LUISS

Department of Business and Management

Bachelor's Degree in Management and Computer Science
Course of Algorithms

Mapping Tariff Discourse: Comparing Relational and Semantic Networks of World Leaders

SUPERVISOR

Prof. Irene Finocchi

CANDIDATE

Michele Turco

ID: 285 251

Academic Year 2024/2025

Abstract

In recent years, tariff policies have re-emerged as central tools of geopolitical and economic strategy. No longer confined to classical protectionism, tariffs are now used both to address trade imbalances and to signal political intent. This dual function, economic and symbolical, has intensified global debate over their strategic use.

The announcement of the 2025 U.S. tariff policy on “Liberation Day” (April 2) sparked widespread public and political reaction. Conversations unfolded not only among consumers and citizens, but also among economic stakeholders, political leaders, and international commentators. To understand how this discourse evolved, this study analyses data from X (formerly Twitter), a platform central to real-time political communication due to its openness, rapid posting, and use by public officials and journalists.

The analysis investigates whether global leaders and influential figures shaped or amplified key narratives surrounding the tariff decision. Using unsupervised clustering on text embeddings, it identifies dominant themes and links users to discourse clusters through a bipartite user–discourse graph. From this, a semantic network of online rhetoric is constructed.

This dynamic network is then compared to a static map of geopolitical alignments to assess whether online discursive behaviour reflects real-world institutional relationships. The comparison sheds light on how rhetorical alliances, political signalling, and ideological convergence manifest in the digital public sphere.

Table of Contents

1 Introduction

1.1 Background

In its initial years, the United States collected the majority of its governmental revenues from **tariffs**, levies on imported foreign goods. However, from the beginning of the 20th century, this was reversed with the institution of the federal income tax and the development of a new consensus that recognized tariffs as regressive and burdensome for working individuals while leaving most of the incomes of the wealthy untaxed (Clausing and Lovely, 2024). Despite this historical shift away from protectionist taxation, on April 2, 2025, during a White House Rose Garden ceremony, the 47th United States President, Donald Trump, signed **Executive Order 14257** to implement his plan of a standard 10% tariff on all imported goods. Certain countries were imposed even higher rates: 46% for Vietnam, 32% for Taiwan, 20% for the European Union, 24% for Japan, 49% for Cambodia, 36% for Thailand, and 54% (later raised to 245%) for China. The move marked a significant departure from decades of liberal trade policy and provoked sharp responses from affected countries, raising the risk of retaliatory tariffs and escalating tensions in the global trading system.

The plan, however, aligned with the *America First* trade policy, aiming to address the USA's deficit and pressure other countries to renegotiate deals and treatments considered unfair and harmful to U.S. interests. Trump indeed denoted April 2 as the *Liberation Day* and described it *as one of the most important days in American History*, explaining why these tariffs are sometimes referred to as *Liberation Day Tariffs*. Trump explicitly stated that

> "April 2, 2025, will forever be remembered as the day American industry was reborn, the day America's destiny was reclaimed and the day that we began to make America wealthy again"

The general tariffs were announced to take immediate effect on April 5th, while country-specific ones were to become effective on April 9. In addition, there were tariffs related to specific sectors, such as a 25% tariff on imports of steel and aluminium (implemented on March 12), often justified by concerns over national security and the protection of domestic industries, and a further 25% tariff on auto-mobiles or certain auto parts.

The decision sparked an **immediate and strong debate**, in particular with regard to whether the justifications for these decisions were sufficient to outweigh the drawbacks that emerged in the following weeks. The USA's economy, as often occurs with such a tariff policy, experienced a drastic slowdown, which naturally caused a shock of the stock market. The S&P 500 index lost more than 10% between April 2nd and 8th, marking its worst performance since the Covid-20 pandemic. Additionally, U.S. consumers began to worry about a highly likely increase in prices. Retailers anticipated struggles with supply chains and cost management, while farmers braced for the impact of retaliatory tariffs. On the other hand, U.S. manufacturers welcomed the decision, as key objectives were to revive this sector, correct trade imbalances, and address unfair trade practices.

On 9 April 2025, with the signature of Executive Order: "Modifying Reciprocal Tariff Rates to Reflect Trading Partner Retaliation and Alignment", President Trump declared a 90-day halt to the imposition of these tariffs in full, except for China, lowering the tariff rate in the EU from 20% to the baseline 10%. The 25% tariffs on steel, aluminium, and cars are still in effect. The pause would serve to enable negotiations with trading partners while keeping pressure, especially on China, whose tariffs were raised to 145%.

> "China probably will eat those tariffs. But at 145, they basically can't do much business with the United States. And they were making from us a trillion dollars a year. They were ripping us off like nobody's ever ripped us off"

In response, on April 10, **European Commission** President Ursula von der Leyen announced that while the EU had voted to impose retaliatory action on US steel and aluminium tariffs, it would suspend their implementation for 90 days. On X @vonderleyen (2024):

> "I welcome President Trump's announcement to pause reciprocal tariffs. It's an important step towards stabilizing the global economy. Clear, predictable conditions are essential for trade and supply chains to function."

However, according to the Kiel Institute, the EU could still see a GDP loss of 0.2%-0.8% depending on the introduction of retaliatory measures. The impact is moderate relative to other recent shocks, such as COVID-19 pandemic (-5.6%) or the latest energy crisis (-2.4%). The reduction in tariffs on a partial basis may cancel out part of its consequences, but there are still robust economic headwinds. The impact on the financial markets has been achieved in the form of increased volatility, with sector-specific risks to EU economies significantly dependent on exports to the US, such as automotive, pharmaceuticals, and machinery.

Date	**Event**
12 March 2025	25% tariffs imposed on imports of steel, aluminium, and certain products from the EU and other partners.
2 April 2025	US administration announces "reciprocal tariffs" with a baseline 10% tariff on all imports and additional country-specific tariffs.
8 April 2025	White House confirms plans to impose an additional 50% tariff on Chinese goods, effective 9 April.
9 April 2025	Trump announces a 90-day pause on additional country-specific tariffs (excluding China). EU tariff reduced from 20% to 10%. Tariffs on China rise to 125%. Steel, aluminium, and auto-mobile tariffs remain.
11 April 2025	White House issues presidential memorandum exempting certain semiconductor-related goods from the 2 April tariff list.

Table 1.1: Timeline of Key US Tariff Actions in March–April 2025

On the other hand, China has also imposed an 84% duty on U.S. imports, in retaliation, which is a significant ratcheting up of the ongoing trade war. The retaliations have raised alarm worldwide, with most warning that the conflict could drive the world economy into recession. The latest tariff hike amounts to an estimated $400 billion cost boost for U.S. consumers and businesses, prior to even factoring in substitution effects. For **China**, the currency is now a potential tool for offsetting the impact. The Chinese yuan has already declined by about 1.6% since mid-March to 7.34 per dollar, and there is growing speculation that officials will allow further devaluation in response to the higher tariffs. Due to these trends, China's economic growth for the full year of 2025 is now predicted to slow down to 4.4%, downgrading from initial expectations. Weakening of the global outlook, coupled with slowing exports to the U.S., will likely take around 0.3 percentage points off growth via trade. Additionally, more palatable consumption and investment in sectors with export orientation ought to reduce growth another 0.4

percentage points. But at least part of this drag is supposed to be cushioned by the scheduled release of an additional 1 trillion yuan of central government bonds, to be announced probably in the third quarter of 2025 (J.P. Morgan, 2025).

Subsequently, on April 11, the White House released a presidential memorandum: "Clarification of Exceptions Under Executive" stating that certain products that contains semiconductors, including smartphones and other consumer electronics would have not not been subject to the tariffs associated with the 2 April executive order (Trump, 2025).

1.2 Purpose

The overall objective of this thesis is to examine whether public declarations, such as official press releases, conferences and social media posts issued by global leaders and influencers are aligned with their self-proclaimed political positions. The statements are examined both from a semantic and relational perspective, with the former measuring proximity of vector representations and the latter based on whether tweets in a cluster are salient within the same user communities. Based on this, the study constructs a bipartite **user–discourse network**, linking each actor to the clusters most closely associated with their rhetorical activity. This structure enables the retrieval of a layered semantic network, where connections between leaders are inferred from shared semantic behaviours and ideological proximity. This approach allows for the identification of discursive alignment and facilitates its comparison with geopolitical orientation, understood as the formal, policy-based relationships between states. The resulting contrast provides a novel analytical lens for understanding transnational signalling, **rhetorical convergence** and digital diplomacy in the tariff announcement landscape.

In parallel, the research employs unsupervised clustering on embedded tweet representations to uncover the dominant semantic clusters emerging from the social media discourse of ordinary users, with the goal of mapping the broader landscape of public opinion surrounding the debated topic. Elite political discourse indeed operates at a different semantic and rhetorical register compared to the general tweet corpus. Clusters derived with this methodology can reflect sub-debates being conducted, community polarization, or thematic trends in specific geopolitical blocks. As discussed in the background, different aspects of the overall theme can be addressed simultaneously, e.g., political and diplomatic implications, economic and financial consequences and sectoral reactions. Such clusters can also reflect concerns of specific countries, sentiments of certain communities or ideologies.

By the integration of natural language processing methods, sentiment detection and social network analysis, this research belongs to a new generation of computational political discourse analysis. It aims to answer if X discourse reflects geopolitical blocs, signals shifts in ideological camps, or discovers emergent rhetorical blocs that disrupt classical blocs.

2 Literature Review

The analysis makes use of a wide range of computer science disciplines that are increasingly being applied in conjunction. The theoretical background of the thesis has roots in word embeddings, sentiment analysis, clustering algorithms, and social network analysis. As a first step, each field will be described individually. Then, recent literature will be reviewed to illustrate how these techniques have been integrated in prior works relevant to this study.

2.1 Foundational Knowledge

2.1.1 Vector Space Representations of Lexical Semantics

Word embeddings are techniques that convert textual content into numerical vector representations, enabling machine learning algorithms to process linguistic data more efficiently. One of the earliest approaches to this was *Latent Semantic Analysis* (LSA) (Deerwester et al., 1990), a foundational contribution to the field that provided a method for capturing semantic relationships between words and documents through dimensionality reduction. The output of LSA consists of dense vector representations for both terms and documents in a lower-dimensional semantic space, allowing similar words and documents to be grouped based on latent patterns of co-occurrence. Subsequently, *Latent Dirichlet Allocation* (LDA) (Blei et al., 2003) was developed to discover topics within a collection of documents. It functions as a generative probabilistic model: groups of words that frequently co-occur are automatically identified across a large corpus and are returned as distinct topics by the algorithm. This marked a significant step in the development of topic modelling and detection.

Later on, a pivotal contribution to this field was the introduction of *Word2Vec* by Mikolov et al. (2013), which represented words as continuous vectors in a high-dimensional space, capturing semantic similarity based on local context within a fixed window.This study had a significant impact as it was among the first methods to enable the quantification of semantic distance between individual words, thereby allowing for the measurement of their proximity in meaning. Word2Vec consists of two architectures, Continuous Bag of Words (CBOW) and Skip-gram, that learn embeddings by predicting a word from its context or vice versa. Despite its strengths, Word2Vec has two main limitations. In the first place, it does not work on a document level and is consequently not capable of capturing the global message conveyed. Moreover, it produces *static embeddings*, meaning each word has a single, context-independent vector. As a result, it cannot handle polysemy or adapt to context-dependent shifts in meaning.

This limitation was addressed with the development of *contextualized word embeddings*, beginning with *ELMo* (Embeddings from Language Models) (Peters et al., 2018) and significantly advanced by transformer-based models such as *BERT* (Devlin et al., 2019), which are grounded in the Transformer architecture introduced by Vaswani et al. (2017). These models generate word vectors dynamically based on surrounding context, allowing the same word to take on different representations depending on its usage.

ELMo introduced the idea of context-dependent embeddings by employing a deep, bidirectional LSTM trained on a language modelling objective. While it significantly improved performance on many downstream NLP tasks, its reliance on sequential processing limited scalability and training efficiency. BERT, on the other hand, is built entirely on self-attention mechanisms, as proposed in the Transformer model. It processes text bidirectionally using masked language modelling and next sentence prediction objectives, enabling it to capture rich contextual information. This shift to attention-based architectures marked a turning point in NLP, leading to substantial gains across a range of tasks including sentence classification, question answering, and semantic similarity.

Building on this, *Sentence-BERT* (SBERT) (Reimers and Gurevych, 2019) was introduced to compute sentence-level semantic similarity efficiently, using cosine distance between sentence embeddings. Unlike BERT, which is optimized for token-level classification, SBERT modifies the architecture by adding a pooling operation to generate fixed-size sentence vectors, making it suitable for downstream tasks that require meaningful comparisons between entire texts. This adjustment enables significantly faster inference in pairwise similarity and clustering tasks, as it eliminates the need to recompute full cross-encodings for each pair. SBERT has proven to be especially effective in clustering, classification, and semantic similarity tasks on short-form text such as tweets, where traditional models often struggle with noise and lack of context.

Upon these advancements, OpenAI introduced `text-embedding-3-large` (OpenAI, 2024), a **state-of-the-art** embedding model capable of generating 3072-dimensional dense vector representations. This model leverages a transformer-based architecture optimized for semantic similarity and retrieval tasks across diverse textual inputs. In this study, it was employed to obtain high-quality sentence embeddings for each tweet, allowing for accurate comparisons in the subsequent clustering and network analysis phases.

2.1.2 Computational Approaches to Sentiment and Opinion Detection

Sentiment analysis, or **opinion mining**, is a natural language processing (NLP) technique whose purpose is to determine the emotional tone or evaluative stance conveyed in textual documents. Generally, the overall aim is to classify sentiment as one of discrete classes, which are most often positive, neutral, or negative. Early approaches were lexicon- or rule-based, employing hand-curated sentiment lexicons where each word received a polarity score or affective label. The foundational study by Pang et al. introduced supervised techniques to accomplish the task, establishing it as central to computational linguistics. Popular tools such as *SentiWordNet* (Baccianella et al., 2010) provided the foundation for sentiment labelling with the accumulation or averaging scores of the words in a sentence or document. Rule-based systems sometimes incorporated syntactic properties to model negation (e.g., ”not good” vs. ”good”) or intensifiers (e.g., ”very bad” vs. ”bad”), but such systems were typically brittle and domain-specific. One notable study among them is *VADER* (Hutto and Gilbert, 2014), which was designed specially for social media and contains heuristics for capitalization, punctuation, and emoticons to achieve surprisingly strong performance on short, informal text.Although these methods were intuitively straightforward and interpretable, they struggled with ambiguity, contextual change, sarcasm, and domain drift—particularly when dealing with noisy, informal text such as social media.

Word embedding development was a significant step forward in this direction: models learned to represent semantic similarity of words and thereby make sentiment classifiers more generalizable, and detect more subtle emotional phrases. Tang et al. (2014) introduced *Sentiment-Specific Word Embeddings* (SSWE), which extended the Word2Vec model by incorporating sentiment supervision in the training process. Their model was aligned not only on context but also on sentiment polarity and thus performed better in downstream sentiment classification tasks. This work was an important step in showing that sentiment information could and should be integrated into the learning process of vectors itself. Later, Severyn and Moschitti (2015) employed pre-trained

Word2Vec representations as input to a Convolutional Neural Network (CNN) for sentiment classification on individual tweets. Their work demonstrated that beginning with high-quality, unsupervised representation and fine-tuning on sentiment data leads to consistent performance even in data-poor conditions like Twitter sentiment tasks. Their architecture performed the best for the SemEval-2015 sentiment analysis task, establishing the role of embeddings as the baseline layer in deep sentiment models

Contemporary sentiment analysis is dominated by transformer-based models. These can be fine-tuned on labeled datasets for sentiment classification tasks. For example, the model `nlptown/bert-base-multilingual-uncased-sentiment` predicts sentiment on a 1–5 star scale, while `textattack/bert-base-uncased-imdb` (Morris et al., 2020) performs binary classification. The `cardiffnlp/twitter-roberta-base-sentiment` (Barbieri et al., 2020) model is trained on over 124 million English tweets and finetuned using a benchmark dataset called **TweetEval**. It is specifically optimized for social media sentiment, classifying text as positive, negative, or neutral.

