



The Hague Centre
for Strategic Studies

GINA Diplomatic Methodological Notes

HCSS Datalab

Jesse Kommandeur, Rens van Dam, Emma Bokel, Alessandro Dell'Orto
December 2025

GINA Diplomatic Methodological Notes

Authors: Jesse Kommandeur, Rens Van Dam, Emma Bokel, Alessandro Dell'Orto

Contributors: Marit Weurding

Created: July 2024

Updated: December 2025

Cover photo: Canva

This methodological note has been constructed by the HCSS Datalab. The conclusions and recommendations presented in this report are the result of in-house independent research and development. Responsibility for the content rests with the authors and the authors alone.

© The Hague Centre for Strategic Studies. All rights reserved. No part of this report may be reproduced and/or published in any form by print, photo print, microfilm or any other means without prior written permission from HCSS. All images are subject to the licenses of their respective owners.

HCSS
Lange Voorhout 1
2514 EA The Hague

Follow us on social media:
@hcssnl

The Hague Centre for Strategic Studies
Email: info@hcss.nl
Website: www.hcss.nl

GINA Diplomatic in short	5
1. Introduction	6
2. Conceptual Foundation	9
2.1. Speech Mentions	9
2.2. Voting Similarities	13
2.3. State Visits	16
3. Typology and Data Sources	19
3.1. United Nations General Debate Corpus	19
3.2. UN General Assembly Voting Data	19
3.3. COLT Leader Travel Data	20
3.4. Countries Standardization	21
4. Modelling	22
4.1. Speech Sentiment Analysis	22
4.2. Voting Similarity Scoring	23
5. Network Construction	24
5.1. Speeches Network	24
5.2. Voting Network	25
5.3. Visits Network	26
5.4. Visual Representation	27
6. Analysis and Visualization	29
6.1. Voting behaviour	29
6.2. Speech Analysis	30
6.3. Visiting States	31
7. Limitations	32
7.1. Data Limitations	32
7.2. Methodological Limitations	32

GINA Diplomatic in short

The Geopolitical Interactive Network Analysis (GINA) Diplomatic dashboard, developed by the HCSS Datalab in cooperation with Strategic Analysts, applies network science to explore global patterns of diplomatic interactions. Drawing on open-source and systematically maintained datasets, it visualizes how countries are interconnected through rhetorical sentiment, policy alignment, and high-level engagement, providing a structured, data-driven view of the diplomatic dimension of international relations.

Based on data from the UN General Debate Corpus (UNGDC), UN General Assembly Voting Data, and the COLT Leader Travel dataset, the tool integrates decades of records on political discourse, voting behaviour, and state visits. It applies Natural Language Processing (NLP) to speech transcripts to quantify sentiment and utilizes voting similarity scores to model political alignment. By linking these datasets through a network-based framework, GINA Diplomatic reveals the relational structures underlying global diplomatic manoeuvring and alliance formation.

The dashboard is organized into three analytical layers. The Speeches layer visualizes the sentiment and rhetorical focus between states, revealing how countries perceive and address one another. The Voting layer maps policy alignment within the UN General Assembly, highlighting the formation of voting blocs and coalitions. The Visits layer tracks high-level diplomatic travel, distinguishing between occasional meetings and sustained bilateral engagement. Together, these layers offer a multi-dimensional overview of state relations from the rhetorical to the practical level.

While GINA Diplomatic provides a consistent and transparent methodological foundation, its insights depend on the availability and quality of open-source data. The nuances of diplomatic language may be flattened by automated sentiment analysis, while the exclusion of consensus votes and non-publicized state visits may limit the visibility of tacit cooperation or secret diplomacy. As such, the dashboard serves as a conservative but robust analytical baseline for exploring how rhetoric, voting, and visitation patterns intersect within the global diplomatic network.

1. Introduction

The current geopolitical landscape is characterized by fluctuating international tensions and shifting alliances. To understand how these interstate occurrences shape global affairs, analytical tools are helpful for gaining insight into diplomatic interactions between state actors. The Hague Centre for Strategic Studies (HCSS) has developed the **GINA Series**, an initiative by the HCSS Datalab in collaboration with Strategic Analysts, to apply network science frameworks to multiple dimensions of international relations. Drawing on bespoke HCSS- and publicly available datasets as shown in figure 1, GINA examines patterns, interdependencies, and power dynamics between states across four domains: Diplomatic, Military, Economic, and Information. The following methodological note provides context for **GINA Diplomatic** and grants transparency in the creation of the dashboard.



Figure 1 - GINA Suite overview

Each interactive dashboard allows users to explore data through dynamic maps, networks, and graphs across varying time frames and geographic scales, from individual countries to the global level. While the series provides a consistent methodological foundation and employs techniques such as natural language processing and similarity modelling, its insights depend on the quality and completeness of open-source data, positioning it as a conservative yet robust starting point for understanding international dynamics.

As a result of globalization, states have become increasingly dependent on another. Therefore, state governments have a strong interest in maintaining stable relationships within the international community.¹ Displays of collaboration or alignment in bi- and multilateral initiatives are often a result of diplomatic actions, the greatest example being the United Nations General

¹ Jochen Prantl and Evelyn Goh, 'Rethinking Strategy and Statecraft for the Twenty-First Century of Complexity: A Case for Strategic Diplomacy', *International Affairs* 98, no. 2 (2022): 443–69, <https://doi.org/10.1093/ia/iiab212>.

