Argument Networks: Structural and Semantic Exploration of Argumentative Discourse

Bachelorarbeit zur Erlangung des akademischen Grades

Bachelor of Science

im Studiengang Code & Context

an der Fakultät für Informatik und Ingenieurswissenschaften

der Technischen Hochschule Köln

vorgelegt von:	Anthony Pierre Aflatoun
Matrikel-Nr.:	111 500 27
Adresse:	Grüner Hof 30
	50739 Köln
	anthony_pierre.aflatoun@smail.th-koeln.de
eingereicht bei:	Prof. Dr. Stefan Bente
Zweitgutachterin:	Prof. Dr. Ivonne Preusser

Köln, 13.01.2025

Abstract

This thesis investigates the integration of structural and semantic elements in modeling argumentative discourse using enriched argumentation graphs. Traditional frameworks often prioritize either structural relationships or semantic content, rarely combining the two effectively. By employing graph embeddings and advanced textual embeddings from large language models, this study proposes a framework to capture the nuanced interplay of arguments in complex discourse.

The methodology involves constructing enriched argumentation graphs that integrate structural and semantic insights, offering a comprehensive representation of argumentative interactions. Techniques such as dimensionality reduction and clustering were used to evaluate the effectiveness of these enriched embeddings. Results show that combining structural and semantic dimensions enhances the clarity and interpretability of argumentation models, outperforming traditional approaches in representing nuanced roles and relationships.

The research contributes to computational argumentation and discourse analysis, providing a foundation for practical applications such as discourse evaluation, debate summarization, and semantic search. Future work could focus on automating argumentation graph construction, integrating multimodal data, and incorporating temporal dynamics to analyze discourse evolution. While limitations include reliance on synthetic datasets and computational demands, this study offers a meaningful step toward bridging structural and semantic insights in computational argumentation.

Keywords: *Computational Linguistics, Computational Argumentation, Natural Language Processing, Network Science, Abstract Argumentation Framework, Argumentation Graph, Representation Learning, Discourse Analysis, Collective Sensemaking, […]*

Table of Contents

List of Tables

List of Figures

1 Introduction

In the contemporary digital landscape, public communication platforms have become instrumental in shaping discourse and facilitating collective decision-making processes. For instance, recent studies have shown that social media platforms like Twitter and Facebook influence political engagement, while online forums often serve as critical venues for community decision-making, highlighting the significant role these platforms play in public discourse [1], [2]. However, as these platforms have expanded in both scale and complexity, traditional methods of discourse analysis have encountered substantial challenges in effectively capturing the full scope of argumentative interactions that take place within them. The rapid, dynamic nature of digital communication introduces levels of complexity — such as speed, scale, fragmentation and diversity of participants — that challenge conventional models due to the need for rapid processing, scalability across diverse conversations, and capturing varied participant perspectives. The transformation of public spaces in the digital age echoes themes critically examined by Habermas in *The Structural Transformation of the Public Sphere* [3, pp. 141–175], where he explores shifts in discourse resulting from changes in communication structures.

Modern digital platforms, including social media networks and online forums, prioritize engagement-driven content, often due to advertising revenue and user retention metrics, which frequently amplifies sensational or polarizing messages, as highlighted by the findings that algorithmic content ranking and user choices limit exposure to diverse viewpoints [4]. Pariser highlights how these algorithms orchestrate not only the advertisements we see but also our entire digital experience, creating a 'filter bubble' that alters how we encounter ideas and information [5, p. 10].

Empirical research further substantiates these claims, demonstrating the extent to which social media can exacerbate polarization within public discourse [6, p. 9216]. Sunstein similarly argues that echo chambers, intensified by social media, undermine effective governance, making it increasingly difficult to converge on sensible solutions [7, p. 10]. The algorithmic mechanisms underpinning content curation inherently favor posts that provoke strong emotional responses — whether positive or negative — leading to the amplification of extreme viewpoints and the marginalization of more moderate or opposing perspectives. Consequently, this design dynamic contributes to a progressively divided public sphere. Similarly, Tufekci's *Twitter and Tear Gas* highlights that social media platforms, driven by ad-financed algorithms, often filter content in ways that favor advertiser-friendly material, which can drown out activist messages, undermine constructive dialogue, and contribute to growing polarization [8, p. xxix]. The structural attributes of digital platforms inherently challenge the facilitation of nuanced and in-depth discussions. Instead, they often replace these discussions with brief, reactionary exchanges that fail to delve into substantive issues.

1.1 Background

Argumentation theory and associated frameworks, such as argumentation graphs, offer a foundational basis for modeling argumentative interactions [9, Ch. 3]. Argumentation graphs serve as visual representations that delineate the relationships between different arguments, including support, attack, and counter-argumentative dynamics. However, these traditional models frequently fall short in capturing the full semantic depth and complexity characteristic of real-world discourse. The intricacies of human communication — such as implicit meanings, rhetorical strategies, and emotional undertones — are often not adequately represented in these models. For example, arguments may involve *unexpressed premises* that remain implicit but are central to the argumentative process [10, p. 4]. Traditional argumentation graphs typically map the structural and logical relationships between arguments but do not account for the more profound layers of meaning, context, or the underlying intentions shaping discourse.

To address these limitations, recent advancements in both graph and textual embeddings present new opportunities for enhancing argumentation models. Graph embeddings enable the representation of structural elements in a continuous vector space, making it easier to capture complex argumentative relationships. These embeddings are particularly useful in preserving the structural properties of arguments while reducing computational complexity [11, pp. 833–834]. For instance, they have been applied to visualize argument clusters and relationships, revealing deeper structural insights in argumentative graphs [12, pp. 3–4].

Conversely, textual embeddings leverage deep learning based natural language models like BERT, which excels at capturing the semantic richness of discourse, including implicit premises and emotional tones. Such embeddings enhance the ability to model long-range dependencies and latent themes within arguments [13, pp. 1–2]. The integration of these embeddings with graph structures allows for a comprehensive understanding of both the structural and semantic layers of discourse [14, pp. 1–2]. This combined approach supports more robust analysis, unveiling underlying drivers and themes in complex argumentative contexts [15, p. 4158].

Beyond its contributions to discourse analysis, the argumentation model developed in this thesis has the potential to serve as a foundational element within a computational engine designed to support diverse applications, including debate evaluation, argument summarization, argument quality assessment, retrieval-augmented generation (RAG) systems, and conversational systems. By enriching argumentation graphs with both structural and semantic insights, this model could power tools that analyze, retrieve, and generate discourse content with high relevance to specific user queries or contexts. Although this thesis concentrates on generating and the evaluation the embeddings, the broader implications suggest potential utility across public communication platforms and decision-support environments, enabling users to navigate complex debates and access relevant discourse effectively while making argument-based insights accessible and actionable.

1.2 Problem Statement

Traditional argumentation models exhibit significant limitations in their capacity to capture and analyze the semantic nuances and contextual dimensions inherent in complex debates. While these models effectively represent the logical structure of arguments — such as identifying which statements support or contradict one another — they fall short in representing the more subtle dimensions of discourse, such as shifts in tone, implicit argumentative strategies, and emotional fluctuations. These shortcomings are particularly pronounced in structured debates, where participants may employ rhetorical devices, sarcasm, or emotionally charged appeals to bolster their arguments, elements which traditional graph models struggle to effectively encapsulate.

To address these deficiencies, this thesis aims to enhance argumentation graphs through the integration of graph and textual embeddings. By combining these two methods, the resulting enriched model will provide a more comprehensive representation of arguments, encompassing both structural relationships and semantic content. This enriched approach could facilitate a more thorough discourse analysis, potentially revealing semantic patterns, rhetorical strategies, and argumentative nuances that traditional models often overlook. Such an approach aims to elucidate how arguments evolve over time. A deeper understanding of these factors is critical for meaningful discourse analysis, particularly within complex, real-world argumentative environments.

1.3 Research Objectives

The principal objectives of this research are as follows:

(1) Develop an argumentation model that encapsulates complex argumentative structures, addressing traditional limitations in capturing nuanced relationships and contradictions.

(2) Generate a synthetic debate dataset using a Large Language Model (LLM) to construct argumentation graphs, establishing a robust framework for modeling argumentative dynamics.

(3) Enrich the semantic embeddings with contextual embeddings based on the graph structure to enhance the graph's representational capacity, enabling detailed analysis of discourse elements.

(4) Compare the semantic, structural, and aggregated semantic embedding spaces by visualizing their clustering patterns using dimensionality reduction techniques like t-SNE and assessing clustering quality with metrics such as the Silhouette Score. This analysis aims to evaluate how effectively each embedding space separates argumentative components and captures semantic and structural information, providing insights into their relative strengths and weaknesses.

1.4 Research Questions

This thesis seeks to answer the following research question:

How can argumentative discourse be effectively modeled as an argumentation graph that incorporates both structural and semantic elements?

1.5 Thesis Structure

Chapter 2: Literature Review – This chapter will provide a comprehensive review of relevant literature on argumentation theory, argumentation graphs, graph embeddings, textual embeddings, and implication of such models for computational argumentation. It will also examine previous studies that have utilized computational methods to analyze discourse and debates. Furthermore, the chapter will identify existing gaps in the literature and demonstrate how the proposed research intends to bridge these gaps.

Chapter 3: Research Methodology – This chapter outlines the research design, including the construction of argumentation graphs using an LLM-generated synthetic dataset, representation learning of both structure and semantics and the aggregation of semantic embeddings. It details methods for comparing embedding spaces through visualization with dimensionality reduction (e.g., t-SNE), clustering analysis using k-means and Silhouette Score, and a combined evaluation of visual and quantitative metrics to assess argumentative roles and relationships.

Chapter 4: Results and Analysis – This chapter will present the findings of the study, focusing on the comparative analysis of embedding spaces and their effectiveness in modeling argumentative roles and relationships. It will include visualizations of embedding spaces reduced using techniques like t-SNE, accompanied by clustering analysis to evaluate the separation and grouping of argumentative elements (e.g., claims, premises). Quantitative metrics, such as Silhouette Scores, will be used to assess clustering quality of semantic and aggregated semantic embeddings spaces.

Chapter 5: Conclusion – This chapter will synthesize the key findings, provide conclusive answers to the research questions, and propose future research directions. The conclusion will reflect on the overall contributions of the thesis, emphasizing potential advancements in modeling argumentative discourse and considering the potential theoretical and practical impacts of these contributions.

2 Literature Review

To establish a rigorous theoretical and methodological foundation for this thesis, this literature review critically examines existing research and frameworks relevant to the analysis of argumentative discourse in digital environments. Given the inherently dynamic, often fragmented nature of online communication, it is imperative to evaluate how traditional argumentation theories and contemporary computational advancements can be leveraged and enhanced to effectively capture the complexities of discourse in the digital era. This chapter systematically reviews foundational concepts in argumentation theory, the application of argumentation graphs, and recent advancements in computational modeling, such as textual and graph embeddings.

The literature review commences by exploring the foundational tenets of argumentation theory, encompassing classical frameworks that have profoundly shaped our understanding of reasoning and argument structures. Subsequently, it critically evaluates the application of these traditional models, elucidating their inherent limitations in the context of digital discourse, and emphasizes the increasing relevance of network analysis and computational methodologies in refining argument mapping and discourse analysis. Furthermore, the review delves into the integration of textual and graph embeddings, emphasizing their potential to transcend some of the limitations of earlier models and to offer a more nuanced and comprehensive representation of both the structural and semantic dimensions of discourse.

Through the synthesis of these diverse research domains, this chapter seeks to identify existing gaps in the literature and to elucidate how the proposed research aims to bridge these gaps by developing an enriched argumentation model. Ultimately, this review positions the present study within the broader landscape of discourse analysis, underscoring its contributions to advancing our understanding of argumentative interactions in the intricate and evolving context of digital communication.

2.1 Argumentation Theory

Argumentation theory represents an interdisciplinary exploration into the principles and mechanisms that underpin the formulation, exchange, and evaluation of arguments across diverse domains. These include philosophical, computational, and practical applications. It plays an instrumental role in deconstructing the underpinnings of discourse, especially in formalized settings such as debates, legal reasoning, and policy-making processes.

The theoretical foundations of argumentation theory trace back to classical rhetoric and dialectics, notably Aristotle's foundational articulation of **ethos**, **pathos**, and **logos** as pillars of persuasion. Aristotle argues that persuasion depends on the character of the speaker (ethos), the emotional state of the audience (pathos), and the logical consistency of the argument (logos) [16, p. 38]. These classical concepts evolved into sophisticated frameworks for both formal and informal argument structures, providing the analytical basis for assessing the logical coherence and persuasive power of arguments. The emphasis on ethos, pathos, and logos continues to serve as a crucial tool for understanding how arguments can be crafted to appeal to both the rational and emotional faculties of the audience, making Aristotle's contributions central to the development of modern argumentation theory.

This section of the literature review is not intended to be exhaustive but rather focuses on the most pertinent models and theories that directly inform the descriptive modeling of argumentative discourse. While there are many significant frameworks within argumentation theory, this review selectively highlights those that provide the greatest insight into understanding and enhancing the structural and semantic analysis of discourse. By concentrating on these targeted models and theories, the review aligns closely with the thesis's objective of developing an enriched argumentation graph model. This focused approach ensures that the discussion remains both relevant and impactful, facilitating a more streamlined and contextually significant analysis that directly supports the thesis's research goals.

2.1.1 Toulmin's Model of Argumentation

Stephen Toulmin's *The Uses of Argument* first published in 1958 presents a practical framework for understanding arguments, particularly in contexts that extend beyond formal deductive reasoning. His model identifies six core components of any argument:

1. **Claim** (C) - the conclusion being argued,

2. **Data** (D) - evidence supporting the claim,
3. **Warrant** (W) - the reasoning that connects the data to the claim,
4. **Backing** (B) - further support for the warrant,
5. **Qualifier** (Q) - the strength of the claim,
6. **Rebuttal** (R) - acknowledgment of counterarguments [9, pp. 89–100].

These elements provide a structured way to dissect arguments, making the model especially useful for analyzing the informal, fragmented discourse typical in digital environments.

The general layout of this model can be illustrated as follows:

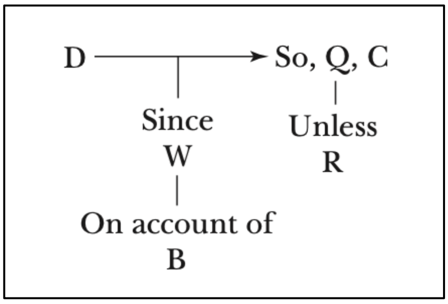

Figure 1 - Diagram illustrating the basic structure of Toulmin's model of argumentation [9, p. 97]

<u>Example of Toulmin's Argument Layout</u>

Toulmin's model of argumentation can be illustrated using the example of determining Harry's nationality. In this case, the **claim** (C) is that *Harry is a British subject.*" The **data** (D) supporting this claim is that "*Harry was born in Bermuda.*" The **warrant** links these two elements by establishing the general principle that "*A man born in Bermuda may be taken to be a British subject."* Since issues of nationality are subject to qualifications, the claim is modified with **presumably** (Q). The **rebuttal** (R) addresses possible challenges, such as if "*both his parents were aliens, or he has since become a naturalized American.*" Finally, the **backing** consists of statutes and legal provisions governing British nationality [9, p. 97].

This argument structure is visually represented in the following diagram:

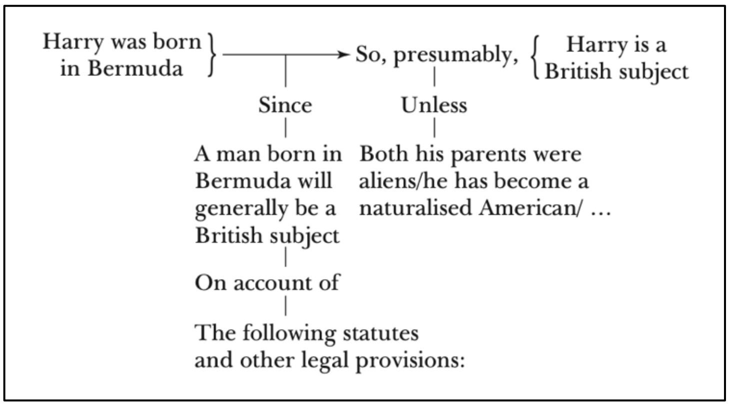

Figure 2 - Visual representation of Harry's nationality argument following Toulmin's model [9, p. 97]

A key contribution of Toulmin's model is its differentiation between **field-invariant** and **field-dependent** components. **Field-invariant** elements — such as the presence of claims, data, and warrants — are consistent across different contexts. However, **field-dependent** elements, such as the type of evidence or the standards for validity, change depending on the domain (e.g., legal reasoning vs. scientific discourse) [9, pp. 14–16]. This adaptability makes the model highly relevant to digital discourse analysis, where arguments often pull from multiple fields and must accommodate different norms of reasoning.

Toulmin also critiques the rigid structures of formal logic, arguing that real-world reasoning rarely conforms to such idealized forms [9, Ch. 4]. His model provides a more flexible approach, accommodating both the structure and the variability of everyday argumentation. The warrant, in particular, serves as an essential link between data and claim, and in digital discourse, this connection is often implicit or unclear. Toulmin's model helps clarify these connections, making arguments more transparent and robust. Additionally, backing reinforces the warrant, ensuring that arguments are well-supported — an increasingly important feature in online debates where misinformation can thrive [9, pp. 95–100].

Soundness and Argumentative Validity

When analyzing the structure and effectiveness of arguments, it is essential to distinguish between *soundness* and *argumentative validity*, concepts pivotal to Toulmin's critique of traditional logic. Toulmin argues that while **formal validity** — ensuring that conclusions logically follow from premises — is crucial, it is insufficient for real-world applications that require more than structural consistency

[9, pp. 110–114]. He emphasizes that practical arguments often demand an evaluation of both the structural form and the contextual backing of the argument's premises.

In formal logic, validity pertains solely to the logical coherence of the argument. Toulmin notes that in structured fields, such as mathematics or formal deduction, this form of validity can be maintained simply by aligning data, warrants, and conclusions. However, Toulmin critiques this as overly restrictive for practical fields, where "field-dependent" standards for validity are often more relevant and adaptable to specific domains, like legal or ethical reasoning [9, pp. 118–123, 136–138].

Soundness, in Toulmin's model, involves both validity and the truthfulness of the premises — meaning an argument is only sound if it is valid and its premises are true. Toulmin's framework, particularly his emphasis on warrants and backing, shifts the focus beyond logical form to include the credibility and applicability of these foundational elements, making soundness a contextually enriched evaluation rather than a purely formal assessment [9, pp. 110–114, 136–138]. For instance, in digital discourse, where claims often span multiple fields, Toulmin's model enables a nuanced analysis of whether arguments are not only logically coherent but also substantively reliable.

Implications for Digital Discourse Analysis

Toulmin's approach to soundness and validity is particularly valuable in analyzing online discourse. His model allows the identification of gaps in reasoning, especially where field-dependent warrants and backing are implicit or ambiguous. This is a critical feature for computational models of discourse, which aim to capture both logical and contextual nuances. In summary, Toulmin's model expands the assessment of arguments from rigid structural validity to a more adaptable and context-sensitive framework, accommodating the layered complexity of digital communication.

2.1.2 Walton's Argument Schemes

Douglas Walton's contributions to argumentation theory revolve around a series of **argumentation schemes** that represent typical forms of **presumptive reasoning**. These schemes are particularly suited for understanding arguments that operate under uncertainty and involve everyday reasoning rather than purely deductive logic. Walton’s approach is useful in analyzing how arguments function within real-world contexts, especially those involving incomplete information and dynamic discussion settings [17, Ch. 3].

Walton categorizes a range of schemes such as **Argument from Expert Opinion**, **Argument from Analogy**, and **Argument from Consequences**. Each scheme provides a generalized structure to understand how conclusions are drawn from premises. For instance, in the **Argument from Expert Opinion**, the reasoning takes the form: “*Expert E asserts that statement A is true, and E is an expert in the relevant field, therefore A is likely true.*” The validity of such an argument is evaluated through **critical questions** — such as assessing the credibility of the expert, the expert’s reliability, and the presence of any biases [17, pp. 14–16]. These critical questions serve as a mechanism to safeguard against misuse and fallacies, which is crucial for maintaining the quality of argumentative discourse.

An example of a common scheme is the **Argument from Sign**, where an observed indicator is presumed to suggest a particular conclusion. For instance, “*John’s hat is not on the peg; therefore, John has left the house*” relies on the assumption that John habitually wears his hat when leaving. Walton emphasizes that inferences are generally **defeasible** — meaning they can be rebutted if further evidence arises, such as finding John still inside the house [17, pp. 13–14]. The concept of defeasibility is key to Walton's argumentation schemes, emphasizing their tentative nature.

Walton’s work is also relevant in the context of digital discourse where arguments often rely on signals of popularity, such as likes or shares on social media, rather than rigorous logical analysis. The **Argument from Popularity**, which is a type of Argument from Appeal to the People (ad populum), suggests that if many people believe or endorse a proposition, it is likely true. Walton points out that such arguments are inherently presumptive and must be evaluated carefully. Critical questions, such as whether the popularity is based on a genuine rational consensus or simply emotional appeal, are essential for assessing their validity in a digital context [17, pp. 83–85].

In conclusion, Walton's argumentation schemes provide a structured framework for understanding and assessing informal arguments. By incorporating critical questions and emphasizing the provisional nature of presumptive reasoning, Walton's tools enable the evaluating of both the **logical structure** and **pragmatic validity** of arguments in various real-world contexts.

2.1.3 Pragma-Dialectics in Argumentation Theory

The pragma-dialectical approach, developed by Eemeren and Grootendorst, presents a systematic model for analyzing argumentative discourse that integrates pragmatic and dialectical elements, emphasizing the role of structured, rule-governed discourse in resolving differences of opinion. Unlike classical rhetorical approaches that emphasize persuasion, pragma-dialectics focuses on achieving resolution through collaborative critical discussion. This is articulated through the concept of a "critical discussion," a structured dialogue where participants test opposing views by adhering to a set of norms aimed at fostering rational exchange [10, pp. 42–44, 51 ff.].

The model establishes four key stages in a critical discussion: *confrontation*, *opening*, *argumentation*, and *concluding*. In the confrontation stage, participants identify and clarify their disagreements. In the opening stage, procedural rules and mutual commitments are established, laying the groundwork for a cooperative exchange. During the argumentation stage, participants present and challenge arguments, advancing their respective standpoints. Finally, the concluding stage determines the discussion's outcome, ideally leading to consensus or a reasoned acknowledgment of remaining differences. This structure helps manage argumentative complexity, maintaining coherence and fostering resolution [10, pp. 57–62].