2.1.3 Clustering Algorithms and the Notion of Data Proximity

Data clustering, also denoted as **cluster analysis**, is a sequence of mathematical operations, namely, an algorithm, whose goal is to discover natural groupings within a set of patterns, objects, or observations represented as points in a high-dimensional space. The thousands of published algorithms and the wide range of scientific domains and applications that have employed clustering techniques are too numerous to enumerate here (Jain, 2010). Clustering is particularly valuable in unsupervised settings, such as social media discourse analysis, where label information is absent and the goal is to identify emergent semantic or ideological groupings (Vakulenko et al., 2017).

Its relevance has only increased with the growing use of vector-based representations of language, where clustering techniques can be used to identify latent structures or communities in semantic space. Moreover, the flexibility of clustering models makes them particularly suited for contexts where the number or shape of clusters is not known a priori, a common condition in open-domain discourse analysis.

At the core of all clustering algorithms lies the concept of distance or similarity between points. The most common is the **Euclidean distance**, the direct-line distance between two points in n-dimensional space. For two vectors $\vec{x}, \vec{y} \in \mathbb{R}^n$, the Euclidean distance is defined as:

$$d_E(\vec{x}, \vec{y}) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}$$

This metric assumes isotropic, continuous feature spaces and works best if all dimensions are on the same scales. In high-dimensional space, however, the Euclidean distance can become ever more meaningless due to the so-called ”curse of dimensionality, since the contrast between the nearest and farthest points diminishes making it harder to distinguish clusters based on distance. (Aggarwal et al., 2001).

Otherwise, especially with text data that is expressed in embeddings, **cosine similarity** is commonly used. Rather than absolute position, it defines the angle between two vectors, and is especially helpful when vector magnitudes have a wide range of values. The cosine similarity is expressed as:

$$\cos(\theta) = \frac{\vec{x} \cdot \vec{y}}{||\vec{x}||\,||\vec{y}||}$$

where $\vec{x} \cdot \vec{y}$ is the dot product and $||\vec{x}||$ is the norm (magnitude) of vector $\vec{x}$. In practice, this metric gives a more intuitive sense of the semantic similarity between sentences or documents in vector space, regardless of length or intensity (Singhal et al., 2001; Huang et al., 2008).

Clustering approaches can be broadly categorized into several methodological families. **Partition-based** methods attempt to directly assign each point to a group such that intra-cluster similarity is maximized and inter-cluster similarity is minimized. These methods may be deterministic, yielding the same output for the same input, or non-deterministic, where different initializations may lead to different results — a property especially common in algorithms that rely on random seeds or centroid selection.

Another prominent category consists of **density-based** clustering methods, which define clusters as regions of high point density separated by regions of low density. These approaches are particularly useful when dealing with clusters of arbitrary shapes or noisy data distributions. Finally, **hierarchical clustering methods** operate by recursively merging or splitting clusters to form a tree-like structure known as a dendrogram.

This framework allows for multi-resolution analysis and does not require pre-specifying the number of clusters, making it advantageous in exploratory settings.

While these paradigms vary in assumptions and computational strategies, they are all grounded in the goal of revealing hidden structure within data. In the context of language data and discourse analysis, they offer powerful tools for capturing implicit thematic or ideological divisions that may not be explicitly labelled or predefined.

2.1.4 Social Network Analysis

Social Network Analysis (SNA) is a methodological framework for studying the structure of social systems by modelling relationships as networks. Rather than focusing solely on individual attributes, SNA emphasizes the relational patterns that emerge among actors (Otte and Rousseau, 2002) An influential early definition was provided by Wetherell et al., who described it in the following terms:

> "Most broadly, social network analysis conceptualises social structure as a network with ties connecting members and channelling resources, focuses on the characteristics of ties rather than on the characteristics of the individual members, and views communities as 'personal communities', that is, as networks of individual relations that people foster, maintain, and use in the course of their daily lives."

At its core, SNA draws on the mathematical foundations of graph theory. A graph is formally defined as $G = (V, E)$, where V is a set of nodes (also called vertices) and $E \subseteq V \times V$ is a set of edges (or links) connecting pairs of nodes (Diestel, 2005). Graphs can be classified in various ways depending on the analytical focus:

- **Directed**: if edges have orientation, meaning a connection from A to B does not imply one from B to A.
- **Undirected**: if edges represent mutual association or presence.
- **Weighted**: if edges carry an intensity or frequency.

From this foundation, researchers quantify relational patterns using measures such as degree, centrality, clustering coefficient, and modularity. These metrics form the analytical backbone of network studies across sociology, political science, and communication research.

- **Degree**: The degree of a node is defined as the number of edges connected to it. In directed graphs, this is often divided into:
  - *In-degree*: the number of incoming edges (e.g., followers).
  - *Out-degree*: the number of outgoing edges (e.g., followees).

  High-degree nodes often represent actors with high visibility or influence within the network.

- **Centrality**: Centrality measures the importance or influence of a node. Several types exist:
  - *Betweenness centrality*: measures how often a node lies on the shortest paths between other nodes; it indicates the node's role as a bridge or broker.
  - *Closeness centrality*: captures how close a node is to all other nodes in the network, based on shortest path distances.
  - *Eigenvector centrality*: assigns importance to a node based on the importance of its neighbours. This is the basis for algorithms like Page Rank.

- **Clustering Coefficient**: The local clustering coefficient quantifies how connected a node's neighbours are to each other. It is calculated as:

  $$C_i = \frac{2 \times \text{number of closed triplets}}{\text{number of connected triples of vertices}}$$

  This metric reflects the level of local cohesion and is useful for identifying tightly connected communities or echo chambers.

- **Modularity**: Modularity measures the strength of division of a network into modules or communities. A high modularity score indicates that nodes within a group are densely connected, while connections between groups are sparse. For undirected networks, modularity is typically computed as:

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(c_i, c_j)$$

where A_{ij} is the adjacency matrix, k_i and k_j are the degrees of nodes i and j, m is the total number of edges, and $\delta(c_i, c_j)$ is 1 if i and j are in the same community, and 0 otherwise.

Considering the type of nodes that constitute the network, we can distinguish two relevant subcategories. In *relational networks*, nodes represent individuals and edges represent observable interactions such as follows, replies, or retweets. These structures are essential for analysing patterns of influence, information diffusion, and community formation.

Conversely, *semantic networks* focus on the content and meaning of communication. Nodes may consist of keywords or clusters of thematically related messages, while edges indicate co-occurrence, conceptual association, or semantic similarity. These networks capture *what* is being said, by *whom*, and in what discursive proximity. This approach provides a complementary lens to traditional relational SNA, particularly in the analysis of ideological alignment, narrative framing, and rhetorical structures.

Importantly, semantic and relational networks are not mutually exclusive. Combining them, such as in bipartite user, discourse graphs—enables a more integrated analysis of how rhetorical strategies align with networked social behaviour.

2.2 Related Works

Even if there are few attempts that are similar to the nature of the whole thesis, there are many studies that address parts of the latter, especially with the widespread adoption of social media and the consequent development of discourse there. In particular, X (formerly Twitter) has become a central platform for examining political and social discourse due to its diverse interaction modalities.

2.2.1 Evolving Models of Twitter Discourse: From Interaction to Meaning

Twitter (X) has been recognized for a long time as an exceptionally sociologically and politically significant platform. Due to its magnitude, openness, and real-time nature, it is not only a medium of communication but also a proxy for broader societal behaviour. Ahmad (2010), examining journalism and media practice, proves that Twitter makes it possible for researchers to track emerging public narratives and issue salience, vital elements needed to identify live ideological discourse shifts. (Kokkinogenis et al., 2015) demonstrate instead Twitter's use as a participatory sensing system, where users' input provides real-time readings of on-the-ground conditions such as traffic or urban mood. These researches lend credence to the argument that Twitter is more than a platform of speeches, but an active stream of information with the ability to reflect both socio-political mood as well as behavioural patterns in near real time.

Considering its powerful ability to capture socially relevant concerns, the majority of researches began to study Twitter trough the lens of social network analysis. The entire X network can be represented as a directed graph where every user is a node and each "follow" link is a directed edge. However, the model reflects latent structure of online communities rather than active interaction or influence. As a result, researchers began to ask more conceptual questions, not just who follows whom, but what is the element that actually causes an individual to be central in a network. One of the earliest breakthroughs was *TwitterRank* (Weng et al., 2010), which adapted Google's PageRank algorithm to measure influence with greater nuance. Unlike previous approaches that simply counted followers or retweets, TwitterRank addressed the reality that Twitter influence is **topic-sensitive**: a technology influencer doesn't necessarily need to be particularly important in politics or sports. To manage this, topic modelling was integrated with graph structure, weighing user links by theme. This procedure is particularly fitting in the context of this thesis, since it is one of the initial attempts to place semantic content within a naturally relational structure, offering conceptual groundwork for more sophisticated combinations of meaning and structure in discourse analysis.

Subsequently, Conover et al. (2011) introduced, in one of the most influential works in the field, two alternative representations of interaction beyond the follower network:

- **Retweet Network**: an edge from user A to user B exists if user B retweets content originally authored by A, indicating information flow from A to B.
- **Mention Network**: an edge from user A to user B exists if user A mentions

user B in a tweet, modeling direct address.

The contribution has been essential to understand how political talk evolves on the site, and it is the reason why it is worth mentioning in this study. Conover's research steered the attention of analysis away from fixed network structure toward dynamic discursive exchange and opened the door for a generation of researchers interested in how political polarization develops and shifts through digital activity.

Later efforts extended early interaction-based studies by creating **hashtag co-occurrence networks**, where nodes represent hashtags and edges connect those of them that occur together in the same tweet or are used by the same user. This was among the first attempts of a Twitter-based semantic network construction, shifting the focus of analysis from interpersonal relationships to the linguistic overlap among them. Of special note is the study on the 2018 Italian election (Radicioni et al., 2021) that had combined co-occurrence networks of hashtags with user interactions to reveal ideological segmentation and communication asymmetry across political parties. This experiment demonstrated the way in which markers within language could be employed to chart political alignment and boundaries of discourse. In the COVID-19 pandemic, the same approaches were utilized to model the ways in which public health narratives, institutional trust, and political framing altered in times of crisis. Hashtag semantic networks under such conditions made available tools for researchers to chart changes in agenda-setting and aggregate sentiment over time (Mattei et al., 2021). Very recently, for the 2022 Italian general elections, Federico et al. (2024) carried out a multi-scale analysis of reply and retweet networks, proving that network structure strongly followed party coalition and campaign strategy. The authors' analysis distinguished between "top-down" communities, led by core leaders, and "broad discussion" communities characterized by decentralized participation and participatory debate.

At the same time, a full-fledged methodology was introduced by Jess and Bayhan (2024), which incorporates text-based features, sentiment scores and keyword occurrence—into dynamic Twitter networks. Although they do not embed tweets into semantic vector spaces, they do incorporate sentiment (using TextBlob and VADER) and keyword flags on nodes, and use these features to animate change in discourse over time. Their case study of the ChatGPT AI discussion demonstrated how keyword trends and emotional tone evolve in concert with the development of the network structure. This hybrid approach is a step towards semantically enriched network analysis, but it remains grounded in rule-based sentiment models and hand-selected keywords rather than learned representations.

Together, such studies reflect a symbiosis between social and semantic organizations in political communication even if, they operate on finite proxies such as co-occurrence or symbolic tags.

2.2.2 Joint Topic–Sentiment Modeling

This thesis takes the aforementioned methodological trajectory further, by embedding tweets into real-valued vector spaces that reflect both semantic and sentiment. Traditionally, these two analysis have been conducted separately, with semantic embeddings providing a strong foundation for modelling thematic, even if they generally lacked sensitivity to emotional valence. In complex domains like political communication, where tone often distinguishes support from critique within discussions on the same matter, joint modelling of topic and sentiment provides richer interpretive value. It has been an active area of research for over a decade and is directly relevant to this study's aim of constructing tweet-level representations that integrate both semantic content and polarity framing. Rather than isolating topic and sentiment as distinct stages, this approach acknowledges the necessity of including them both, as emotional tone frequently signals ideological alignment or opposition within the same topical field.

One of the earliest formal models to jointly capture topic and sentiment was the *Joint Sentiment/Topic* (JST) model proposed by Lin and He, which extended Latent Dirichlet Allocation (LDA) by introducing a sentiment layer into the generative process. In JST, documents are modelled as mixtures over sentiment–topic pairs, allowing words to be associated not only with thematic topics but also with sentiment labels. Although JST and its successors rely on discrete latent variables, the conceptual premise, that sentiment and topic co-occur and condition each other, remains foundational for more recent neural and embedding-based models.

A subsequent shift occurred with the introduction of neural representation learning. Maas et al. (2011) proposed a hybrid model that combined unsupervised word vector learning with sentiment supervision using document-level polarity annotations. This enabled the learned embedding space to reflect both semantic similarity and emotional valence. Building on this, Tang et al. (2014) introduced *Sentiment-Specific Word Embeddings* (SSWE), trained under distant supervision using emoticons, to address the common challenge of separating semantically close but sentimentally opposite words. Rouvier and Favre (2016a) extended these ideas by developing a fusion architecture that integrates lexical, syntactic, and sentiment embeddings within a neural network, yielding robust classification performance. Similarly, Liu et al. (2021) introduced cross-domain sentiment-aware embeddings that preserved polarity distinctions even across thematically diverse contexts such as product reviews.

Whereas earlier approaches like JST operated through probabilistic generative modelling, a growing body of work now adopts a two-step architecture: first deriving continuous vector representations of texts (semantic and/or sentiment-aware), and then applying clustering algorithms to organize these embeddings into topic–sentiment groupings. This decoupled pipeline, often denoted as "Semantic Space Partitioning for Topic Discovery", will be the focus of the next section.

2.2.3 Semantic Space Partitioning for Topic Discovery

Recent progress in representation learning introduced a new strand of activity in topic modelling: instead of inferring topics through probabilistic inference over discrete distributions over words, these models induce topics by clustering high-dimensional semantic spaces constructed by pre-trained language models. This paradigm shifts the analytical focus from generative modelling to **spatial clustering**. It has three basic advantages: (1) it maintains word order and semantic specificity, (2) it disentangles feature acquisition from topic induction, and (3) it facilitates easy incorporation of sentiment and other discursive signals.