Assembly (UNGA) with its 193 member states.² Cooperation within an organisation of this size, in which a variety of interests are represented, requires strategic diplomatic movements. Diplomatic activity can take various forms. For example, prior research has demonstrated how aspects such as language can reflect the general sentiment of and its impact on relationships between states.³ By providing visualizations of interstate sentiment networks, based on UNGA data of speech mentions, GINA Diplomatic proves to be a useful tool for analysing diplomatic relationships.

General Debate speeches at the UNGA offer insight into the rhetorical dimension of diplomacy that voting records alone cannot capture. While votes reveal positions on specific resolutions, speeches demonstrate how states frame international issues, which actors they choose to mention or ignore, and the sentiment they express toward other members of the international community. This linguistic data of general sentiment among states exposes strategies, diplomatic signalling, and normative positioning that may diverge from voting patterns. For example, a state might vote alongside a partner while maintaining critical rhetoric, or conversely, express support while abstaining on key resolutions. Speech analysis therefore reveals the subjective dimension of interstate relations: how states perceive and publicly characterize one another, which relationships they prioritize through speech mention, and whether diplomatic tone is indicating cooperation or confrontation.

UNGA voting patterns provide a quantifiable measure of policy alignment between states across a comprehensive range of international issues. Unlike declarations or bilateral statements, voting records represent observable positions that states adopt publicly, making them particularly valuable for tracking diplomatic alignment over time. The two-thirds majority requirement for the passing of UNGA resolution causes continuous incentives for coalition building and strategic positioning. Therefore, voting behaviour is also a reflection of diplomatic relationships and influence. By analysing voting similarities, researchers can identify enduring alliances, detect emerging partnerships, and measure the extent to which converging general sentiment turns into coordinated positioning on global issues. This dimension captures the preferences of states: how they position themselves when abstract principles must be turned into practical stances.

State visits represent the most direct form of bilateral diplomatic engagement, reflecting conscious decisions by governments about where to invest high-level political attention and resources. The frequency and level of visits signal relationship prioritization in ways that speeches and votes cannot: a head of state personally traveling to another capital demonstrates commitment beyond rhetorical alignment or procedural voting. Visit patterns reveal which partnerships states actively maintain versus those they only visit sporadically or situationally. Shifts in state visit frequency often occur when states plan on or start discussing formal agreements. Therefore, this dimension captures the intensity and directionality of bilateral engagement, enabling analysis of whether stated partnerships translate into consistent interaction and whether diplomatic shifts manifest in altered visit patterns before becoming noticeable in other dimensions.

² Sascha Werthes, 'The United Nations: In between International and Global Governance', in *Handbook on Global Governance and Regionalism* (Edward Elgar Publishing, 2022), <https://www.elgaronline.com/edcollchap/book/9781800377561/book-part-9781800377561-18.xml>.

³ Fisher, Scott, Graig R Klein, and Juste Codjo. 'Focusdata: Foreign Policy through Language and Sentiment'. *Foreign Policy Analysis* 18, no. 2 (2022): orac002. <https://doi.org/10.1093/fpa/orac002>. Scott Fisher et al., 'Focusdata: Foreign Policy through Language and Sentiment', *Foreign Policy Analysis* 18, no. 2 (2022): orac002, <https://doi.org/10.1093/fpa/orac002>.

GINA Diplomatic creates an interactive analytical environment that applies network analysis to three major open-source datasets on UNGA speech mention and general sentiment, voting similarities, and state visits between member states. The dashboard allows for the visualization and analysis of how states reflect general sentiment of other UNGA members through use of language, how they display alliances via UNGA voting mechanisms and the frequency of state visits. What follows are insights into the interstate relationships created in the international arena, possibly also supporting research on how (economic and military) power influence global diplomatic networks. Therefore, the tool enables users to systematically explore complex interstate relationships through a range of configurable lenses and focus layers, supporting more nuanced insights into the relational dimensions of diplomacy.

The methodological note accompanying the dashboard aims to guide the reader from the conceptual foundations of the GINA Diplomatic framework to its practical application within the dashboard. Each chapter builds on the previous one, moving from theory to data to network construction and finally to analytical interpretation. The structure mirrors the logic of the GINA series and aims to make both the reasoning behind the model and its operational choices transparent.

Section 2 introduces and explains the conceptual foundation of GINA Diplomatic. By defining the concepts of diplomacy and its connection to general interstate sentiment expressed in UNGA speeches, voting similarities in UNGA mechanisms, and state visits, the section provides a basis of understanding the various forms, causes and drivers, and dynamics that influence diplomatic relationships.

Section 3 presents the typology and data sources underpinning the dashboard. It describes the categories used to classify diplomatic relationships via speech mentions, voting behaviour and visits records, outlining how these categories relate to their respective datasets. The section also details the harmonisation of country entities across sources and concludes with a data summary table, offering a concise overview of definitions, coverage and operationalisation.

Section 4 explains how the network layers are constructed from these datasets. It outlines how actors become nodes, how interactions become edges, how weights are assigned and how temporal filters are applied. A dedicated subsection describes the visual logic shared across both layers – layout, colouring and scaling – to ensure users understand how to interpret the structures displayed in the dashboard. The section ends with a summary table of all modelling decisions.