Meta-Theoretical Principles: Functionalization, Externalization, Socialization, and Dialectification

Pragma-dialectics is structured around four core meta-theoretical principles that guide its framework: *functionalization*, *externalization*, *socialization*, and *dialectification*. These principles help distinguish pragma-dialectics from other models by focusing on both normative and practical dimensions [10, p. 52].

1. **Functionalization** interprets each linguistic move in argumentation as a purposeful act directed at resolving the dispute, emphasizing the role of each statement in advancing the discussion [10, pp. 52–54].
2. **Externalization** focuses on publicly observable commitments, analyzing the actual claims participants make rather than speculating about intentions, thus promoting transparency [10, pp. 52–55].
3. **Socialization** highlights argumentation's interactive nature, framing arguments as part of a collaborative process where participants respond and adapt to each other's contributions [10, pp. 53–56].
4. **Dialectification** formalizes argumentation as a rule-governed exchange, ensuring contributions are structured for critical evaluation rather than persuasion alone, reinforcing procedural standards of reasonableness [10, pp. 52–53, 57].

Types of Speech Acts

In pragma-dialectics, speech acts are essential tools that shape the direction and effectiveness of argumentative discourse. Each type of speech act contributes uniquely to the process of critical discussion, aligning statements and responses with the objective of resolving differences rationally. By categorizing speech acts, pragma-dialectics provides a framework that clarifies how each communicative move advances the discourse, either by establishing claims, requesting elaboration, committing to a course, expressing attitudes, or signaling shifts in the debate's status.

Assertives are foundational, conveying claims or information that serve as premises for argumentation. For instance, a statement like "The data supports this conclusion" sets the groundwork for further discussion by presenting a factual basis for a viewpoint [10, pp. 63–64, 67–68]. Meanwhile, ***directives*** aim to prompt a response or action from the other party. Examples include requests for evidence or challenges to clarify a position, as in, "Can you provide proof?" which pushes the discussion toward greater depth and justification [10, p. 64, pp. 67–68].

Commissives are acts through which a speaker commits to a stance or course of action, thereby setting mutual expectations in the dialogue. Statements like "I will defend this point" help establish reliability and focus, reinforcing participants' commitments to constructive engagement [10, pp. 64–65, 67–68]. ***Expressives***, on the other hand, convey the speaker's attitudes, such as agreement or respect, often influencing the cooperative tone of the interaction. For example, saying “I appreciate your point” fosters goodwill and a collaborative atmosphere, which can be crucial for productive discourse [10, p. 65, pp. 67–68].

Finally, ***declarations*** are statements that enact a change in the state of the discussion, such as formally concluding or redirecting it. By stating, “This debate is concluded,” a participant may signal the end of a discussion, marking an essential transition in the discourse [10, pp. 66–68].

Each type of speech act serves a specific purpose, keeping contributions purposeful and aligned with the goal of resolving the discussion rationally. In pragma-dialectics, this categorization enables a nuanced analysis of argumentative interactions, making it possible to see how each move in the dialogue contributes to or detracts from the critical discussion’s progress.

Problem Validity and Intersubjective Validity

In addition to the four principles above, pragma-dialectics incorporates the meta-theoretical criteria of *problem validity* and *intersubjective validity*, which are essential for evaluating the quality of argumentative discourse.

Problem Validity assesses whether the contributions in a discourse effectively advance the resolution of the issue at hand. Arguments are deemed problem-valid when they adhere to the procedural rules of the pragma-dialectical model, facilitating a genuine move toward resolving the central disagreement. This ensures that each step in the discourse aids in clarifying or addressing the main point of contention, preventing fallacies that might derail or distort the argument [10, pp. 17, 22, 132]. Problem validity is integral to maintaining the quality of discourse, as it reflects adherence to a methodical code of conduct that precludes discussion violations, ensuring that the argument remains focused on resolution [10, pp. 57, 187].

Intersubjective Validity, on the other hand, emphasizes the shared acceptance among discussion participants. For an argument to be intersubjectively valid, it must be deemed reasonable by all parties involved within the discussion framework [10, pp. 17, 22, 129]. This criterion reinforces mutual recognition of argumentative moves as rational, fostering an environment of trust and cooperation [10, p. 57]. By ensuring that participants assess arguments based on mutually agreed-upon standards, intersubjective validity allows the discourse to advance constructively, even when disagreements persist [10, p. 187].

Relevance to modelling Argumentative Discourse

While pragma-dialectics does not directly inform the technical design of the enriched argumentation graph central to this thesis, it provides essential guidance for conducting a structured debate as a case study. The four meta-theoretical principles and the criteria of problem and intersubjective validity offer a framework for ensuring that discourse within the study adheres to high standards of rationality and fairness. By structuring the debate in accordance with these principles, a case study establishes a disciplined environment, facilitating a robust evaluation of how effectively the enriched model captures discourse patterns based on structured, rational exchanges.

2.1.4 Key Developments in Computational Applications of Argumentation Theory

The integration of computational methods into argumentation theory has significantly advanced the analysis and modeling of complex argumentative structures. This evolution has led to the development of sophisticated tools and frameworks that enhance our understanding of argumentation in various contexts.

One notable advancement is the emergence of **argument mining**, a subfield of natural language processing (NLP) focused on automatically identifying and extracting argumentative components from textual data. This process involves detecting premises, conclusions (or claims), and other argumentative elements and the relationships between them, facilitating large-scale analysis of argumentative discourse especially within the realm of digital communication. Machine learning techniques have been instrumental in this domain, enabling the development of models that can process and analyze vast amounts of text to identify argumentative structures [18, pp. 99–100]. Recent developments employ pre-trained transformer-based deep-learning models, such as BERT, GPT and T5, which can be fine-tuned for specific argumentation mining tasks [19], [20]. These models are also versatile enough to be trained for additional tasks like sentiment analysis, summarization, classification, and similarity analysis [21, Sec. VIII]. These developments opened the door to new applications of NLP and deep learning in a diverse range of computational linguistic tasks, which will be further explored in later chapters.

Another significant development is the application of **formal argumentation frameworks** in various domains. These frameworks provide structured representations of arguments and their interrelations, allowing for the systematic evaluation of argument validity and strength. Dung's foundational work [22] established the groundwork for abstract argumentation frameworks used in AI.

Furthermore, the development of computational models of argumentation has led to practical applications across various domains. Rahwan and Simari's provide comprehensive coverage of argumentation in AI, highlighting various computational techniques [23].

In summary, the computational turn in linguistics has introduced powerful tools and methodologies that enhance the analysis, modeling, and application of argumentative discourse. These developments have paved the way for more nuanced and scalable analyses of arguments, particularly in digital environments where the volume and complexity of discourse continue to grow.

Building upon these advancements, the subsequent section will delve into **Abstract Argumentation Frameworks**, a pivotal framework that represents argumentative structures as networks. This approach facilitates the visualization and analysis of complex argumentation patterns, offering deeper insights into the interplay between different argumentative components.

2.2 Abstract Argumentation Frameworks

Argumentation graphs serve as visual and computational representations of arguments, capturing relationships such as support, attack, and counterarguments among different claims. **Dung's abstract argumentation framework** [22] has revolutionized the formalization of argumentation by providing a versatile and robust methodology for analyzing interactions between arguments. This approach finds utility across domains such as artificial intelligence, social reasoning, and decision-making.

2.2.1 Dung’s Abstract Argumentation Framework

In this framework, arguments are treated as abstract entities, disconnected from their internal structure, to focus solely on their interactions — specifically attacks and defenses. This abstraction enables the analysis of arguments in diverse domains, from artificial intelligence to social reasoning.

Dung’s framework formalizes argumentation by representing arguments and their interactions as a pair:

$$\boldsymbol{AF} = (\boldsymbol{AR}, \boldsymbol{attacks}),$$

where AR is a set of arguments, and $attacks \subseteq AR \times AR$ specifies which arguments attack others.

A set $S \subseteq AR$ is **conflict-free** if no two arguments in S attack each other.

(1) An argument A is **acceptable** with respect to S if every argument attacking A is counter-attacked by some argument in S .

(2) A set S is **admissible** if it is conflict-free and all its arguments are acceptable with respect to S itself [22, p. 326].

Dung’s framework also introduces **argumentation semantics** to evaluate the status of arguments in the graph:

- **Grounded Extensions**: These represent the minimal fixed point of the framework’s characteristic function F, offering a conservative baseline of acceptable arguments [22, p. 329].
- **Preferred Extensions**: Maximal admissible sets, representing the most inclusive defensible arguments [22, p. 327].
- **Stable Extensions:** Sets that are conflict-free and attack every argument not included in the set, representing a stricter notion of argument acceptability [22, p. 328].

Grounded extensions provide a cautious approach, ensuring universal acceptability, while preferred extensions emphasize inclusivity. Stable extensions, the strictest of the three, guarantee that all external arguments are attacked, reflecting a robust standard for argument acceptability.

Example: Consider a debate over climate action. Let A_1 represent the claim "*Reducing carbon emissions will mitigate climate change*," while A_2 states "*Mitigation efforts are economically unsustainable*," attacking A_1. Meanwhile, A_3, asserting "*Long-term economic stability depends on climate action*," counters A_2. If A_3 successfully undermines A_2, then A_1 becomes defensible, highlighting A_3's pivotal role in validating A_1. This interplay is visually represented as a graph where nodes A_1, A_2, and A_3 denote arguments, and directed edges represent attacks and counterattacks.

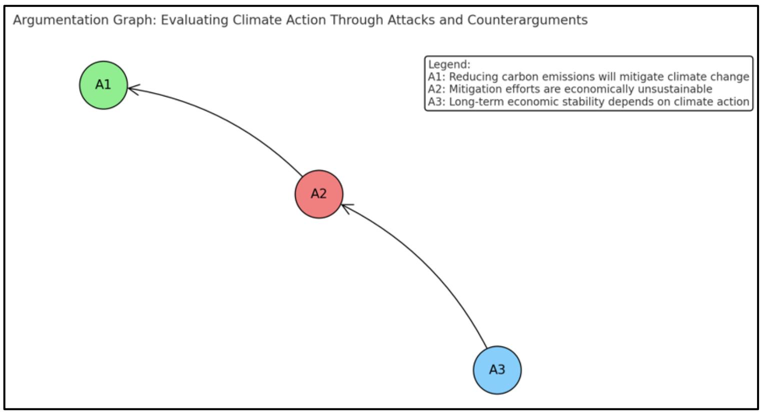

Figure 3 - Argumentation graph illustrating the interplay of attacks and counterarguments in a climate action debate, with A_1, A_2, and A_3 representing key claims

The power of Dung's framework lies in its generality and adaptability. For example, the *Nixon diamond problem*, a classic test case, involves conflicting arguments: "*Nixon is a pacifist because he is a Quaker*" (A_1) versus "*Nixon is anti-pacifist since he is a republican*" (A_2) [22, p. 327]. Here, both arguments attack each other, resulting in two preferred extensions ($\{\mathrm{A}_1\}$ and $\{A_2\}$) but no stable consensus under grounded semantics. This illustrates how different semantics lead to distinct outcomes, enabling nuanced analysis of argument structures.

Dung's framework has proven invaluable in analyzing conflicts, providing tools for diverse applications, from automated reasoning to ethical deliberation. The abstract nature of Dung's framework ensures broad applicability, offering a foundation for automated reasoning systems, ethical deliberation, legal analysis, and multi-agent systems. By resolving conflicts and evaluating the acceptability of arguments, it has become a cornerstone in computational argumentation and decision support.

Furthermore, Dung's theory bridges the gap to broader computational frameworks, such as logic programming and nonmonotonic reasoning. Stable semantics align with stable models in logic programming, facilitating diverse practical applications [22, Sec. 4]. This connection enhances its utility in real-world contexts, ensuring its relevance in fields as diverse as law, negotiation theory, and multi-agent systems, where resolving conflicting perspectives is critical.

2.2.2 Extensions of Abstract Argumentation Frameworks

Dung's foundational abstract argumentation framework provides a robust structure for analyzing argument interactions, but its focus on binary attack relationships limits its capacity to model more nuanced scenarios. Over the years, various extensions of the original framework have been developed to address specific challenges and expand its applicability. These extensions introduce additional constructs and relationships to represent more complex argumentative interactions.

Bipolar Argumentation Framework

Cayrol and Lagasquie-Schiex extended Dung's foundational framework by introducing the **Bipolar Argumentation Framework (BAF)**, which incorporates both **defeat (attack)** and **support** relations among arguments [24, pp. 378–379]. While Dung's framework focused solely on conflict, the BAF acknowledges the dual nature of argumentation by explicitly representing support as an independent relation, enabling more nuanced modeling of real-world interactions where arguments are not only contested but also reinforced [24, pp. 379–382].

In the BAF, arguments and their interactions are defined as a triplet $(\mathrm{A}, \mathrm{R_{def}}, \mathrm{R_{sup}})$, where $\mathrm{R_{sup}}$ represents positive interactions distinct from the defeat relation $\mathrm{R_{def}}$ [24, p. 382]. This framework introduces the concepts of **supported defeat** and **indirect defeat**, allowing for richer representations of layered interactions among arguments, such as chains of supportive relations culminating in a defeat [24, p. 383]. For instance, an argument A_1 might support A_2, which then defeats A_3, creating a nuanced interdependence. The capabilities of BAFs can be illustrated by revisiting Dung's climate change example (Figure 3). The BAF enriches this scenario by introducing $\mathrm{A_4}$, a supporting argument for A_3, such as "Economic studies show climate action boosts innovation." Here, $\mathrm{A_4}$ strengthens A_3, which in turn counters A_2, effectively bolstering A_1. This multi-layered interaction highlights BAF's ability to integrate both conflict and reinforcement, providing a more comprehensive view of argument dynamics.

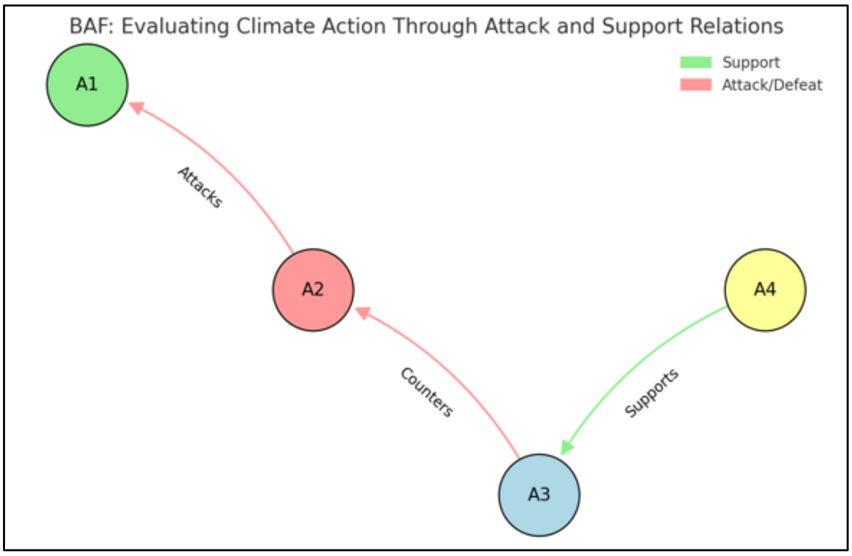

Figure 4 - Bipolar Argumentation Framework illustrating climate action arguments: A_2 attacks A_1, A_3 counters A_2, and A4 supports A_3, reinforcing A_1 indirectly.

The Bipolar Argumentation Framework further extends Dung's original framework by introducing the notions of **set-defeat** and **set-support**, which generalize the defeat and support relations to apply to sets of arguments. Formally, a set $S \subseteq A$ is said to **set-defeat** an argument $A \in A$ if S contains an argument that establishes either a supported or indirect defeat for A [24, p. 383]. Similarly, S **set-supports** A if there exists a sequence of support relations emanating from arguments in S that ultimately reinforce A. Building on these relations, the concept of **defense by a set of arguments** is introduced [24, p. 383]. A set S defends an argument A collectively if, for every argument B that set-defeats A, there exists an argument $C \in S$ that in turn set-defeats B. These generalizations allow for a more robust evaluation of argument interactions, particularly in cases where layered or group-based dynamics play a significant role, thus enhancing the framework's applicability to real-world scenarios.

To evaluate argument sets under the Bipolar Argumentation Framework, the concepts of **conflict-free sets** and **safe sets** provide foundational tools for managing interactions. A set is defined as **conflict-free** if no argument within it set-defeats another, ensuring *internal* coherence and aligning with Dung's original framework [24, pp. 384–385]. However, the introduction of the support relation necessitates an additional layer of rigor: a **safe set** satisfies both *internal* and *external* coherence by disallowing situations where the same argument is simultaneously set-supported and set-defeated by the set [24, p. 385]. These concepts form the groundwork for revised acceptability semantics that account for both defeat and support.

Building on this foundation, Cayrol and Lagasquie-Schiex refined Dung's stable and preferred semantics to address the dual nature of argument interactions. **Stable extensions** remain conflict-free but must defeat all arguments outside the set, considering supportive effects [24, pp. 385–386]. **Preferred extensions** are similarly adapted using criteria like **s-admissibility** (safe admissibility), which ensures *external* coherence, and **c-admissibility** (closed admissibility), which requires sets to be closed under the support relation, encompassing all arguments that reinforce members of the set [24, pp. 386–387]. These refinements address the complexities of real-world argumentation, where positive and negative interactions coexist.

By capturing the dual nature of argumentative interactions, BAFs enable more precise analysis of competing and supportive claims, enabling potential applications in domains like legal reasoning, negotiation, and multi-agent systems.

Value-Based Argumentation Framework

Another significant advancement in modeling argumentation is **Bench-Capon's Value-Based Argumentation Framework (VAF)** [25]. VAFs extend Dung's Argumentation Framework (AF) by associating arguments with underlying **values**, which influence the resolution of conflicts between arguments based on the relative importance of these values. Unlike Dung's AF, where the success of an attack is determined purely by the structural relationships among arguments, VAFs introduce a preference relation over values, allowing the evaluation of arguments to reflect the ***subjective*** **priorities** of individuals or groups [25, pp. 1–2].

To illustrate, consider a policy debate over whether to implement universal free college education. Suppose there are three main arguments:

(1) A_1: "Free college education promotes social equality" (value: equality),

(2) A_2: "Universal free education places an unsustainable financial burden on taxpayers" (value: economic efficiency),

(3) A_3: "Access to higher education improves long-term economic productivity" (value: innovation).

In this example, A_2 attacks A_1 because the financial burden undermines equality goals, and A_3 supports A_1 by emphasizing long-term economic benefits.

In a VAF, the framework is defined as a five-tuple $(AR, attacks, V, val, valpref)$, where:

- AR is the set of arguments (A_1, A_2, A_3),
- $attacks$ represents the binary relation denoting attacks (A_2 attacks A_1),
- V is a set of values ($V = \{equality,\ economic\ efficiency,\ innovation\}$),
- val is a mapping that associates arguments with values ($val(A_1) = equality$, etc.),
- $valpref$ is a transitive, irreflexive, and asymmetric preference relation on V (e.g., $equality > economic\ efficiency$) [25, pp. 2–3].

This formulation allows arguments to be evaluated not only based on their ability to counteract others but also on the precedence of the values they promote.

VAFs are particularly useful in contexts like legal reasoning, ethical debates, and public policy discussions, where value judgments play a central role. Continuing with the hyptothetical example, if the audience prioritizes **equality** over **economic efficiency**, A_1 is likely to be preferred over A_2, despite A_2's valid critique. However, if **economic efficiency** is prioritized, A_2 might outweigh A_1. Similarly, A_3's support for A_1 becomes more significant if **innovation** is also highly valued.

A crucial contribution of VAFs is the concept of **objectively acceptable arguments** – arguments that are included in the preferred extension of the framework regardless of the value ordering. For example, if A_3 demonstrates that innovation unequivocally supports equality without significantly impacting economic efficiency, A_3 may be objectively acceptable, ensuring that its inclusion in the debate transcends subjective value hierarchies [25, pp. 3–4]. This property ensures that some arguments can be deemed rationally compelling across varying subjective preferences, providing a stable foundation for resolving disputes even in highly polarized scenarios.

Additionally, Bench-Capon introduces **strategic heuristics** for extending VAFs, which can alter the status of arguments by modifying their context within the framework. These heuristics allow participants to strategically influence disputes by introducing new arguments, adjusting existing attacks, or altering the value preferences that guide the resolution process. For instance, consider introducing a new argument A_4: "*Subsidizing education fosters global competitiveness*" (value: innovation). This hyptothetical argument counters A_2's financial concerns by linking efficiency to long-term benefits. By carefully analyzing the position of arguments within chains and the parity of attacking chains, it is possible to shift the defensibility of arguments, making them objectively acceptable, subjectively acceptable, or indefensible [25, pp. 9–10].

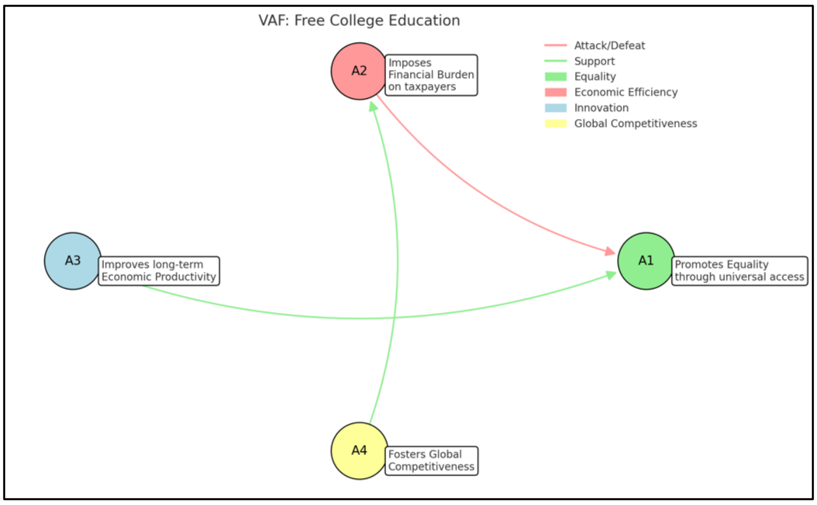

Figure 5 - VAF for Free College Education: The graph illustrates the relationships among arguments (A_1–A_4) in the debate on universal free college education, highlighting their interactions and underlying values.

In summary, VAFs enhance Dung's AF by integrating value-based reasoning, enabling a nuanced analysis of argumentation that accounts for subjective preferences. The hyptothetical example of universal free college education demonstrates how VAFs allow stakeholders to navigate complex debates, weighing arguments not just on logical grounds but also on the competing values they represent. This innovation bridges the gap between abstract formal frameworks and practical decision-making scenarios, proving indispensable for understanding complex argumentative dynamics in domains where values play a pivotal role.