Part of the elegance of this approach is the leveraging of semantic embeddings, typically acquired from algorithms like Word2Vec, Doc2Vec, or Sentence-BERT, to represent texts in a vector space where proximity means thematic congruity. Clustering algorithms such as HDBSCAN or k-means are then applied to these vectors to extract emergent topic clusters. Since these reflect local densities in the semantic space, they will represent more coherent and fine-grained discourse structures than the output of generative models such as LDA or JST. A second strand of research concurrent with embedding-based clustering explores deep generative models of document semantics with neural variational inference. One such significant work in this direction is the

Neural Variational Document Model (NVDM) (Miao et al., 2016), which learns continuous latent representations of documents with a variational autoencoder learned for text input reconstruction. Even though NVDM does not cluster explicitly, the resultant document embeddings can be conceptualized or grouped post-hoc and serve as a conceptual bridge between LDA-type models and modern representation-based models. *Top2Vec* (Angelov, 2020) is arguably the most successful implementation of this idea, which learns word and document representations simultaneously with Doc2Vec and discovers topics as centroids of dense clusters within this representation. By employing UMAP dimension reduction and HDBSCAN density-based clustering, Top2Vec discovers the number of topics automatically and infers highly informative topic labels by nearest-neighbour word vectors. Central to this is that it requires no manual preprocessing (e.g., removal of stop words or stemming) and outperforms LDA on interpretability and information gain. Building on similar concepts, *BERTopic* (Grootendorst, 2022) incorporates Sentence-BERT embeddings to continue this pipeline as long as a class-based TF-IDF module to further refine topic descriptors. Sentiment is not explicitly modeled, but the modular nature of BERTopic has seen it become the go-to option for unsupervised short text analysis, most notably across platforms like Twitter.

Sia et al. (2020) propose an allied strategy by clustering word embeddings directly and reranking cluster keywords based on coherence and semantic uniqueness, hypothesizing that clusters of pretrained vectors can substitute traditional topic models.

Culminating this trajectory, Hanny and Resch (2024) propose a clustering-based joint topic–sentiment (JTS) framework specifically tailored to short-form social media discourse. Their model first generates SBERT-based semantic embeddings and applies *UMAP* to reduce them to a five-dimensional space. Sentiment vectors are then extracted using a transformer classifier (*twitter-roberta-base-sentiment*) and concatenated with the semantic vectors. The combined representation is clustered using HDBSCAN, yielding discourse groupings that are semantically coherent and affectively distinct. This architecture represents a methodological synthesis that aligns most closely with the present study: it integrates both semantic proximity and affective valence into a unified spatial representation, allowing ideological and emotional fault lines to emerge through clustering.

3 Methodology

3.1 Dataset

In order to retrive data, the current study utilized the **Apify** platform, an automation tool that is designed for web scraping and browser-based data extraction. Established in 2015 by Jan Čurn and Jakub Balada, Apify employs headless browser technology to simulate user action and scrape structured information off of web interfaces. Its nature lets developers build scalable crawlers in JavaScript, so it’s extremely powerful on dynamic or JavaScript-rich websites such as Twitter.

3.1.1 Tweets

The dataset considering social network discourse initially comprised 50,748 tweets in multiple languages, collected using the actor apidojo/tweet-scraper (2024). Search terms applied to scrape these tweets ranged from general keywords such as ”tariffs,” ”*Trump tariffs*” and ”*stock market crash*” to more specific phrases which evoked stances pertaining to policy instruments, geopolitical agents, or economic indicators. The desire was to sample a large range of discursive positions, including potentially marginalized or fringe opinions which might remain at least rhetorically active in framing the debate.

The period for collecting data extended for 2 weeks, from the time the imposition tariff was declared on April 2nd until shortly after the suspension announcement, April 16th, as illustrated in Figure 3.1.

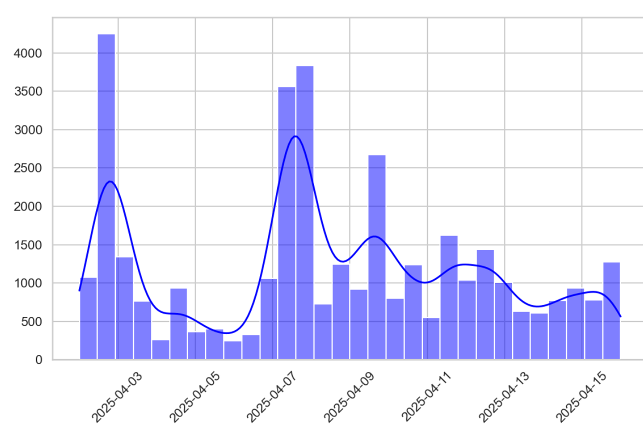

Figure 3.1: Distribution by Date

This temporal structure verifies that the debate did not unfold smoothly but instead formed around punctuated moments of political and media significance, reaffirming the need for a temporally conscious analysis of rhetorical alignment and influence.

35 categories (columns), with some of them even being complex data, such as Json strings, were presented in the dataset. These encompassed tweet identifiers, full and raw text, language, creation time, retweet and like counts, reply metadata and client source (e.g., Twitter Web App, iPhone, Android). Although some fields were sparsely populated or redundant, others could have been useful for secondary tasks like language filtering, temporal segmentation, or source-based audience profiling. For example, over 52% of the tweets were English, followed by Portuguese and Spanish (12% each), and 21% other languages. The research retrieved instead only few Italian tweets, in spite of the intense activity of the Italian Prime Minister. The detailed distribution is shown in Figure 3.2

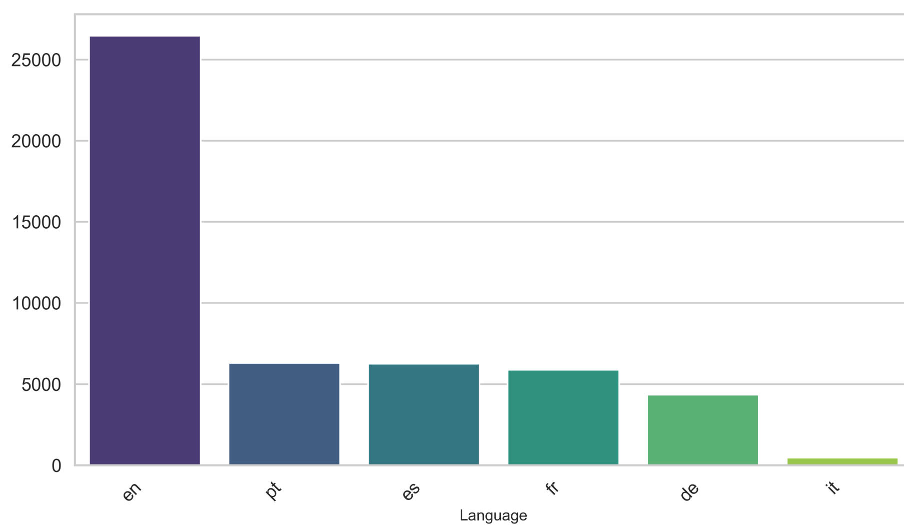

Figure 3.2: Language distribution across tweets

It is notable that a vast number of tweets do not contain **hashtags** (Figure 3.3). This underuse confirms that they are not a overpowering discursive expression device in this corpus. Thus, any semantic or thematic analysis relying exclusively on their co-occurrence would be partial and even misleading. While sometimes informative, hashtags are in this cases too sparsely and erratically distributed to constitute a solid basis for discourse segmentation. This result supports the methodological choice of relying alternately on embedding-based representations of full tweet text, more suitable for implicit stance, rhetoric, and topic composition capture.

Hyperlinks, on the other hand, are more uniformly distributed, with more than half of all tweets containing at least one link, as shown in Figure 3.3 This is an important finding, as there was some preliminary consideration of excluding link-containing tweets from the analysis, based on the suspicion that they might constitute automated, promotional, or non-discursive content (e.g., news headlines or promotional tweets by influencers). But their substantial share of the dataset and their potential rhetorical importance (e.g., to offer evidence, boost arguments, or frame discourse) justified their inclusion. Removing them would have introduced a bias that disproportionately mutes some voices and communication approaches.

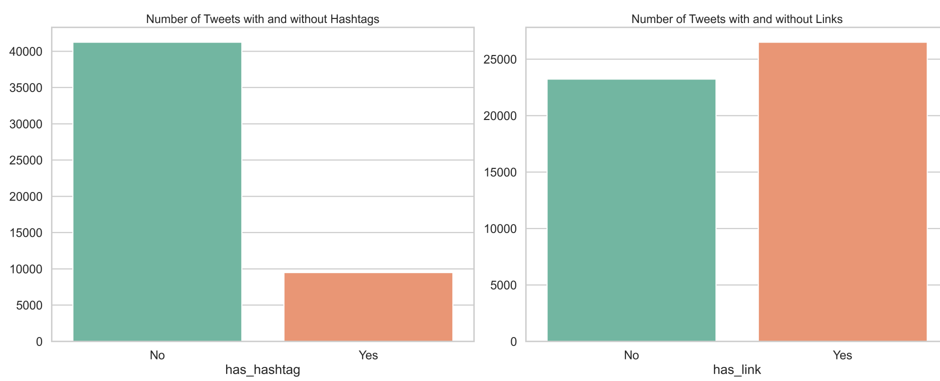

Figure 3.3: Proportion of tweets containing hyperlinks and hashtags

This metadata richness supported the dataset's multi-dimension observation prior to text embedding, and it laid the groundwork for additional filtering and preprocessing.

3.1.2 Users

The primary objective of this study was to derive quantitative insights through social network analysis, focusing on the semantic relationships among influential figures in the discourse. As a consequence, public statements from selected relevant individuals were also considered in the study. All of them were retrieved from official sources, such as governmental websites or uploaded videos of press conferences.

To ensure that the sample of selected individuals was well-balanced and capable of representing diverse viewpoints on the debate, coherent principles guided their selection. A key consideration was that these figures should have been in positions enabling their statements to shape, reflect, or amplify different facets of the discussion. Firstly, a primary selection criterion identified individuals holding top-level executive positions. Presidents, Prime Ministers, or Chancellors in countries with significant influence on global trade dynamics, including members of the G7, G20 and other leading economies, were consequently selected (Table 3.1.2). Clearly, those countries whose leaders did not posses an X profile could have not been taken into account.

Table 3.1: List of Selected Public Figures and Their Twitter Handles

Name	**Title**	**X (Twitter) Handle**
Ursula von der Leyen	President of the European Commission	`@vonderleyen`
Emmanuel Macron	President of France	`@EmmanuelMacron`
Donald Trump	Former President of the United States	`@realDonaldTrump`
Giorgia Meloni	Prime Minister of Italy	`@GiorgiaMeloni`
Olaf Scholz	Chancellor of Germany	`@OlafScholz`
Viktor Orbán	Prime Minister of Hungary	`@PM_ViktorOrban`
Javier Milei	President of Argentina	`@JMilei`
Narendra Modi	Prime Minister of India	`@narendramodi`
Recep Tayyip Erdoğan	President of Turkey	`@RTErdogan`
Shigeru Ishiba	Prime Minister of Japan	`@shigeruishiba`
Luiz Inácio Lula da Silva	President of Brazil	`@LulaOficial`
Cyril Ramaphosa	President of South Africa	`@CyrilRamaphosa`
Keir Starmer	Prime Minister of the UK	`@Keir_Starmer`
Mark Carney	Prime Minister of Canada	`@MarkJCarney`
Anthony Albanese	Prime Minister of Australia	`@AlboMP`
Prabowo Subianto	Prime Minister of Indonesia	`@prabowo`
Claudia Sheinbaum	Prime Minister of Mexico	`@Claudiashein`
Pedro Sánchez	Prime Minister of Spain	`@sanchezcastejon`

Subsequently, public figures whose commentary demonstrated a clear impact on how tariff related issues are framed, discussed, or perceived in broader public and policy arenas were taken into consideration. Considering his global influence and often unpredictable public commentary, **Elon Musk** constitutes a pertinent example of this class. Moreover, his significant activity on X further justified his inclusion in the analysis.

Additionally, to provide comprehensive coverage of the U.S. position, the main character of this discussion, Vice President **JD Vance** and former Vice President **Kamala Harris** were included. These figures, while not current heads of state, possess considerable influential power and are representative of different clusters of opinion.

Table 3.2: List of Selected Public Figures and Their Twitter Handles

Name	**Title**	**X (Twitter) Handle**
JD Vance	vice-President of the United States	`@JDVance`
Kamala Harris	Former Vice-President of the United States	`@KamalaHarris`
Elon Musk	CEO of SpaceX, Tesla; owner of Twitter	`@elonmusk`

Moreover the Apify Actor `apidojo/twitter-user-scraper` was employed to retrieve the *follower relationships* on the social media platform X among the selected public figures, enabling the construction of the underlying relational graph in figure 3.4 used in subsequent network analysis.

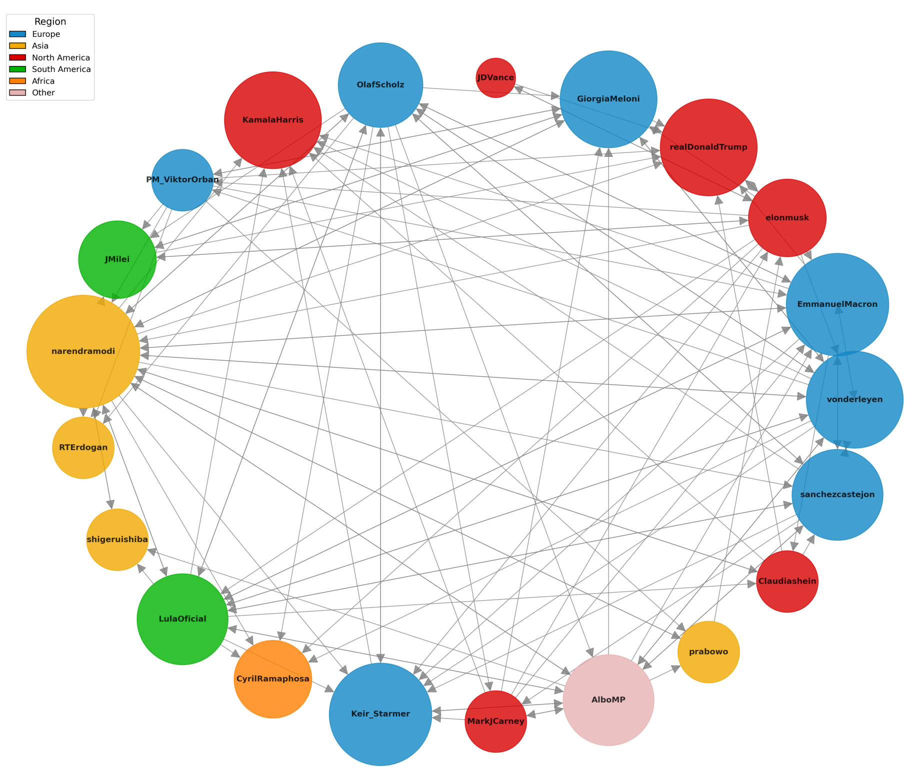

Figure 3.4: Relational Network

The latter will be examined in greater detail in chapter 4, particularly in the discussion of findings from the graph analysis, including centrality measures and community detection.

3.2 Preprocessing

Raw text data can be particularly difficult to embed into dense vector representations due to the fact that it is plagued by a range of inconsistencies that can be damaging to the embedding process. This is doubly so for tweets, as they tend to be authored in informal, platform-specific language and often forego conventional spelling or punctuation. In addition, while emojis, user mentions, hashtags, and hyperlinks may be are informative for certain tasks, they can potentially confound the ability to capture a tweet's underlying semantic message content (Eisenstein, 2013; Baldwin et al., 2013).