Section 5 introduces the analytical outputs available through the dashboard. It presents examples of insights that can be drawn from the global and state-level networks and discusses the complementary visualisations that support broader interpretation, including maps, quantitative summaries and temporal patterns. This section illustrates how users can interpret the visualization of abstract relationships to meaningful analysis by exploring variation across actors, regions and time.

Section 6 outlines key data and methodological limitations, clarifying how the results should be interpreted and where future extensions are most needed. Readers can either follow the full sequence or consult individual sections as standalone references, depending on whether their primary interest is conceptual framing, data transparency, modelling choices or analytical use.

Taken together, the methodological note is structured to offer both a contextual foundation and a technical roadmap for using the GINA Diplomatic dashboard. Readers can read it to increase their overall understanding of the tool or consult each section as a standalone reference depending on whether their focus is on conceptual clarity, data transparency, methodological choices or analytical interpretation.⁴

2. Conceptual Foundation

This section establishes the conceptual basis for the GINA Diplomatic framework by clarifying how diplomatic activities are understood and applied in this project. Clear definitions of speech mention and general sentiment, voting similarities, and state visits are essential for ensuring that these interactions can be measured consistently and represented accurately within a network-science approach. Section 2.1 outlines the core elements, forms, and drivers of diplomacy through speech mention, while Section 2.2 and 2.3 develop an equivalent foundation for bilateral voting similarities and state visits. Together, these concepts provide the groundwork for the typology in Section 3 and for an overall understanding of how diplomatic behaviour can shape the structure and evolution of interstate relationships.

2.1. Speech Mentions

2.1.1. Definition

Overall, the term diplomacy refers to “the whole of relations between states in today’s international conjuncture”.⁵ State governments deploy diplomatic activities to advance their countries’ interests in the international arena via institutions like embassies, consulates and global initiatives, such as the United Nations (UN). The current manifestation of diplomacy has previously been characterized as preventive, often taking place in the form of consultations and the spread of norms and values through international institutions.⁶

Bi- and multilateral organisations are considered a soft but effective way of advancing diplomatic interests.⁷ Out of all globally operating institutions, the UN has the most member states. A strategically deployed form of soft diplomacy includes openly sharing one’s countries’ position regarding global issues or disapproving another member states’ behaviour within international organisations. During UNGA meetings, where each country gets a chance to speak, representatives chosen to advance domestic interests can make use of diplomatic language, strategically strengthening a countries international image.

By visualizing data concerning general sentiment between states reflected in UNGA speeches, GINA Diplomatic maps the frequency and manner to which one member state refers to other

⁴ Portions of this methodological note were created with support from generative AI tools. In line with the HCSS Maxims for Responsible Use of Generative AI, the use of these tools served only to augment the analytical process, not to substitute expert judgement. All AI generated suggestions were verified for accuracy, authenticity, and integrity. Sweijts, T., Kommandeur, J., and de Ruijter, A. (2024). Augmentation, Not Substitution. HCSS Manual for the Responsible Use of Generative AI.

⁵ Abdurahmanli, Elvin. ‘Definition of Diplomacy and Types of Diplomacy Used between States’. *Anadolu Akademi Sosyal Bilimler Dergisi* 3, no. 3 (2021): 580–603.

⁶ Sayfullayeva, Marjona. ‘The Formation of Modern Global Balance and International Diplomatic Relations: Challenges and Opportunities’. *Академические исследования в современной науке* 4, no. 19 (2025): 79–85.

⁷ Surendra Pokharel, ‘Diplomatic Language: An Analysis of Salutations from Speeches Used in International Diplomacy’, *Journal of International Affairs* 3, no. 1 (2020): 180–93, <https://doi.org/10.3126/joia.v3i1.29094>.

member states. This is conceptualized as the term “speech mention”, overall referring to the general sentiment expressed between UNGA member states. A “positive mention” refers to a reference with positive associations, possibly indicating a good bilateral diplomatic relationship. Vice versa, a “negative mention” includes a state mentioning another state during a speech with negative associations, indicating less stable diplomatic ties.

Diplomatic efforts through language can thus take a variety of forms. State representatives can share a positive attitude toward another state’s behaviour but may also use large scale events to disregard or disapprove other’s behaviour that does not reflect the first states’ norms and values. Another possibility is for states to remain neutral. For example, a country may not want to discuss and/or mention activities of states neglecting international treaties if they still have an interest in maintaining beneficial bilateral ties with this actor.

There are various reasons for a country to deploy diplomatic measures. Usually, motivations are rooted in geo- political or economic spheres. For example, contextual drivers like conflicts or the digital evolution drive countries to improve or establish strong bilateral ties with states that provide the materials critical for such developments.

Since geo-political and economic contexts are constantly prone to change, diplomacy remains a dynamic domain subjected to external drivers that can alter the trajectory of state alliances. Therefore, fragmentation of relationships and the formation of new diplomatic ties are not novel characteristics of the international sphere. It is important to understand the context in which diplomatic relations are formed and what types of disputes could possibly drive them apart. **GINA Diplomatic** gives insight into the frequency of positive, neutral, and negative bilateral mention of UNGA member states across a long period of time, providing a tool to monitor the possibility of shifting alliances.

Previously discussed conceptualizations of diplomacy and “speech mention” form the foundation for the analysis of the various forms, causes and drivers, and dynamics of state referencing during UNGA speeches. The following paragraphs expand on these aspects of the GINA Diplomatic framework while focusing on differing economic and material contexts, as well as global power dynamics and its impact on interstate relationships.

