4) = \tau(\alpha_4) = 0.8$, whereas, under the T-norm-p semantics, $\sigma_{T_p}^G(\alpha_4) = 0$. The DC semantics respects the stability property, defaulting to the weight for α_4 's strength. The T-norm-p semantics, however, rejects α_4 due to its lack of a CST.

Stability then works in tandem with the following properties such that the behaviour of semantics can be proven to be desirable iteratively over any state under any completeness.

First, we consider the effect of adding a statement that has been assigned a zero strength, i.e., it should not affect the strength of any statement it attacks or supports, as in *neutrality* (Amgoud and Ben-Naim 2018) for abstract GS.

Property 6. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X} \cup \{\alpha_1\}, \mathcal{A}', \mathcal{S}', \tau' \rangle$, a GS σ satisfies neutrality iff for any $\alpha_2 \in \mathcal{X}$, where $\mathcal{A}'(\alpha_2) \cup \mathcal{S}'(\alpha_2) = \mathcal{A}(\alpha_2) \cup \mathcal{S}(\alpha_2) \cup \{\alpha_1\}$, $\sigma^{G'}(\alpha_3) = \sigma^G(\alpha_3) \forall \alpha_3 \in \mathcal{A}(\alpha_2) \cup \mathcal{S}(\alpha_2)$, $\tau'(\alpha_2) = \tau(\alpha_2)$ and $\sigma^{G'}(\alpha_1) = 0$, $\sigma^{G'}(\alpha_2) = \sigma^G(\alpha_2)$.

This property, as with stability, takes a notably different approach to the T-norm semantics given their focus on CSTs, namely that the strengths of statements are calculated recursively, but should be unaffected by bottom-strength statements, which must be unsubstantiated, as illustrated below.

Example 6. Let us revisit our running example and let $\alpha_7 = \langle \top, d \rangle$ (“It is a fact that CO₂ measurements are unreliable”) with $\tau(\alpha_7) = 0$. Then, under the DC and T-norm-p semantics, $\sigma_{\wedge D}^G(\alpha_7) = \sigma_{T_p}^G(\alpha_7) = 0$, i.e., α_7 is bottom-strength. However, the strength of α_4 under DC is $\sigma_{\wedge D}^G(\alpha_4) = 0.8$ (unchanged). However, under the T-norm-p semantics, it is still $\sigma_{T_p}^G(\alpha_4) = 0$ (bottom-strength). Under the DC semantics, α_7 is bottom-strength and, thus, has no influence on the strength of α_4 . In contrast, the bottom-strength of α_7 under T-norm-p leads to the bottom-strength of α_4 due to the multiplicative property, thus violating neutrality.

The next two properties consider the effect of an attacking or supporting statement being added which is not bottom-strength, taking inspiration from *bi-variate monotony* (Amgoud and Ben-Naim 2018), *monotonicity* (Baroni, Rago, and Toni 2019) and *P3* in (Skiba, Thimm, and Wallner 2023). The first requires that adding an attacker should not strengthen a statement.

Property 7. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X} \cup \{\alpha_1\}, \mathcal{A}', \mathcal{S}', \tau' \rangle$, a GS σ satisfies attacked premise iff for any $\alpha_2 \in \mathcal{X}$, where $\mathcal{A}'(\alpha_2) = \mathcal{A}(\alpha_2) \cup \{\alpha_1\}$, $\mathcal{S}'(\alpha_2) = \mathcal{S}(\alpha_2)$, $\sigma^{G'}(\alpha_3) = \sigma^G(\alpha_3) \forall \alpha_3 \in \mathcal{A}(\alpha_2) \cup \mathcal{S}(\alpha_2)$ and $\tau'(\alpha_2) = \tau(\alpha_2)$, $\sigma^{G'}(\alpha_2) \leq \sigma^G(\alpha_2)$.

Next, adding a supporter should not weaken a statement.

Property 8. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X} \cup \{\alpha_1\}, \mathcal{A}', \mathcal{S}', \tau' \rangle$, a GS σ satisfies supported premise iff for any $\alpha_2 \in \mathcal{X}$, where $\mathcal{A}'(\alpha_2) = \mathcal{A}(\alpha_2)$, $\mathcal{S}'(\alpha_2) = \mathcal{S}(\alpha_2) \cup \{\alpha_1\}$, $\sigma^{G'}(\alpha_3) = \sigma^G(\alpha_3) \forall \alpha_3 \in \mathcal{A}(\alpha_2) \cup \mathcal{S}(\alpha_2)$ and $\tau'(\alpha_2) = \tau(\alpha_2)$, $\sigma^{G'}(\alpha_2) \geq \sigma^G(\alpha_2)$.