2.2.3 Limitations of traditional Argumentation Frameworks

Abstract Semantics

Traditional argumentation graphs excel at illustrating structural relationships between arguments, such as attacks and defenses. However, their abstraction from the logical content of arguments imposes significant limitations, particularly when applied to real-world scenarios requiring semantic depth and logical consistency. Caminada and Wu highlight these challenges, emphasizing the disconnect introduced by abstract evaluation methodologies when handling nuanced argumentative contexts [26].

A key issue arises from the **three-step process** outlined by Caminada and Wu [26, pp. 2–3]. The first step involves constructing arguments from a knowledge base and identifying attack relations, followed by an abstract evaluation of argument acceptability (Step 2). During this phase, the internal logical content of arguments is disregarded in favor of a purely topological analysis of the argument graph. This abstraction introduces critical risks to the coherence of conclusions derived in Step 3, where logical entailments are determined based on the arguments deemed acceptable. As noted by Caminada and Wu, this approach can result in accepted arguments whose conclusions fail to satisfy logical consistency or closure requirements [26, pp. 4–7].

For instance, abstract semantics often operate "blindly", without regard for the logical foundations or internal connections among arguments. Caminada and Wu illustrate this problem by questioning how a set of accepted arguments can ensure consistency of their conclusions or guarantee closure under logical entailment [26, p. 5]. They argue that abstract frameworks may inadvertently admit conclusions that violate these principles, particularly when purely topological methods, such as those defining conflict-free or admissible sets, are employed [26, p. 6].

Caminada and Wu's critique extends to the uncritical reliance on abstract semantics in contexts where logical and semantic integrity are paramount. They note that while abstract argumentation provides a useful conceptual simplification, it risks becoming disconnected from the real-world argumentative contexts it aims to model. This detachment undermines the applicability of traditional frameworks in domains requiring rigorous logical validation [26, p. 6].

Challenges in expressiveness and applicability

In the preceding sections, we discussed the structure and capabilities of abstract argumentation frameworks (AFs), along with several key extensions that introduce support relations and value-based reasoning. However, traditional AFs and their immediate extensions often face limitations in expressiveness and applicability when dealing with complex, real-world scenarios.

One critical shortcoming lies in the inability of classical AFs to naturally represent varying degrees of uncertainty, priority, or evidential strength among arguments. The basic binary attack relation typically oversimplifies the nuanced interplay of competing claims, which can be more subtly captured by **weighted argumentation frameworks**. Weighted AFs assign numerical strengths or preferences to arguments, providing a more fine-grained analysis of argument acceptability. Such quantitative extensions allow for distinguishing between stronger and weaker attacks, as well as for aggregating multiple sources of support or opposition [27].

Another avenue of enhancement involves **probabilistic argumentation frameworks**, which integrate uncertainty into the modeling process. By associating probabilities with arguments or attacks, these frameworks enable reasoning under incomplete or ambiguous information [28]. This probabilistic dimension is particularly relevant in domains like legal reasoning, medical diagnosis, and data-driven decision-making, where the truth value or reliability of premises is often not definitive.

Temporal argumentation frameworks introduce a dynamic perspective, accounting for the fact that arguments and their relationships may evolve over time. Such temporal extensions facilitate the modeling of scenarios where the availability or relevance of evidence fluctuates, and the strength of an argument may increase or decrease as new information emerges [29], [29]. This temporal dimension is crucial in domains like policymaking, where evolving social, economic, or environmental conditions continuously reshape the argumentative landscape.

Beyond weighting, uncertainty, and temporality, **evidential argumentation frameworks** incorporate evidence management and source reliability into their reasoning processes. By systematically tracking the origin, credibility, and corroboration of evidence, these frameworks offer a richer understanding of why certain arguments prevail and how trust or skepticism in sources may shift outcomes [30].

In summary, while Dung's original framework and its immediate extensions (e.g., bipolar and value-based AFs) offer robust foundations for modeling argumentative interactions, these classical models fall short in terms of expressiveness and applicability when confronted with intricate, real-world debates. The subsequent waves of research—weighted, probabilistic, temporal and evidential frameworks—address these shortcomings by introducing mechanisms that better capture the complexity, uncertainty, and dynamism of practical argumentation. These developments ensure that computational models remain both theoretically sound and operationally useful, bridging the gap between abstract theoretical constructs and the nuanced demands of real-world decision-making processes.

These limitations underscore the importance of advancing argumentation graphs beyond their traditional design. Future developments should focus on integrating deeper semantic insights, potentially through hybrid models that combine structural graph-based representations with semantic embeddings derived from natural language processing techniques. Such enhancements are essential for paving the way for more robust and context-aware analytical tools.

2.3 Representation Learning for Computational Argumentation

In recent years, the study of computational argumentation has advanced significantly, driven by the increasing availability of large datasets, powerful machine learning models, and novel representation learning techniques. At the heart of these developments lies the need to effectively represent arguments and their relationships in a manner conducive to automated processing, analysis, and inference. Traditional frameworks, such as those introduced by Dung and its extensions, provide robust formal structures for capturing the interplay of arguments, but they often lack the flexibility and nuance needed to handle complex, large-scale, real-world scenarios.

Representation learning addresses these limitations by enabling the encoding of arguments into machine-readable formats that preserve critical semantic, structural, and contextual information. Techniques such as textual embeddings, graph embeddings, and their integration offer innovative ways to capture the richness of arguments. These representations underpin key computational tasks, including argument mining, stance detection, relation classification, and the evaluation of argument strength or persuasiveness.

This section explores the role of representation learning in computational argumentation, focusing on its three primary facets: **textual representations**, which model the semantic and contextual nuances of argument components; **graph-based representations**, which encapsulate the structural relationships among arguments in graph form; and **integrated representations**, which combine the strengths of both textual and graph-based embeddings to provide a holistic view of argumentative frameworks. By analyzing these approaches, this section aims to highlight the transformative potential of representation learning for advancing argumentation research and its applications across diverse domains.

2.3.1 Textual Representations

In computational argumentation, analyzing and processing arguments effectively relies on transforming raw text into structured representations that encode semantic meaning and contextual nuance. This transformation, known as **textual representation learning**, is foundational for tasks like argument mining, stance detection, and sentiment analysis. These tasks often involve uncovering implicit logical structures and rhetorical strategies embedded in argumentative discourse.

The Role of BERT in Textual Representation Learning

A significant breakthrough in textual representation learning is **BERT (Bidirectional Encoder Representations from Transformers)**, developed by Devlin et al. BERT fundamentally reshaped natural language processing (NLP) by introducing a model that captures deep, bidirectional representations of text. Unlike traditional models that process text in a unidirectional manner, BERT incorporates context from both preceding and succeeding words, enabling it to generate embeddings that reflect the full context of a word or phrase within its textual environment [31].

This capability makes BERT particularly effective for computational argumentation. Arguments are often composed of interdependent claims and counterclaims that require an understanding of how components interact contextually. For instance, in a debate on environmental policy, understanding the nuanced relationship between the premises "reducing emissions protects biodiversity" and "biodiversity loss impacts human well-being" requires bidirectional contextual embeddings. BERT excels at capturing such relationships, enabling the precise identification and classification of argumentative structures.

One of BERT's strengths lies in its **pretraining objectives**: *masked language modeling* (MLM) and *next sentence prediction* (NSP). These objectives prepare BERT to understand both local and global text coherence. MLM involves predicting missing words in a sentence, equipping the model to handle incomplete or noisy arguments often found in real-world discourse. NSP allows BERT to assess the coherence of two sentences, a skill that aligns well with tasks like detecting premise-conclusion relationships in argumentative texts [31].

Applications of BERT in Argumentation

BERT has become an indispensable tool in computational argumentation, driving advancements in tasks that require a nuanced understanding of text. Its bidirectional understanding of text allows for nuanced interpretation of arguments, enabling applications across a wide spectrum of use cases. In **argument mining**, BERT-based systems are adept at identifying and classifying argumentative components such as claims and premises from diverse text sources. These capabilities are particularly useful in domains where arguments are complex and nuanced. For instance, a study demonstrated that BERT, when fine-tuned, outperforms traditional machine learning models like GloVe and ELMo in extracting argumentative components such as premises and conclusions from legal texts, highlighting its effectiveness in domains where argument structures are complex and nuanced [32].

BERT’s contextual embeddings are highly effective for **stance detection**, particularly in tasks requiring an understanding of relationships between textual components such as headlines and article bodies. Karande et al. demonstrated this in a study on fake news detection, where stance detection was integrated with BERT-based embeddings to measure the similarity between article elements. This approach achieved a state-of-the-art accuracy of 95.32% by leveraging cosine similarity as an additional feature. The results underscore BERT’s capacity to capture nuanced linguistic relationships, making it an invaluable tool for tasks like credibility analysis and the detection of misinformation [33].

BERT's contextual embeddings have significantly advanced **sentiment analysis** by effectively capturing emotional nuances within text. Alaparthi and Mishra's comparative study evaluated four sentiment analysis techniques – SentiWordNet, logistic regression, LSTM, and BERT – using a dataset of 50,000 IMDB movie reviews. Their findings demonstrated BERT's superior performance across accuracy, precision, recall, and F1 scores [34]. Similarly, Wu et al. explored BERT's architecture and optimization strategies in sentiment analysis, confirming its robust performance, especially after fine-tuning [35]. These studies underscore BERT's effectiveness in capturing emotional undertones, thereby enhancing the analysis of persuasive strategies in various texts.

Implications for Argumentation Research

The integration of BERT into computational argumentation research has significantly advanced the field, enabling models to process text with a level of depth and precision that was previously unattainable. By capturing bidirectional context and nuanced meaning, BERT-based systems provide insights into the structure, tone, and persuasiveness of arguments. These capabilities support applications across diverse domains, from legal analysis to educational tools and policy-making frameworks. As BERT continues to inform the development of optimized models, its influence on computational argumentation remains profound.

2.3.2 Graph Representations

Graph-based representations play a pivotal role in computational modeling, particularly for encoding complex relational data into structured forms suitable for machine learning and inference. By organizing entities as nodes and their interrelations as edges, these representations facilitate the analysis of both local and global dependencies within data. Such models are instrumental in diverse domains, including social networks, biological systems, and knowledge graphs.

Representation Learning with Graph Embeddings

Graph embeddings transform graph-structured data into continuous vector spaces, preserving the relational and structural properties of the original graphs.

This transformation enables the application of standard machine learning techniques to data that would otherwise require specialized algorithms.

A significant approach in this domain involves **Graph Neural Networks (GNNs)**, which employ message-passing algorithms to iteratively aggregate information from neighboring nodes and edges. GNNs are well-suited to capturing both local and global graph patterns, making them versatile tools for representation learning [36]. Among GNNs, **Graph Convolutional Networks (GCNs)** are particularly noteworthy for their ability to perform convolutional operations over graph structures, enabling efficient learning of node embeddings that reflect both the node's features and its structural context [37, pp. 5–10].

An important advancement within this framework is the integration of attention mechanisms, as seen in **Graph Attention Networks (GATs)**. By assigning varying weights to neighboring nodes during the aggregation process, GATs focus on the most relevant nodes and edges. This selective attention is particularly advantageous in contexts where specific relationships are more influential, such as hierarchical data or layered networks [38, pp. 2–5].

Graph embeddings are utilized in various graph analysis tasks, including node classification, node clustering, node recommendation, link prediction, and graph classification. These embeddings transform graph structures into low-dimensional spaces, preserving their relational and structural properties while enabling efficient computation. Concrete applications include **community detection in social networks**, **recommendation systems**, and **knowledge graph completion**, where embeddings are used to infer missing links or relationships [39, Sec. 5].

Integrated Representations of Textual and Graph Embeddings

Graph embeddings effectively capture relational data structures but often fail to encapsulate the semantic richness of individual nodes. On the other hand, textual embeddings from pre-trained language models excel in providing deep semantic representations but lack the contextual relational depth offered by graph embeddings. Integrating these two paradigms yields a comprehensive representation that synergistically leverages both structural and semantic information.

One approach involves encoding textual data as node features within graph-based models. By embedding textual attributes of nodes using pre-trained language models, dense vector representations enriched with semantic information are generated. These enriched graphs are then processed using algorithms like GNNs or GATs, which produce representations that unify semantic and structural characteristics. This methodology, as exemplified by the **STAGE framework**, simplifies such integration by employing large language models to generate robust text embeddings, which are subsequently input into ensemble GNN architectures. The STAGE approach demonstrates competitive performance on node classification benchmarks while minimizing computational complexity and training overhead [40].

The combination of textual and graph embeddings is particularly impactful for complex tasks requiring both content understanding and relational reasoning. Models like **GraphFormers** exemplify this synergy by iteratively fusing textual and graph-based information, enhancing representation quality while maintaining computational efficiency [14].

Implications for Computational Argumentation

The integration of textual and graph embeddings offers a nuanced way to combine semantic richness with structural insights, facilitating advancements in computational argumentation. A thorough examination of the provided sources highlights significant contributions to this interdisciplinary domain.

Argument Quality Assessment

Marro et al. introduce a hybrid neural framework explicitly combining textual and graph embeddings for argument quality assessment. Their approach uses graph embeddings to capture the relational context among arguments and BERT-based textual embeddings to model the linguistic properties of natural language arguments. This methodology allows for the evaluation of quality dimensions such as cogency, rhetoric, and reasonableness in a corpus of student persuasive essays. The results indicate that this integrated model significantly outperforms state-of-the-art baselines, underscoring the value of combining textual and structural features in tasks like argument evaluation [15].

Debate Evaluation and Stance Prediction

Ruiz-Dolz et al. propose a hybrid method for evaluating complete argumentative debates by integrating formal argumentation theory with NLP techniques. Although this work does not directly combine textual embeddings with graph representations, it employs Transformer-based embeddings to model the natural language properties of arguments alongside graph-based computations of argument acceptability. This combination enhances the evaluation of debates, allowing the model to predict the winning stance with high accuracy by leveraging both logical and linguistic insights [41].

Argument Graph Construction

Lenz et al. present a modular pipeline for transforming natural language texts into argument graphs. Their work emphasizes the extraction of Argumentative Discourse Units (ADUs) and the prediction of relations such as support and attack to construct structured argument graphs. Although textual embeddings are not explicitly integrated, the modularity of the pipeline allows for potential enhancements using pre-trained language models to enrich the semantic representation of nodes in the graph [42].

The reviewed studies demonstrate that the integration of textual and graph embeddings is not only feasible but also advantageous for computational argumentation. While some works explicitly adopt hybrid methods, others provide foundational insights that can be extended with integrated approaches. Future research should focus on expanding these hybrid methodologies across diverse argumentation tasks, with an emphasis on scalability and interpretability.

2.3.3 Conclusion: Advancing Computational Argumentation through Representation Learning

Representation learning has significantly advanced the field of computational argumentation, offering powerful tools to analyze and model arguments in complex, large-scale, real-world contexts. Textual embeddings, exemplified by models such as BERT, have proven indispensable in capturing the semantic richness and contextual nuances of argumentative components, facilitating applications like argument mining, stance detection, and quality assessment. On the other hand, graph embeddings provide a robust means of encapsulating the structural relationships among arguments, enabling the analysis of interconnected argumentative structures.

The integration of textual and graph embeddings represents a promising avenue, offering a holistic approach to computational argumentation. By leveraging the complementary strengths of semantic and structural representations, hybrid methods have demonstrated their efficacy in tasks such as argument quality assessment, debate evaluation, and the construction of argument graphs. These studies highlight the potential of integrated approaches to bridge the gap between formal argumentation theories and practical applications in diverse domains.

As computational argumentation continues to evolve, future research should prioritize expanding the use of hybrid methodologies, refining their scalability, and improving interpretability. This focus will ensure that representation learning remains a cornerstone of computational argumentation, driving its applicability to increasingly complex and impactful real-world scenarios.

2.4 Gaps in the Literature and Justification for the Present Study

The evolution of computational argumentation has provided powerful tools for analyzing discourse, yet significant gaps remain in addressing the full complexity of real-world argumentative interactions. Traditional models often focus on either the structural relationships between arguments or their semantic content but rarely integrate both dimensions comprehensively. This chapter synthesizes the limitations identified in existing research and establishes the motivation for the present study. By bridging structural and semantic insights through the integration of graph and textual embeddings, this thesis explores a novel avenue to argumentation modeling. The chapter also highlights the broader significance of this work, positioning it within the field and outlining its potential applications for advancing discourse analysis in diverse contexts.

2.4.1 Summary of Key Limitations in Existing Research

As detailed in the preceding sections, traditional argumentation frameworks — even when extended with support relations or value-based considerations — often rely on highly abstracted structural models. While these models provide valuable insights into attack, defense, and preference dynamics, they nonetheless exhibit critical shortcomings in capturing the full complexity of real-world discourse. Specifically, they frequently neglect or oversimplify.

Semantic Richness

Abstract semantics largely ignore the nuanced linguistic features of arguments, making it difficult to represent the emotional tone, rhetorical strategies, implicit premises, and domain-specific contextual cues that can profoundly affect how arguments are perceived and evaluated.

Contextual and Pragmatic Dimensions

Many computational models omit the interactive, evolving nature of discourse, overlooking how arguments can shift in tone, content, or relevance over time. This limitation is particularly acute in digital settings, where rapid exchanges and constant feedback loops shape argumentative trajectories.

Integrated Representations - Contextual and Pragmatic Dimensions

While advances in graph embeddings capture structural information and language models like BERT excel at revealing semantic contexts, few studies combine these two dimensions in a unified model. As a result, opportunities for more robust and holistic discourse analysis remain largely untapped.

Scalability and Practical Utility

Existing methods that do integrate semantic and structural insights often involve highly specialized or computationally expensive approaches, limiting their feasibility for large-scale digital discourse analysis or real-time applications (e.g., automated debate moderation, policy-making platforms).

These gaps underscore the need for a more comprehensive modeling framework that can accommodate the complexities of digital discourse without sacrificing the analytical power of formal argumentation or the semantic depth afforded by modern language models.

2.4.2 Positioning the Current Research

This thesis is positioned at the intersection of argumentation theory, computational linguistics, and representation learning. By integrating graph embeddings with advanced textual embeddings, the proposed model aims to provide an enriched argumentation graph that captures both the relational structure of arguments and their semantic content. Building on frameworks such as Dung's abstract argumentation [22] and transformer-based language models (e.g., BERT [31]), this research specifically addresses the limitations outlined above.

Semantic–Structural Fusion

Rather than treating semantic analysis as a post hoc supplement, this work weaves semantic embeddings directly into graph construction, enabling a more cohesive understanding of argumentation that accounts for both logical interactions and the subtleties of language.

Enhanced Nuance

The model is designed to accommodate implicit premises, emotional appeal, and rhetorical devices, moving beyond the binary attack-and-support paradigm to incorporate richer modes of argumentative engagement.

Scalable Methodology

By leveraging established machine learning pipelines and well-documented frameworks (e.g., GNN architectures), the proposed approach aims to be both robust and efficient. This emphasis on scalability is pivotal for handling large, dynamic corpora typical of modern digital platforms.

2.4.3 Relevance and Contribution

By bridging formal argumentation structures with sophisticated semantic representations, this thesis endeavors to advance computational discourse analysis in several keyways.

Potential Methodological Innovation

Introducing an integrated framework that unifies structural and semantic insights provides a versatile tool for modeling, visualizing, and interpreting complex argumentative interactions. This hybrid perspective has the potential to offer more fine-grained analyses and to detect hidden or implicit argumentative strategies.

Practical Impact

From debate evaluation to policymaking and decision-support systems, a richer understanding of how arguments are interlinked, supported, and contested can facilitate more informed and collaborative public discourse. The proposed model could underpin applications that automatically highlight contradictory viewpoints, expose logical gaps, or summarize sprawling debates.

Future Extensions

While the immediate focus is on demonstrating the model's efficacy using a small dataset, the broader implications extend to diverse domains where robust argument analysis is critical — including legal, political, and educational contexts. The groundwork established here paves the way for deeper explorations into real-time conflict detection, sentiment-aware argument mapping, and retrieval-augmented argument generation.

In sum, the thesis seeks to experiment within underexplored areas in argumentation research with the goal of potentially developing a richer, more holistic model of how we argue and deliberate in debates.

3 Research Methodology

3.1 Research Design

This study aims to evaluate and compare the effectiveness of different embedding spaces—semantic, topological, and aggregated semantics—in representing argumentative discourse. The methodology focuses on embedding generation, visualization, clustering, and comparative analysis to assess their individual and combined capabilities.

3.2 Graph Schema

This research adopts a refined version of Bipolar Argumentation Frameworks to model argumentative discourse, integrating key advancements inspired by Toulmin's model. These refinements address the abstract semantic limitations observed in traditional frameworks. By incorporating distinct node types—such as Premises, Claims, Questions, and Authors—the schema delivers a nuanced and comprehensive representation of argumentative interactions, bridging gaps in logical structure and contextual subtleties.

3.2.1 Types of Nodes

1. **Question**: Questions function as critical elements within the discourse, framing the central topics or issues under debate. These nodes are categorized into:
    - **Initiating Questions**: Standalone, thematic entry points that do not target specific elements, enabling the exploration of broad dimensions within the discourse.
    - **Follow-up Questions**: Targeted inquiries directed at specific Claims or Premises, fostering detailed scrutiny of individual argumentative components.
2. **Claim**: Claims represent the argumentative positions or assertions presented in the discourse. These nodes may either support or oppose other Claims or Premises, forming the backbone of the argumentation structure.
3. **Premise**: Premises provide foundational evidence or reasoning that substantiates Claims. They supply the logical or empirical underpinnings necessary for validating the discourse.

4. **Author**: Authors identify the entities (individuals or groups) responsible for generating Questions, Claims, or Premises. This node type introduces an additional layer of contextual depth, allowing for a detailed analysis of authorship's influence on the discourse.

3.2.2 Types of Relationships

1. **Questions**: A directed relationship from a Question node to a Claim or Premise node, representing the act of inquiry and explicitly linking questions to their targets.
2. **Answers**: A directed relationship from a Claim or Premise node to a Question node, clarifying the connections between argumentative responses and their corresponding inquiries.
3. **Support**: A directed relationship indicating reinforcement or agreement between nodes, such as a Premise supporting a Claim. This relationship is pivotal for modeling constructive argumentative dynamics.
4. **Attack**: A directed relationship denoting contradiction or opposition between nodes, such as a Claim refuting another Claim or Premise. This relationship captures adversarial dynamics critical to argumentative discourse.
5. **Authored_by**: A directed relationship linking argumentative nodes (e.g., Questions, Claims, Premises) to their respective Author nodes. This connection enhances accountability and provides insights into the origins of arguments.