2.1.2. Forms of Speech Mention

Diplomacy exercised via the expressing of general sentiment, conceptualized in speech mention, can be shaped differently depending on the analysed interstate relationship. For analytical clarity, three broad forms are distinguished by the HCSS Datalab: positive, neutral, and negative interstate references. Because of the differing nature of these forms, deployment looks different for each and impacts diplomatic spheres in respectively different manners.

Positive references of one state to another during UNGA speeches involves the consolidation or strengthening of interstate relationships by sharing agreement and alignment. It contributes to positive diplomatic relationships between states and signals a sign of strength in numbers to other nations possibly willing to diverge from commonly shared norms and values.

Neutral references can manifest in the form of refrainment by referring directly to another state when mentioning disapproved activities, as well as overall refrainment from engaging with the practical matter for which a state could be reprimanded. Even when the practical matter contrasts a state’s institutional norms and values, the latter can occur due to a fear of possible consequences. For example, small middle powers may prefer to stay neutral when a strong geo-

economic actor is involved, such as the U.S. or China, because they fear economic consequences after sharing disagreement with one of the global economic powerhouses.

Negative references during UNGA speeches include stating disagreement with other member states or even reprimanding another state's behaviour. While these types of speech mention can worsen diplomatic relationships between the involved actors, it can positively impact other diplomatic ties. If multiple countries reprimand another state actor's actions, common understanding and strength of their interstate relationships can increase. When state representatives have beforehand already discussed and aligned on an issue, chances of sharing negative sentiment during UNGA speeches can be incentivized because of belief to be a part of a majority. Additionally, the reprimanding of a member's behaviour can be motivated by absolute divergence from a state's norms and values.

In practice, states mix and can shift between different forms of diplomatic language. Countries that adhered to a position based on structural drivers may adjust their statements if situational aspects alter internationally operating domains. The distinctions nevertheless provide a coherent typology that aligns with how general sentiment between countries is recorded and visualized. Overall, general sentiment of state actors toward one another can be influenced by structural, situational, and relational drivers.

2.1.3. Causes and Drivers of Speech Mention

General sentiment is shaped by the interaction of structural, situational, and relational drivers. These factors do not determine general sentiment on their own, but together they influence how states speak about each other and how this general sentiment might change over time.

Structural drivers refer to international treaties and historic aspects that tie countries together by shared conditions imposed by international organisations or long-term cooperation shaping internal networks within global ecosystems. For example, strongly integrated and well-respected bi- and multilateral relationships could cause positive general sentiment between states. Therefore, positive referencing to establish and display diplomatic partnerships on a global level can be a consequence of partnerships structurally embedded in global networks via shared treatments and agreements.

Situational drivers are short-term shocks or changes in the global ecosystem that can strengthen diplomatic ties, lead to sudden shifts in alliances, or further diverge states from one another. Examples of situational drivers include military conflicts that worsen already existing geopolitical tensions (e.g., Russia-Ukraine war), and economic measures (e.g., tariffs) imposed to disadvantage another country their position in a competitive market, simultaneously affecting multiple states operating in the same global value chain. In the latter example, negatively impacted states could turn toward one another and share disapproval of such measures in UNGA meetings, increasing negative mentioning of aggressive states.

Relational drivers reflect on how the interconnectedness and interdependence of states in, for example, global production networks can influence general sentiment between states. Diplomatic relationships do not occur in isolation but are often results and parts of a global network of governmental connections, which makes these bonds continuously subjected to change.

Together, these drivers show that general sentiment between state actors reflects both integrated structural conditions and dynamic changes in the global sphere, as well as the

relational architectures in which governments operate. Understanding these factors provides the foundation for analysing how general sentiment evolves over time and how it shapes the broader domain of diplomacy.

2.1.4. Dynamics and Trajectories

The global network of interstate relationships develops and progresses over time, possibly in scale and composition, as well as the strategic incentives that drive countries to seek diplomatic ties. General sentiment between countries can impact the evolution or creation of their interstate relationship. These sentiments can be subjected to change by external and internal drivers. For example, external drivers such as changing behaviours of partner states, or domestically oriented drivers like leadership changes, can impact diplomatic ties between aligned countries.

Cooperation efforts may end or expand due to other states expressing interest for or showcasing capabilities for membership of certain multilateral organisations. An example is the G7, a western oriented forum for economic consultations, that has experienced membership changes over time. Originally, the collective started off with six members in 1975, but within a year, Canada entered as a seventh member.⁸ The purpose of the collaborative was to discuss economic challenges that all members faced in the era of oil crises and currency fluctuations. Therefore, the creation of the G7 illustrates how situational drivers, like economic shocks, can create or strengthen alliances. Simultaneously, the collective demonstrates how situational and relational drivers can negatively influence general sentiment and the trajectory of interstate relationships. Between 1996, shortly after the ending of the cold war, and 2014, Russia had also been a part of the then called G8.⁹ However, Russia's membership was withdrawn after the country illegally annexed the Crimea, resulting in an overall increase in negative sentiment toward Russia by countries diplomatically tied to Ukraine.