These properties offer an iterative view of how the effects of added attackers and supporters should be governed, which does not depend on CSTs and thus considers cases with partially-complete or incomplete statements. For instance, in Example 6, adding statement α_7 , which is an attacker of α_1 , leads to a decrease in α_1 's overall strength, despite α_7 missing a support for its premise g .

Next, we consider the effect of increasing the strengths of statements that attack or support other statements, as in *bi-variate reinforcement* (Amgoud and Ben-Naim 2018) and *monotonicity* (Baroni, Rago, and Toni 2019). First, strengthening an attacker should not strengthen a statement.

Property 9. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X}', \mathcal{A}', \mathcal{S}', \tau' \rangle$, a GS σ satisfies weakened premise iff for any $\alpha_1 \in \mathcal{X}$, where $\mathcal{A}'(\alpha_1) = \mathcal{A}(\alpha_1)$, $\mathcal{S}'(\alpha_1) = \mathcal{S}(\alpha_1)$, $\tau'(\alpha_1) = \tau(\alpha_1)$, and $\sigma^{G'}(\alpha_2) = \sigma^G(\alpha_2) \forall \alpha_2 \in \mathcal{A}(\alpha_1) \cup \mathcal{S}(\alpha_1) \setminus \{\alpha_3\}$, where $\alpha_3 \in \mathcal{A}(\alpha_1)$ and $\sigma^{G'}(\alpha_3) > \sigma^G(\alpha_3)$, $\sigma^{G'}(\alpha_1) \leq \sigma^G(\alpha_1)$.

Analogously strengthening a supporter should not weaken a statement.

Property 10. Given two SGs $\mathcal{G}$ and $\mathcal{G}' = \langle \mathcal{X}', \mathcal{A}', \mathcal{S}', \tau' \rangle$, a GS σ satisfies strengthened premise iff for any $\alpha_1 \in \mathcal{X}$, where $\mathcal{A}'(\alpha_1) = \mathcal{A}(\alpha_1)$, $\mathcal{S}'(\alpha_1) = \mathcal{S}(\alpha_1)$, $\tau'(\alpha_1) = \tau(\alpha_1)$, and $\sigma^{G'}(\alpha_2) = \sigma^G(\alpha_2) \forall \alpha_2 \in \mathcal{A}(\alpha_1) \cup \mathcal{S}(\alpha_1) \setminus \{\alpha_3\}$, where $\alpha_3 \in \mathcal{S}(\alpha_1)$ and $\sigma^{G'}(\alpha_3) > \sigma^G(\alpha_3)$, $\sigma^{G'}(\alpha_1) \geq \sigma^G(\alpha_1)$.

These properties assess added statements' effects, conditioning on their strengths rather than weights, and direct attackers and supporters rather than CSTs, thus accommodating statements of all three levels of completeness.

Note that, for Properties 6-10, the opposite effect on the statement's strength holds for the reverted modification, i.e., for removing or reducing the strengths of statements.

Next, a bottom-strength premise causes a statement to be bottom-strength. This property is related to *argument death* (Spaans 2021).

Property 11. A structured GS σ satisfies bottom-strength premise iff for any $\alpha_1 \in \mathcal{X}$, where $\exists x \in \text{Prem}(\alpha_1)$ such that $\exists \alpha_2 = \langle \Phi_2, \Psi_2 \rangle \in \mathcal{A}(\alpha_1)$ with $\Psi_2 \equiv \neg x$ and $\sigma^G(\alpha_2) = 1$ and $\nexists \alpha_3 = \langle \Phi_3, \Psi_3 \rangle \in \mathcal{S}(\alpha_1)$ with $\Psi_3 \equiv x$ and $\sigma^G(\alpha_3) > 0$, $\sigma^G(\alpha_1) = 0$.

For partially-complete or incomplete statements especially, this property provides an alternative reason for requiring rejection, i.e., there is maximal evidence to the contrary of one of its premises since they are, critically, part of a conjunction, as opposed to provability, where this is the case

unless every premise in the statement is supported. Note that bottom-strength premise is also desirable for complete statements. We illustrate this property in the following.