3.2.3 Graph Representation

The enhanced framework seamlessly integrates semantic and structural elements, enabling advanced analysis and visualization of argumentative discourse. By incorporating specialized node and relationship types, the schema captures the multidimensional nature of argumentation. This dual-layered approach facilitates precise modeling of both logical coherence and contextual nuances, thereby enhancing the analytical depth and practical applicability of the framework.

3.2.4 Future Considerations

Future refinements of this schema could incorporate additional elements to augment its representational capacity. For example:

- **Evidence Nodes**: Introducing nodes explicitly representing factual backing would enhance the robustness of argumentation.
- **Nuanced Relationships**: Expanding the range of relationship types, such as conditional dependencies (e.g., arguments valid only in specific contexts), would allow for more sophisticated modeling of interdependencies and argumentative strategies.
- **Weighted Relationships**: Assigning weights to relationships could reflect the relative strength or significance of connections, facilitating a more granular analysis of argument dynamics.

These advancements would broaden the schema’s utility, potentially enabling its application in increasingly complex discourse environments and fostering its relevance in both theoretical research and practical implementations.

3.3 Dataset

This study uses a **synthetic dataset** generated by a Large Language Model (LLM) to ensure flexibility and control over the structure and characteristics of argumentative discourse. The dataset is designed to simulate real-world debate scenarios while maintaining consistency in annotations and relationships to facilitate robust experimentation. It

3.3.1 Dataset Characteristics

The dataset consists of debates generated with annotated argumentative components, including claims, premises, and questions, providing a structured foundation for analysis. Relationships between these components are explicitly annotated with support and attack edges, enabling the construction of detailed and informative argumentation graphs. As a synthetic dataset, it offers the flexibility to tailor argument structures precisely, accommodating scenarios of varying complexity to meet the specific needs of the study.

The dataset centers around a debate on **Universal Basic Income (UBI),** with three fictional participants representing pro, contra, and moderate positions on the topic. It comprises approximately 100 argumentative nodes, each representing a distinct argument, and the relationships between them, which illustrate the interplay of supporting and opposing ideas. This structure allows for a nuanced exploration of the debate, capturing the dynamics of argumentation and the connections between differing perspectives.

3.3.2 Dataset Generation Process

Prompt Design

The dataset is generated using a carefully crafted prompt for a Large Language Model (LLM) to simulate argumentative discourse. The prompt is designed to elicit a range of argumentative components (questions, claims, premises) and relationships (question, answer, support, attack) to ensure diversity and relevance. Few-shot prompting is utilized by providing the LLM with multiple annotated examples of argumentative structures within the prompt. These examples illustrate the desired components (e.g., claim, premise) and relationships (e.g., support, attack), improving the quality and consistency of the generated outputs.

The full prompt can be found on GitHub [43].

Generation Procedure

A specialized Large Language Model (LLM) from OpenAI, known as GPT-o1 and designed for reasoning tasks, is used to generate the argumentative text based on carefully designed prompts. To ensure sufficient variability and richness in the argument structures, multiple rounds of generation are conducted. After generation, the outputs are meticulously reviewed to maintain the overall quality and relevance of the dataset.

3.3.3 Preprocessing

Preprocessing plays a crucial role in transforming raw data into a structured format suitable for analysis within this study. This step involves multiple processes to ensure the dataset is correctly represented as an argumentation graph, enabling the seamless application of embedding techniques and subsequent evaluations.

Graph Construction

The construction of the argumentation graph begins with importing the synthetic dataset, which is provided in a structured CSV format. This dataset consists of two primary files: *nodes.csv* and *edges.csv*. These files contain essential information about the argumentative discourse units (ADUs) and their interconnections, respectively. The preprocessing pipeline follows a systematic approach to transform these files into a usable graph format:

1. **Node Representation:**
    - Each row in the *nodes.csv* file corresponds to a specific ADU, including elements such as claims, premises, questions, and authors.
    - Nodes are created for every unique ADU, ensuring that the graph captures all relevant argumentative components.
2. **Edge Connectivity:**
    - Relationships between nodes, such as "support" and "attack," are derived from the *edges.csv* file.
    - Each row in this file specifies a directed edge between two nodes, defining the nature of their relationship.

Tools and Libraries

Preprocessing relies on state-of-the-art graph tools, such as neo4j, to construct and manipulate the graph. These tools provide robust support for importing CSV data, creating graph structures, and encoding additional metadata as node and edge attributes. By leveraging these libraries, the preprocessing phase ensures consistency and scalability.

Output

The final output of the preprocessing stage is a fully constructed argumentation graph where:

- Nodes represent ADUs (e.g., claims, premises, questions, authors).
- Edges encapsulate the relationships between these nodes (e.g., support, attack).
- Additional metadata, such as authorship and node attributes, is integrated into the graph.

This graph serves as the foundation for generating embeddings and conducting comparative analyses, bridging the gap between raw data and meaningful insights.

3.4 Embedding Generation

The generation of embeddings forms a critical component of this research, enabling the comparative analysis of semantic, structural, and integrated embedding spaces. This section outlines the methods and rationale behind each embedding approach, emphasizing the objectives, processes, and implications for modeling argumentative discourse.

3.4.1 Semantic Embeddings

Semantic embeddings capture the linguistic and contextual nuances of arguments, providing a dense vector representation of textual content. Capturing these nuances is crucial for argumentative discourse analysis as it enables a deeper understanding of the underlying intent, meaning, and rhetorical strategies employed in arguments. By preserving context and subtle linguistic variations, semantic embeddings help in accurately representing the interplay of ideas, which is essential for modeling complex argumentative interactions.

Choice of Model

Pre-trained models, such as *all-mpnet-base-v2*, are utilized to generate these embeddings due to their specialization in sentence-level representations. Unlike traditional models such as BERT, *all-mpnet-base-v2* offers enhanced performance for sentence embeddings, making it particularly suitable for capturing the semantics of argumentative discourse components.

Sentence-level embeddings excel in maintaining the contextual coherence of entire statements, a critical factor for argumentative discourse analysis, where the focus is on the intent and meaning of full arguments rather than individual words. Conversely, token-level embeddings, like those from BERT, provide finer granularity, which is beneficial for tasks that require detailed word-level interpretation. However, this granularity often necessitates additional aggregation steps to derive sentence-level meaning, which can introduce noise and reduce consistency. As a result, *all-mpnet-base-v2* is preferred for its ability to directly generate coherent, high-quality sentence representations with minimal preprocessing.

This model is specifically chosen for its ability to maintain high accuracy across diverse datasets, outperforming other sentence-level models like Universal Sentence Encoder (USE) or Sentence-BERT in capturing fine-grained contextual meaning. Additionally, *all-mpnet-base-v2* integrates masked language modeling and permuted language modeling techniques, which enable a deeper understanding of sentence semantics and contextual relationships. These capabilities make it particularly well-suited for tasks like argumentative discourse analysis, where capturing nuanced interactions between ideas is critical.

Embedding Process

The embedding process involves generating 768-dimensional vectors for each argumentative node (e.g., claims, premises, and questions). These embeddings encapsulate the semantic content of the nodes, providing a rich representation that facilitates downstream analyses.

To ensure uniformity and improve the comparability of results, the generated embeddings are normalized. This process involves scaling the embeddings to have unit norm, which ensures that all vectors are represented in the same magnitude range. By normalizing, the model prevents certain nodes from disproportionately influencing analysis due to differences in vector magnitude, thereby enabling a fair comparison across all nodes in the embedding space. This approach ensures that the semantic characteristics of each node are accurately and consistently represented within the embedding space.

3.4.2 Structural Embeddings

Structural embeddings focus on encoding the relational and topological properties of the argumentation graph. This study employs **Fast Random Projection (FastRN)** embeddings to generate vector representations based solely on the graph's structural attributes—nodes and edges—without incorporating node-specific features. By isolating these structural aspects, the embeddings provide a clear and unbiased representation of the argumentation framework's architecture. This approach enables a deeper exploration of its organizational patterns and topological characteristics, highlighting key structural dynamics effectively. The emphasis on structure rather than content makes this approach invaluable for identifying key relational dynamics within the argumentative discourse.

Choice of Method

FastRN is chosen for its efficiency and effectiveness in capturing graph topology. At its core, FastRN operates by initializing each node in the graph with a random vector, which is then iteratively refined based on the node's connections within the graph. During this refinement process, each node updates its embedding by aggregating information from its neighbors' embeddings, weighted by their respective degrees. This iterative propagation process captures the local and global structural context of each node, ensuring that nodes with similar roles or connectivity patterns in the graph converge to similar embeddings.

Unlike alternative methods that often require extensive computational resources or additional data features, FastRN provides a streamlined solution focused purely on structural properties. The method computes embeddings by analyzing the connections and roles of nodes within the graph, which helps highlight relationships such as hubs, central nodes, or hierarchical positioning. This choice ensures scalability for large graphs while maintaining the precision necessary for accurate analysis. Furthermore, its adaptability allows the method to be applied to diverse graph structures, making it a robust tool for argumentation analysis.

Nodes and Relationships

Nodes included in this process are claims, premises, authors, and questions. Each node type represents a distinct argumentative component, contributing unique insights into the discourse structure. Relationships between these nodes are categorized as follows:

- **Attack** and **Support**: These are treated as undirected edges, capturing the essence of argumentative connections that define agreements or contradictions.
- **Questions** and **Answers**: These relationships are similarly treated as undirected edges, providing a generalized view of the inquiry-response dynamics inherent in the graph.
- **Authored_by**: This relationship, linking argumentative nodes to their respective authors, is reversed in direction to mitigate its dominance in the graph. This reversal ensures that no single relationship type disproportionately influences the overall embedding representation, thereby preserving a balanced view of the structural dynamics.

The inclusion of these nodes and relationships ensures that the graph comprehensively captures the complexity of argumentative interactions. By treating certain edges as undirected, the model abstracts away unnecessary directional biases, allowing the relationships to be analyzed without assuming an inherent hierarchy or causality. This decision is particularly impactful in contexts where the directionality of relationships may obscure broader patterns of interaction, such as reciprocal argumentation or mutual support. By simplifying these edges, the analysis emphasizes the overall connectivity and clustering within the graph, leading to a more holistic understanding of its structural properties.

Hyperparameters

Key hyperparameters for FastRN include an embedding dimension of 768 and a normalization strength of 0.5. The embedding dimension defines the vector size used to represent each node, balancing the need for expressive power with computational efficiency. The normalization strength parameter scales the initial random vector assigned to each node by its degree raised to the power of normalization strength. This scaling process ensures that nodes with higher connectivity are appropriately weighted, reflecting their centrality or importance within the graph. Additionally, this approach prevents overemphasis on isolated or low-degree nodes, fostering a balanced representation of the entire graph's topology.

These hyperparameter settings were chosen after iterative experimentation to optimize the model's performance while maintaining computational feasibility.

Implications

Structural embeddings are crucial for understanding the relational dynamics within the argumentation graph. By relying solely on topological properties, these embeddings reveal the organizational structure and interactions between argumentative components. This approach allows researchers to uncover patterns such as clustering, node centrality, and relational symmetry, which are not immediately evident from the raw data.

Moreover, structural embeddings provide a foundation for cross-comparative analyses. By abstracting away semantic content, they enable an unbiased examination of how argumentative structures evolve across different contexts or datasets. This abstraction is particularly useful for tasks such as identifying common argumentation strategies, detecting structural anomalies, or assessing the robustness of argumentative frameworks in diverse discourse environments.

Overall, the reliance on purely structural features ensures a focused and rigorous analysis, offering insights that complement those derived from semantic embeddings.

3.4.3 Aggregated Semantic Embeddings

Aggregated semantic embeddings offer a powerful means of capturing the nuanced interplay between arguments, questions, and premises in argumentative discourse. By leveraging GraphSAGE, this approach enriches each node's representation with **contextual information** from its local neighborhood, allowing the model to incorporate semantic nuances from surrounding nodes. This enhancement captures inter-node relationships, emphasizing how individual arguments relate to their broader discourse context, thus improving the overall representation quality. This method emphasizes the importance of both individual node features and their connections within the discourse graph, creating embeddings that are well-suited for sophisticated analyses of argumentation patterns and relationships. By leveraging the semantic context captured by node features, this method enriches each node's embedding with information from its neighbors, enabling a more nuanced representation of the argumentative discourse while preserving its semantic integrity. This approach ensures that the embeddings not only represent the individual arguments but also encapsulate their interaction within the discourse network, creating a robust and dynamic model suitable for a range of analytical tasks.

Choice of Method

GraphSAGE, which stands for **Graph Sample and Aggregate**, is employed as the aggregation method due to its ability to generate node representations by sampling and aggregating information from neighboring nodes. Under the hood, GraphSAGE operates by iteratively sampling a fixed number of neighbors for each node and applying an aggregation function, such as mean, max-pooling, or an LSTM-based aggregator, to combine their features. The resulting aggregated features are then concatenated with the node's own features and passed through a neural network to generate an updated embedding. This process is repeated over multiple layers, allowing the model to incorporate information from increasingly distant neighbors in the graph, capturing both local and extended contexts effectively. This capability is particularly effective for argumentative graphs, where the semantic connections between arguments, questions, and premises are crucial. Furthermore, GraphSAGE's flexibility allows for iterative refinement of embeddings, ensuring that each representation captures both immediate and extended contexts within the graph.

By aggregating information from sampled neighbors, GraphSAGE mitigates the computational overhead associated with processing entire neighborhoods, making it scalable for larger datasets. Additionally, its aggregation strategies, such as mean or pooling functions, provide mechanisms to balance the influence of densely and sparsely connected nodes, resulting in embeddings that reflect the unique semantic properties of each argument while integrating the broader discourse dynamics.

Input Features and Relationships

Input Features

To ensure that each node's embedding captures relevant contextual semantic information, preprocessed semantic embeddings derived from PCA-reduced *all-mpnet-base-v2* vectors are utilized. These embeddings encapsulate the critical semantic properties of claims, premises, and questions. Principal Component Analysis (PCA) transforms the original high-dimensional features into orthogonal components ranked by variance, reducing data complexity while retaining the most informative aspects. This approach minimizes computational resource requirements, such as memory and runtime, while enhancing robustness by focusing on essential features and preventing overfitting during training. However, this dimensionality reduction may involve trade-offs, such as the potential removal of less prominent features that could carry nuanced semantic information in certain contexts.

Relationships

Relationships in the graph are configured to ensure that meaningful connections between nodes are captured:

- **Attack** and **Support**: Maintained in their natural orientation to preserve the inherent argumentative dynamics. These relationships emphasize the adversarial and supportive interactions that form the backbone of argumentative discourse.

- **Questions** and **Answers**: Treated as undirected edges to generalize inquiry-response dynamics. This configuration abstracts the directional constraints, focusing on the semantic exchange between nodes rather than the explicit flow of inquiry and response.

By carefully curating input features and relationship orientations, the model ensures that both the semantic content and the relational structure of the discourse are effectively represented, enabling deeper and more accurate analyses.

Training and Hyperparameters

Key hyperparameters for the GraphSAGE model include:

- **Embedding Dimension**: 768, matching the dimensionality of the semantic embeddings for consistency. This alignment ensures that downstream analyses can seamlessly integrate embeddings from different layers.
- **Sample Sizes**: [10, 5], representing the number of neighbors sampled at each layer to balance depth and breadth in the graph. Sampling reduces computational complexity while maintaining sufficient context to enrich node embeddings.
- **Learning Rate**: 0.01, optimized for stable convergence. The chosen rate prevents abrupt changes in weights while ensuring steady progress during training.
- **Epochs**: 40, providing sufficient iterations for the model to learn meaningful embeddings. This balance avoids underfitting while preventing overfitting.
- **Aggregator**: “Mean”, chosen to prevent over-smoothing and dimensionality collapse, ensuring that each node retains unique semantic features. This strategy maintains the individuality of nodes while integrating information from their neighborhoods.

Implications

Aggregating semantic embeddings using GraphSAGE enhances the representational power of the model by incorporating contextual information from neighboring nodes. This method preserves the semantic richness of individual arguments while contextualizing them within the discourse network. The result is a set of embeddings that capture not only the content of individual nodes but also the broader semantic relationships within the graph. By incorporating neighbor information, these embeddings reveal intricate patterns of argumentation, such as recurring themes, implicit dependencies, and contextual nuances.

This approach facilitates tasks such as identifying argumentation strategies, assessing discourse coherence, and analyzing the interplay between different argumentative roles. For instance, by examining clusters of nodes with similar embeddings, researchers can uncover thematic groupings or argumentative alignments that provide insights into the discourse's structure and intent.

By focusing exclusively on semantic aggregation, the embeddings provide a robust framework for evaluating the semantic structure of discourse while avoiding the complexities of integrating structural embeddings. This approach also ensures scalability, as the reliance on preprocessed semantic features and GraphSAGE's efficient sampling mechanisms allow the model to handle larger argumentative datasets. Consequently, the aggregated embeddings serve as a versatile tool for both theoretical research and practical applications, bridging the gap between detailed semantic analysis and holistic discourse modeling.

3.5 Comparison of Embedding Spaces

This section conducts a preliminary evaluation of semantic, structural, and aggregated semantic embedding spaces to explore their respective capabilities in representing argumentative discourse. By combining visualization, clustering evaluation, and quantitative metrics, the analysis ensures that both qualitative and quantitative dimensions of the embeddings are examined. Visualization reveals clustering tendencies and thematic patterns, clustering evaluation quantifies group cohesion and separation, and quantitative metrics provide precise measurements of embedding performance. This combination is crucial for capturing the diverse aspects of how embedding spaces represent argumentative discourse.

3.5.1 Visualization

Visualization serves as an indispensable tool for uncovering the latent structures within embedding spaces, providing intuitive insights into clustering patterns and distributional properties of argumentative components. This study utilizes PCA and **t-Distributed Stochastic Neighbor Embedding (t-SNE)** to project high-dimensional embeddings into a more interpretable three-dimensional space.

Process

Initially, PCA is employed to reduce noise and enhance computational efficiency by retaining only those components that collectively explain 99% of the data variance. This threshold balances the need for simplification with the preservation of meaningful information. By focusing on components that capture the majority of the variance, PCA ensures that the core patterns and structures within the data are retained, while irrelevant or redundant variations are discarded. This approach enables downstream analyses to operate efficiently without compromising the integrity of the original data. This step ensures that the essential characteristics of the data are preserved while discarding redundant information.

Subsequently, t-SNE is applied to map the embeddings into a three-dimensional space. t-SNE operates by constructing pairwise probability distributions for points in the high-dimensional space, which reflect their similarities, and then optimizing a corresponding distribution in the lower-dimensional space to minimize divergence. By focusing on local neighborhoods, t-SNE preserves small-scale relationships while sacrificing some global structure, meaning that while local clusters of similar points are well represented, the overall spatial arrangement of distant clusters may not reflect their true global relationships. For example, in a dataset where distant groups represent separate thematic topics, t-SNE may emphasize internal cohesion within each topic but distort their relative positions in the overall visualization. This method effectively highlights subtle clustering tendencies and reveals intricate patterns in the embeddings.

The resulting visualizations are scatterplots where nodes are color-coded based on argumentative roles (e.g., claims, premises), authorship, or debate sections to emphasize thematic patterns.

Configuration Details

Semantic Embeddings: PCA preprocessing is applied to the embeddings before t-SNE. The t-SNE parameters include a perplexity of 15, a maximum of 100,000 iterations, and a **cosine similarity** metric. Author-related metadata is excluded, allowing the analysis to focus exclusively on semantic content.

Structural Embeddings: Similarly, PCA preprocessing is used before t-SNE application. Parameters remain consistent (perplexity: 15, iterations: 100,000), but the **Euclidean distance** metric is employed. Author information is included to capture relational dynamics effectively.

Aggregated Semantic Embeddings: Unlike the other two spaces, PCA preprocessing is omitted to preserve the complete semantic richness of the embeddings. t-SNE parameters include a perplexity of 20, a maximum of 100,000 iterations, and a **cosine similarity** metric. Author information is excluded to maintain focus on aggregated semantic properties.

Parameter Implications

Perplexity: This parameter influences how t-SNE balances local and global aspects of the data. A perplexity of 15 reflects a focus on preserving local neighborhoods, ideal for datasets where local clustering is critical. Aggregated semantic embeddings, with a perplexity of 20, capture slightly broader contexts.

Distance Metric: The choice of **cosine** similarity for semantic and aggregated embeddings emphasizes the directional relationship between vectors, which is crucial for capturing semantic nuances. Conversely, **Euclidean** distance for structural embeddings reflects spatial relationships, aligning with the graph topology.

Implications

The scatterplots derived from these visualizations are valuable for gaining a preliminary understanding of how well each embedding space captures the semantic and structural characteristics of the discourse. By observing the distribution and clustering of nodes based on argumentative roles and thematic contexts, researchers can identify initial strengths and limitations of each embedding space. Notable differences in clustering patterns across semantic, structural, and aggregated semantic embeddings provide insights into their potential efficacy in modeling argumentative interactions.

3.5.2 Clustering Analysis

Clustering analysis provides a **quantitative** complement to visualization, enabling an initial evaluation of grouping quality within each embedding space. This study employs **k-means clustering**, a robust unsupervised learning method, to partition the nodes and assess their separation into meaningful clusters.

Process

The k-means algorithm partitions data into a pre-defined number of clusters by iteratively minimizing the variance within each cluster. The process begins by initializing **cluster centroids**, often chosen randomly, and assigning data points to the nearest centroid based on a specified distance metric. The centroids are then updated as the mean of all points assigned to their cluster. This process repeats until convergence, where data point assignments stabilize, or a maximum number of iterations is reached. K-means effectively groups nodes by *minimizing* intra-cluster variance while *maximizing* inter-cluster separation.

Understanding Silhouette Score

The **Silhouette Score** is a quantitative metric used to assess the quality of clustering. It measures how similar a data point is to its assigned cluster compared to other clusters. For each data point, the Silhouette Score is calculated as the difference between the mean distance to points in the same cluster (**cohesion**) and the mean distance to points in the nearest neighboring cluster (**separation**), normalized by the maximum of the two values.

Scores range from -1 to 1, where higher values indicate well-defined clusters with high cohesion and separation. A score near 0 suggests overlapping clusters, and negative scores imply misclassified points. For example, a Silhouette Score of 0.8 might indicate that most data points are clearly associated with their respective clusters, with minimal overlap between clusters. Conversely, a score of -0.2 could suggest significant misclassification, where points are assigned to clusters they are not closely related to, reflecting poor clustering quality.

This metric is particularly valuable for evaluating how effectively different embedding spaces distinguish between semantic differences.