2.1.5. Analytical Purpose

Taken together, the conceptual foundations as discussed in this chapter clarify what diplomacy via the expression of general sentiment entails. Moreover, a framework including the forms, causes and drivers, and dynamics and trajectories of general sentiment between state actors has been introduced. This framework provides the analytical basis for understanding international relationships because of globally integrated networks rather than as isolated diplomatic interactions. The following section will lay out a similar explanation for a framework concerning diplomacy via voting similarities in the UNGA.

⁸ Peter I. Hajnal, *The G8 System and the G20: Evolution, Role and Documentation* (Routledge, 2016), <https://doi.org/10.4324/9781315557861>.

⁹ Rafał Matera, 'G8 – Suspension or Fall? From Cold War to Cold War. A Short History of Russian Participation in Major Industrial Democracies', *Humanities and Social Sciences Quarterly*, ahead of print, 2017, <https://doi.org/10.7862/rz.2017.hss.29>.

2.2. Voting Similarities

2.2.1. Definition

Voting similarities as reflected by GINA Diplomatic are conceptualized as con- or divergences in states' positions regarding topics subjected to UNGA voting mechanisms.¹⁰ It concerns the final position of states on voting matters and uses an average of voting results from a variety of resolution types, illustrating the extent of bilateral alignment between member states for these specific types of resolutions. Considering that UNGA general votes require a two-thirds majority for a resolution to pass, creates the need for strategic diplomatic movements between countries, the UNGA is a significant place for diplomats to pursue their countries' interests by interacting with other delegates.¹¹

The UNGA voting mechanism's two-thirds majority requirement ensures that diplomatic efforts carry tangible consequences. States cannot achieve desired outcomes through unilateral positioning alone; they must actively cultivate coalitions and negotiate compromises to secure passage or rejection of resolutions. This structural feature transforms the UNGA into a venue where diplomatic effectiveness becomes empirically observable through voting outcomes. Consequently, shifts in voting similarity patterns serve as indicators of successful relationship-building, revealing which states have strengthened their capacity to influence others' positions and which partnerships have become productive over time.

2.2.2. Forms of Voting and their Similarities

Results of UNGA voting mechanisms can reflect different types of interstate relationships. Voting results of UNGA resolutions can indicate alliances and rivalries between states by displaying (dis)agreements on globally relevant subjects.¹² Moreover, voting results can provide knowledge for states on which other member states might be worth seeking out diplomatic relationships with. Additionally, member states of the UNGA can use the option to abstain from voting on resolutions, which is often considered a statement of neutrality. HCSS conceptualizes these different forms of UNGA voting results as follows:

Alliances and agreement in the context of UNGA voting mechanisms can be interpreted as states creating strength in numbers by using strong diplomatic relationships to ensure a resolution gets either accepted or rejected by the UNGA. Since the mechanism relies on a majority vote, interstate alliances can be effective diplomatic tools to prosper the integration of a state's norms and values into international systems. By showcasing a high number of voting similarities, state actors seemingly have similar norms and values or are intertwined in a bilateral dependency dynamic.

Abstention of voting occurs expressed through 'abstain' in occasions of neutrality. However, countries can also abstain from voting on specific UNGA resolutions when the subject matter is controversial or countries that are diplomatically relevant for the voting state are involved. A member state that is highly dependent on another state for the provision of a specific good or

¹⁰ In UN General Assembly voting, yay (or yes/for) indicates explicit support for the resolution or proposal, nay (or no/against) indicates explicit opposition, and abstain means a state is present and voting but chooses not to take a formal position for or against. Abstentions are often used to signal reservations, maintain neutrality, or balance competing diplomatic interests without directly blocking or endorsing the text.

¹¹ Soo Yeon Kim and Bruce Russett, 'The New Politics of Voting Alignments in the United Nations General Assembly', *International Organization* 50, no. 4 (1996): 629–52, <https://doi.org/10.1017/S0020818300033531>.

¹² Ibid.

service, might not want to vote against a resolution that the other state is in favour of. When a resolution strongly diverges from the voting state's norms and values, but voting behaviour might influence diplomatic ties, a country can choose to abstain from voting to mediate possible consequences.

Rivalry and rejection are reflected in a low number of voting similarities between states. When countries are often on opposite sides regarding voting outcomes, there is the general idea that they have diverging or competing priorities. If this occurs more frequently between specific states, they might start to consider the other as a rival in the international arena, which complicates or prevents the formation of diplomatic relationships. In UNGA meetings, states that are competitors in specific global dimensions can, regardless of domestic values, still consciously choose to disadvantage another state by voting against a resolution that the opposite state is in favour of.

Together, these forms provide a structured way to interpret voting con- and divergences between UNGA member states. They clarify how voting can reflect, and possibly create or change, relationships, providing a foundation for analysing how states position themselves within multilateral diplomatic systems.

2.2.3. Causes and Drivers of Voting Similarities

Voting similarities in UNGA mechanisms are shaped by the interaction of structural, situational, and relational drivers. These factors do not determine voting outcomes on their own, but can influence how states choose and shift positions, consequently impacting how diplomatic relationships develop over time. HCSS conceptualizes these drivers as follows:

Structural drivers that can cause voting similarities are overlap between states in, for example, government organisation, a state's capability to be independent from other states in terms of welfare and safety, or economic drivers like alignments in trade market structure (e.g., open trade). Other structural, integrated drivers that cause voting similarities are shared ideologies, norms and values, or historical alliances.