To counteract these issues, an extensive preprocessing pipeline was performed before creating semantic embedding and assigning sentiment vectors. Duplicate tweets, as found through matching tweet IDs between overlapping keyword queries, were removed to prevent redundancy. Tweets exceeding 280 characters were excluded since these are typically media sources or influencers and tend to contain summaries or fact-based reporting rather than spontaneous opinion-based tweets. Similarly, AI-generated content, particularly responses by tools like **Grok** or **Perplexity**, and tweets referring to the use of these tools were excluded since it was found that they were adding system-based conversation rather than actual human responses.

Following the removal of noisy or non-organic content, the text was sanitized in a uniform way. URLs, user mentions, and hashtags were removed from the tweet body to improve semantic lucidity, though hashtags were retained in an auxiliary field for potential subsequent analysis. Non-standard characters (including emojis) and redundant white spaces were also removed for maintaining textual consistency and improving the compatibility of the data with former-based embedding models. The tweets that were not in English were translated using the *Gemini 2.0 Flash* model, a light and fast generative model created by Google DeepMind for high-throughput natural language applications (DeepMind, 2024),

The filtering pipeline ensured that raw tweets fetched were all treated identically before being represented. This resulted in a dataset of **32633** filtered texts (including 69 statements from influencial individuals), each ready for semantic and sentiment-level embedding.

Additionally, only the following fields were retained for further analysis:

Table 3.3: Description of columns retained in the cleaned tweet dataset

Column Name	**Description**
`id`	Unique identifier of the tweet
`createdAt`	Timestamp when the tweet was posted
`author`	Username or handle of the tweet author
`clean text`	Preprocessed text of the tweet with URLs, mentions, and emojis removed
`hashtags`	List of hashtags originally included in the tweet
`replyCount`	Number of replies received by the tweet
`retweetCount`	Number of retweets the tweet received
`likeCount`	Number of likes the tweet received
`quoteCount`	Number of times the tweet was quoted
`lang`	Language of the tweet after translation (typically English)
`has link`	Binary indicator for presence of a hyperlink
`has hashtags`	Binary indicator for presence of at least one hashtag
`original lang`	Original language of the tweet before translation

3.3 Semantic–Sentiment Fusion Pipeline

The pipeline employed in this thesis to generate tweet-level vector representations utilizes UMAP in an similar manner to the **Joint Topic–Sentiment** (JTS) framework (Hanny and Resch, 2024). In the JTS model, the dimensionality reduction technique was applied to semantic vector prior to their combination with the sentiment ones, producing low-dimensional representations (5 dimensions). This approach allows for a compact input to clustering algorithms and also balances the contributions of sentiment and semantic features by reducing dimensionality of the latter. In this thesis, instead, dimensionality reduction is used purely for computational efficiency and not as a trade-off for feature importance. Indeed, before the semantic and sentiment vectors are concatenated, their respective dimensions have been weighted explicitly. The following sections describe this pipeline (embedding method, dimensionality reduction, sentiment modeling and the fusion process) in detail.

3.3.1 Semantic Representation via Tweet Embeddings

The embedding step proved to be a critical phase for obtaining consistent and interpretable results across the analysis pipeline. After testing multiple approaches, including open-source models such as *Sentence-BERT* (Reimers and Gurevych, 2019), the model that yielded the most stable and semantically coherent results was OpenAI's *text-embedding-3-large* (OpenAI, 2024). This transformer-based model encodes textual data into high-dimensional vectors, projecting each input into a 3072-dimensional embedding space by default. The model also accommodates variable output dimensionality via a dimensions parameter so that it can be tuned according to task complexity or resource availability. For the work, the full 3072-dimensional representation was preserved in order to maximize semantic opulence, especially considering the stylistic variety and brevity of tweet text.

Compared to earlier models (Sentence-BERT variants), text-embedding-3-large performed superiorly in fine-grained semantic variation and topical coherence at the higher level. These contrasts were especially dramatic when applied to noisy, stylistically diverse short-form text like political tweets. Empirical evaluation included visual inspection of UMAP plots and qualitative checking of cluster coherence, which confirmed the model's ability to more dependably cluster semantically coherent content. These traits rendered it particularly well-suited to the discourse clustering and rhetorical analysis aims of this thesis.

3.3.2 Uniform Manifold Approximation and Projection

The semantic embeddings generated were particularly space-consuming, being projected in a 3072-dimensional space. To enable efficient clustering and analysis without sacrificing the underlying data structure, dimensionality reduction was applied to the embedding matrix.

Through various tests of techniques, including *Principal Component Analysis* (PCA) and *t-Distributed Stochastic Neighbors Embedding* (t-SNE), the best interpretable and meaningful results were obtained using **Uniform Manifold Approximation and Projection** (UMAP). UMAP (McInnes et al., 2020) is a non-linear manifold learning algorithm that projects high-dimensional data into a lower-dimensional space (e.g., 2D or 80D) while preserving both local proximity and global topology. Unlike PCA (Jolliffe, 2002), a linear approach, and t-SNE (Van der Maaten and Hinton, 2008), focusing on local structure in the cost of global relationships, UMAP offers more balanced representation, especially appropriate for semantic embeddings. Because those capture meaning in terms of angular, not absolute, distance, cosine similarity is better than Euclidean distance to compare them in high-dimensional space. This is in line with the common practice in text representation literature where comparisons based on cosine always perform better than those based on Euclidean for semantic tasks involving sentence-level embeddings (Reimers and Gurevych, 2019; Aggarwal et al., 2001). Consequently, in the process here UMAP plays a supporting role by projecting the embeddings in a space where Euclidean distance is more suitable for downstream processing such as fusion with sentiment vectors and clustering by euclidean distance-based algorithms. For this comparison, UMAP was initialized with $neighbors = 15$ and output dimensionality to 80. The number of neighbors parameter dictates the balance between holding on to fine-grained clusters and more generic topological relationships; a middle-of-the-road value of 15 was used to hold onto the nuances of discourse within a topic-congruent but widely rhetorically and ideologically diverse dataset.

The **80 output dimensions** were chosen to be consistent with existing work like BERTopic (Grootendorst, 2022), which shows that this number of dimensions is enough to preserve semantic fidelity without sacrificing computational efficiency. This reducing dimensionality step was therefore essential to the balance between model tractability and representational depth.

3.3.3 Sentiment Vectors

As the introduction reads, the central aim of this analysis is to design a topic detection technique that is capable of extracting not only semantic meaning but also sentiment direction from text data. Recognizing that affective framing is a significant variable that shapes public discourse, especially in the context of politicized or polarized issues, the study incorporates sentiment as a structural element of the tweet representation. To achieve this, all of the statements or tweets in the dataset were labeled with sentiment scores along the categories negative, neutral, and positive, which reflected the model-assigned probabilities that the text expresses each of these emotional stances towards the topic material that it is discussing.

For this exercise, we employed the *twitter-roBERTa-base-sentiment model* (Loureiro et al., 2022), a fine-tuned variant of RoBERTa-base pre-trained on approximately 124 million tweets from January 2018 to December 2021. The model was then further fine-tuned for sentiment classification on the **TweetEval** benchmark, a collection of standardized classification tasks developed for evaluating Twitter data. The model outputs a **three-dimensional probability vector**, of which each dimension is the predicted probability that the tweet is negative, neutral, or positive, respectively.

$$\mathbf{v_{sentiment}} = \begin{bmatrix} p_{\text{positive}} \\ p_{\text{neutral}} \\ p_{\text{negative}} \end{bmatrix}$$

where p_{positive} denotes the model-assigned probability that the tweet expresses a positive sentiment, p_{neutral} denotes the probability that the tweet is emotionally neutral with respect to its topic and p_{negative} denotes the probability that the tweet expresses a negative sentiment.

These were subsequently employed as structured sentiment features and appended to the overall representation of every tweet such that downstream clustering algorithms could consider not only what is being discussed but how it is emotionally described.

3.3.4 Concatenation of Semantic and Sentiment Vectors

To keep sentiment features on the same scale as semantic embeddings, both sets of vectors were standardized using the **StandardScaler** function (Pedregosa et al., 2011). This transformation works by relocating each feature to zero mean and unit variance without causing any distortions in the geometry of the resulting combined space, a vital operation for distance-based algorithms like clustering Hastie et al. (2009). After this step, semantic and sentiment vectors were merged to produce a **83-dimensional representation** for every tweet (80-dimensional semantic vector + 3-dimensional sentiment vector). However, because of the uneven dimensionalities of these sub-vectors, naive concatenation risked overpowering the higher-dimensional semantic vector without taking into account the contribution of the sentiment one. To avoid this, a scaling procedure was employed to weight the magnitude of each sub-vector depending on both on its dimensionality and on an empirically defined value. In practice this step is a trivial yet effective method to ensure that, when the Euclidean distance between two data points is computed, certain classes of dimensions are given a heavier weight than others. Each group's scaling factor was determined as the square root of the division of its weight by its dimensionality.

$$\text{scale}_{\text{group}} = \sqrt{\frac{w_{\text{group}}}{d_{\text{group}}}}$$

In comparison to prior models such as the one proposed by Rouvier and Favre (2016b), that performs hidden-layer fusion in deep architectures, this approach takes a linear and interpretable re-weighting of the embedding dimensions. Aside from easing implementation, this permits explicit control over the distance metric used in clustering. Whereas more recent architectures like SKEAFN (Liu et al., 2021) and CDSAWE (Zhu et al., 2023) learn weights of feature fusion dynamically by employing attention or adapting objectives to the domain, ours brings modularity and interpretability—benefits especially useful with unsupervised methods based on distance-based segmentation.

Figure 3.5: Umap Projections by Semantic and Sentiment Weight

As evident from the UMAP representations (Figure 3.5) obtained by reducing the number of dimensions to 2, changing the **relative weighting** between sentiment and semantic vectors matters drastically to the separability and topology of the embedded discourse space. As the sentiment is given more weight, the resultant form comes to be more diffuse and striated with elongated and stacked bands instead of dense, separate groupings. These representations encode the inherent smoothness of sentiment as a space of features: compared to semantic content, which moves in discontinuous steps between theme or topic, sentiment shifts more gradually and smoothly across texts, often along valence gradients. Sentiment embeddings therefore spread less tightly and instead fill out space more homogeneously, leading to dense overlapping regions in the projection.

This is because sentiment vectors, typically encoded as soft probability distributions within affective categories, do not perform as well in creating clusters that are well-separated (Mohammad and Turney, 2013; Rouvier and Favre, 2016a). Their local density, being very high, causes semantic variation to collapse and blurs the topic boundaries. This becomes most egregious in bigger data sets, in which sentiment features dominate local neighbourhoods and cause semantically distinct tweets to cluster in close vicinity by virtue of affective similarity alone.

In contrast, semantic features capture discrete conceptual boundaries such as political actors, issues, rhetorical frames that are more conducive to clean cluster formation (Mikolov et al., 2013). This is easily observable in the most left plots of the figure, where there is higher semantic weighting with tighter, interpretable cluster geometries forming. The clusters are tidy spatially, often with visible inter-cluster margins, which are more desirable for the subsequent process.

Because of this, semantic features were assigned greater weight (0.85 compared to 0.15 for sentiment). This weighting captures the more categorical, discontinuous structure of semantic variation without sacrificing affective subtlety as a secondary, continuous organizing principle. The interplay between them enables affective tone to subtly modulate thematic structure without overwhelming it, producing a projection space both topologically stable and highly interpretable.

3.4 Clustering Social Media Discourse

The first part of the clustering process was aimed at classifying public opinions in the discourse. Different algorithms were initially implemented for such task. However, it is worth to report only two of them,*K-means* and *Hierarchical*, as those are the ones that performed better. On the contrary, density based approaches such as DBSCAN or HDBSCAN did not yield optimal results, probably because of the nature of the data itself that does not present highly dense regions, as visible from the projection of the data in the two dimensional space.

3.4.1 K nearest neighbours

The k-means algorithm, probably one of the most widely used in the clustering field, was employed to group tweet embeddings into distinct sets of similar semantic meaning. K-means operates by iteratively labelling each point with the nearest centroid and subsequently updating centroids to be the mean of the labelled points with the aim of minimizing overall variance in the clusters (MacQueen, 1967). To make structure discovery meaningful,the selected similarity measure was the **Euclidean distance**, since in this setting all the tweet embeddings were already normalized and projected onto a common space using UMAP and then concatenated with sentiment vectors.

The first and most crucial step in the clustering process was determining k, the optimal number of clusters. The **within-cluster sum of squares** (WCSS) and **silhouette score** (Rousseeuw, 1987) were the two measures employed for this purpose. WCSS captures compactness within clusters as it sums up the squared distances between every point and the centroid of the corresponding group. Because this value decreases monotonically with increasing numbers of clusters, the "**elbow method**" is used to identify the point at which further subdivision yields slower decreases. The silhouette score gives a different view by quantifying how well-clustered the groups are. For each point, it calculates the average distance to points of the same and of the closest other cluster, normalized by the larger of the two. The resulting score will range between -1 and 1, with higher values indicating tighter separation and better organized cluster structure. In this study, WCSS and silhouette scores were computed for a value of k ranging from 5 to 14, as larger values are not optimal to track social media discourse. Looking at the joint behaviour, it was possible to identify the best trade off between intra-cluster compactness and inter-cluster separateness.

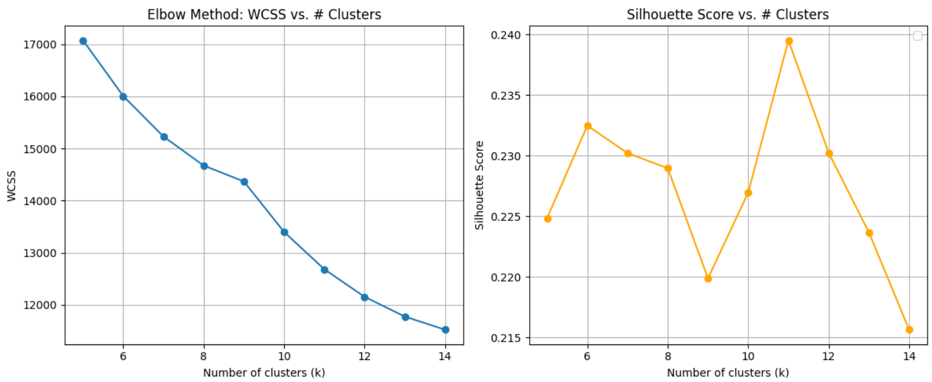

Figure 3.6: WCSS and Silhouette Score by k

Figure 3.6 shows how the WCSS gradually decreases with more number of clusters. A very gentle inflection point is observed around k=10 where the rate of reduction starts to become less prominent indicating a possible "elbow" in the curve. The right-hand plot presents the Silhouette Score on the same scale. The global scores are very low, which is typical in high-dimensional spaces, but local maxima at $k = 6$ and $k = 11$ indicate stronger inter-cluster separation at those values. A sharp drop at $k = 9$ is an indicator of weak cohesion, while further from $k = 11$ the silhouette score started to decrease. Collectively, these diagnostic plots imply $k = 7$ and $k = 10$ are strong candidates. Clustering at $k = 7$ offers an acceptable compromise between separation and tightness. However, qualitatively examined, this configuration was felt to be too coarse. The thematic boundaries tended to blur dissimilar narratives, reducing distinctive ideological framing or sector-specific issues into broader, less distinct groupings. Conversely, $k = 11$ gave a more comprehensive and interpretable clustering that captured a broader scope of visions (economic, strategic, ideological, and sectoral). Such clarification was particularly needed given the richness of the discourse, which involves multiple sectors (e.g., manufacturing, trade sovereignty, international retaliation, domestic populism) and is highly divergent in opinion.