Example 7. In our running example, let $\alpha_8 = \langle e \wedge f, \neg b \rangle$ (“Extreme weather events are increasing and solar activity is high, thus global temperatures are not increasing”) with $\tau(\alpha_8) = 0.5$. Now, $\sigma_{\wedge D}^G(\alpha_8) = 0.707$ and $\sigma_{T_p}^G(\alpha_8) = 0$. Let us now add a new statement $\alpha_9 = \langle \top, \neg f \rangle$ (“It is a fact that solar activity is not high”) with $\tau(\alpha_9) = 1$. Then, the strength of α_8 is $\sigma_{\wedge D}^G(\alpha_7) = \sigma_D^{G^1}(e) \times \sigma_D^{G^1}(g) = 0.707 \times 0 = 0$ and $\sigma_{T_p}^G(\alpha_8) = 0$. Both semantics now reject α_8 , but for different reasons. The DC semantics rejects it because one of its premises (f) is fully contradicted by α_9 , while the T-norm-p semantics already gives it bottom-strength due to incomplete support.

Next, statements with universally top-strength premises should also be top-strength.

Property 12. A structured GS σ satisfies top-strength premises iff for any $\alpha_1 \in \mathcal{X}$, where $\forall x \in \text{Prem}(\alpha_1)$, $\nexists \alpha_2 = \langle \Phi_2, \Psi_2 \rangle \in \mathcal{A}(\alpha_1)$ such that $\Psi_2 \equiv \neg x$ and $\sigma^G(\alpha_2) > 0$ and $\exists \alpha_3 = \langle \Phi_3, \Psi_3 \rangle \in \mathcal{S}(\alpha_1)$ such that $\Psi_3 \equiv x$ and $\sigma^G(\alpha_3) = 1$, $\sigma^G(\alpha_1) = 1$.

This property somewhat overrides a statement’s weight when its premises are maximally supported, as in many approaches in argumentation that do not include weights. This seems intuitive particularly in our SGs, where the premises of statements (of any completeness) are conjunctions.

Example 8. In our running example, consider $\alpha_5 = \langle a, e \rangle$, $\alpha_2 = \langle \top, e \rangle$, with $\tau(\alpha_5) = 0.8$, $\tau(\alpha_2) = 1$. Then, $\sigma_{\wedge D}^G(\alpha_5) = 1$, $\sigma_{T_p}^G(\alpha_5) = 0.8$. Despite α_2 fully supporting the premise of α_5 (as the weight of α_2 is 1), the strength of α_5 under the T-norm-p semantics is not 1. In contrast, the DC semantics assigns full strength to α_5 when its premise is fully supported.

Finally, attackers and supporters may affect statements symmetrically, i.e., two statements’ premises that negate one another have “opposite” strengths, as in *duality* (Potyka 2018) and *franklin* (Amgoud and Ben-Naim 2018).⁹

Property 13. A structured GS σ satisfies mirroring iff for any $\alpha_1, \alpha_2 \in \mathcal{X}$, where $\alpha_1 = \langle \Phi_1, \Psi_1 \rangle$, $\alpha_2 = \langle \Phi_2, \Psi_2 \rangle$, $|\text{Prem}(\alpha_1)| = |\text{Prem}(\alpha_2)|$, $\Phi_1 \equiv \neg \Phi_2$ and $\tau(\alpha_1) = \tau(\alpha_2) = 0.5$. Then, $\sigma^G(\alpha_1) = 1 - \sigma^G(\alpha_2)$.

Mirroring seems intuitive and reflects a fairness in how statements are evaluated, potentially inspiring trust in users.

Example 9. In our running example, let $\alpha_{10} = \langle \neg a, \neg e \rangle$ (“CO₂ emissions are not rising, thus extreme weather events are not increasing”) with $\tau(\alpha_5) = \tau(\alpha_{10}) = 0.5$. Then, $\sigma_{\wedge D}^G(\alpha_{10}) = 1 - \sigma_{\wedge D}^G(\alpha_5) = 1 - 0.55 = 0.45$, and $\sigma_{T_p}^G(\alpha_5) = 0$ and $\sigma_{T_p}^G(\alpha_{10}) = 0.45$. Despite α_5 and α_{10} having premises that negate one another and equal base weights, their strengths are not complementary (do not sum to 1) under

⁹For the DC semantics, mirroring’s conditions are triggered only for statements with single literals in their premises. However, this may not be the case for other modular structured GS.