Situational drivers consist of states lobbying for the passing or rejection of a resolution that would (dis)advantage their countries interests.¹³ Other examples of situational drivers include powerful states leveraging their position in, for example, global value chains, by imposing restrictions in critical sectors, to pressure other countries into aligning with their point of view. Furthermore, sudden global events can create or change national interests, consequently leading to new alliances. For example, states that face similar threats in terms of military conflicts because another actor is showing increased aggression, might find each other in shared opinions regarding arms resolutions.¹⁴

Relational drivers that can lead to voting similarities between UNGA member states include bi- and multilateral agreements that bind countries to act in accordance with certain norms and values. Additionally, states are often in alliances with neighbouring nations, creating regional blocs that share interests and therefore are more likely to show similar voting behaviour.

¹³ Samuel Brazys and Diana Panke, 'Analysing Voting Inconsistency in the United Nations General Assembly', *Diplomacy & Statecraft* 28, no. 3 (2017): 538–60, <https://doi.org/10.1080/09592296.2017.1347450>.

¹⁴ Matthew DiLorenzo and Bryan Rooney, 'International Constraints, Political Turnover, and Voting Consistency in the United Nations General Assembly', *Foreign Policy Analysis* 17, no. 1 (2021): oraa017, <https://doi.org/10.1093/fpa/oraa017>.

Dependency dynamics in multi- and bilateral relations can also cause states to shift positions to secure the continuation of, for example, trade networks and military alliances.

Together, these drivers show that voting similarities within the UNGA can be a consequence of a variety of factors. Understanding these factors provide a valuable contribution to understanding the underlying dynamics of diplomatic relationships and alliances within international networks.

2.2.4. Dynamics and Trajectories

Voting similarity patterns in the UNGA reflect the multifaceted nature of diplomatic ties, shaped by the interplay of structural, situational, and relational drivers. Structural drivers – such as shared political systems, overlapping ideological orientations, or similar positions within global trade networks – create baseline alignment that usually evolves gradually. Situational drivers introduce volatility: crisis events, economic coercion, or shifting threat perceptions can rapidly realign voting behaviour as states recalculate their interests. Relational drivers, including bilateral agreements and regional bloc memberships, create both opportunities for coordination and constraints that bind states to collective positions. By tracking voting similarity trajectories over time, researchers can distinguish between sustainable alignments rooted in structural convergence and temporary coordination driven by situational pressures, thereby revealing which diplomatic relationships rely on deeply integrated commonalities versus strategic opportunities.

The divergences in European responses to China since 2016 illustrates how these drivers interact to shape voting patterns. Central and Eastern European states that joined China's Belt and Road Initiative could have been expected to show increased voting similarity with Beijing on resolutions, suggesting that relational drivers like economic dependency through BRI participation would outcompete structural drivers like shared liberal democratic governance norms. However, research has indicated that these states showed continuous, or even increased, alignment with the EU and U.S., while still maintaining economic ties with Beijing.¹⁵ This indicates that structural and situational drivers (e.g., geopolitical pressure) can counteract relational ones. In this case, the Eastern European countries balanced economic benefits with maintenance of long-term diplomatic relationships.¹⁶ Such situations display how voting trajectories are subjected to strategic choices that try to balance domestic interests with diplomatic relations.

2.2.5. Analytical Purpose

The concepts discussed in this section clarify how voting similarities can reflect diplomatic relationships and operates as a system that connect states through varying forms, drivers, and trajectories. By distinguishing between voting results representing alliances and rivalries, this framework provides a basis for interpreting UNGA voting behaviour as a representation of interstate relationships rather than a collection of isolated transactions. These distinctions make it possible to identify which alliances are seemingly close-knit, reflected by a high number of voting similarities. Simultaneously, divergences indicate which countries may have complicated diplomatic relationships. Such alliances and rivalries may be created and diversified by instances occurring on the global stage, resulting in and reflecting concentrated or strategically significant

¹⁵ Christoph V Steinert and David Weyrauch, 'Belt and Road Initiative Membership and Voting Patterns in the United Nations General Assembly', *Research & Politics* 11, no. 1 (2024): 20531680241233784, <https://doi.org/10.1177/20531680241233784>.

¹⁶ *Ibid.*

dependencies that influence the need for alignment with specific UNGA member states. In the network approach used in this study, each voting interaction becomes a directed and weighted connection between states. This translation allows patterns of erosion, solidification and interdependence to be represented as part of a wider international diplomatic structure that can be compared across actors and time periods.

2.3. State Visits

2.3.1. Definition

The concept of state visits, as visualized by GINA Diplomatic, can be understood as the visiting of a (representative of) head of government and state (HOGS) to another state actor as a form of diplomatic engagement. State visits are overall considered high level events and therefore, depending on each state's significance, receive global attention.¹⁷

Since state visits are often linked to trade negotiations, conflict mediation, alliance signalling, and institutional cooperation, fluctuations in their frequency, direction, or intensity can indicate emerging political realignments or strategic interests.¹⁸ As a result, reviewing the quantity of state visits of various HOGS across different time periods, enables a richer understanding of how evolving diplomatic patterns may influence, or be driven by, changes in global dynamics.

2.3.2. Forms of State Visits

As mentioned, state visits can be motivated by different incentives, causing the content of such visits to change in accordance with the goal of the two HOGS. To gain a better insight into the range of reasons for diplomatic interaction, HCSS distinguishes the following forms of state visits:

Trade negotiations include formal discussions between two nations that are seeking to create or advance their economic partnership. During state visits in the form of trade negotiations, HOGS can for example discuss bilateral rules for trade and possibly create a bilateral trade agreement.