Therefore, the final selection, $k = 11$, was guided not only by these statistical metrics, but also by considering whether clusters were associated with coherent storylines and rhetorically differentiated positions. The combined application of WCSS trends, silhouette stability, and thematic coherence ensured that resulting clustering organization was both mathematically sound and analytically accurate, allowing it to serve as a solid foundation for follow-on analyses.

3.4.2 Hierarchical Clustering

Hierarchical clustering (Ward Jr, 1963) was also employed as an alternative approach. Unlike centroid-based procedures that require pre-specifying the number of clusters and iterative optimization (e.g., k-means), hierarchical clustering constructs a nested group hierarchy by successively merging or splitting observations based on their pairwise distances. This results in a **dendrogram**, a tree-like structure that embodies the whole multi-scale organization of the data and allows to explore at which levels of detail clusters form and evolve. An agglomerative (bottom-up) approach was used in this research, starting with each tweet as an individual cluster and merging over time ever-more similar clusters together until there was a global structure in place. The **linkage criterion** employed was average pairwise distances, intermediary between the sensitivity to chaining of single linkage and the bias towards compactness of complete linkage.

A prominent advance is that hierarchical clustering is especially suited to exploratory social media discourse analysis because it doesn't assume clusters to be spherical, all of the same size, or easily discriminable in advance. Instead, this algorithm outlines the process by which clusters form on scales from fine-grained differences of rhetoric to broad dissimilar of topics and enables flexible post hoc choice of clustering levels.

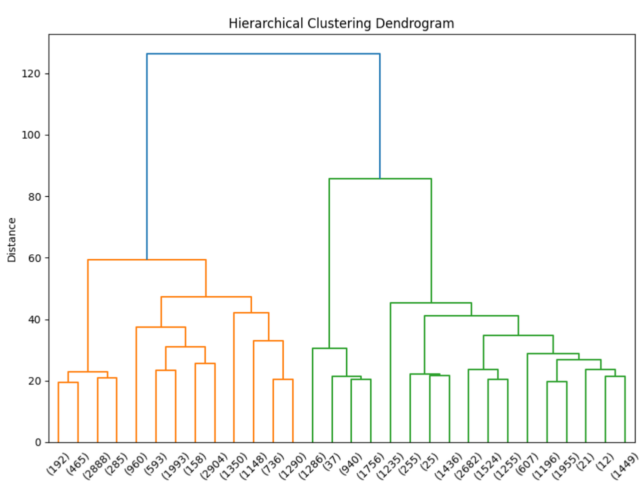

Figure 3.7: Agglomerative Clustering

Considering Figure 3.7 following the branches from top to bottom, it is possible to observe how the smaller subgroups smoothly merge into broader clusters, and two large branches can be observed at the top. These are the two principal points of bifurcation in the space of discourse, and they likely are the two opposing ideological groupings or divergent perspectives. For achieving a finite number of clusters from the dendrogram, a horizontal cut is made at a particular distance measure. In this segmentation, the cut-off was selected to yield **11 clusters**, the choice informed by the prior quantitative diagnostics (elbow method and silhouette analysis) and maintained through qualitative evaluation. Specifically, 11 groups had already been discovered to appropriately capture a balanced proportion of sectoral, ideological, and geopolitical narratives in the dataset.

Having similarity with K-Means-influenced structure at k=11 also allows for *cross-comparison* among cluster algorithms, facilitating triangulation of findings. The dendrogram thus serves not only as a diagnostic tool for identifying topology of discourse, but also as a validation process that confirms the analytical coherence and interpretive comprehensiveness of the ten-cluster solution employed throughout the study.

3.4.3 Characterization and Thematic Labeling

To produce concise and evocative thematic labels, an instruction-refined large language model (LLM)—*Gemini 2.5*, was utilized. The model was seeded with each cluster's set of tweets and then prompted to produce consistent, stance-sensitive labels:

> **Prompt Used for Cluster Labeling (Gemini 2.5)**
>
> You are analyzing a group of tweets that all belong to the same thematic cluster. These tweets are reactions to Trump's dramatic 2025 tariff announcement. The global reaction was intense — that context is already known.
>
> **Your task:** Write a title (4–5 words) that captures what this cluster *argues*, *believes*, or *signals* about the decision. If it is Pro Trump or against him. Pro Europe or else.
>
> **Guidelines:**
>
> - Use clear, direct language.

- Make it punchy, emotionally charged, or rhetorically sharp — not neutral or merely descriptive.
- Reflect the cluster's dominant *stance* or *rationale* (e.g., outrage, approval, strategic logic, economic fear, ideological framing), and the main sector mentioned.
- Avoid hashtags, slogans, or explanatory phrases.
- Do **not** use unclear pronouns: the subject must be specific.
- Do **not** restate that tariffs were announced — assume this is understood.
- Do **not** use the phrase 'Trump's tariffs' — that's already implicit.

Return just the title. No explanations. No formatting. Just plain text.

The prompt was the product of iterative testing by trial and error under the need to extract not just thematic synopses, but also rhetorically evocative and ideologically rooted terms. In its earliest versions, the model consistently produced overly general titles, synopsizing blocs in neutral, descriptive ways (e.g., "Trade Policy Debate", "Global Reactions to Tariffs"). Although factually accurate, such outputs missed the discursive acuity, stance marking, and political construal required to reflect the implicit argumentative organization of each cluster. To address this drawback, a number of significant modifications were made. To begin with, the prompt was grounded in **contextual knowledge**: it made an explicit mention of Trump's 2025 tariff declaration as an established geopolitical context. This dissuaded the model from regurgitating background facts and refocused it on interpretation rather than narration. Second, the model was asked specifically to label what the cluster "argues," "believes," or "signals." This moved the task from value-neutral summarization to **rhetorical analysis**. The prompt also included examples of ideological and emotional positions (e.g., "outrage, approval, economic fear"), already preparing the model to pick up on strategic framings and sector-specific issues, rather than depending on surface-level keywords. Third, the prompt included negative constraints to exclude **unwanted behaviour**, such as using vague pronouns, repeating parrot-like the tariff announcement, or falling back on generic catchphrases such as "Trump's tariffs". These exclusions were extremely useful in enhancing the accuracy and correctness of the output, forcing the model to focus on specific topics and views. Finally, the prompt was subjected to a multi-model review process, which involved GPT-4 and Claude 3 testing, to see that it generated consistent, stance-aware, and ideologically well-informed titles for a broad variety of cluster types.

Because of all these advances, model output quality highly improved. The generated titles began exhibiting clear instances, and sectoral framing that enabled each discourse cluster to be situated within a broader ideological and geopolitical context.

3.5 Clustering Elite-Level Discourse

Although it was theoretically possible to trace **high-order discourse** back to clusters established by public debate on X, this strategy proved to be unviable for analysis. Preliminary attempts at categorizing the statements of influential individuals into common clusters encountered considerable limitations. Existing thematic categories were not ideologically specific or semantically defined enough to satisfactorily pinpoint the rhetorical agendas of political actors. Their work was integrated seamlessly into larger narrative frameworks, making it challenging to detect or examine their communication strategies with precision. Second, the discourse emanating from elites differs in a qualitative way from that of the masses. Unlike spontaneous or off-the-cuff social media postings, public figures' commentary, presented through press conferences, official statements, or other routinised communication channels, is well considered, scripted with care, and in line with institutional norms. Thus, they neglect to convey the rhetorical tone, idiom, or thematic agenda. Consequently, the cluster analysis was conducted once more exclusively in relation to the statements obtained from elite-level individuals alone. Because the underlying processes of semantic-sentiment fusion, dimensionality reduction, and clustering were identical, it would be repetitive to describe it thoroughly in the sections that follow. Instead, this subsection is specifically committed to illustrate the process for the selection of the number of clusters and for the construction of a user-to-user semantic networkNewman (2001).

3.5.1 Number of Clusters Selection

As with the tweet-level analysis, multiple clustering algorithms were applied to the embeddings of statements made by political leaders. However, unlike the public discourse data, where both algorithms converged on a similar structure, the clustering of elite statements produced divergent results across methods. After empirical comparison, hierarchical clustering was selected as the basis for further analysis, as it yielded more interpretable and thematically coherent groupings.

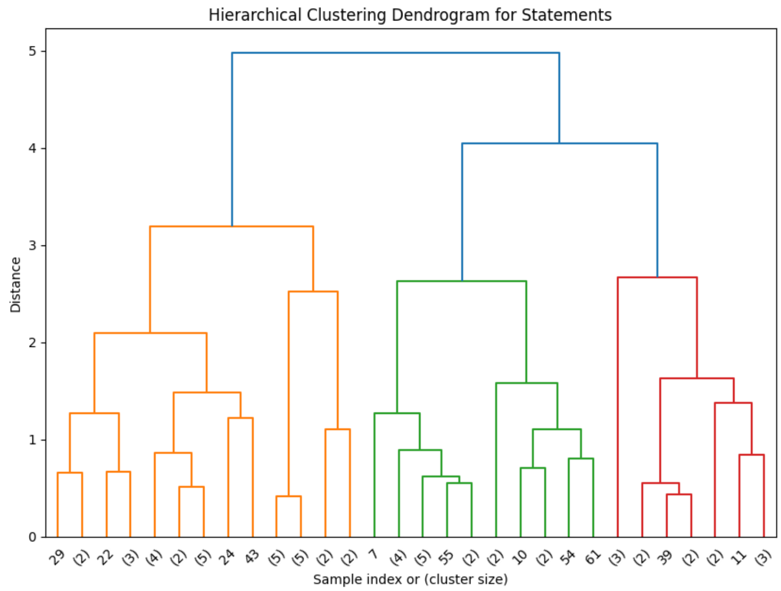

Figure 3.8: Dendogram

The dendrogram in Figure 3.8 illustrates the hierarchical structure of the statements. Based on the distribution of linkages and distance thresholds, seven clusters were identified as the optimal solution. This parameter suggests a reasonable trade-off: it prevents over generalizations and, simultaneously, the decomposition into too narrow categories. Additionally, this value of $k = 7$ yielded the most useful segmentation also from an analytical point of view, with clusters capturing distinct rhetorical and geopolitical narratives.

3.5.2 Semantic Network Construction

To examine rhetorical alignment among powerful political figures, a **bipartite graph** was constructed first. One set of nodes in the graph is users, and the other is the discourse clusters of their statements. An edge between a user and a cluster was created only if more than 30% of the user's total statements belonged to the cluster. This **threshold** ensures that each edge corresponds to a salient rhetorical relation, rather than incidental or marginal interaction with a topic and also ensures that sporadic misidentifications do not have an impact on the general model.

More formally, let:

- N_u be the total number of statements made by user u,
- $N_{u,c}$ be the number of statements made by user u that belong to cluster c.

An edge between user u and cluster c was created if:

$$\frac{N_{u,c}}{N_u} > 0.30$$

From this bipartite network, a **semantic user-to-user** graph was derived. In this projection, each node corresponds to a user, and an edge is established between two users if they were connected to at least one common cluster in the bipartite structure. The intuition behind this projection is that **shared rhetorical affiliations**, expressed through engagement with the same themes, indicate a form of discursive proximity between actors.

The most delicate modelling decision concerned the threshold for edge formation: specifically, how many shared clusters should be required for two users to be considered discursively linked. While a minimum of one shared cluster ensures maximum connectivity, it also risks inflating weak or incidental ties. Conversely, setting the threshold too high may eliminate meaningful but infrequent overlaps. Given the limited size of the elite discourse dataset, and the selective nature of leaders' public statements, a conservative threshold of one cluster was deemed appropriate, especially since most of the noise was dealt with the way the bipartite network was built in the first place.

4 Results

This chapter addresses the main **research questions** guiding the study. First, it analyses the **dominant narratives** that developed during the debate over the 2025 tariff announcement. Specifically, it examines whether the storylines centred on particular economic sectors (e.g., automotive, agriculture) or reflected particular ideological or geopolitical orientations.

Then, the research explores the formation of **communities** within the Relational Network of world leaders on X. This part studies whether clusters that have formed are aligned with the political inclinations of these leaders and how they connect to public knowledge regarding their participation in global trade talks.

Finally, the third part of the chapter illustrates the model's findings when linking themes to specific individuals. It empirically tests cluster coherence by examining whether highly publicized utterances by public figures were appropriately clustered according to **thematic similarity**.

4.1 Clusters Analysis

This section presents the findings derived from the unsupervised clustering approach applied to tweet embeddings, with the aim of revealing prevailing themes in public discourse. The first part discusses the nature of the clusters found, with labels provided by the large language model (LLM). Then, a visual examination of the clustering results is presented to enable intuitive verification of cluster compactness and separation. As expected, the clustering method which yielded the best results was the hierarchical one, considering both a visual interpretation and empirical evidence. As a consequence, only the results obtained with this approach will be illustrated.

4.1.1 Interpretation of Clusters via LLM Analysis

The titles were generated to best capture the rhetorical attitude and thematic content of each cluster. The LLM-assigned labels to the eleven discourse clusters reveal broad thematic and geopolitical scope, including the nuanced and multifaceted reaction to the 2025 tariff announcement. The findings are depicted in the table below.

Table 4.1: LLM-Generated Titles for Discourse Clusters

Cluster #	**Generated Title**
1	American Farmers, Consumers Bleed
2	Tech Rallies, America Prospers
3	Trade War Fuels Market Crash
4	Tech Oligarchs Sway Policy
5	Global Auto Industry Crippled
6	China Trade War: Economic Warfare
7	EU Threatens US Tech Giants
8	China Weaponizes Rare Earths
9	Global Retaliation Targets Boeing
10	Global Trade War: Trump's Folly
11	Europe Warns Trade Disaster

Many clusters articulate the problem in the context of **domestic economic harm**. For instance, "*American Farmers, Consumers Bleed*" (Cluster 1) presents the perspective of a sector heavily hit by the imposition of tariffs, while "*Trade War Fuels Market Crash*" (Cluster 3) depicts the situation from a financial perspective, with a focus on negative stock market performance. Together, the stories portray public anxieties around inflation, consumer expense, and economic instability, nudging towards a general perception of self-destruction due to the policy shift.

Other groups centre on the **tech industry**, which is a constant topic of discussion in the debate. Headlines such as "*Tech Rallies, America Prospers*" (Cluster 2) and "*Tech Oligarchs Sway Policy*"(Cluster 4) indicate a discursive contest between narratives of techno-nationalist hope and stories of corporate influence and regulatory resistance, both domestic and transatlantic. The decision to exempt the technology sector from penalty tariffs was indeed very controversial, and these clusters represent the two dominant interpretations of the act fairly accurately. In the first case, the exemption was interpreted as a prudent and strategic move, maintaining the competitiveness of an essential industry while preserving the symbolic strength of the tariff policy. Conversely, the second framing conveys concern that the ruling suggested excessive political involvement by major tech firms, which some viewed as imposing national policy to advance private interests.

A third set of clusters pertains to **China's geopolitical role**, with titles such as "China

Trade War: Economic Warfare" (Cluster 6) and "China Weaponizes Rare Earths" (Cluster 8), both capturing strategic and rhetorical applications of trade as economic warfare. China is indeed arguably the second most salient actor in this geopolitical framework. The initial, more general cluster represents China as a unilateral power able to increase the confrontation and, by extension, push against Trump's vision of unilateral pressure. In contrast, the second cluster highlights a specific strategic response, China's potential use of rare earths as leverage. These components, essential to the US high-tech industry and predominantly sourced from China, are a strong symbol of the imbalances inherent in global supply chains and raise concerns about the long-term sustainability of US economic resilience.