Comparative Insights

The analysis focuses on comparing the **grouping quality** within semantic and aggregated semantic embeddings, assessing their relative strengths in capturing argumentative nuances. By examining Silhouette Scores across varying cluster configurations, researchers can identify the optimal cluster count and evaluate which embedding type achieves superior separation of argumentative components. This quantitative approach provides a starting point for assessing the representational power and utility of each embedding space.

3.5.3 Evaluation Metrics

To ensure a balanced comparison, the analysis incorporates both qualitative and quantitative evaluation metrics. These complementary approaches provide a foundational understanding of the embedding spaces' effectiveness.

Silhouette Score: This metric **quantitatively** assesses clustering quality, offering insights into the degree of separation and cohesiveness among clusters. By comparing Silhouette Scores across different embedding spaces, researchers can identify initial trends in their relative efficacy in modeling argumentative structures.

Visual Insights: Observational patterns derived from scatterplots provide **qualitative** insights that numerical metrics may overlook. For instance, scatterplots can reveal tightly clustered groups of nodes corresponding to specific argumentative roles, such as claims or premises, or expose anomalies like outlier nodes that do not fit into expected clusters. These visual cues help identify trends or inconsistencies in the data that might not be evident through metrics alone, such as thematic overlaps or hierarchical relationships within the argumentation structure. These observations help reveal thematic groupings, latent relationships, and clustering tendencies within the data, adding depth to the analysis.

Synthesis

The combination of Silhouette Scores and visual insights ensures a preliminary evaluation of the embedding spaces. While the Silhouette Score provides a quantitative measure, visual analysis captures contextual nuances and practical implications. Together, these metrics offer an exploratory understanding of how each embedding space represents the complexity of argumentative discourse, setting the stage for further validation and refinement.

3.6 Code and Reproducibility

Ensuring the reproducibility of experiments and results is a cornerstone of rigorous scientific research. This section outlines the tools, workflows, and best practices employed in this study to facilitate reproducibility and provide transparency in the analysis of embedding spaces and argumentative discourse.

3.6.1 Project Workflow

The project follows a modular and well-documented workflow to ensure that each step can be replicated independently. Moreover, the complete workflow can be executed as a single **data pipeline** with one command, streamlining the process and minimizing user intervention. The workflow involves:

(1) **Data Preparation:**

- Combining multiple CSV files into unified *nodes.csv* and *edges.csv* files using the *merge_data.py* script.
- Ensuring consistency and correctness in data format and structure.

(2) **Graph Construction:**

- Using the *main.py* script to construct the argumentation graph in Neo4j.
- Loading the combined *nodes.csv* and *edges.csv* files into the database, ensuring seamless integration of all argumentative components.

(3) **Embedding Generation:**

- Generating semantic embeddings (*all-mpnet-base-v2*), structural embeddings (*FastRN*), and aggregated semantic embeddings (*GraphSAGE*).
- Employing dedicated scripts for each type of embedding, with hyperparameters documented and easily adjustable within their respective Python files.

(4) **Visualization and Analysis:**

- Utilizing Jupyter notebooks to visualize embeddings in three-dimensional space using t-SNE.
- Performing clustering and comparative analysis using k-means and Silhouette Scores.

3.6.2 Tools and Technologies

The following tools and technologies are employed to ensure accuracy and reproducibility:

- **Python** for scripting and analysis, leveraging libraries such as *numpy, pandas, scikit-learn,* and *matplotlib.*
- **Neo4j** as the graph database platform for managing and querying the argumentation graph.
- **Jupyter Notebooks** for interactive visualizations and exploratory analysis.
- **Pre-trained Models** like *all-mpnet-base-v2* for semantic embeddings and *GraphSAGE* for aggregated embeddings.
- **Version Control** using Git to track changes in scripts, datasets, and configurations.

3.6.3 Reproducibility Features

To further ensure consistency and reproducibility across stochastic methods, a random seed of 42 is set for all applicable processes, such as embedding generation, clustering, and visualization. This guarantees that results are replicable under identical configurations. It is important to note that this does not apply to the synthetic dataset generation, as its variability is intentionally preserved to simulate real-world scenarios.

To enhance reproducibility, the following practices are implemented:

1. **Documentation:**
    - Each script includes detailed comments explaining its functionality, expected inputs, and outputs.
    - The *README.md* file provides clear instructions for installation, usage, and troubleshooting.
2. **Hyperparameter Configuration:**
    - All hyperparameters for embedding generation and analysis are centralized within easily accessible Python files.
    - Clear guidelines on tuning hyperparameters are provided in the documentation.

3. **Dataset Versioning:**

   - The synthetic dataset used in this study is versioned and accompanied by the prompt used for its generation to ensure traceability.

4. **Reproducible Pipelines:**

   - The entire workflow, from data preparation to visualization, can be executed using modular scripts and notebooks.
   - The ability to execute the complete pipeline with a single command ensures streamlined reproducibility and minimizes user error.
   - Output files, such as embeddings and visualizations, are consistently named and organized for easy reference.

3.6.4 Challenges and Limitations

Despite the rigorous workflow, certain limitations should be acknowledged:

- **Synthetic Dataset:** While the synthetic dataset ensures control over structure and annotations, it may not fully capture the complexity of real-world argumentative discourse.
- **Computational Requirements:** Depending on the size of the datasets, the embedding generation process, particularly for GraphSAGE, is resource-intensive and may require high-performance computing resources.

By addressing these challenges and adhering to the outlined practices, this study provides a robust foundation for reproducible and transparent research in the modeling of argumentative discourse.

3.7 Methodological Summary

This section synthesizes the diverse methodologies employed in this study to systematically generate, visualize, and compare embedding spaces for argumentative graphs. It underscores the balance between methodological rigor and interpretability, ensuring that the analytical outcomes align closely with the research objectives.

The methodological framework prioritizes efficiency and clarity, offering a streamlined yet robust approach to understanding argumentative discourse. Embedding spaces are evaluated using visualization and clustering techniques, which provide complementary insights into their representational strengths and limitations. These methodologies emphasize the dual importance of qualitative interpretation and quantitative rigor in examining semantic and structural nuances within the discourse.

Scatterplots and clustering metrics emerge as core analytical deliverables, bridging the gap between abstract representations and tangible insights. Through techniques such as t-SNE visualizations and Silhouette Score assessments, the study highlights patterns, thematic groupings, and structural dynamics intrinsic to argumentative discourse. This comprehensive approach ensures that the chosen methodologies not only fulfill theoretical objectives but also offer practical applicability for advancing research in discourse modeling.

The results presented in the subsequent chapter build on this methodological foundation, offering an in-depth exploration of the comparative performance of semantic, structural, and aggregated semantic embeddings. These analyses illuminate the potential of embedding spaces to capture the intricate interplay of arguments, fostering a deeper understanding of their strengths and limitations within varied discourse environments.

4 Results and Analysis

4.1 Overview of Results

This section provides an in-depth examination of the results obtained through the methodologies described in the previous chapter. By addressing the limitations of traditional argumentation models highlighted in the problem statement, the analysis evaluates the extent to which semantic, structural, and aggregated semantic embedding spaces capture nuanced argumentative dynamics. Quantitative metrics, visualization insights, and clustering analysis are employed to assess their representational strengths and limitations, with a focus on modeling the interplay between structural relationships and semantic content.

4.2 Structural Embedding Analysis

4.2.1 Visualization Insights

The structural embedding space is visualized through three distinct perspectives to explore the representational capacity of the embeddings and highlight key clustering patterns. A 3D visualization is employed for all perspectives, providing a detailed spatial representation of the embedding space. For enhanced readability, an interactive digital version is made available online, allowing for deeper exploration of the clustering patterns and relationships.

These visualizations collectively provide a comprehensive view of the structural embedding space, uncovering its capacity to represent argumentative dynamics from multiple dimensions. Readers are encouraged to explore the digital version for enhanced interaction with the embedding space and deeper insights [44]. Detailed observations and implications from these perspectives will be explored in the interpretative analysis.

Separation of Argumentative Roles

Nodes are color-coded based on their type, distinguishing between Claims, Premises, Questions, and Authors. This perspective offers insights into how well the structural embeddings separate argumentative components, emphasizing their relational roles and interactions.

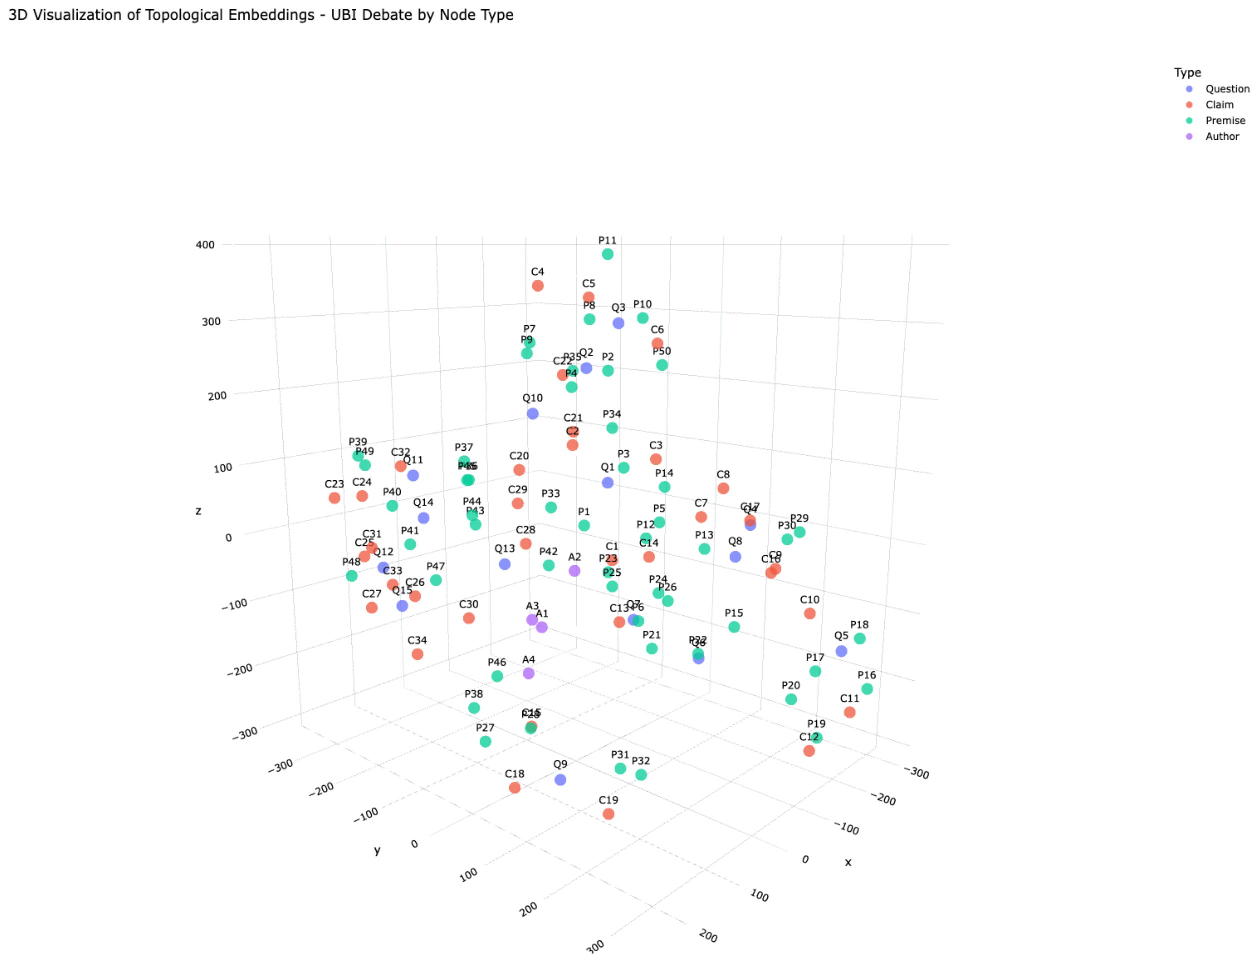

Figure 6 - Structural Embeddings by Node Type

Thematic Sections

Nodes are color-coded according to their thematic section, such as Introduction, Ethical and social dimensions, Macroeconomic Impacts, etc. This perspective examines whether the embeddings capture thematic groupings within the graph, reflecting logical segmentations of the discourse.

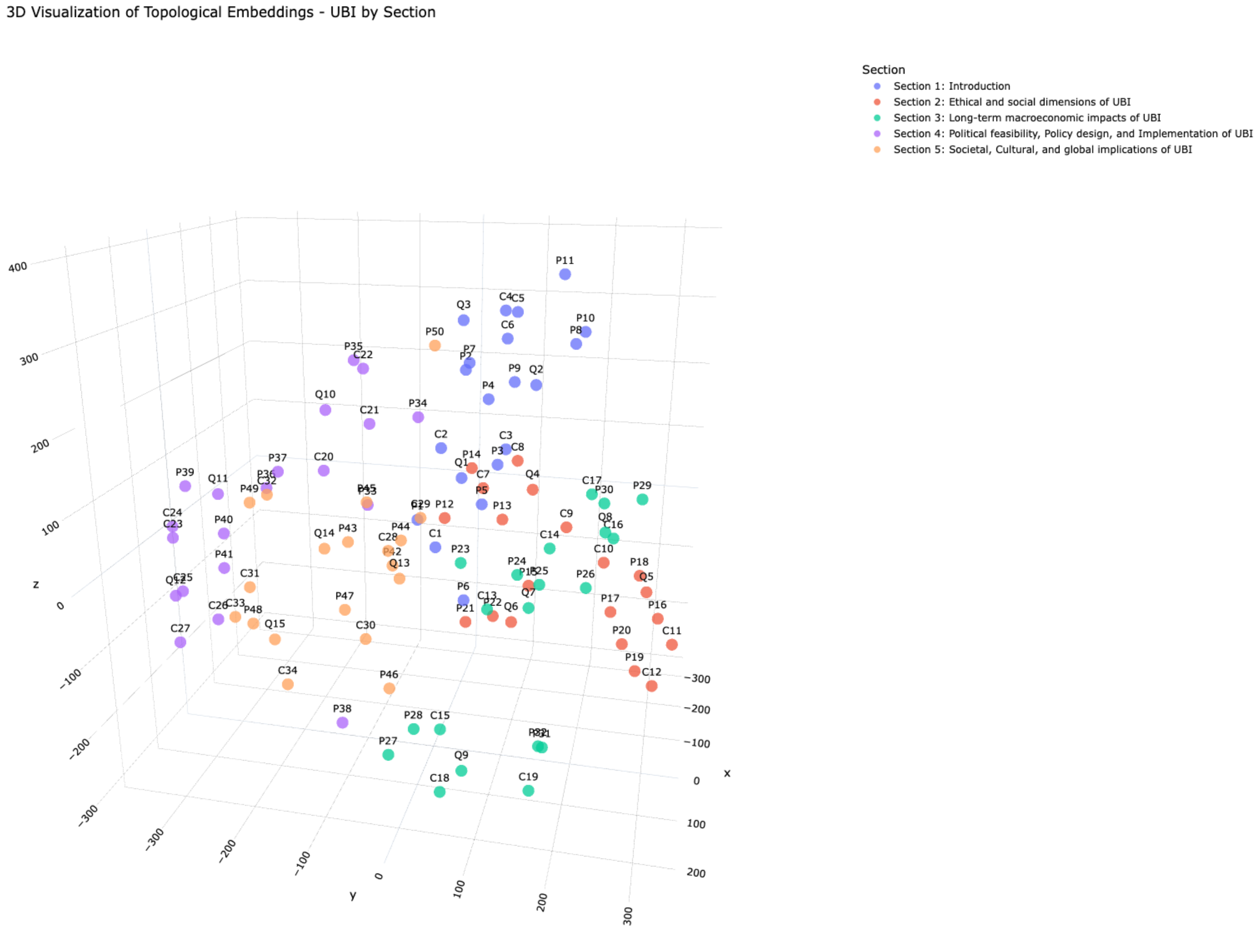

Figure 7 - Structural Embeddings by Thematic Sections

Author Attribution

Each node is color-coded based on its author. This visualization explores the impact of authorship on structural organization, identifying whether individual authors' contributions exhibit distinct patterns or overlap significantly with others.

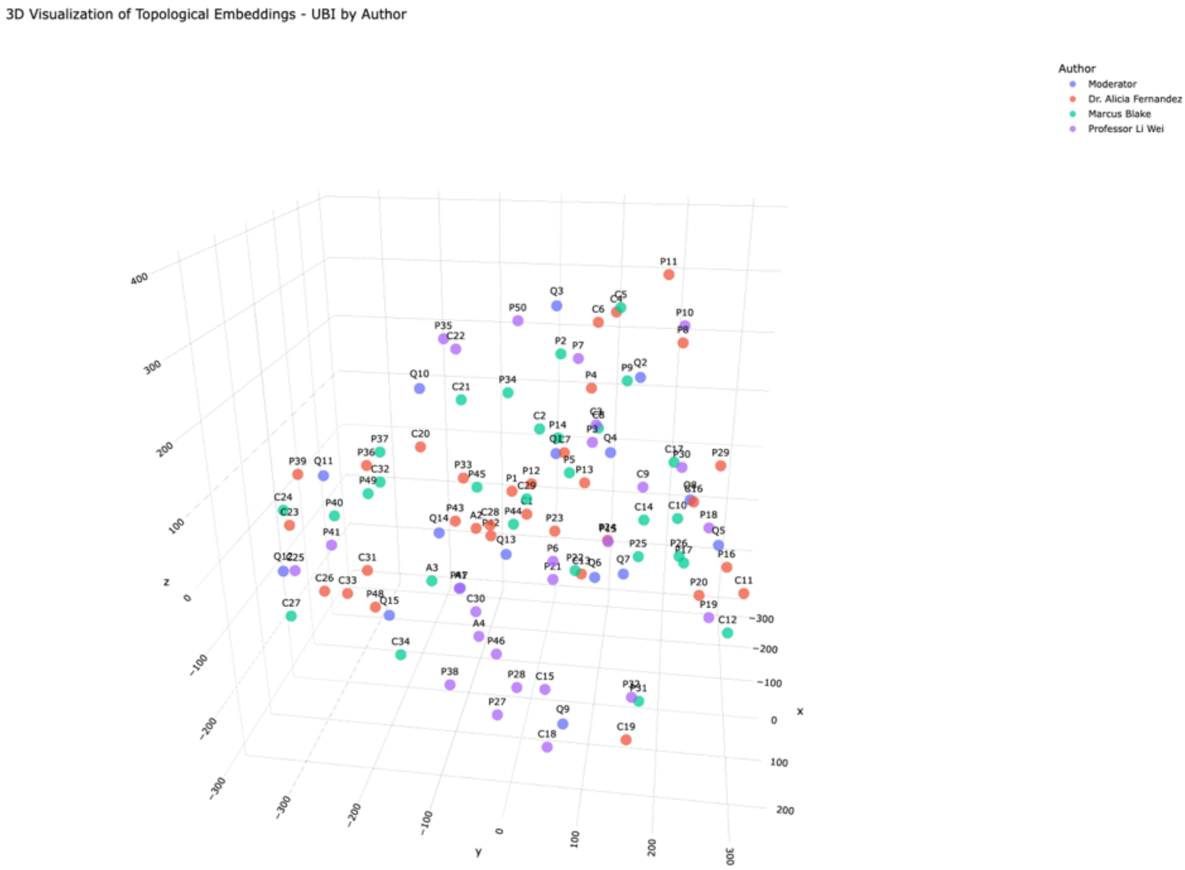

Figure 8 - Structural Embeddings by Author Attribution

4.2.2 Interpretative Analysis

The structural embeddings reveal significant insights into the relational and organizational dynamics of the argumentation graph. Observations from the visualizations are detailed below:

Separation of Node Types

The embeddings demonstrate a nuanced clustering of elements, affirming their ability to represent the logical architecture of the argumentation graph. Elements closely connected within the graph, either directly or indirectly, tend to form cohesive clusters:

- **Argumentative Elements and Authors**: Claims, Premises, and Questions often cluster near their respective authors, reflecting the direct relationship between argumentative contributions and their originators.
- **Supporting and Attacking Premises**: Premises providing support or opposition are consistently situated near the Claims they reinforce or contradict, indicating the embeddings' ability to capture argumentative dynamics effectively.
- **Claims and Questions**: Claims cluster closely with the Questions they address or are challenged by, highlighting the embeddings' capacity to model inquiry-response relationships within the discourse.

This clear alignment of elements underscores the structural embeddings' capacity to capture both explicit and implicit connections within the argumentative framework, while maintaining logical cohesion.

Thematic Groupings

Nodes associated with distinct thematic sections, such as Economics, Ethics, and Politics, exhibit localized clustering. This pattern indicates that structural embeddings successfully capture thematic coherence within argumentative discourse, even in the absence of explicit semantic content. These thematic clusters often reflect logical segmentations in the discourse, where arguments and premises related to a specific topic naturally converge within the embedding space. For example, discussions centered on "Economics" form a distinct cluster, likely encompassing claims and premises discussing financial policies, market behaviors, and economic theories.

However, overlaps between themes suggest interdependencies where discussions on specific topics influence the other topic and vice versa. Such overlaps highlight the embeddings' sensitivity to nuanced relationships between themes, capturing the interplay of ideas that span multiple discourse categories. Additionally, these patterns provide insight into areas where the discourse transitions between topics, often indicative of multifaceted arguments that leverage concepts from different thematic domains. This ability to capture both distinct clusters and inter-thematic overlaps enhances the utility of structural embeddings in modeling complex argumentative environments.

Authorship Dynamics

The author-based visualization reveals nuanced patterns due to the directional reversal of the "AUTHORED_BY" relationship, which minimizes its dominance in the embedding space. As a result, distinct authorship clusters are less pronounced. Instead, argumentative elements tend to align more closely based on their relational roles rather than strict authorship boundaries.

Notably, the contributions of the Moderator and Professor Li Wei, a neutral debate participant, are more distributed across the embedding space. This distribution reflects their balanced engagement with both pro and contra participants, emphasizing their neutral and connective roles within the debate. For other participants, clustering remains partially visible, particularly where arguments strongly align with specific perspectives or recurring themes. These observations suggest that while authorship contributes to the structural organization, it is secondary to the relationships and argumentative dynamics captured by the embeddings.

Collectively, these findings underscore the structural embeddings' capacity to model the topological characteristics of argumentative graphs, capturing logical relationships, thematic segmentation, and authorial impact effectively. These observations lay the groundwork for exploring how structural embeddings contribute to broader discourse modeling and identifying areas for further refinement.

4.3 Semantic Embedding Analysis

4.3.1 Quantitative Performance

This section quantitative evaluates the clustering quality within the semantic embeddings space. Clustering quality is assessed using **Silhouette Scores** across a range of cluster numbers to identify optimal configurations for capturing nuanced argumentative dynamics. The analysis focuses on how these embeddings address limitations in traditional models, such as modeling semantic content and interconnections within the argumentation graph.