Conflict mediation concerns a HOGS participating as a mediator in discussions to resolve specific disputes between other states. During such state visits, HOGS can help negotiate terms for conflict resolutions by using peaceful negotiation tools.

Alliance signalling are state visits that are meant to strengthen the bilateral ties of two countries, usually during high-profile events. Visits like this do not only send a signal to the two parties involved but also display a strong sign of bilateral alliance and cooperation to the rest of the global community.

¹⁷ Ali Balci et al., 'High-Level Leader Visits: A Promising Area of Study in IR', *International Studies Perspectives*, 19 May 2025, ekaf008, <https://doi.org/10.1093/isp/ekaf008>.

¹⁸ Matt Malis and Alastair Smith, 'State Visits and Leader Survival', *American Journal of Political Science* 65, no. 1 (2021): 241–56, <https://doi.org/10.1111/ajps.12520>; Faradj Koliev and Magnus Lundgren, 'Visiting the Hegemon: Explaining Diplomatic Visits to the United States', *Research & Politics* 8, no. 4 (2021): 20531680211066767, <https://doi.org/10.1177/20531680211066767>.

Institutional cooperation refers to state visits that are focused on the bilateral cooperation between legislative or governmental institutions to advance diplomatic relationships between two states.

Together, these forms provide a structured way to interpret data on state visits between sovereign nations. They clarify how diplomatic visits can reflect, and possibly create or change, interstate relationships, providing a foundation for analysing how states position themselves within the global system.

2.3.3. Causes and Drivers of State Visits

State visits are shaped by the interaction of structural, situational, and relational drivers. While these factors could individually motivate state visits, they could also co-contribute to a HOGS' incentive for visiting another state's government. Overall, state visits can have a great influence on how diplomatic relationships develop over time. HCSS conceptualizes these drivers as follows:

Structural drivers of organising state visits include the solidification of bilateral relationships.¹⁹ For example, countries could aim to strategically enhance political alignments to further embed their diplomatic relationship to one another into the global system. Structural drivers could also take the shape of bilateral agreements and treaties regarding institutional or economic cooperation, making bilateral commitments legally binding.

Situational drivers of state visits are sudden occurrences within the international community that require HOGS of different states to come together and consult one another. Examples of such occurrences can include geopolitical events like the start of interstate military conflicts, calling for negotiation efforts, or economic shocks on the global market to which states want to coordinate responses.²⁰

Relational drivers refer to drivers based on the varying aspects bilateral relationships are built upon. A HOGS may need to visit another state to discuss conditions of formerly created agreements concerning bilateral economic partnerships. Such relational drivers can, for example, lead to a state visit in the form of trade negotiations.²¹

Together, these drivers show that state visits can be driven by a variety of causes, reflecting enduring structural conditions, dynamic shifts in global affairs, and the broader network architectures in which states are embedded. Understanding these factors provides the foundation for analysing why state visits occur and how they can influence or be influenced by global dynamics and trajectories.

2.3.4. Dynamics and Trajectories

Diplomacy exercised in the form of state visits can be shaped and motivated by a broad scale of international developments. Therefore, the frequency and general quantity of diplomatic visits also fluctuate over time. It is therefore important to consider that state visits are not just isolated prompts of international relations but rather are a part of the global interactive network because

¹⁹ Malis and Smith, 'State Visits and Leader Survival'.

²⁰ Koliev and Lundgren, 'Visiting the Hegemon'.

²¹ 'State Visits and International Trade - Nitsch - 2007 - The World Economy - Wiley Online Library', accessed 5 December 2025, <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9701.2007.01062.x>.

of the occurrences that motivate them as well as the general influence they can have on global events.

State visits happen within the dynamic context of international relations. Therefore, motivation for states to initiate diplomatic visits are often a result of interstate disputes or opportunities arising from international collaboration. The interconnectedness and therefore the interdependence between state in the global economy has for example contributed to states seeking out stronger bilateral ties with countries whom they are reliant on for economic prosperity.

The trajectory of bilateral state visits then is impacted by the dynamics that motivate them in the first place, and the results of the interaction between the HOGS. Visits that include constructive and efficient consultations will lead to stronger bilateral ties and a probable increase or continuing quantity of interstate visits. When diplomatic visits result in disputes about, for example, trade conditions, general sentiment between states can decline. This in turn can cause a decrease in interstate visits because HOGS decide to no longer pursue bilateral agreements with one another. Yet again, because of the various events that happen over time, bilateral relationships are constantly subjected to challenges or opportunities of enhancement, explaining why analysis of data concerning quantity of state visits can show fluctuations rather than a consistency in numbers.

2.3.5. Analytical Purpose

The concepts discussed in this section clarify how state visits can reflect interstate relationships and serve as a diplomatic tool within the international system. As explained, analysis of frequency of state visits can be enriched by understanding the varying forms they can take, and the drivers, dynamics and trajectories that impact interstate relationships. Therefore, the framework provides a basis for understanding state visits as a representation of interstate relationships rather than a collection of isolated transactions. In the network approach used in this study, each state visit becomes a directed and weighted connection between states. Translating the data this way exposes patterns of in- and decreases of the importance and strength regarding bilateral relationships across varying time periods.