Aside from broad geopolitical problems, several clusters highlight fairly **sector-specific issues**. For example, "*Global Retaliation Targets Boeing*" (Cluster 9) and "*Global Auto Industry Crippled*" (Cluster 5) highlight the vulnerability of specific industries. The latter, in reality, is predominantly a European issue, as the auto-mobile sector was specifically affected and large producers suffered huge market losses following the announcement. These articles highlight the varied impact of the tariff policy on strategic segments of the economy.

There are also clusters serving as **direct criticism** of the policy's reasoning and implications. "*Global Trade War: Trump's Folly*" (Cluster 10) is arguably one of the most overtly critical, conceivably reflective of a generalized sentiment, especially given that nearly half of the dataset contains non-English tweets. This general cluster serves as an indicator of the model's performance. Indeed, although the majority of clusters express negative sentiments regarding the tariffs, the relatively significant contribution of the model lies in its capacity to separate these responses by sector and theme, distinguishing them from critiques that are not deeply articulated (Cluster 10).

In addition, the model performs optimally in segmenting between the **three big global players**, China, the U.S., and the European Union. The EU common reaction is encapsulated in the cluster "*Europe Warns Trade Disaster*" (Cluster 11) showing global concern about the implications not just for Europe, but even for the United States. Lastly, "*EU Threatens US Tech Giants*" (Cluster 7) encapsulates mounting transatlantic tensions, particularly over the EU's regulatory approach against hegemonic US tech giants. The clusters altogether depict the multi-dimensionality of the discourse and support the model in separate thematic layers within a complex, politically charged environment.

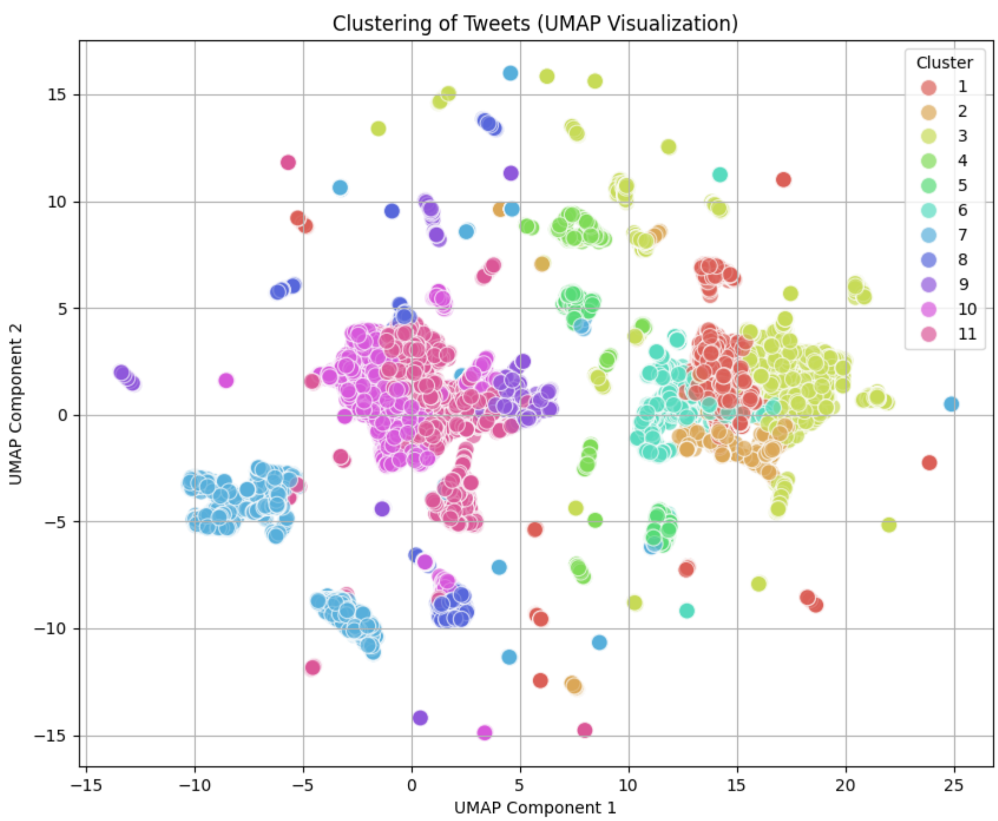

Figure 4.1: 2D Visualizations of Clusters

More generally, the titles register a discursively polarized space, constructed from different perspectives of national interest, global responsibility, and economic consequence. The use of high-emotional-charge, ideologically charged lexis by the LLM (e.g., "crippled," "sway," "folly") underscores the rhetorical extremity of the debate and constitutes a potent interpretive tool with which to regard the clustering output.

4.1.2 Analytical Analysis

Besides their rhetorical significance, how the tweets are distributed across clusters is additional evidence of the relative salience and visibility of each theme in the overall public discourse.

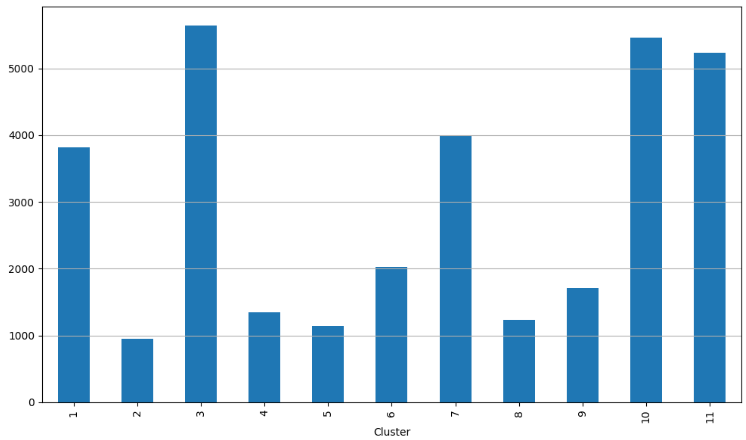

Figure 4.2: Number of tweets in each cluster

The number of tweets is not evenly distributed across clusters, as displayed in Figure 4.2. Groupings 3, 10, and 11 ("*Trade War Fuels Market Crash*", "*Global Trade War: Trump's Folly*", "*Europe Warns Trade Disaster*") are indeed high-density areas of discourse, indicating that public reaction was largely skewed against criticism of the tariff policy. Economic fear and geopolitics of turmoil themes characterized the internet conversation, with the model correctly distinguishing three general framings: European alarm, financial markets' fear, and domestic outrage. Such interpretations are even supported by average sentiment scores per cluster, which all exhibit an high negative score in the largest population groups.

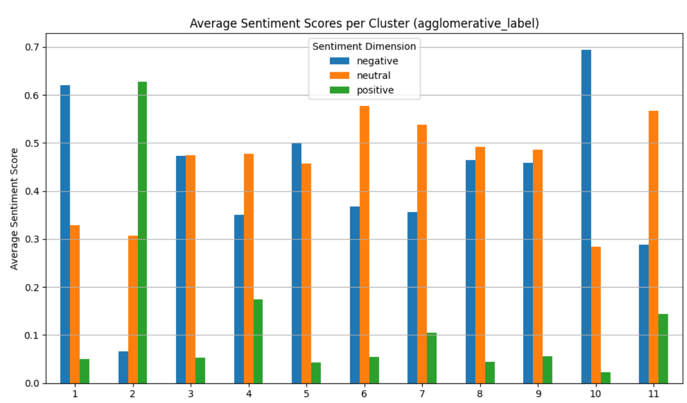

Figure 4.3: Average sentiment in each cluster

Not unexpectedly, the most numerous classes also correspond to the strongest **negative sentiment scores**, reported in Figure 4.3. Clusters 1, 3, and 10, identified with titles like "*American Farmers, Consumers Bleed*" and "*Global Trade War: Trump's Folly*" are characteristic of dominant disapproval, economic, and opposition-to-tariff-policy narratives. These emotions, ubiquitous across the dataset, register discontent from home economic harm to global chaos. Their high volume and tone indicate the public's overall critical response to the action.

Cluster 2, nonetheless, while smallest in tweet quantity, is the only group boasting a positive sentiment profile. The title "*Tech Rallies, America Prospers*" reflects a techno-optimistic sentiment, a speciality account presenting the exclusion of the tech sector as a strategic win. While marginal, the cluster demonstrates a tiny but apparent enclave of favour for Trump's policy, that is, from individuals who see protection of American innovation as justified and a good move. The same, if rather more inconclusive, dynamic is seen in Cluster 11, "*Europe Warns Trade Disaster*," which shows isolated pockets of favourable sentiment. This may signal either a degree of sympathy with the European Union's diplomatic initiative or a latent desire for negotiation and resolution—its framing as a "warning" and not flat rejection keeps possibilities for open-ended constructive interaction.

Other low-volume clusters, such as 5 and 8 reflect sector-specific concerns (e.g., automotive and rare earths), and tend to register neutral or mildly negative sentiment, suggesting technical commentary rather than emotionally polarized rhetoric. Meanwhile, neutral sentiment also dominates Clusters 6, 7, and 9, which, despite higher activity, function more as informational or strategic discursive nodes rather than sites of emotive expression.

Overall, this sentiment-layered analysis reinforces earlier findings: high-volume clusters concentrate around negative or neutral tones, especially in discussions of economic fallout, international retaliation, and elite-driven policymaking. In contrast, smaller or more specialized clusters are where sentiment polarity, either positive or interpretively open, tends to surface. The interplay of sentiment, volume, and thematic framing reveals a discourse environment stratified not just by topic, but also by tone, ideology, and rhetorical function.

4.2 Relational Network

As stated in the previous chapter, follower relations among prominent global political leaders were revealed to facilitate analysis and visualization of their relational network structure. The graph indicates potential rhetorical influence trajectories, ideological clusters and strategic digital alignments.

4.2.1 Graph Visualization

The resultant directed graph is presented below, where node size reflects follower in-degree and node colour represents the regional affiliation.

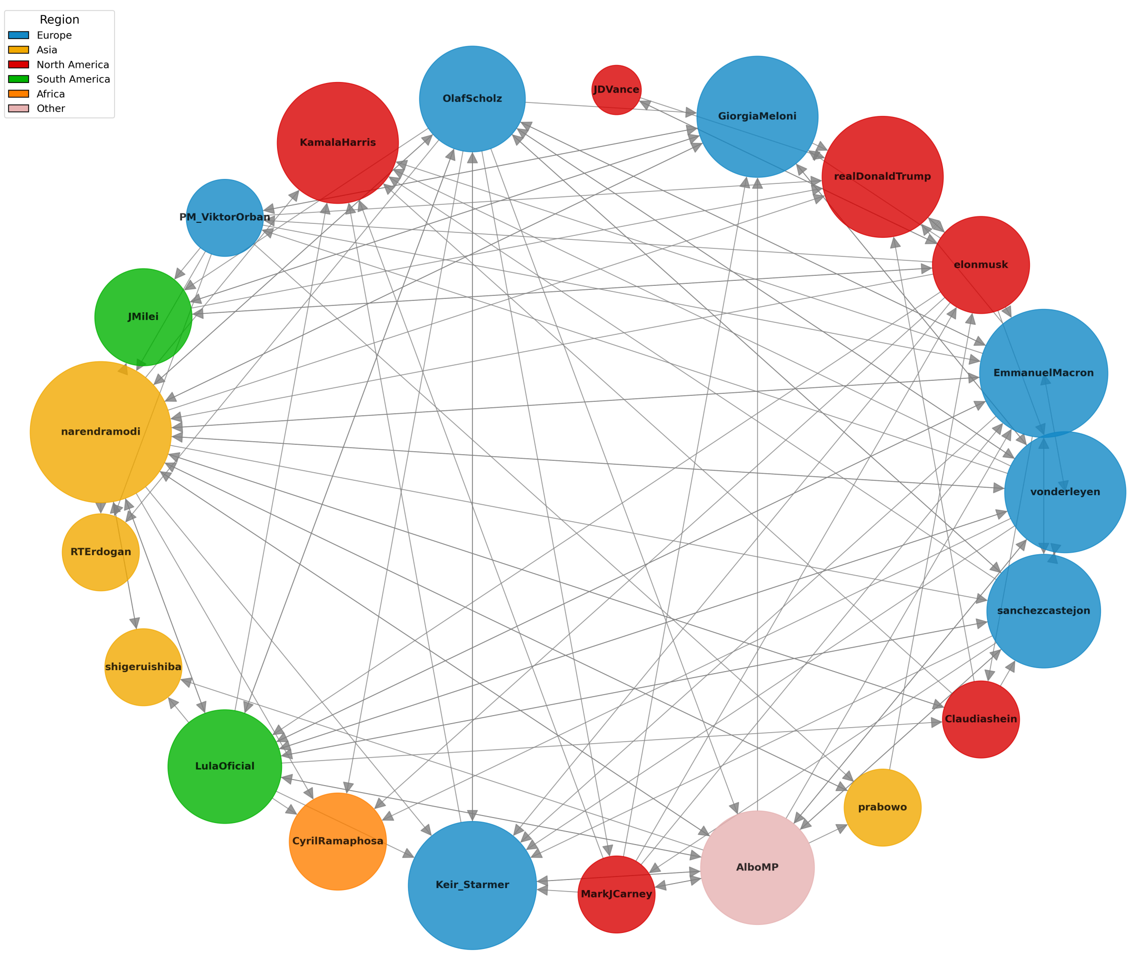

Figure 4.4: Relational Network

The visualization indicates various interesting structural characteristics. The Prime

Minister of India, **Narendra Modi**, is the most followed individual in this network, having the highest in-degree centrality. This result may be unexpected; however, it is consistent with his high activity on the site and international recognition. Donald Trump, on the other hand, is quite disconnected: though he has a high profile, he is not followed very frequently by other leaders in the dataset. Surprisingly, **J.D. Vance** is being followed by **Elon Musk** alone and not by **Trump** himself, even though the current Vice President is Vance. To be precise, the sole individual being followed by the President of the White House is Musk.

It is possible to identify a clearly recognizable European cluster (marked in blue), among which the leaders such as **Von Der Leyen**, **Macron**, **Scholz** and **Sánchez** are very well connected. This pattern is consistent with institutional proximity and policy similarity in the EU setting. **Kamala Harris**, by contrast, has numerous in-links, particularly from European actors, but no out-links, suggesting a role as a symbolic or institutional oppositional node to Trump-aligned actors. **Cyril Ramaphosa** and **Recep Tayyip Erdoğan** similarly lack out-links, which is consistent with little two-way interaction in the virtual space. Erdoğan is only followed by Modi and Scholz, though, maintaining a perception of regional isolation.

4.2.2 Centrality Analysis

Following the topological analysis, a set of **centrality measures** was computed to assess influence, embeddedness, and structural prominence within the leaders' follower network.