Property	σ_{T_p}	σ_{T_m}	$\sigma_{\wedge D}$	$\sigma_{\wedge Q}$	σ_D	σ_Q
1: Directionality	✓	✓	✓	✓	✓	✓
2: Rewriting	✓	✓	✗	✗	–	–
3: Provability	✓	✓	✗	✗	–	–
4: Weak Provability	✓	✓	✓	✓	–	–
5: Stability	✗	✗	✓	✓	✓	✓
6: Neutrality	✗	✗	✓	✓	✓	✓
7: Attacked Premise	✓	✓	✓	✓	✓	✓
8: Supported Premise	✓	✓	✓	✓	✓	✓
9: Weakened Premise	✗	✗	✓	✓	✓	✓
10: Strengthened Premise	✓	✓	✓	✓	✓	✓
11: Bottom-Strength Premise	✓	✓	✓	✗	–	–
12: Top-Strength Premises	✗	✗	✓	✗	–	–
13: Mirroring	✗	✗	✓	✓	–	–

Table 1: Structured (σ_{T_p} , σ_{T_m} , $\sigma_{\wedge D}$ and $\sigma_{\wedge Q}$) and abstract (σ_D and σ_Q) GS’ satisfaction (denoted with ✓) or violation (denoted with ✗) of the assessed properties, where incompatibility of a GS with a property is denoted with –.

T-norm-p semantics. In contrast, the DC semantics shows a desirable duality in attackers’ and supporters’ effects.

4.3 Theoretical Analysis

We now assess whether instantiations of the DC semantics, as well as existing structured and abstract GS, satisfy the properties from §4.2. Table 1 summarises the results.

First, we assess the **T-norm semantics** in §3, σ_{T_p} and σ_{T_m} .

Proposition 1. σ_{T_p} and σ_{T_m} satisfy Properties 1-4, 7-8, and 10-11, but violate Properties 5-6, 9, and 12-13.

Both GS satisfy identical properties across our collection, performing poorly when we consider the novel properties. Even stability and neutrality, two basic properties in the realm of abstract GS, are violated due to the fact that they do not consider the weight as a starting point, nor the strength as a recursive function, respectively. Attacked and supported premises are both satisfied, while only weakened, and not strengthened, premise is satisfied due to the asymmetry in the way the attackers and supporters are evaluated, which leads to mirroring being violated also. Similarly, bottom-strength premises is satisfied, but top-strength premises is violated due to the (potentially non-maximal) weight of the statement playing a role in the calculation. Given top-strength premises’ fundamental role across statements of any completeness, this represents a significant weakness.

Next, we consider the **DC semantics**, supported by either the DF-QuAD semantics ($\sigma_{\wedge D}$) or the QEM semantics ($\sigma_{\wedge Q}$), as defined in §4.1.

Proposition 2. $\sigma_{\wedge D}$ satisfies Properties 1 and 4-13, but violates Properties 2-3. $\sigma_{\wedge Q}$ satisfies Properties 1, 4-10, and 13, but violates Properties 2-3 and 11-12.

Both DC semantics only satisfy directionality of the existing properties. The violation of the other existing properties (rewriting and provability) can be justified by their incompatibility with the novel properties which we introduce.

Proposition 3. Rewriting, stability and top-strength premises are incompatible.

Proposition 4. Provability and stability are incompatible.

This means that a user selecting a semantics must choose between its satisfaction of one set of properties or another, so our approach represents a viable alternative to the existing work, particularly in cases with incomplete information. Examples 3 and 4 demonstrate what we believe is further justification for this. For example, rewriting is not satisfied by either semantics since more evidence, even if the statements constituting it are not complete, will strengthen a statement, which seems to be reasonable particularly for our settings with partially-complete or incomplete statements. Moreover, provability is violated by design, since we wanted GS that do not necessarily reject a statement when its support is missing, only when the statement’s weight is also minimal, hence the GS’ satisfaction of weak provability instead. Both GS align with stability in treating the weight as a “starting point”, and from there they satisfy attacked and supported premises, as well as weakened and strengthened premises, due to the recursive way they handle attackers and supporters. Interestingly, the DC semantics, when supported by the DF-QuAD semantics, satisfies bottom-strength premise and top-strength premises, but it does not when supported by the QEM semantics. This is due to the saturation effect that is seen at the extremities of the strength scale in the DF-QuAD semantics, giving rise to this intuitive behaviour for this specific type of SG (with conjunctions as premises). Finally, mirroring is satisfied by both GS.