Key Observations

A graph of Silhouette Scores versus the number of clusters reveals peaks indicating the most coherent cluster configurations.

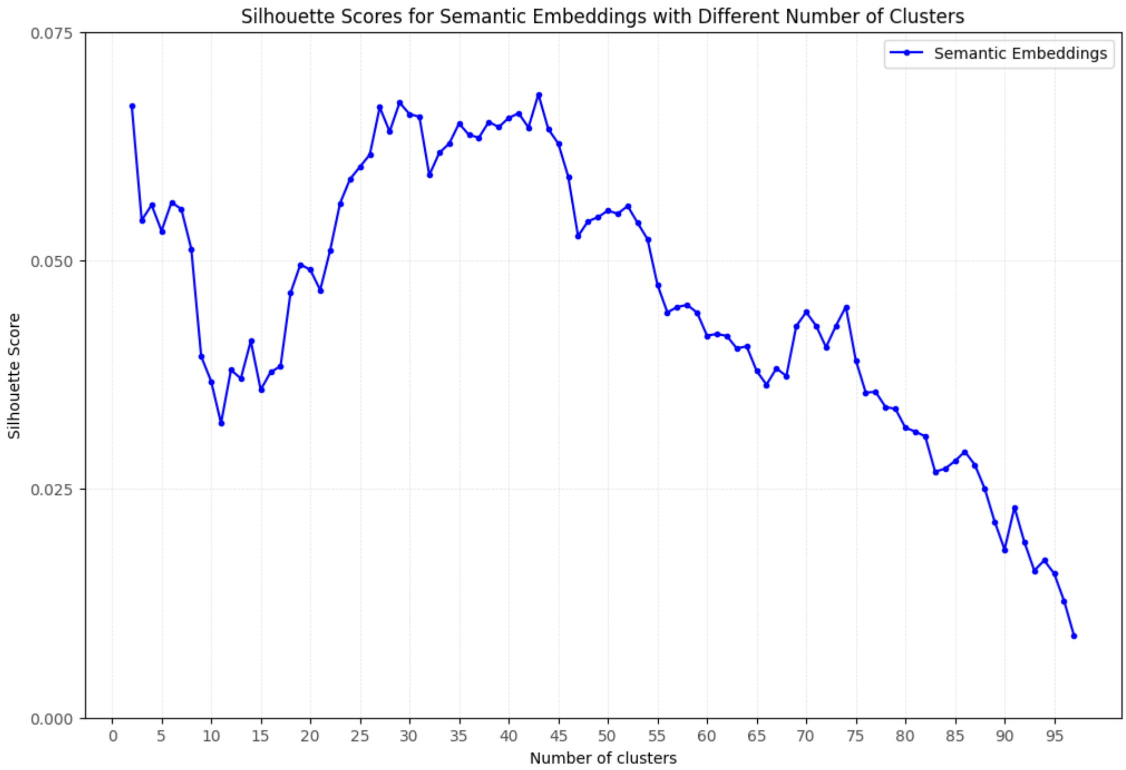

Figure 9 - Semantic Embeddings: Silhouette Score vs the number of clusters (k-means)

For semantic embeddings, the following cluster numbers yielded the highest scores:

Table 1 - Semantic Embeddings: Peak Silhouette Scores against Number of Clusters

Number of Clusters	Silhouette Score
43	0.0682
29	0.0673
27	0.0668
41	0.0661
38	0.0652

These scores highlight the embeddings' capability to delineate clusters corresponding to semantic groupings. Aggregated semantic embeddings are expected to demonstrate improved clustering due to their integration of structural and semantic information, enabling more nuanced representations.

Interpretation

The clustering performance of semantic embeddings indicates that the representation captures meaningful patterns in argumentative components, though the Silhouette Scores suggest a moderate separation of clusters. The variation in peak cluster numbers reflects the complexity of discourse elements, with larger cluster counts likely corresponding to finer-grained distinctions between roles, themes, and relationships.

Aggregated semantic embeddings, by combining structural and semantic information, are hypothesized to exhibit improved cluster cohesion and separation. Their performance will be discussed in comparison to semantic embeddings to evaluate the added value of structural integration in capturing argumentative dynamics. Further analysis and visualization insights are provided in subsequent sections.

4.3.2 Visualization Insights

As mentioned earlier, visualizing the embedding space provides critical insights into its representational capacity and the clustering patterns it produces. Here, the semantic embedding space is explored through a 3D visualization, offering a spatial representation of how these embeddings encode argumentative components and their relationships. To enhance accessibility and interactivity, an online digital version of the visualization has been made available, allowing readers to further examine the clustering patterns in detail.

These visualizations collectively reveal a nuanced perspective on the semantic embedding space, demonstrating its ability to capture argumentative dynamics from multiple angles. By leveraging the interactive version, readers can delve deeper into specific clusters and relationships [45]. The following sections outline detailed observations and their implications.

Separation of Argumentative Roles

Nodes are color-coded based on their type, distinguishing between Claims, Premises, and Questions. This perspective offers insights into how well the semantic embeddings separate argumentative components, emphasizing their relational roles and interactions. Observing these clusters helps assess the semantic embeddings' ability to preserve logical distinctions among these key argumentative roles.

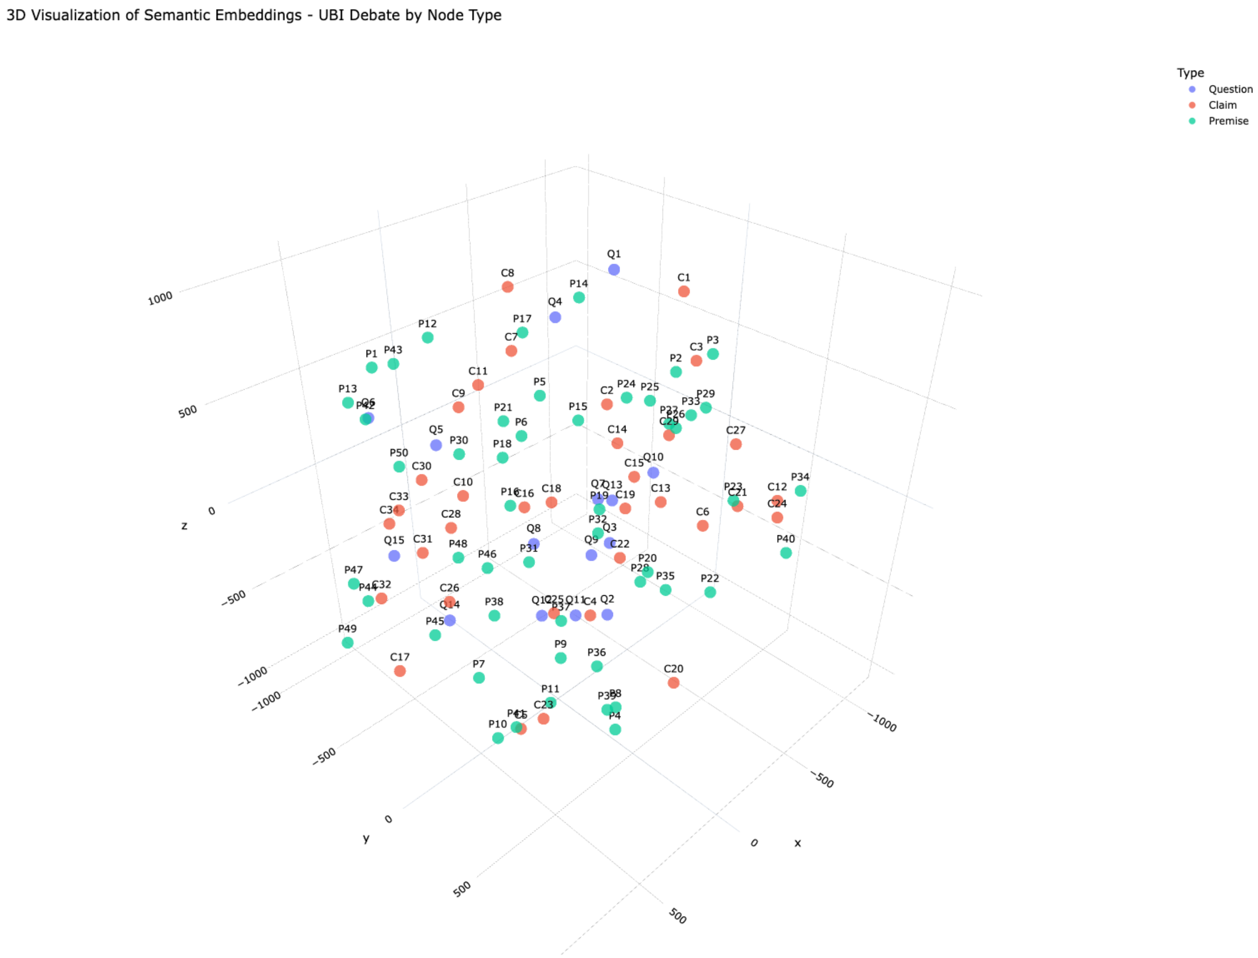

Figure 10 - Semantic Embeddings by Node Type

Thematic Sections

Nodes are color-coded according to their thematic section, such as "Introduction," "Ethical and Social Dimensions," "Macroeconomic Impacts," etc. This perspective examines whether the embeddings capture thematic groupings within the graph, reflecting logical segmentations of the discourse. Clear thematic clusters indicate the embeddings' capacity to encode contextual nuances effectively, while overlaps may reveal areas of discourse where themes are interconnected.

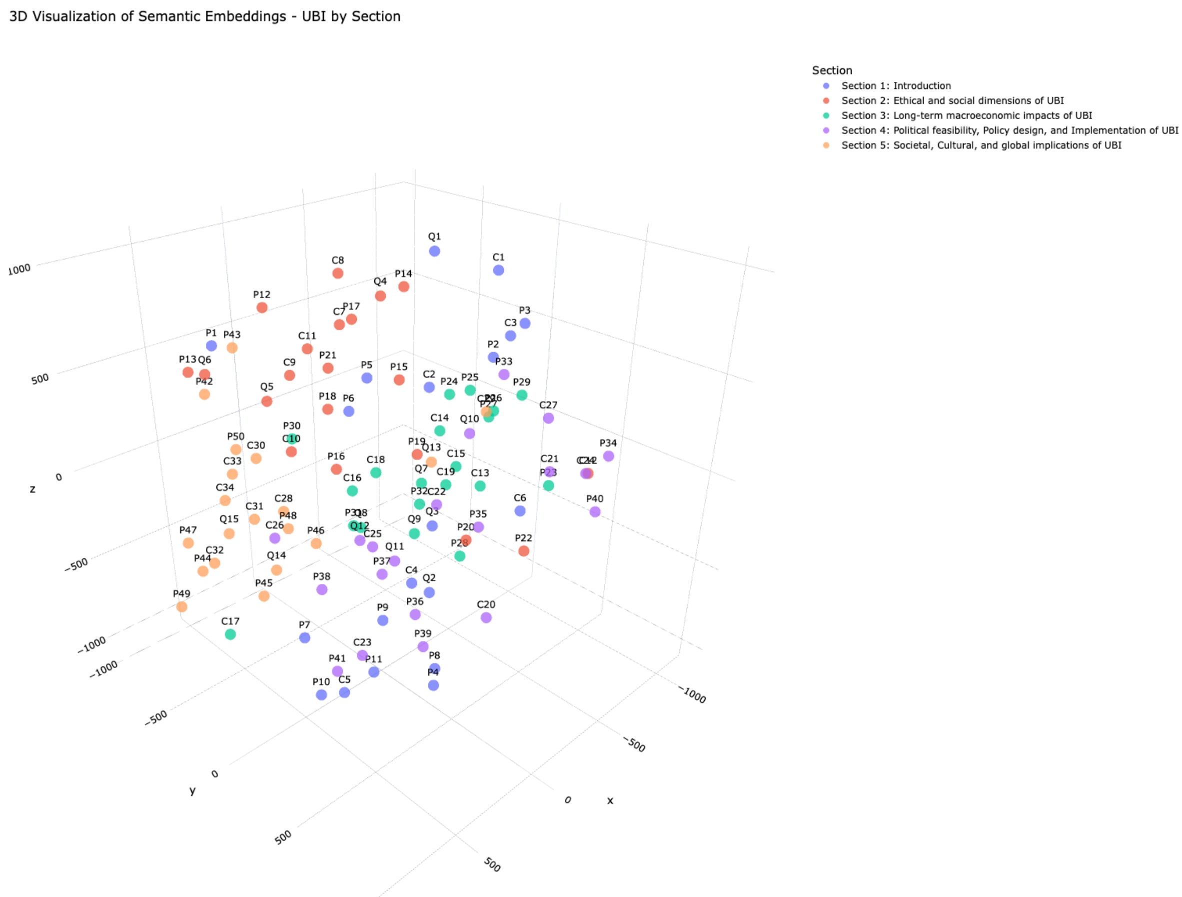

Figure 11 - Semantic Embeddings by Thematic Section

Author Attribution

Each node is color-coded based on its author. This visualization explores the impact of authorship on the semantic organization of the embeddings, identifying whether individual authors' contributions exhibit distinct patterns or overlap significantly with others. Such patterns can provide insights into the semantic consistency or diversity in the arguments presented by different authors.

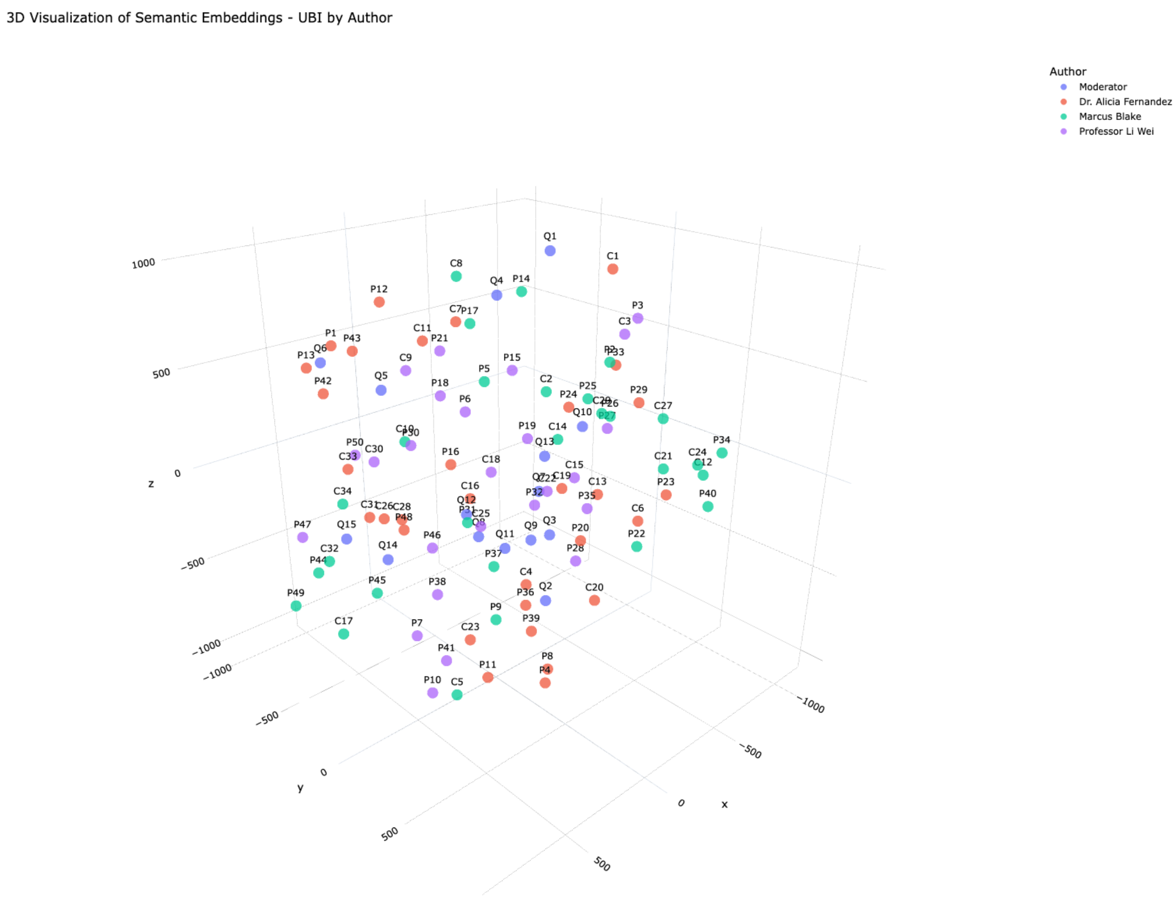

Figure 12 - Semantic Embeddings by Author

4.3.3 Interpretative Analysis

This section provides an interpretative analysis of the clustering patterns observed in the semantic embedding space. By examining the separation of node types, thematic groupings, and authorship dynamics, this analysis highlights the representational strengths and limitations of semantic embeddings in capturing argumentative discourse.

Separation of Node Types

As expected, semantic embedding space does not exhibit observable clustering based on node types such as Claims, Premises, or Questions. This lack of distinction suggests that semantic embeddings prioritize contextual and topical information over the explicit argumentative roles of nodes. While this is a limitation in differentiating roles, it aligns with the model's focus on capturing semantic relationships and content.

Thematic Groupings

Clear and distinct thematic clustering is observed in the semantic embedding space, which is expected for a semantic model as each section of the discourse handles semantically distinct topics. For example, sections such as "Ethical and Social Dimensions" and "Macroeconomic Impacts" form coherent clusters, reflecting the embeddings' ability to group arguments by topic. This indicates that semantic embeddings effectively encode contextual nuances and thematic separations within the discourse.

However, the "Introduction" section exhibits a more distributed clustering compared to other focused sections of the debate. This distribution is anticipated, as the introduction typically serves as a broad overview, addressing multiple themes and setting the stage for the subsequent focused discussions. In contrast, sections that focus on politics and economics demonstrate tighter, more localized clusters due to their concentrated focus on specific topics.

Some overlaps are evident between thematic clusters, particularly where arguments span multiple domains or involve interconnected discussions. This suggests that some arguments bridge these themes, highlighting areas where discourse transitions or integrates different perspectives. These overlaps provide valuable insights into the complexity of the arguments and the interconnected nature of real-world debates.

Authorship Dynamics

The authorship dynamics in the semantic embedding space closely resemble those observed in the structural embedding space. Distinct and separated clusters are formed by participants with opposing stances:

- **Pro (Dr. Alicia Fernandez)** and **Contra (Markus Blake)** participants exhibit clear separation, as their contributions strongly oppose each other. This pattern reflects the embeddings' ability to encode the semantic divergence inherent in their arguments.
- **Neutral Participants (Moderator and Professor Li Wei)** act as bridges between the pro and contra stances. This bridging role is expected since their contributions are more balanced, addressing both sides of the debate and often mediating discussions. Their distributed positioning highlights their role in facilitating or contextualizing the arguments of other participants.

These patterns illustrate that semantic embeddings effectively capture the semantic context and stance-specific dynamics of authors, further emphasizing their utility in representing argumentative discourse. However, the reliance on semantic content over structural roles limits their capacity to differentiate node types, suggesting the need for complementary approaches to enhance representational depth.

This analysis underscores the strengths of semantic embeddings in thematic representation and authorship-based clustering, while also identifying areas for improvement in role differentiation. These findings form a foundation for evaluating the integration of structural and semantic information in subsequent analyses.

4.4 Aggregated Semantics Embedding Analysis

4.4.1 Quantitative Performance

The quantitative analysis of aggregated semantic embeddings focuses on their clustering quality and ability to represent nuanced argumentative dynamics. Using Silhouette Scores as a metric, the study evaluates how effectively these embeddings delineate clusters of argumentative components by integrating both semantic and structural information. The combination of GraphSAGE's neighborhood aggregation of semantic embeddings enables the model to capture both content and contextual relationships, offering a holistic representation of the discourse.

Key Observations

A graph of Silhouette Scores versus the number of clusters reveals peaks indicating the most coherent cluster configurations.

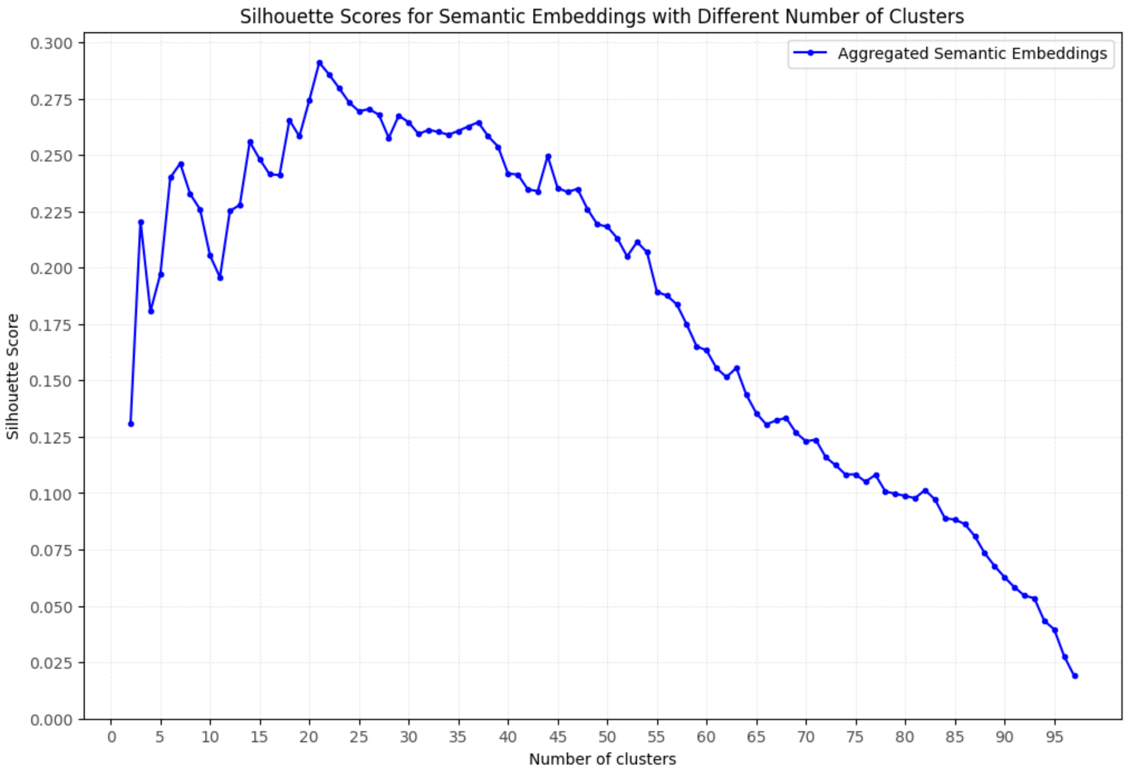

Figure 13 - Aggregated Semantic Embeddings: Silhouette Score vs the number of clusters (k-means)

A plot of Silhouette Scores against the number of clusters reveals notable peaks, indicating the most coherent cluster configurations. For aggregated semantic embeddings, the following cluster numbers yielded the highest scores:

Table 2 - Aggregated Semantic Embeddings: Peak Silhouette Scores against Number of Clusters

Number of Clusters	Silhouette Score
21	0.2910
26	0.2704
29	0.2674
18	0.2654
37	0.2645

These scores surpass those observed for semantic embeddings alone, reflecting the added value of integrating structural information. The improvement in cluster cohesion and separation underscores the ability of aggregated embeddings to model the interplay between argumentative roles, thematic contexts, and structural relationships within the discourse.