Table 4.2: Centrality Measures of Political Figures in the Discourse Network

User	Degree	Betweenness	Closeness	PageRank
`vonderleyen`	0.9500	0.0470	0.5558	0.0561
`EmmanuelMacron`	0.9000	0.0438	0.5780	0.0533
`elonmusk`	0.8000	0.1896	0.4379	0.1066
`realDonaldTrump`	0.4500	0.0222	0.5352	0.0702
`GiorgiaMeloni`	0.6500	0.0315	0.5558	0.0571
`JDVance`	0.1500	0.0000	0.2949	0.0202
`OlafScholz`	0.9000	0.0547	0.4661	0.0499
`KamalaHarris`	0.4000	0.0000	0.5226	0.0588
`PM_ViktorOrban`	0.5000	0.0189	0.4250	0.0343

User	Degree	Betweenness	Closeness	PageRank
`JMilei`	0.4000	0.0154	0.4983	0.0426
`narendramodi`	1.3500	0.2133	0.6283	0.0926
`RTErdogan`	0.1500	0.0000	0.4378	0.0246
`shigeruishiba`	0.2000	0.0000	0.4250	0.0250
`LulaOficial`	0.9000	0.0329	0.5352	0.0471
`CyrilRamaphosa`	0.2500	0.0000	0.5226	0.0367
`Keir_Starmer`	0.6000	0.0111	0.5780	0.0542
`MarkJCarney`	0.4500	0.0173	0.3440	0.0245
`AlboMP`	0.8500	0.0519	0.4817	0.0519
`prabowo`	0.2500	0.0114	0.4250	0.0255
`Claudiashein`	0.3500	0.0014	0.4129	0.0256
`sanchezcastejon`	0.7500	0.0246	0.4817	0.0433

As the centrality table reveals, Narendra Modi is the most central actor on almost all measures: he has the highest degree centrality (normalized by dividing by the maximum possible degree in a simple graph n-1 where n is the number of nodes in G) (Hagberg et al., 2008) reflecting that he is followed by more actors than anyone else, and also ranks first on betweenness and closeness centrality. This indicates that Modi not only has high visibility, but also plays a bridging role, acting as a strategic connector between actors that would otherwise be disconnected. Although this role may be counter-intuitive for him, it is plausibly accounted for by Modi's **intensive and strategic platform use**, which enhances his visibility. European leaders like von der Leyen, Macron, and Scholz also score highly across the board, reflecting both their structural embeddedness in the dense European subnetwork and their symbolic salience in transnational digital diplomacy. By contrast, actors like JD Vance, Cyril Ramaphosa, and Recep Tayyip Erdoğan have consistently low values for centrality, identifying them as marginal or **peripheral actors**, either as a function of low activity, sparse reciprocated connections, or deliberate digital restraint. Interestingly, if we consider PageRank centrality, which controls not only for the number but also the prestige of incoming links, Elon Musk ranks first. Although followed by fewer actors than Modi or the European leaders, Musk is followed by several highly central actors, including Trump, Macron, and Von Der Leyen. amplifying his structural influence.

4.2.3 Community Detection

The application of the **Louvain** algorithm (Blondel et al., 2008) to the relational follower network yields two primary communities, each reflecting distinct ideological and geopolitical alignments. However, several notable deviations from expected groupings, particularly involving European figures, offer important insights into how discursive trust and strategic affinity play out in digital space.

The first community primarily includes centrist to centre-left political figures closely aligned with Western institutions, including the European Union, NATO, and broader rules-based international frameworks. Leaders such as von der Leyen, Macron, Scholz, Sánchez and Keir Starmer, form the structural core of this community.

What is particularly striking is the presence of Kamala Harris within this European-dominated cluster. She receives more connections from European leaders than Donald Trump, in spite of the latter's executive role and ongoing political influence. This asymmetry strongly suggests a widespread loss of symbolic trust in Trump among traditional allies, reinforcing the idea that his rhetorical posture, particularly his anti-EU, anti-NATO, and unilateralist framing, has not only alienated institutional Europe but actively repositioned him outside their digital-political trust networks.

The second cluster consists of figures more sceptical of multilateralism and more supportive of economic sovereignty, protectionism, or non-aligned diplomacy. Here we find Donald Trump, Elon Musk, JD Vance, Viktor Orbán, Giorgia Meloni, Narendra Modi, Javier Milei, and others—actors frequently associated with nationalist rhetoric or challenges to the global liberal order.

The inclusion of Giorgia Meloni and Viktor Orbán in this group is particularly revealing. Both are EU member state leaders, yet their follower alignments reflect a rhetorical and ideological distance from the EU core. Their placement in this cluster suggests that despite institutional belonging, their symbolic alignment is closer to nationalist or unilateralist leaders—especially Trump, whose anti-EU stance resonates with their own scepticism toward Brussels-based governance.

This divergence between institutional affiliation and rhetorical affinity underscores the usefulness of the follower network in capturing alliances that may not be visible through formal diplomacy alone. In particular, Meloni's earlier classification in the semantic network as rhetorically aligned with European leaders contrasts with her structural alignment here—suggesting a complex double position: institutionally embedded in Europe, but digitally proximate to post-nationalist actors.

4.3 Semantic Network

This section shifts the analytical lens from the general public discourse to the rhetorical strategies employed by political leaders themselves. While earlier analyses focused on clustering the broader population's responses, here the goal is to map the semantic structure of elite-level discourse and assess how it aligns with geopolitical positions and digital network formations.

4.3.1 Semantic Clustering of Political Statements

The same LLM-based approach used for labelling the public tweet clusters was applied to the statements made by political leaders. However, the resulting titles demonstrate a qualitatively distinct discursive register. Unlike the broader dataset, which includes emotionally charged and sector-specific reactions, the clusters extracted from elite-level statements exhibit a more strategic and diplomatically framed tone. This reflects the nature of **official communications**, where rhetorical restraint and political signalling are more prevalent.

After that, the cluster were linked to the relevant characters, as described in the previous chapter.

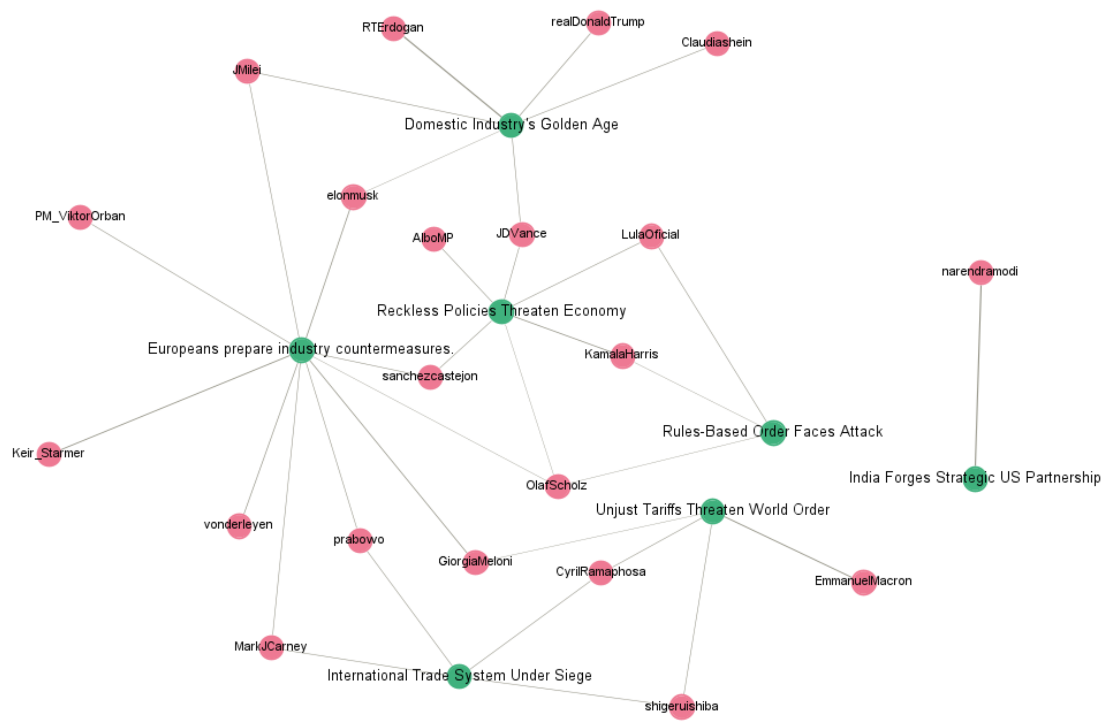

Figure 4.5: Bipartite Graph of Users and Opinions

In the graph in Figure 4.5 each edge represents a strong relationship between a leader (red nodes) and a **discourse cluster** (green nodes). The results confirm the model's validity: the alignments are substantively consistent with the geopolitical stance and rhetorical behaviour of each figure, as publicly registered during the post-tariff announcement period.

Cluster 1, ”*European Prepares Industry Countermeasures*”, with the highest number of connections (10), captures a broad rhetorical alignment among political leaders in response to the U.S. tariff announcement. The cluster includes prominent European leaders like von der Leyen, Macron, Scholz, Sanchez, and Orban, along with non-European leaders like Milei, Carney, and Prabowo. It mostly captures an aligned rhetorical stance, especially among European leaders, in announcing institutional resistance to Trump's policy. In particular, it points to the EU's aspiration to reassert its normative leadership in global trade governance in the context of increasing American economic nationalism.

> "We are already finalising a first package of countermeasures in response to tariffs on steel. And we are now preparing for further countermeasures, to protect our interests and our businesses if negotiations fail."
>
> — Ursula Von Der Leyen, 03/04/2025

But this example also shows a **deficiency in the LLM-generated labelling process**. As the model correctly sensed the usual contrarian tone, the labelling "Europe Defends, Prepares Countermeasures" misleads regarding the reaction as exclusively European. Actually, the presence of some non-European leaders confirms that opposition to the tariff policy reached far beyond the EU, emphasizing the need for human checking of machine-generated results.

Cluster 7, "*Domestic Industry's Golden Age*" also marks an unambiguous ideological boundary, uniting actors who were broadly sympathetic with protectionist and nationalist economic rhetoric. The cluster includes Donald Trump, Elon Musk, and JD Vance, all of whom openly located the 2025 tariffs as forces of industrial revival and national prosperity.

> "With today's action, we are finally going to be able to make America great again, greater than ever before. Jobs and factories will come roaring back into our country and you see it happening already."
>
> — Donal Trump, 02/04/2025

The inclusion of Recep Tayyip Erdogan and Javier Milei is in line with their respective discursive positions. Milei's personal alignment with Trump and his framing of the policy as part of a broader anti-globalist agenda support his placement here. Erdogan's inclusion is consistent with Turkey's favourable treatment under the tariff regime, which was interpreted domestically as a strategic win. The presence of Mexican Prime Minister Claudia Sheinbaum further reinforces the cluster's internal consistency. Her statements delivered realistic endorsement or at least reduced the negative impacts of the policy, projecting an open diplomatic position and not one of opposition.

> "In Mexico's case, there are no additional tariffs and that's good for the country."
>
> — Claudia Sheinbaum, 20/04/2025

Most significantly, it is the only cluster in which **sentiment is strongly positive**, confirming the under-representation of these opinions within total discourse. Its relatively modest size indicates the asymmetry of elite response, in which the majority of leaders were negative or neutral in their evaluations and only a minority were actually supportive of the policy.

Cluster 2, "*Unjust Tariffs Threaten World Order*", includes Macron, Meloni, Ramaphosa, and Shigeru Ishiba, leaders who took diplomatically phrased but substantively critical positions. Their remarks emphasized the erosion of international trade norms, the unilateral nature of the U.S. move, and its dangers to rules-based economic collaboration.

> "The decision announced last night was brutal and unfounded. It's unfounded because you can't correct trade imbalances by imposing tariffs, and the fundamentals of economic theory show the opposite, especially when imbalances do not take into account services, including digital services to cite just one, not others."
>
> —Emmanuel Macron, 03/04/2025

The cluster indicates a stance of multilateral alignment, in which criticism is expressed through appeals to international legitimacy instead of direct ideological confrontation.

Cluster 4, titled "*Reckless Policies Threaten Economy*" contains an inconsistent but cohesive list: Kamala Harris, Olaf Scholz, Lula da Silva, Pedro Sánchez, and Anthony

Albanese. As a group, the leaders defined the U.S. tariff action as irresponsible and destabilizing, often using language that questioned the strategic rationale of the action and warned against unintended global consequences. The cluster suggests both internal and external opposition to Trump's rationale for the action, framing the decision as reckless and fiscally perilous.

> "So look, some people are describing what's been happening in recent months as absolute chaos. And, of course, I understand why, and it's certainly true of those tariffs, tariffs that, as I predicted, are clearly inviting a recession."
>
> — Kamala Harris, 30/04/2025

It is important to note that despite his inclusion in Cluster 4, JD Vance himself has actually strongly reaffirmed his approval for the new tariff stance. Still, due to an unstated rebuke of the previous administration's failure in trade, the model appears to have misinterpreted the tone, classifying the statement as overall negative, when in fact it is actually consistent with protectionist objectives.

> "For 40 years, we have gone down that pathway. We've seen closing factories, we've seen rising inflation. We've seen the cost of housing so high that most Americans can't afford to buy a home right now. "President Trump is taking this economy in a different direction."
>
> — JD Vance, 03/04/2025

Cluster 5, titled *International Trade System Under Siege*, comprises Cyril Ramaphosa, Shigeru Ishiba, and Mark Carney. The latter three focused more on world power rivalry and system weakness rather than bilateral conflict.

> "Currently, we are witnessing how an economic superpower is shaking the world by imposing high tariffs on many countries."
>
> — Prabowo Subianto, 08/04/2025

Their rhetoric placed the U.S. policy not only as an American action, but as a tectonic shift in global trade that put America at the center of the world economy and exposed the weakness of existing supply chains and diplomatic arrangements.

Cluster 3, titled "*Rule-Based Order Faces Attack*", links Kamala Harris, Olaf Scholz, and Lula da Silva, reinforcing a shared framing of the tariff announcement as a threat to global economic well-being.

> "This is an attack on a trade order that has created prosperity all over the globe, a trade order that is essentially the result of American efforts."
>
> — Olaf Scholz, 03/04/2025

Rather than focusing on national or regional grievances, their rhetoric emphasized the systemic harm of protectionist escalation, casting the policy as a dangerous catalyst for global slowdown, trade fragmentation, and macroeconomic instability. This framing positions the tariff move as a collective loss, not a bilateral dispute, an outlook rooted in multilateral economic cooperation and concern for institutional fallout.

Cluster 6, "*India Forges Strategic US Partnership*", goes to Narendra Modi alone. The result is not really surprising since he's the only word leader who did not emerge with clear words on tariff policy.

4.3.2 Projected Graph

Having viewed the clusters of themes in separation, the bipartite network is utilized to construct the anticipated user-to-user network such that semantic proximity between leaders can be explained more effectively.