We now assess the two abstract GS, the **DF-QuAD semantics** (σ_D) and the **QEM semantics** (σ_Q), as defined in §3, along the compatible properties.

Proposition 5. σ_D and σ_Q satisfy Properties 1 and 5-10. (They are incompatible with the other properties.)

Although both of the abstract GS satisfy all of the properties with which they are compatible, they are incompatible with many others, showing the advantages of structured GS in considering statements’ logical structure.

While both T-norm semantics perform well in the existing properties which require complete information, under incomplete information, i.e., with partially-complete or incomplete statements, where the novel properties are particularly powerful, there are clear weaknesses. This highlights how the DC semantics, and in particular that supported by DF-QuAD, offer clear advantages, especially given that all of the properties cannot be satisfied by a given semantics. Thus, in settings with varying completeness of information, our approach fills a gap in the literature.

Finally, we give a theorem that demonstrates the conditions under which a DC semantics, supported by a specific abstract GS, aligns with the abstract GS itself, namely when the premises of all statements consist of at most one literal (meaning SGs effectively become QBAFs).

Theorem 1. Let $\mathcal{X}$ be such that $\forall \alpha \in \mathcal{X}, |\text{Prem}(\alpha)| \leq 1$. Then, for any given abstract GS σ_i , $\sigma_{\wedge i} = \sigma_i$.

5 Conclusions and Future Work

In this paper, we introduced a novel methodology for obtaining GS in SGs that naturally accommodate incomplete information. Our modular methodology dialectically evaluates the literals in a statement’s premise before aggregat-

ing these evaluations based on the statement’s construction. This separation allows our semantics to handle incomplete information by leveraging any existing GS for QBAFs to evaluate the available evidence without requiring complete support. We then demonstrated how our DC semantics can effectively leverage abstract GS for structured argumentation. Furthermore, we discussed existing and novel properties for GS in SGs, revealing some incompatibilities between existing properties, e.g., rewriting and provability, and our novel properties, e.g., stability and top-strength premises, showing that users must select between them given a particular contextual setting, e.g., based on the completeness notions we define. Our theoretical analysis demonstrated that the DC semantics supported by the DF-QuAD abstract GS satisfies all of our novel properties, while existing structured GS struggle with cases without complete information.

While this paper establishes the foundations of our novel approach using a toy example for pedagogical reasons, we see significant potential for real-world applications, e.g., in the analysis of debates mapped onto SGs through argument mining (Lawrence and Reed 2019), possibly enhanced by large language models (Cabessa, Hernault, and Mushtaq 2025; Gorur, Rago, and Toni 2025). It is well known that mining structured arguments is a complex problem, not least due to the presence of enthymemes causing incompleteness (Hunter 2022; Stahl et al. 2023). A methodology leveraging GS that can handle incomplete information could alleviate some of the burden that structured argumentation’s requirements for completeness place on these methods.

Our methodology represents an early step forward in GS for structured argumentation, and opens various potentially fruitful directions for future work, in particular targeting more general settings of structured argumentation, e.g., in ABA or ASPIC⁺. First, we would like to broaden our analysis, both within our restricted form of SGs, exploring various other instantiations of our DC semantics, e.g., using other abstract GS such as the *exponent-based restricted semantics* (Amgoud and Ben-Naim 2018), as well as extending to different forms of SGs with more complex logical premises. Specifically, our ability to deal with partially-complete or incomplete statements may be related to the possibility of deriving assumptions in non-flat ABA (Cyras et al. 2017). We would also like to consider whether other properties for abstract GS are desirable, e.g., *open-mindedness* (Potyka 2019) (the satisfaction of which gives a clear advantage of σ_Q over σ_D), *attainability* (Cocarascu, Rago, and Toni 2019), whose suitability in our SGs was not obvious due to the conjunctive premise. It would also be interesting to formalise the interplay between properties, e.g., between provability and stability, which are incompatible without restrictions. Further, it would be interesting to investigate whether our analysis could be applicable to probabilistic argumentation (Kohlas 2003; Dung and Thang 2010; Hunter and Thimm 2014; Gabbay and Rodrigues 2015b; Fazzinga, Flesca, and Furfaro 2018). Finally, we plan to explore how our methodology translates to real-world applications, e.g., in multi-agent model reconciliation (Vasileiou et al. 2024) given GS’ capability for modelling human or machine reasoning (Rago, Li, and Toni 2023).