The analysis also highlights the dynamic range of cluster sizes, with larger clusters likely representing broader thematic segments, while smaller clusters capture fine-grained distinctions between argumentative components. This versatility demonstrates the efficacy of aggregated embeddings in addressing the diverse representational requirements of argumentative discourse.

Interpretation

The significant improvement in Silhouette Scores for aggregated semantic embeddings compared to purely semantic embeddings can be attributed to the integration of structural information alongside semantic content. By leveraging GraphSAGE's neighborhood aggregation, these embeddings capture not only the intrinsic properties of individual nodes but also the contextual relationships defined by the argumentation graph's topology. This dual-layer representation enables the model to account for both the semantic nuances and the structural dynamics inherent in argumentative discourse.

Several factors contribute to these results. First, the structural information enhances the embeddings' ability to model inter-node relationships, such as support, attack, and inquiry-response dynamics. This added layer of context ensures that nodes with similar roles or connectivity patterns converge in the embedding space, leading to more coherent clusters. Second, the incorporation of structural attributes allows the embeddings to distinguish between nodes that may share similar semantic features but differ in their relational roles within the graph. For example, a Premise and a Claim with comparable semantic content may occupy distinct positions in the aggregated embedding space due to their differing connections and argumentative functions.

However, the reliance on structural integration also introduces potential limitations. The improved clustering quality observed in aggregated embeddings could reflect a **bias towards structural cohesion**, potentially overshadowing finer-grained semantic distinctions. This trade-off underscores the need for careful balancing between semantic and structural features to ensure comprehensive representation. Additionally, the sensitivity of aggregated embeddings to graph topology may lead to variability in clustering performance across datasets with differing structural characteristics. These factors highlight the importance of dataset-specific parameter tuning and further validation to generalize these findings across diverse argumentative contexts.

4.4.2 Visualization Insights

The visualization of aggregated semantic embeddings provides an intuitive understanding of how these representations encode argumentative components and their relationships. By projecting the high-dimensional embedding space into a three-dimensional scatterplot using t-SNE, key patterns and clustering tendencies are revealed.

Again, the online versions of these visualization are available in an online repository for dynamic exploration of the embedding space [46].

The visualization explores three perspectives: separation of argumentative roles, thematic sections, and author attribution.

Separation of Argumentative Roles

Nodes are color-coded based on their argumentative type, distinguishing between Claims, Premises, and Questions. This perspective offers insights into how well the aggregated semantic embeddings separate argumentative components, emphasizing their relational roles and interactions. Observing these clusters helps assess the semantic embeddings' ability to preserve logical distinctions among these key argumentative roles.

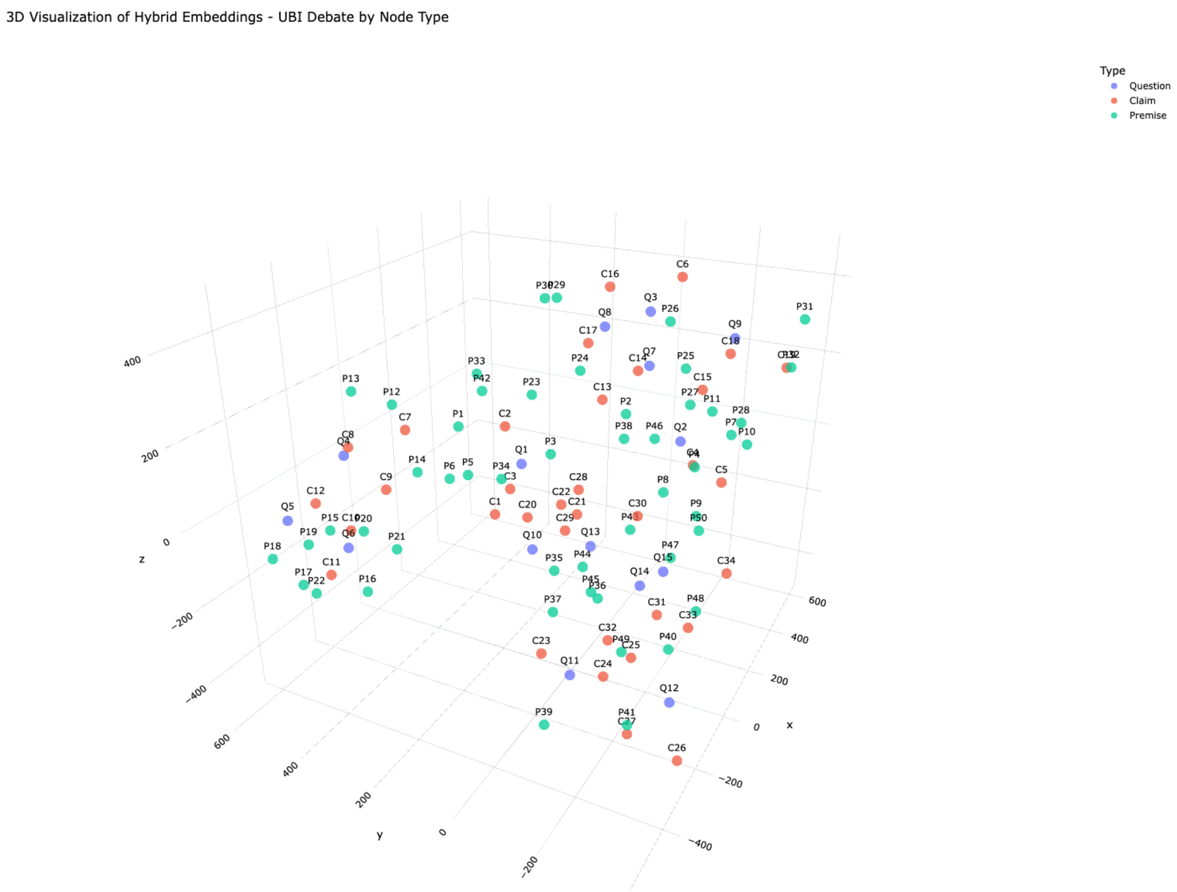

Figure 14 - Aggregated Semantic Embeddings by Argumentative Role

Thematic Sections

Nodes are color-coded according to their thematic section, such as "Introduction," "Ethical and Social Dimensions," "Macroeconomic Impacts," etc. This perspective examines whether the embeddings capture thematic groupings within the graph, reflecting semantic segmentations of the discourse. Clear thematic clusters indicate the embeddings’ capacity to encode contextual nuances effectively, while overlaps may reveal areas of discourse where themes are interconnected.

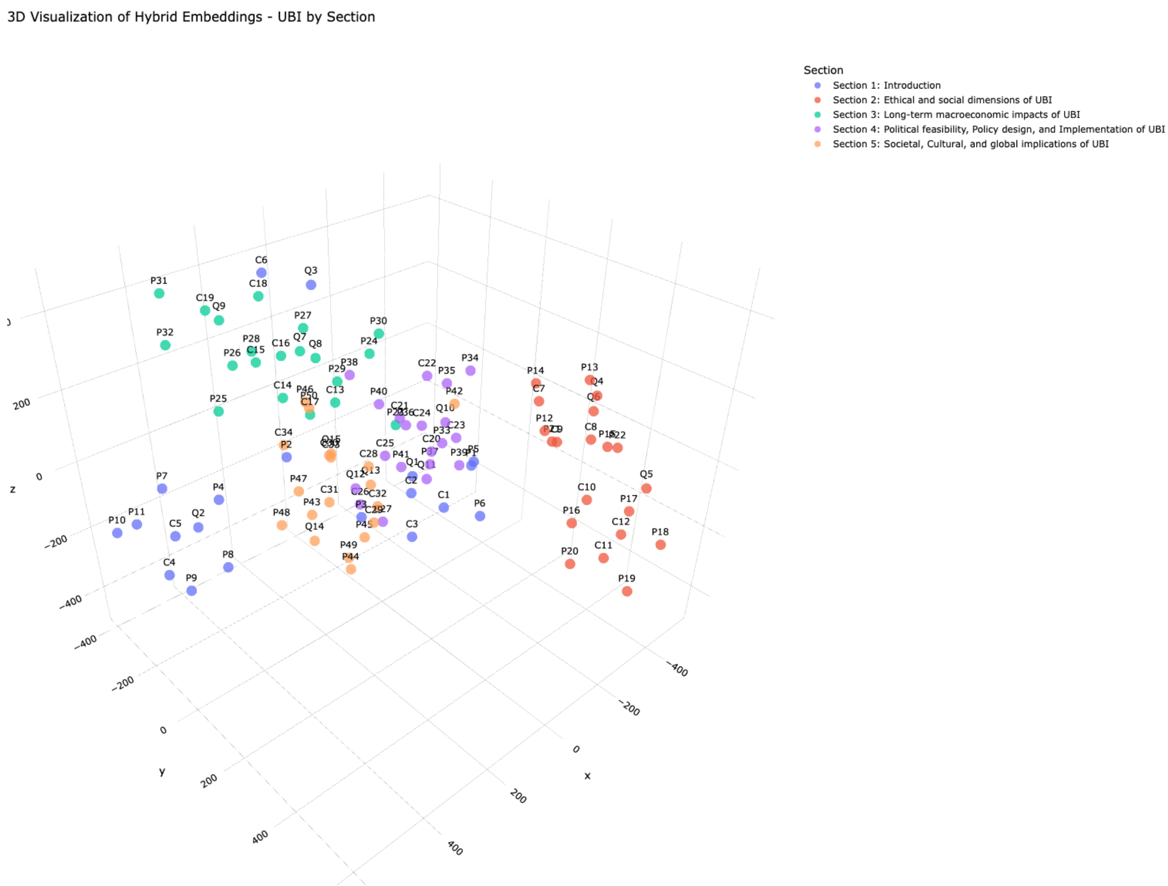

Figure 15 - Aggregated Semantic Embeddings by Thematic Sections

Author Attribution

Each node is color-coded based on its author. This visualization explores the impact of authorship on the semantic organization of the embeddings, identifying whether individual authors' contributions exhibit distinct patterns or overlap significantly with others. Such patterns can provide insights into the semantic consistency or diversity in the arguments presented by different authors.

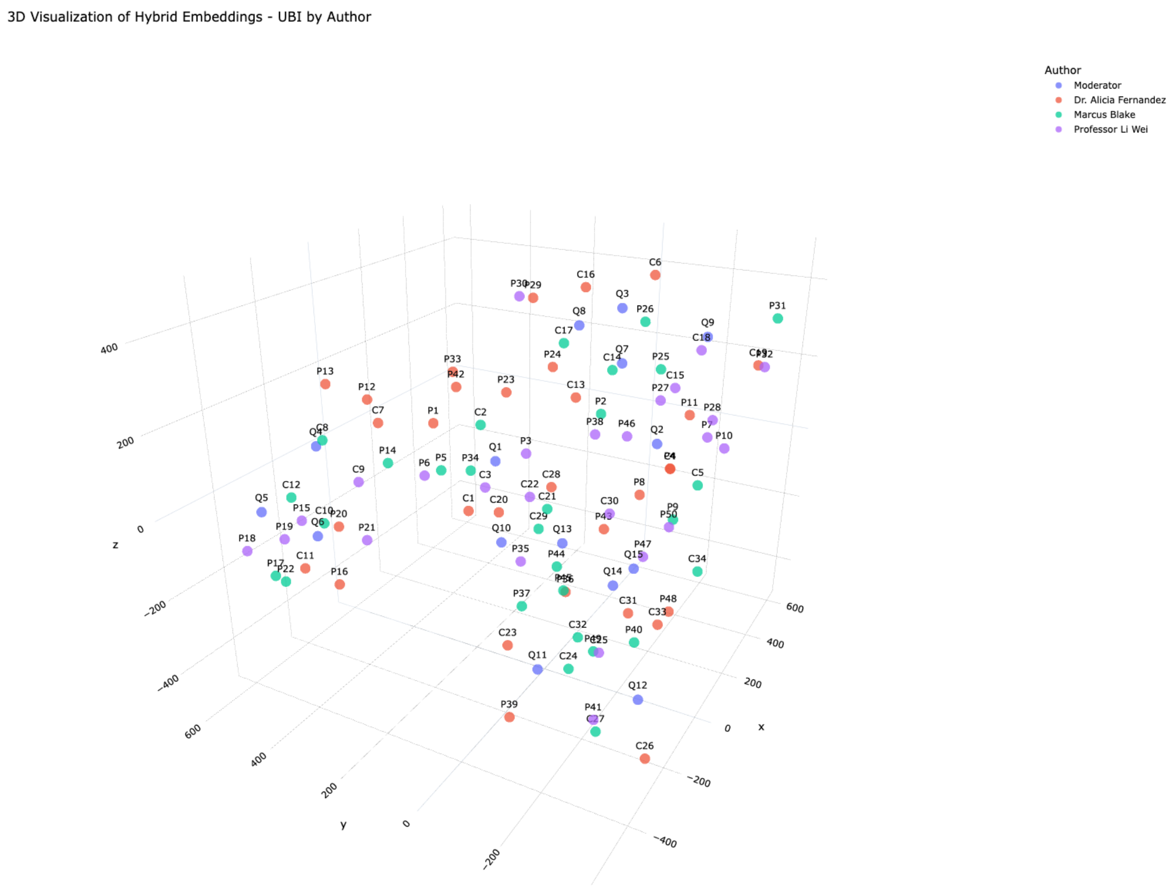

Figure 16 - Aggregated Semantic Embeddings by Author

4.4.3 Interpretative Analysis

Aggregated semantic embeddings offer a comprehensive representation of argumentative discourse by combining the strengths of semantic and structural approaches. This section interprets the clustering patterns observed in the visualization and evaluates the embeddings' capacity to capture argumentative roles, thematic coherence, and authorship dynamics.

Separation of Node Types

Aggregated embeddings do not demonstrate superior differentiation between argumentative roles compared to semantic embeddings. However, they excel in maintaining logical proximity to related elements, such as Claims, Premises, and Questions. Questions cluster around their associated Claims, effectively capturing inquiry-response relationships. This ability to preserve logical connections affirms the model's capacity to represent both content and context, potentially addressing a key limitation of purely semantic approaches.

Thematic Groupings

Thematic clusters are better defined compared to those generated by semantic embeddings, with minimal overlap between sections. Aggregated embeddings preserve the semantic richness of arguments while integrating structural relationships, resulting in coherent and contextually accurate thematic representations. Overlaps observed between interdisciplinary sections reflect genuine discourse complexities, such as arguments spanning multiple domains. These patterns highlight the embeddings' superior capacity to model nuanced argumentative dynamics while maintaining thematic clarity.

Authorship Dynamics

The embeddings reveal diminished authorial influences compared to both semantic-only and structure-only embeddings. This reduced clarity may be attributed to the omission of direct authorship relationships and authors from the embedding space, a limitation shared with semantic-only embeddings. However, the aggregation process appears to dilute the isolated semantic individuality of each node with the features of its neighbors, thereby obscuring clear links to their respective authors. While this trade-off results in improved thematic grouping through aggregation, it comes at the expense of semantic specificity and identifiable authorship. The dispersed positioning of neutral participants still emphasizes their bridging role, but the reduced clarity in authorial clusters suggests a loss of semantic granularity in favor of relational and thematic cohesion.

In summary, aggregated semantic embeddings offer a nuanced framework for modeling argumentative discourse by integrating semantic and structural information. While they excel in capturing thematic coherence and logical proximity among related elements, they face challenges in preserving semantic individuality and authorial clarity. These insights underscore their potential utility for advanced discourse analysis, albeit with limitations that highlight the trade-offs inherent in aggregating diverse informational layers.

4.5 Comparative Evaluation of Semantic Embedding Spaces

This section provides a comparative evaluation of **semantic** and **aggregated semantic embedding spaces**, focusing on their quantitative clustering results, including the number of clusters and Silhouette Scores. By analyzing these metrics, the study assesses their respective capacities to represent argumentative discourse effectively.

4.5.1 Quantitative Analysis

Quantitative evaluation of the embedding spaces focuses on their clustering quality, as measured by Silhouette Scores, across varying cluster configurations. This approach ensures an objective comparison of their capacity to delineate distinct argumentative components.

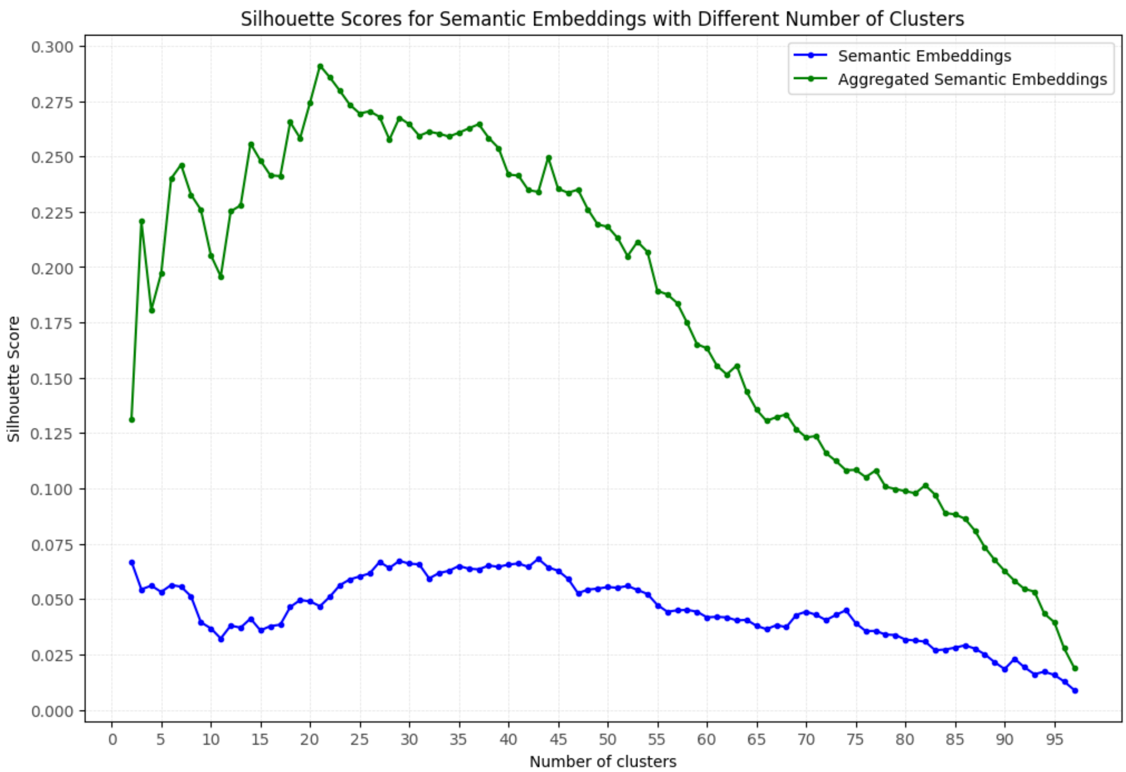

Figure 17 - Silhouette Score per Cluster: Semantic Embeddings vs. Aggregated Semantic Embeddings

Semantic Embeddings

The clustering performance of semantic embeddings reveals moderate Silhouette Scores, with the highest peak observed at 0.0682 for 43 clusters. These results indicate that semantic embeddings effectively group nodes based on contextual and topical similarities but lack precision in delineating argumentative roles and relationships. The variability in peak cluster counts reflects the embeddings' sensitivity to the complex interplay of semantic nuances within the discourse. While these embeddings excel in capturing thematic distinctions, their inability to incorporate structural features limits their capacity to model inter-node relationships comprehensively.

Aggregated Semantic Embeddings

Aggregated semantic embeddings demonstrate the highest clustering quality among the two spaces, achieving a peak Silhouette Score of 0.291 for 21 clusters. By integrating semantic and structural information through GraphSAGE's neighborhood aggregation, these embeddings provide a holistic representation of the discourse. This integration enhances cluster cohesion and separation, allowing for a nuanced depiction of both argumentative roles and thematic contexts. The improved clustering performance underscores the added value of combining content and context, although this approach may occasionally obscure fine-grained semantic distinctions in favor of structural coherence.

4.5.2 Comparative Interpretation

The comparative evaluation highlights the trade-offs inherent in semantic and aggregated semantic embedding spaces. Semantic embeddings excel in capturing contextual and thematic nuances, making them ideal for analyses focused on semantic content. However, their inability to represent structural relationships limits their applicability for tasks requiring a detailed understanding of argumentation dynamics.

Aggregated semantic embeddings offer a more comprehensive representation by combining the strengths of semantic and structural dimensions. Their superior clustering performance and balanced representation of content and context make them well-suited for complex discourse analyses. Nonetheless, the integration of structural features may occasionally obscure fine-grained semantic distinctions, necessitating careful consideration of specific analytical goals.

This comparative analysis underscores the importance of selecting embedding spaces based on the research objectives and the specific aspects of argumentative discourse being studied. By leveraging the unique strengths of each approach, researchers can tailor their methodologies to achieve a deeper and more nuanced understanding of argumentative interactions.

4.6 Challenges and Limitations

4.6.1 Bias in Synthetic Datasets

One significant challenge was the inherent bias within the synthetic dataset used for evaluation. Although the dataset was designed to emulate real-world argumentative discourse, its artificial nature may not fully capture the complexity and diversity of genuine debates. For instance, it fails to account for cultural variability in argumentation styles, which can significantly influence the framing and progression of discourse. Additionally, the dataset does not adequately reflect incomplete arguments or evolving debates, both of which are common in real-world contexts. These aspects are critical for modeling authentic argumentative dynamics and highlight the need for more representative datasets in future analyses. These biases could potentially influence the clustering results, leading to findings that are less representative of broader, more nuanced argumentative contexts. Furthermore, the dataset's structure and content may have inadvertently favored certain embedding methods over others, introducing skewed performance assessments. This bias underscores the need for datasets that better simulate the diversity and unpredictability of real-world discourse, encompassing a wider range of argumentative styles, themes, and contexts.