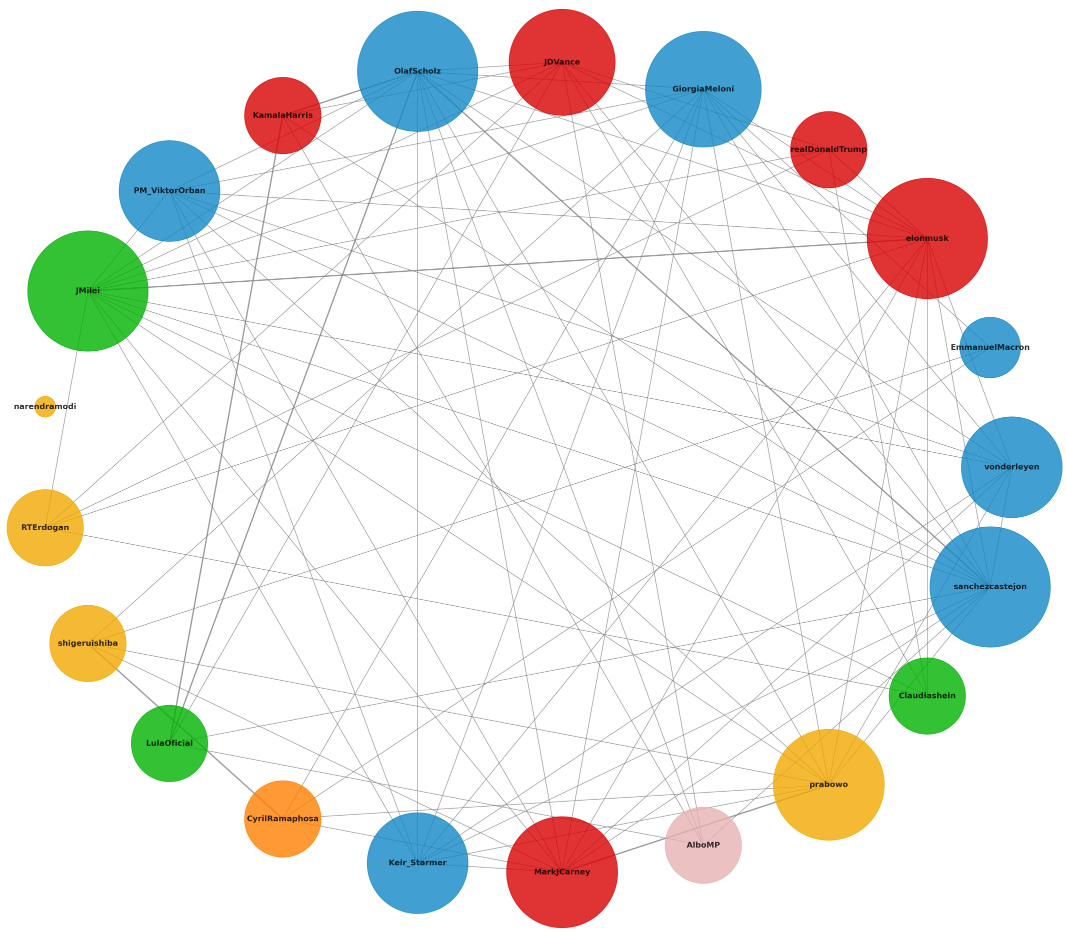

Figure 4.6: Semantic Network

The graph in Figure 4.6 is able to capture not only similarities of an ideological nature, but also discursive alignments that at times diverge from regional or institutional conventions.

Significantly, **Giorgia Meloni** falls in the European grouping, emphasizing her rhetorical solidarity with leaders such as Scholz, von der Leyen, and Sánchez. This is unexpected given her national-conservative image, but she repeatedly denounced the tariff move and said it was harmful to European interests, highlighting the way policy content can overcome ideological differences in framing discursive memberships.

> ”The introduction of tariffs by the United States against the European Union is a measure that I consider wrong and that does not benefit any of the parties involved.”
>
> — Giorgia Meloni, 02/04/2025

On the other hand, Donald Trump is merely linked to those who vehemently supported the tariffs, like JD Vance, Elon Musk, and Claudia Sheinbaum. This narrow rhetorical solidarity demonstrates his **segregated structural position** in the follower network and accentuates the ideological purity of his protectionist discourse.

The case of Elon Musk is less clear-cut. His entanglements with both nationalist actors and EU leaders demonstrate his ambiguity: while making publicly political overtures on the attractiveness of zero tariffs, he also made qualified support for safeguarding American industry. This duplicity mirrors his dual role as both political player and business giant.

> ”At the end of the day, I hope it’s agreed that both Europe and the United States should move ideally, in my view, to a zero tariff situation, effectively creating a free trade zone between Europe and North America.”
>
> — Elon Musk, 05/04/2025

Perhaps most noteworthy is the isolation of Narendra Modi. While the most central node in the network of followers, he stands outside the semantic graph, which means either that he is not engaging substantively with the tariff question or that the global scope of his rhetorical stance is limited. This absence may be explained by the low tariff rate imposed on India, which muted any urgent discursive response (though it is noteworthy that Turkey and Mexico, also lightly impacted, did issue statements). Modi’s silence, therefore, may reflect either strategic restraint or media ecosystem filtering, particularly in Western digital spaces.

5 Conclusions

This thesis explored how political leaders shape and reflect public discourse around tariff policy through both rhetorical content and digital network behavior. By combining semantic clustering and network analysis, it offered a multilayered view of geopolitical communication in the aftermath of the 2025 tariff announcement.

5.1 Limitations of the Study

This study faced three primary limitations: **linguistic**, **computational**, and **interpretive**. First, significant linguistic and accessibility barriers arose when attempting to collect and analyse statements from countries where English is not the dominant language. Many relevant political speeches or public remarks were either not fully translated, infrequently cited by international media, or embedded in difficult-to-navigate platforms, particularly in the cases of Asian, Latin American, and Eastern European leaders. This reduced the representativeness of elite-level discourse in the global sample.

Second, due to computational limitations, it was not feasible to process the full scope of tweet data surrounding the 2025 tariff debate. The dataset had to be constrained in both volume and timeframe, potentially overlooking slower-moving or less visible discourse formations that may have emerged beyond the two-week window.

Finally, while the use of LLM-generated cluster titles and transformer-based sentiment classifiers provided nuanced interpretive tools, these models are not without flaws. The LLM occasionally produced titles that overgeneralized or introduced regional bias, as seen in the mislabelling of non-European actors in the "Europe Defends" cluster. Similarly, the sentiment classifier, although advanced, sometimes failed to capture irony, sarcasm, or layered rhetorical cues, particularly in politically ambivalent statements or those relying on subtext rather than explicit tone.

5.2 Inconsistencies Between Relational and Semantic Networks

Among the primary findings of this study is that there are widespread discrepancies between the structural position of leaders in the relational network and their semantic congruence in rhetorical language. These differences are indicative of serious issues in digital diplomacy, performative politics, and shifting influence dynamics in the platform age.

To cite but one, Giorgia Meloni is relationally placed together with Donald Trump and other nationalist figures in the network, but semantically aligns with the bloc of Europe against the tariffs. This could be a sign of **functional contradiction** – being both a representative of a European institution and constructing transatlantic populist alignment – or strategic placement in order to act as a discursive broker between different parties.

Elon Musk presents a second ambivalence: structurally placed within Trump's network and widely known as one of his strongest backers in the 2024 election, yet semantically less ambivalent about the policy of tariffs itself. Musk's suggestion that "zero tariffs" would be the optimal policy is a statement of business-first pragmatism over against ideological loyalty, a tension between political influence and economic interest.

Narendra Modi, on the other hand, is the most central node of the graph of followers but virtually nowhere to be found in the semantic battle. Perhaps this is to suggest that India's position – insulated from the most acute tariffs – did not demand rhetorical action, or perhaps this is a sign of asymmetries between data availability and access to international media. In either case, it is an argument for interpreting silence as strategically desirable, as opposed to simple absence.

Cumulatively, these cases raise deeper questions. Is there rhetorical inconsistency or duplicity in these leaders' discourse versus actions? Or are they attempting to be located in positions that are no longer institutionally compatible in a sundered global system, requiring simultaneously institutional loyalty and nationalist signalling? Or are they possibly the sole individuals who can pass through polarized blocs precisely because they hold double memberships and adaptable uncertainty?

Surely, it is not possible to answer these questions without a deep understanding of the underlying situation. However, these themes highlights a particularly intriguing area to explore, enabled by the power of mathematical and statistical tools. Quantifying "*how truthful*" a sentence is, based on corresponding actions is not so far from reality.

Bibliography

Aggarwal, C. C., Hinneburg, A., and Keim, D. A. (2001). On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory*, pages 420–434. Springer.

Ahmad, A. N. (2010). Is twitter a useful tool for journalists? *Journal of media practice*, 11(2):145–155.

Angelov, D. (2020). Top2vec: Distributed representations of topics.

apidojo/tweet-scraper (2024). apidojo/tweet-scraper. `https://console.apify.com/actors/61RPP7dywgiy0JPD0/information/latest/readme`.

Baccianella, S., Esuli, A., Sebastiani, F., et al. (2010). Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Lrec*, volume 10, pages 2200–2204. Valletta.

Baldwin, T., Cook, P., Lui, M., MacKinlay, A., and Wang, L. (2013). How noisy social media text, how diffrnt social media sources? In *Proceedings of the sixth international joint conference on natural language processing*, pages 356–364.

Barbieri, F., Camacho-Collados, J., Espinosa Anke, L., and Neves, L. (2020). Tweet-Eval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.

Clausing, K. A. and Lovely, M. E. (2024). Why trump's tariff proposals would harm working americans. *Peterson Institute for International Economics Policy Brief*, pages 24–1.

Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., and Flammini, A. (2011). Political polarization on twitter. In *Proceedings of the international aaai conference on web and social media*, volume 5, pages 89–96.

DeepMind, G. (2024). Gemini 2.0 flash model overview. `https://deepmind.google/technologies/gemini/#gemini-flash`.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding.

Diestel, R. (2005). *Graph Theory*. Springer-Verlag Heidelberg, New York, 3 (electronic edition) edition.

Eisenstein, J. (2013). What to do about bad language on the internet. In Vanderwende, L., Daumé III, H., and Kirchhoff, K., editors, *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 359–369, Atlanta, Georgia. Association for Computational Linguistics.

Federico, L., Mounim, A., Caldarelli, G., and Riotta, G. (2024). Multi-scale analysis of the community structure of the twitter discourse around the italian general elections of september 2022. *Scientific Reports*, 14(1):15980.

Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008). Networkx. `https://networkx.org/`. Version accessed: 2025.

Hanny, D. and Resch, B. (2024). Clustering-based joint topic-sentiment modeling of social media data: a neural networks approach. *Information*, 15(4):200.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition.

Huang, A. et al. (2008). Similarity measures for text document clustering. In *Proceedings of the sixth new zealand computer science research student conference (NZCSRSC2008), Christchurch, New Zealand*, volume 4, pages 9–56.

Hutto, C. and Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 216–225.

Jain, A. K. (2010). Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666. Award winning papers from the 19th International Conference on Pattern Recognition (ICPR).

Jess, N. and Bayhan, H. G. (2024). Visualizing the evolution of twitter (x. com) conversations: A comprehensive methodology applied to ai training discussions on chatgpt. *arXiv preprint arXiv:2407.03484*.

Jolliffe, I. T. (2002). *Principal component analysis for special types of data*. Springer.

J.P. Morgan (2025). Us tariffs: Another round in the trade war? Accessed: 2025-05-28.

Kiel Institute (2025). New u.s. tariffs hit the u.s. itself hardest. Accessed: 2025-05-28.

Kokkinogenis, Z., Filgueiras, J., Carvalho, S., Sarmento, L., and Rossetti, R. (2015). *Mobility Network Evaluation in the User Perspective: Real-Time Sensing of Traffic Information in Twitter Messages*, page 219–234.

Lin, C. and He, Y. (2009). Joint sentiment/topic model for sentiment analysis. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 375–384.

Liu, J., Zheng, S., Xu, G., and Lin, M. (2021). Cross-domain sentiment aware word embeddings for review sentiment analysis. *International Journal of Machine Learning and Cybernetics*, 12:343–354.

Loureiro, D., Barbieri, F., Neves, L., Anke, L. E., and Camacho-Collados, J. (2022). Timelms: Diachronic language models from twitter.

Maas, A., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. (2011). Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting* *of the association for computational linguistics: Human language technologies*, pages 142–150.

MacQueen, J. B. (1967). Some methods for classification and analysis of multivariate observations. In Cam, L. M. L. and Neyman, J., editors, *Proc. of the fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. University of California Press.

Mattei, M., Caldarelli, G., Squartini, T., and Saracco, F. (2021). Italian twitter semantic network during the covid-19 epidemic. *EPJ data science*, 10(1):47.

McInnes, L., Healy, J., and Melville, J. (2020). Umap: Uniform manifold approximation and projection for dimension reduction.

Miao, Y., Yu, L., and Blunsom, P. (2016). Neural variational inference for text processing. In *International conference on machine learning*, pages 1727–1736. PMLR.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space.

Mohammad, S. M. and Turney, P. D. (2013). Nrc emotion lexicon. *National Research Council, Canada*, 2:234.

Morris, J., Lifland, E., Yoo, J. Y., Grigsby, J., Jin, D., and Qi, Y. (2020). Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126.

Newman, M. E. (2001). The structure of scientific collaboration networks. *Proceedings of the national academy of sciences*, 98(2):404–409.

OpenAI (2024). Text embedding models. `https://platform.openai.com/docs/guides/embeddings`.

Otte, E. and Rousseau, R. (2002). Social network analysis: A powerful strategy, also for the information sciences. *Journal of Information Science*, 28:441–453.

Pang, B., Lee, L., and Vaithyanathan, S. (2002). Thumbs up? sentiment classification using machine learning techniques. In *Proceedings of the 2002 Conference on*

Empirical Methods in Natural Language Processing (EMNLP 2002), pages 79–86. Association for Computational Linguistics.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations.

Radicioni, T., Saracco, F., Pavan, E., and Squartini, T. (2021). Analysing twitter semantic networks: the case of 2018 italian elections. *Scientific Reports*, 11(1):13207.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks.

Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65.

Rouvier, M. and Favre, B. (2016a). Sensei-lif at semeval-2016 task 4: Polarity embedding fusion for robust sentiment analysis. In *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, pages 202–208.

Rouvier, M. and Favre, B. (2016b). SENSEI-LIF at SemEval-2016 Task 4: Polarity Embedding Fusion for Robust Sentiment Analysis. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 202–208, San Diego, California, USA. Association for Computational Linguistics.

Severyn, A. and Moschitti, A. (2015). Twitter sentiment analysis with deep convolutional neural networks. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 959–962.

Sia, S., Dalmia, A., and Mielke, S. J. (2020). Tired of topic models? clusters of pretrained word embeddings make for fast and good topics too!

Singhal, A. et al. (2001). Modern information retrieval: A brief overview. *IEEE Data Eng. Bull.*, 24(4):35–43.

Tang, D., Wei, F., Yang, N., Zhou, M., Liu, T., and Qin, B. (2014). Learning sentiment-specific word embedding for twitter sentiment classification. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1555–1565.

Trump, D. J. (2025). Clarification of exceptions under executive order 14257 of april 2, 2025, as amended. `https://www.whitehouse.gov/presidential-actions/2025/04/clarification-of-exceptions-under-executive-order-14257-of-april-2-2025-as-amended/`. Presidential memorandum.

Vakulenko, S., Nixon, L., and Lupu, M. (2017). Character-based neural embeddings for tweet clustering. In Ku, L.-W. and Li, C.-T., editors, *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 36–44, Valencia, Spain. Association for Computational Linguistics.

Van der Maaten, L. and Hinton, G. (2008). Visualizing data using t-sne. *Journal of machine learning research*, 9(11).

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

@vonderleyen (2024). Post on x (formerly twitter). Posted on April 10, 2024. Accessed: 2025-05-28.

Ward Jr, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244.

Weng, J., Lim, E.-P., Jiang, J., and He, Q. (2010). Twitterrank: finding topic-sensitive influential twitterers. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, page 261–270, New York, NY, USA. Association for Computing Machinery.

Wetherell, C., Plakans, A., and Wellman, B. (1994). Social networks, kinship, and community in eastern europe. *The Journal of Interdisciplinary History*, 24(4):639–663.

Zhu, C., Chen, M., Zhang, S., Sun, C., Liang, H., Liu, Y., and Chen, J. (2023). Skeafn: Sentiment knowledge enhanced attention fusion network for multimodal sentiment analysis. *Information Fusion*, 100:101958.