Acknowledgments

Antonio Rago and Stylianos Loukas Vasileiou contributed equally to this work. Antonio Rago and Francesca Toni are partially funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101020934) and by J.P. Morgan and the Royal Academy of Engineering under the Research Chairs and Senior Research Fellowships scheme (grant agreement no. RCSR2021\11\45). Stylianos Loukas Vasileiou and William Yeoh are partially supported by the National Science Foundation (NSF) under award 2232055. Son Tran is partially supported by the NSF grants #1914635, #2151254, and ExpandAI grant #2025-67022-44266. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or governments.

References

- Alarcón, J. R., and Kampik, T. 2024. QBAF-Py: Quantitative bipolar argumentation in C with Python bindings. In *COMMA*, 361–362.
- Alsinet, T.; Chesñevar, C. I.; Godo, L.; and Simari, G. R. 2008. A logic programming framework for possibilistic argumentation: Formalization and logical properties. *Fuzzy Sets Syst.* 159(10):1208–1228.
- Amgoud, L., and Ben-Naim, J. 2015. Argumentation-based ranking logics. In *AAMAS*, 1511–1519.
- Amgoud, L., and Ben-Naim, J. 2018. Evaluation of arguments in weighted bipolar graphs. *Int. J. Approx. Reason.* 99:39–55.
- Amgoud, L.; Doder, D.; and Vesic, S. 2022. Evaluation of argument strength in attack graphs: Foundations and semantics. *Artif. Intell.* 302:103607.
- Anaissy, C. A.; Suntwal, S.; Surdeanu, M.; and Vesic, S. 2024. On learning bipolar gradual argumentation semantics with neural networks. In *ICAART*, 493–499.
- Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G. R.; Thimm, M.; and Villata, S. 2017. Towards artificial argumentation. *AI Magazine* 38(3):25–36.
- Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds. 2018. *Handbook of Formal Argumentation*. College Publications.
- Baroni, P.; Rago, A.; and Toni, F. 2018. How many properties do we need for gradual argumentation? In *AAAI*, 1736–1743.
- Baroni, P.; Rago, A.; and Toni, F. 2019. From fine-grained properties to broad principles for gradual argumentation: A principled spectrum. *Int. J. Approx. Reason.* 105:252–286.
- Besnard, P., and Hunter, A. 2014. Constructing argument graphs with deductive arguments: a tutorial. *Argument Comput.* 5(1):5–30.
- Besnard, P.; García, A. J.; Hunter, A.; Modgil, S.; Prakken, H.; Simari, G. R.; and Toni, F. 2014. Introduction to structured argumentation. *Argument Comput.* 5(1):1–4.
- Bezou-Vrakatseli, E.; Cocarascu, O.; and Modgil, S. 2024. Ethix: A dataset for argument scheme classification in ethical debates. In *ECAI*, 3628–3635.
- Bistarelli, S.; Taticchi, C.; and Meo, M. C. 2023. An interleaving semantics of the timed concurrent language for argumentation to model debates and dialogue games. *Theory Pract. Log. Program.* 23(6):1307–1333.
- Budan, P. D.; Gonzalez, M. G. E.; Budán, M. C. D.; Martinez, M. V.; and Simari, G. R. 2023. Strength in coalitions: Community detection through argument similarity. *Argument Comput.* 14(3):275–325.
- Cabessa, J.; Hernault, H.; and Mushtaq, U. 2025. Argument mining with fine-tuned large language models. In *COLING*, 6624–6635.
- Cayrol, C., and Lagasque-Schiex, M. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In *ECSQARU*, 378–389.
- Cocarascu, O.; Rago, A.; and Toni, F. 2019. Extracting dialogical explanations for review aggregations with argumentative dialogical agents. In *AAMAS*, 1261–1269.
- Cramer, M., and Dauphin, J. 2020. A structured argumentation framework for modeling debates in the formal sciences. *Journal for General Philosophy of Science* 51:219–241.
- Cyras, K.; Fan, X.; Schulz, C.; and Toni, F. 2017. Assumption-based argumentation: Disputes, explanations, preferences. *FLAP* 4(8).
- Cyras, K.; Rago, A.; Albini, E.; Baroni, P.; and Toni, F. 2021. Argumentative XAI: A survey. In *IJCAI*, 4392–4399.
- de Tarlé, L. D.; Bonzon, E.; and Maudet, N. 2022. Multiagent dynamics of gradual argumentation semantics. In *AAMAS*, 363–371.
- Delobelle, J., and Villata, S. 2019. Interpretability of gradual semantics in abstract argumentation. In *ECSQARU*, 27–38.
- Dung, P. M., and Thang, P. M. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *COMMA*, 171–182.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77(2):321–358.
- Fazzinga, B.; Flesca, S.; and Furfaro, F. 2018. Probabilistic bipolar abstract argumentation frameworks: complexity results. In *IJCAI*, 1803–1809.
- Gabbay, D. M., and Rodrigues, O. 2015a. Equilibrium states in numerical argumentation networks. *Logica Universalis* 9(4):411–473.
- Gabbay, D. M., and Rodrigues, O. 2015b. Probabilistic argumentation: An equational approach. *Logica Universalis* 9(3):345–382.
- García, A. J., and Simari, G. R. 2004. Defeasible logic programming: An argumentative approach. *Theory Pract. Log. Program.* 4(1-2):95–138.