4.6.2 Complexity of Real-World Data Acquisition

Real-world data is not neatly structured as the synthetic dataset used in this analysis. Obtaining data from real-world discourse would involve the use of argument mining techniques, which introduce additional layers of complexity. These processes, including preprocessing and graph construction, are far more intricate in real-world scenarios, often requiring extensive manual validation to ensure accuracy. The synthetic dataset's simplified structure fails to reflect the inherent messiness of authentic argumentative discourse, further limiting the applicability of findings.

4.6.3 Limitations of Dataset Size and Content

The small dataset size was another critical limitation, as it consisted of only 100 nodes with very limited content, often reduced to a few sentences per node. This restriction not only limits the generalizability of findings but also reduces the ability of embeddings to capture more complex relationships that emerge in larger and more diverse datasets. A larger dataset with richer content would likely reveal deeper insights into the strengths and weaknesses of different embedding approaches.

4.6.4 Connectivity Challenges in Synthetic Graphs

Additionally, the connectivity of the synthetic data graph posed challenges when compared to real-world argumentative graphs. Real-world discourse often exhibits highly complex and irregular connectivity patterns, such as multi-layered interdependencies and non-linear argumentative structures. The simplified connectivity in the synthetic dataset may fail to represent these complexities, potentially limiting the applicability of findings to more intricate real-world scenarios. Future studies should prioritize the use of datasets that more closely mirror the nuanced connectivity patterns observed in real-world argumentation.

4.6.5 Inadequacies in Graph Schema Design

The graph schema itself was another source of limitation, as it did not model all the complexities inherent in real-world debates and argumentative discourse. For example, the schema may have overlooked nuanced relationships such as temporal aspects, rhetorical strategies, indirect rebuttals, emotional appeals, or audience-specific framing. These types of dynamics play a critical role in real-world argumentation, and their absence restricts the ability of embeddings to fully capture the multi-dimensional nature of debates. Future schema designs should incorporate such relational types to better reflect the intricacies of authentic argumentative discourse.

4.6.6 Hyperparameter Sensitivities

Another challenge involved the hyperparameters of the stochastic methods employed, particularly GraphSAGE and FastRN. These methods rely on various hyperparameters, such as neighborhood sampling sizes and learning rates, which can significantly impact the quality of the resulting embeddings. The process of tuning these hyperparameters is computationally intensive and requires careful consideration to avoid suboptimal configurations that might distort the analysis. For instance, overly small neighborhood sampling sizes in GraphSAGE can result in a lack of contextual information, reducing the model's ability to capture meaningful relationships. Conversely, excessively large sampling sizes might introduce noise by incorporating irrelevant nodes. Similarly, poorly chosen learning rates can lead to slow convergence or unstable training processes, further impacting the quality of the embeddings. These examples highlight the importance of systematic parameter tuning to ensure robust and accurate analysis outcomes. Establishing robust and automated hyperparameter optimization techniques could mitigate this issue and enhance the reliability of these methods.

4.6.7 Limitations in Semantic Embedding Model Choice

The choice of semantic embedding model also introduced limitations. Different semantic models have varying strengths and weaknesses, and the selection of a particular model inherently biases the results. For instance, models prioritizing contextual embedding may underperform in representing explicit argumentative roles, while others may struggle with capturing subtle semantic nuances. A more comprehensive evaluation of multiple semantic embedding models could provide a clearer understanding of their comparative strengths and limitations in the context of argumentative discourse analysis.

4.6.8 Directionality Assumptions in Aggregation Methods

Lastly, the directionality of relationships during aggregation in methods like FastRN and GraphSAGE posed a significant challenge. These methods rely on specific assumptions about the directional flow of information, which may not always align with the underlying dynamics of argumentative discourse. For example, reciprocal rebuttals, where two opposing arguments engage in back-and-forth responses, challenge the assumption of unidirectional flow and require models to capture bidirectional dynamics. Similarly, arguments that incorporate layered inquiries—where a question is followed by multiple sub-questions and responses—may not be adequately represented by methods assuming a simple flow of information. These examples highlight the importance of refining aggregation and discourse modelling techniques to better capture the complex, multi-directional nature of real-world argumentative interactions.

4.6.9 Preliminary Nature of Evaluation

This analysis represents a very preliminary evaluation of the embeddings' quality. Much more work is required to properly assess the aggregated embeddings, particularly in terms of their performance in downstream machine learning tasks.

Evaluating how these embeddings perform in real-world applications, such as automated argument classification or stance detection, will provide a more comprehensive understanding of their utility and limitations. Future research must focus on these aspects to establish the robustness and practicality of aggregated embeddings.

4.6.10 Conclusion

Collectively, these challenges reveal critical areas for improvement in dataset development, visualization methods, embedding design, and computational strategies. Addressing these limitations will be essential to advance the robustness and applicability of embedding-based methodologies in the field of argumentative discourse analysis.

4.7 Implications and Future Directions

The findings presented in this chapter offer significant insights into the representational capacities of embedding spaces in capturing the nuanced dynamics of argumentative discourse. These insights underscore the potential for further advancements in the field of argumentation analysis, particularly in the development of embedding methodologies that integrate semantic and structural dimensions effectively.

One of the key implications of this study lies in the demonstrated benefits of aggregated semantic embeddings. By combining the semantic richness of argument content with the structural relationships encoded within argumentation graphs, these embeddings provide a more holistic representation of discourse. This dual-layered approach enhances clustering coherence and captures interconnections between argumentative components, offering new opportunities for improving discourse modeling frameworks. Future research should build on these findings by exploring additional methods for aggregating diverse data dimensions, such as temporal dynamics or rhetorical strategies, to enrich the representational depth of embedding spaces.

The observed limitations of synthetic datasets highlight the need for more representative and diverse data sources. Synthetic datasets, while useful for controlled evaluations, fail to capture the variability and complexity of real-world argumentative interactions. Future research should prioritize the collection and utilization of real-world datasets, encompassing diverse cultural, thematic, and linguistic contexts. Such datasets will enable a more robust evaluation of embedding methodologies and ensure their applicability to a wider range of discourse scenarios.

Another promising direction involves refining the aggregation techniques used in embedding methodologies. The observed trade-offs between structural and semantic specificity suggest the need for approaches that dynamically balance these dimensions based on the discourse context. Techniques that incorporate adaptive weighting schemes or context-aware aggregation mechanisms could address this challenge, ensuring that embeddings are tailored to the specific characteristics of the discourse being modeled.

Moreover, the integration of advanced machine learning techniques, such as transformer-based models and graph neural networks, offers significant potential for improving embedding quality. These approaches could capture bidirectional and hierarchical relationships within argumentative graphs more effectively, addressing limitations in existing aggregation methods. Researchers should also explore the application of these embeddings in downstream tasks, such as automated argument classification, stance detection, or summarization, to assess their utility in practical applications.

Finally, the findings emphasize the importance of interdisciplinary collaboration in advancing argumentation analysis. Combining insights from fields such as computational linguistics, cognitive science, and rhetoric can lead to the development of richer models that reflect the multifaceted nature of argumentative discourse. Such collaborations could also inform the design of visualization tools that facilitate intuitive exploration of embedding spaces, bridging the gap between technical methodologies and user-friendly interpretability.

4.8 Summary

This chapter has provided a detailed examination of the results obtained through the structural, semantic, and aggregated embedding spaces, shedding light on their strengths, limitations, and potential for modeling argumentative discourse. The structural embeddings excelled in capturing topological relationships, highlighting the logical architecture of argumentation graphs and thematic coherence. Semantic embeddings demonstrated robust thematic clustering but faced challenges in distinguishing argumentative roles. Aggregated semantic embeddings emerged as the most effective representation, combining structural and semantic dimensions to deliver enhanced clustering quality and nuanced modeling of argumentative dynamics.

The comparative evaluation revealed critical trade-offs between these approaches, emphasizing the importance of aligning embedding methodologies with specific analytical objectives. While aggregated embeddings excel in comprehensive discourse modeling, their reliance on structural integration may obscure finer-grained semantic details. The study also identified key challenges, including the limitations of synthetic datasets, the complexity of real-world discourse acquisition, and the sensitivity of embedding methods to hyperparameters and graph schemas.

The chapter concluded with a discussion of the implications of these findings and proposed future directions for advancing embedding methodologies. By addressing current limitations and exploring novel techniques, researchers can refine the tools and approaches used in argumentative discourse analysis, contributing to a deeper understanding of how arguments are structured, interconnected, and contextualized. These advancements hold promise for both theoretical exploration and practical applications, paving the way for more sophisticated and adaptable frameworks in the study of argumentation.

5 Conclusion

This chapter synthesizes the findings, addressing the research question and reflecting on the theoretical and practical aspects of the study. Furthermore, it critically examines the limitations and suggests avenues for future exploration.

5.1 Summary of Research

This study set out to explore how argumentative discourse can be effectively modeled as an argumentation graph integrating both structural and semantic elements. The research aimed to address the limitations in traditional argumentation models by incorporating graph embeddings and advanced textual embeddings into a more comprehensive argumentation graph framework. By combining machine learning techniques for representational learning, including dimensionality reduction and clustering, with theoretical insights from argumentation theory, the study illustrated how such a framework can capture nuanced argumentative roles and relationships. The findings indicate that integrating semantic and structural embeddings enhances the ability to analyze and represent discourse, offering a more detailed approach to discourse analysis.

5.2 Contributions

The study contributes to the fields of argumentation theory and discourse analysis by offering a framework that brings together traditional argumentation graphs with semantic embeddings to better represent argumentative interactions. Integrating large language models, such as Sentence Transformers, into the framework has added depth to the semantic aspects of the graphs. The research provides a basis for developing computational tools that can evaluate and summarize complex debates, which could prove useful in a variety of contexts. While these contributions are promising, they build on and align with existing advancements, extending the work of prior studies to address specific gaps.

5.3 Future Work

Several directions for future research arise from this study. Automating the construction of argumentation graphs from real-world discourse data, such as social media or online forums, through Argument Mining techniques represents one promising area. Such automation would enable the framework to function as a foundational data model for computational systems. Another area of interest lies in adapting the framework for Retrieval-Augmented Generation (RAG) and conversational systems, potentially enabling these systems to produce more contextually relevant and coherent responses. The model could also be refined to support semantic search engines, allowing for more efficient retrieval of specific arguments or perspectives. Future research might also expand the framework to include multimodal data, such as visual or auditory inputs, and explore how arguments evolve over time, adding a temporal dimension to the analysis.

5.4 Broader Computational Applications

The flexibility and adaptability of the framework suggest its potential for broader computational applications. The model's representation of argument structures and semantics could support tools for summarizing and evaluating discourse, providing concise overviews of complex debates. It might inform conversational systems designed to offer meaningful and informed interactions or enhance semantic search engines for locating relevant arguments more efficiently. In decision-making contexts, the framework could aid stakeholders by summarizing argument clusters and offering evidence-based recommendations, which could streamline processes in areas like policy-making and legal reasoning. Additionally, integrating the model into public communication platforms could assist in moderating discourse by detecting, clustering, and retrieving arguments, promoting balanced engagement with complex issues.

5.5 Limitations

This study's findings should be considered in light of several limitations. The reliance on synthetic datasets may not fully reflect the complexities of natural discourse, potentially limiting the generalizability of the results. Moreover, the integration of advanced embeddings and graph structures entails significant computational demands, which could pose challenges for scalability, particularly in real-time applications. Additionally, the framework's performance across varied argumentative environments has not yet been thoroughly validated, indicating the need for further exploration in diverse contexts.

5.6 Final Reflections

This study aimed to explore ways to integrate structural and semantic insights into modeling argumentative discourse, building on established methodologies and addressing both theoretical and practical considerations. The findings highlight the potential for enriched argumentation graphs to support discourse analysis, debate evaluation, and decision-support systems. Future work will focus on refining and expanding the framework, with the hope of contributing to a deeper understanding of argumentative interactions and their applications across various fields.

Bibliography

[1] S. Boulianne, “Social media use and participation: a meta-analysis of current research,” *Information, Communication & Society*, vol. 18, no. 5, pp. 524–538, May 2015, doi: 10.1080/1369118X.2015.1008542.

[2] L. Rainie, S. Verba, H. Brady, A. Smith, and K. L. Schlozman, “Social Media and Political Engagement,” Pew Research Center. Accessed: Oct. 15, 2024. [Online]. Available: https://www.pewresearch.org/internet/2012/10/19/social-media-and-political-engagement/

[3] J. Habermas, *The Structural transformation of the public sphere: an inquiry into a category of bourgeois society*, 10. print. in Studies in contemporary German social thought. Cambridge, Mass: MIT Press, 1999.

[4] E. Bakshy, S. Messing, and L. A. Adamic, “Exposure to ideologically diverse news and opinion on Facebook,” *Science*, vol. 348, no. 6239, pp. 1130–1132, Jun. 2015, doi: 10.1126/science.aaa1160.

[5] E. Pariser, *The filter bubble: what the Internet is hiding from you*. London: Penguin books, 2012.

[6] B. Ca *et al.*, “Exposure to opposing views on social media can increase political polarization,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 115, no. 37, Nov. 2018, doi: 10.1073/pnas.1804840115.

[7] C. R. Sunstein, *#Republic: divided democracy in the age of social media*. Princeton: Princeton university press, 2017.

[8] Z. Tufekci, *Twitter and tear gas: the power and fragility of networked protest*. New Haven London: Yale University Press, 2017.

[9] S. E. Toulmin, *The Uses of Argument*, 2nd ed. Cambridge: Cambridge University Press, 2003. doi: 10.1017/CBO9780511840005.

[10] F. H. van Eemeren and R. Grootendorst, *A Systematic Theory of Argumentation: The pragma-dialectical approach*. Cambridge: Cambridge University Press, 2003. doi: 10.1017/CBO9780511616389.

[11] M. Xu, “Understanding graph embedding methods and their applications,” 2020, *arXiv*. doi: 10.48550/ARXIV.2012.08019.

[12] L. Malmqvist, T. Yuan, and S. Manandhar, “Visualising Argumentation Graphs with Graph Embeddings and t-SNE,” 2021, doi: 10.48550/ARXIV.2107.00528.

[13] H. Cao, “Recent advances in text embedding: A Comprehensive Review of Top-Performing Methods on the MTEB Benchmark,” 2024, *arXiv*. doi: 10.48550/ARXIV.2406.01607.

[14] J. Yang *et al.*, “GraphFormers: GNN-nested Transformers for Representation Learning on Textual Graph,” 2021, *arXiv*. doi: 10.48550/ARXIV.2105.02605.

[15] S. Marro, E. Cabrio, and S. Villata, “Graph Embeddings for Argumentation Quality Assessment,” in *Findings of the Association for Computational Linguistics: EMNLP 2022*, Abu Dhabi, United Arab Emirates: Association for Computational Linguistics, 2022, pp. 4154–4164. doi: 10.18653/v1/2022.findings-emnlp.306.

[16] Aristotle and G. A. Kennedy, *On rhetoric: a theory of civic discourse*, 2nd ed. New York: Oxford University Press, 2007.

[17] D. Walton, *Argumentation Schemes for Presumptive Reasoning*, 0 ed. Routledge, 2013. doi: 10.4324/9780203811160.

[18] R. M. Palau and M.-F. Moens, “Argumentation mining: the detection, classification and structure of arguments in text,” in *Proceedings of the 12th International Conference on Artificial Intelligence and Law*, Barcelona Spain: ACM, Jun. 2009, pp. 98–107. doi: 10.1145/1568234.1568246.

[19] J. Cabessa, H. Hernault, and U. Mushtaq, “In-Context Learning and Fine-Tuning GPT for Argument Mining,” 2024, *arXiv*. doi: 10.48550/ARXIV.2406.06699.

[20] S. Dutta, J. Juneja, D. Das, and T. Chakraborty, “Can Unsupervised Knowledge Transfer from Social Discussions Help Argument Mining?,” 2022, *arXiv*. doi: 10.48550/ARXIV.2203.12881.

[21] R. K. Singh, "Advancements in Natural language Processing: An In-depth Review of Language Transformer Models," *IJRASET*, vol. 12, no. 6, pp. 1719–1732, Jun. 2024, doi: 10.22214/ijraset.2024.63408.
[22] P. M. Dung, "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games," *Artificial Intelligence*, vol. 77, no. 2, pp. 321–357, Sep. 1995, doi: 10.1016/0004-3702(94)00041-X.
[23] G. Simari and I. Rahwan, *Argumentation in Artificial Intelligence*. 2009. doi: 10.1007/978-0-387-98197-0.
[24] C. Cayrol and M. C. Lagasquie-Schiex, "On the Acceptability of Arguments in Bipolar Argumentation Frameworks," in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, vol. 3571, L. Godo, Ed., in Lecture Notes in Computer Science, vol. 3571. , Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 378–389. doi: 10.1007/11518655_33.
[25] T. Bench-Capon, "Value Based Argumentation Frameworks," Jul. 15, 2002, *arXiv*: arXiv:cs/0207059. doi: 10.48550/arXiv.cs/0207059.
[26] M. W. A. Caminada and Y. Wu, "On the Limitations of Abstract Argumentation," 2011. Accessed: Dec. 23, 2024. [Online]. Available: https://www.semanticscholar.org/paper/On-the-Limitations-of-Abstract-Argumentation-Caminada-Wu/e2dbc6a37124c6a7597d768d60ea0104fa27699b
[27] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge, "Weighted argument systems: Basic definitions, algorithms, and complexity results," *Artificial Intelligence*, vol. 175, no. 2, pp. 457–486, Feb. 2011, doi: 10.1016/j.artint.2010.09.005.
[28] A. Hunter, "A probabilistic approach to modelling uncertain logical arguments," *International Journal of Approximate Reasoning*, vol. 54, no. 1, pp. 47–81, Jan. 2013, doi: 10.1016/j.ijar.2012.08.003.
[29] M. L. Cobo, D. C. Martínez, and G. R. Simari, "On Admissibility in Timed Abstract Argumentation Frameworks," in *ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings*, H. Coelho, R. Studer, and M. J. Wooldridge, Eds., in Frontiers in Artificial Intelligence and Applications, vol. 215. IOS Press, 2010, pp. 1007–1008. doi: 10.3233/978-1-60750-606-5-1007.
[30] N. Oren and T. J. Norman, "Semantics for Evidence-Based Argumentation," in *Computational Models of Argument : Proceedings of COMMA 2008*, IOS Press, 2008, pp. 276–284. Accessed: Dec. 09, 2024. [Online]. Available: https://abdn.elsevierpure.com/en/publications/semantics-for-evidence-based-argumentation
[31] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran, and T. Solorio, Eds., Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186. doi: 10.18653/v1/N19-1423.
[32] G. Zhang, P. Nulty, and D. Lillis, "Enhancing Legal Argument Mining with Domain Pre-training and Neural Networks," 2022, *arXiv*. doi: 10.48550/ARXIV.2202.13457.
[33] H. Karande, R. Walambe, V. Benjamin, K. Kotecha, and T. Raghu, "Stance detection with BERT embeddings for credibility analysis of information on social media," *PeerJ Computer Science*, vol. 7, p. e467, Apr. 2021, doi: 10.7717/peerj-cs.467.
[34] S. Alaparthi and M. Mishra, "Bidirectional Encoder Representations from Transformers (BERT): A sentiment analysis odyssey," 2020, *arXiv*. doi: 10.48550/ARXIV.2007.01127.
[35] Y. Wu, Z. Jin, C. Shi, P. Liang, and T. Zhan, "Research on the Application of Deep Learning-based BERT Model in Sentiment Analysis," 2024, *arXiv*. doi: 10.48550/ARXIV.2403.08217.
[36] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A Comprehensive Survey on Graph Neural Networks," 2019, doi: 10.48550/ARXIV.1901.00596.

[37]F. Liang, C. Qian, W. Yu, D. Griffith, and N. Golmie, “Survey of Graph Neural Networks and Applications,” *Wireless Communications and Mobile Computing*, vol. 2022, pp. 1–18, Jul. 2022, doi: 10.1155/2022/9261537.
[38]P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, “Graph Attention Networks,” 2017, *arXiv*. doi: 10.48550/ARXIV.1710.10903.
[39]H. Cai, V. W. Zheng, and K. C.-C. Chang, “A Comprehensive Survey of Graph Embedding: Problems, Techniques and Applications,” 2017, *arXiv*. doi: 10.48550/ARXIV.1709.07604.
[40]A. Zolnai-Lucas, J. Boylan, C. Hokamp, and P. Ghaffari, “STAGE: Simplified Text-Attributed Graph Embeddings Using Pre-trained LLMs,” 2024, *arXiv*. doi: 10.48550/ARXIV.2407.12860.
[41]R. Ruiz-Dolz, S. Heras, and A. García-Fornes, “Automatic Debate Evaluation with Argumentation Semantics and Natural Language Argument Graph Networks,” 2022, *arXiv*. doi: 10.48550/ARXIV.2203.14647.
[42]M. Lenz *et al.*, “Towards an Argument Mining Pipeline Transforming Texts to Argument Graphs,” 2020, doi: 10.48550/ARXIV.2006.04562.
[43]A. Aflatoun, *GitHub - Voidcake/Argument-Networks*. (Jan. 05, 2025). Python. Accessed: Jan. 08, 2025. [Online]. Available: https://github.com/Voidcake/Argument-Networks.git
[44]“Structural Embeddings.” Accessed: Jan. 12, 2025. [Online]. Available: https://github.com/Voidcake/Argument-Networks/tree/main/resources/visualisations/embeddings/structural
[45]“Semantic Embedding · GitHub.” Accessed: Jan. 12, 2025. [Online]. Available: https://github.com/Voidcake/Argument-Networks/tree/main/resources/visualisations/embeddings/semantic
[46]“Aggregated Semantics - GitHub.” Accessed: Jan. 12, 2025. [Online]. Available: https://github.com/Voidcake/Argument-Networks/tree/main/resources/visualisations/embeddings/aggregated_semantics

Erklärung

Ich versichere, die von mir vorgelegte Arbeit selbstständig verfasst zu haben. Alle Stellen, die wörtlich oder sinngemäß aus veröffentlichten oder nicht veröffentlichten Arbeiten anderer oder der Verfasserin/des Verfassers selbst entnommen sind, habe ich als entnommen kenntlich gemacht. Sämtliche Quellen und Hilfsmittel, die ich für die Arbeit benutzt habe, sind angegeben. Die Arbeit hat mit gleichem Inhalt bzw. in wesentlichen Teilen noch keiner anderen Prüfungsbehörde vorgelegen.

Köln, 13.01.2025

Ort, Datum

Unterschrift