- Goffredo, P.; Cabrio, E.; Villata, S.; Haddadan, S.; and Sanchez, J. T. 2023. Disputool 2.0: A modular architecture for multi-layer argumentative analysis of political debates. In *AAAI*, 16431–16433.
- Gonzalez, M. G. E.; Budán, M. C. D.; Simari, G. I.; and Simari, G. R. 2021. Labeled bipolar argumentation frameworks. *J. Artif. Intell. Res.* 70:1557–1636.
- Gorur, D.; Rago, A.; and Toni, F. 2025. Can large language models perform relation-based argument mining? In *COLING*, 8518–8534.
- Hadoux, E.; Hunter, A.; and Polberg, S. 2023. Strategic argumentation dialogues for persuasion: Framework and experiments based on modelling the beliefs and concerns of the persuadee. *Argument Comput.* 14(2):109–161.
- Hecham, A.; Bisquert, P.; and Croitoru, M. 2018. On a flexible representation for defeasible reasoning variants. In *AAMAS*, 1123–1131.
- Heyninck, J.; Raddaoui, B.; and Straßer, C. 2023. Ranking-based argumentation semantics applied to logical argumentation. In *IJCAI*, 3268–3276.
- Hunter, A., and Thimm, M. 2014. Probabilistic argumentation with incomplete information. In *ECAI*, 1033–1034.
- Hunter, A.; Polberg, S.; and Thimm, M. 2020. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artif. Intell.* 281:103236.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *Int. J. Approx. Reason.* 54(1):47–81.
- Hunter, A. 2022. Understanding enthymemes in deductive argumentation using semantic distance measures. In *AAAI*, 5729–5736.
- Irwin, B.; Rago, A.; and Toni, F. 2022. Forecasting argumentation frameworks. In *KR*.
- Jedwabny, M.; Croitoru, M.; and Bisquert, P. 2020. Gradual semantics for logic-based bipolar graphs using t(co)norms. In *ECAI*, 777–783.
- Kampik, T.; Potyka, N.; Yin, X.; Cyraś, K.; and Toni, F. 2024. Contribution functions for quantitative bipolar argumentation graphs: A principle-based analysis. *Int. J. Approx. Reason.* 173:109255.
- Kohlas, J. 2003. Probabilistic argumentation systems: A new way to combine logic with probability. *J. Appl. Log.* 1(3-4):225–253.
- Lawrence, J., and Reed, C. 2019. Argument mining: A survey. *Computational Linguistics* 45(4):765–818.
- Mantadelis, T., and Bistarelli, S. 2020. Probabilistic abstract argumentation frameworks, a possible world view. *Int. J. Approx. Reason.* 119:204–219.
- Mercier, H., and Sperber, D. 2011. Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences* 34(2):57–74.
- Modgil, S., and Prakken, H. 2014. The *ASPIC*⁺ framework for structured argumentation: a tutorial. *Argument Comput.* 5(1):31–62.
- Odekerken, D.; Lehtonen, T.; Borg, A.; Wallner, J. P.; and Järvisalo, M. 2023. Argumentative reasoning in *ASPIC*⁺ under incomplete information. In *KR*, 531–541.
- Oluokun, B.; Paulino-Passos, G.; Rago, A.; and Toni, F. 2024. Predicting human judgement in online debates with argumentation. In *CMNA*, 70–80.
- Oren, N.; Yun, B.; Vesic, S.; and Baptista, M. S. 2022. Inverse problems for gradual semantics. In *IJCAI*, 2719–2725.
- Polberg, S., and Hunter, A. 2018. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *Int. J. Approx. Reason.* 93:487–543.
- Polberg, S.; Hunter, A.; and Thimm, M. 2017. Belief in attacks in epistemic probabilistic argumentation. In *SUM*, 223–236.
- Potyka, N. 2018. Continuous dynamical systems for weighted bipolar argumentation. In *KR*, 148–157.
- Potyka, N. 2019. Open-mindedness of gradual argumentation semantics. In *SUM*, 236–249.
- Potyka, N. 2021. Interpreting neural networks as quantitative argumentation frameworks. In *AAAI*, 6463–6470.
- Potyka, N. 2022. Attractor - A Java library for gradual bipolar argumentation. In *COMMA*, 369–370.
- Prakken, H. 2024. An abstract and structured account of dialectical argument strength. *Artif. Intell.* 335:104193.
- Rago, A.; Toni, F.; Aurisicchio, M.; and Baroni, P. 2016. Discontinuity-free decision support with quantitative argumentation debates. In *KR*, 63–73.
- Rago, A.; Cocarascu, O.; Bechlivanidis, C.; Lagnado, D. A.; and Toni, F. 2021. Argumentative explanations for interactive recommendations. *Artif. Intell.* 296:103506.
- Rago, A.; Cocarascu, O.; Oksanen, J.; and Toni, F. 2025. Argumentative review aggregation and dialogical explanations. *Artif. Intell.* 340:104291.
- Rago, A.; Li, H.; and Toni, F. 2023. Interactive explanations by conflict resolution via argumentative exchanges. In *KR*, 582–592.
- Ruiz-Dolz, R.; Heras, S.; and García-Fornes, A. 2023. Automatic debate evaluation with argumentation semantics and natural language argument graph networks. In *EMNLP*, 6030–6040.
- Skiba, K.; Thimm, M.; and Wallner, J. P. 2023. Ranking-based semantics for assumption-based argumentation. In *FCR*, 44–52.
- Spaans, J. P. 2021. Intrinsic argument strength in structured argumentation: A principled approach. In *CLAR*, 377–396.
- Stahl, M.; Düsterhus, N.; Chen, M.; and Wachsmuth, H. 2023. Mind the gap: Automated corpus creation for enthymeme detection and reconstruction in learner arguments. In *EMNLP*, 4703–4717.
- Thimm, M.; Polberg, S.; and Hunter, A. 2018. Epistemic attack semantics. In *COMMA*, 37–48.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *ECAI*, 750–755.

- Toni, F. 2013. A generalised framework for dispute derivations in assumption-based argumentation. *Artif. Intell.* 195:1–43.
- Vasileiou, S. L.; Kumar, A.; Yeoh, W.; Son, T. C.; and Toni, F. 2024. Dialectical reconciliation via structured argumentative dialogues. In *KR*.
- Vassiliades, A.; Bassiliades, N.; and Patkos, T. 2021. Argumentation and explainable artificial intelligence: a survey. *Knowl. Eng. Rev.* 36:e5.
- Vesic, S.; Yun, B.; and Teovanovic, P. 2022. Graphical representation enhances human compliance with principles for graded argumentation semantics. In *AAMAS*, 1319–1327.
- Visser, J.; Konat, B.; Duthie, R.; Koszowy, M.; Budzynska, K.; and Reed, C. 2020. Argumentation in the 2016 US presidential elections: annotated corpora of television debates and social media reaction. *Lang. Resour. Evaluation* 54(1):123–154.
- Wang, Z., and Shen, Y. 2024. Bilateral gradual semantics for weighted argumentation. In *AAAI*, 10732–10739.
- Yin, X.; Potyka, N.; and Toni, F. 2023. Argument attribution explanations in quantitative bipolar argumentation frameworks. In *ECAI*, 2898–2905.
- Yin, X.; Potyka, N.; and Toni, F. 2024. CE-QArg: Counterfactual explanations for quantitative bipolar argumentation frameworks. In *KR*.
- Yun, B., and Vesic, S. 2021. Gradual semantics for weighted bipolar SETAFs. In *ECSQARU*, 201–214.