Eventually, the combination of this and the previously discussed frameworks allows the translation of these concepts into an operational typology and a network-science approach that identifies diplomatic positions as interconnected nodes and links. This enables GINA Diplomatic to represent several patterns of diplomatic interaction as dynamic networks, allowing for systematic comparison across contexts and time periods.

3. Typology and Data Sources

This section outlines the three core data sources that form the empirical basis of GINA Diplomatic and explains how each captures a distinct dimension of diplomatic behaviour. Section 1.1 introduces the United Nations General Debate Corpus, which provides longitudinal, machine-readable records of states' annual speeches and thus their rhetorical signalling. Section 1.2 describes UN General Assembly Voting Data as a systematic record of foreign policy alignment, detailing how vote types such as yes, no, abstain and non-voting reflect states' formal positions on resolutions. Section 1.3 presents the COLT Leader Travel dataset, which documents high-level diplomatic visits and associated activities, enabling the analysis of physical interaction patterns across states. Together with a standardised country reference system described in Section 1.4, these datasets supply harmonised textual, behavioural and interactional evidence that can be integrated into the network-science framework of GINA Diplomatic.

3.1. United Nations General Debate Corpus

The annual General Debate of the United Nations General Assembly constitutes the most comprehensive global forum for diplomatic communication. Since 1946, representatives from member states have used this platform to articulate their government's positions, signal strategic priorities and persuade the international community. To capture these rhetorical dimensions of diplomacy, GINA Diplomatic utilises the United Nations General Debate Corpus (UNGDC).²² This dataset provides a machine-readable collection of all statements delivered during the General Debate, allowing for the systematic analysis of political text as data.

The dataset covers the period from the first session in 1946 through to the 79th session in 2024. It contains 10,952 individual speeches from 202 distinct countries, including historical states that have since dissolved or merged. As described in section 3.2 for UNGA Voting Data, GINA Diplomatic references each voting session to its primary year of development.

3.2. UN General Assembly Voting Data

Voting behaviour in the United Nations General Assembly (UNGA) constitutes a systematic and universally comparable record of state foreign policy positions. To capture these alignments, this project utilises the **United Nations General Assembly Voting Data**, compiled by the UN Dag Hammarskjöld Library. This dataset provides the authoritative, official record of how Member States have positioned themselves on resolutions adopted by the General Assembly from its inception in 1946 through to the end of the 79th session in September 2025.²³

The dataset covers voting records for resolutions adopted during the General Assembly's regular, special and emergency special sessions. It comprises over 916,000 individual voting entries, with each row representing the decision of a single Member State on a specific resolution. Crucially, the dataset includes only resolutions adopted through a **recorded vote**, excluding those adopted by consensus or without a vote. This distinction limits the data to

²² Slava Jankin et al., 'United Nations General Debate Corpus 1946–2024', version 13.0, 2025, <https://doi.org/10.7910/DVN/OTJX8Y>.

²³ 'United Nations General Assembly Voting Data: Resolutions 1 (11 December 1946) to 79/328 (5 September 2025)', with UN. Dag Hammarskjöld Library, UN, 2024, <https://digitallibrary.un.org/record/4060887>.

instances where states explicitly registered their position, typically occurring when divergent preferences prevent unanimous agreement. Votes on individual paragraphs of draft resolutions or on drafts that failed to be adopted are not included in this collection.

Vote categories are standardised across the entire temporal range of the dataset. Each entry records the specific choice made by a Member State using a four-part typology: Yes (Y), No (N), Abstention (A) and Non-voting (X). The dataset derives these records directly from the UN Digital Library, ensuring that the attribution of votes reflects the official position of the Member State at the time of the resolution.

Temporal alignment is necessary to integrate these records into the annual framework of the GINA series. While the UNGA Voting Data organises information by resolution and session, which often span across calendar years, GINA Diplomatic conducts its analysis on a yearly basis. To reconcile these structures, each voting session is referenced to its primary year of development. This operational choice ensures that voting patterns can be consistently compared with annual indicators from the diplomatic travel and speech domains.

3.3. COLT Leader Travel Data

Diplomatic visits by state representatives serve as a primary mechanism for maintaining international alliances, signalling strategic priorities, and negotiating high-level agreements. To capture the full scope of these physical diplomatic interactions, GINA Diplomatic utilizes the Country and Organisation Leader Travel (COLT) dataset developed by the Frederick S. Pardee Institute for International Futures. This dataset provides a comprehensive and continuously updated record of foreign travel by national leaders and key ministers, enabling the construction of a directed network that visualizes how states project influence and sustain relationships through in-person diplomatic interactions.²⁴

The dataset tracks travel by state representatives from over 200 countries from 1990 to 2024, and is composed by regular consultations of top media outlets as well as stock photo agencies or official government, organisation, and personal social media accounts.

Leader travel is the key variable of interest and it includes a any time a serving head of government or state leaves the recognized borders of the country they represent. This includes travel for official interactions, personal vacations, medical care, and other ad-hoc trips abroad. It also tracks short on-the-ground layovers and stops leaders make in route to their final destination, provided an event takes place that goes beyond stopping for aircraft refuelling, resupply, or brief maintenance. Travel by Foreign and Defence Ministers is coded from 1990 onwards, for the United States, Russia, and China. Finance Ministers, Foreign Ministers, and Defence Ministers travels are coded from 01 January 2018 to present for all G20 countries. Crown Princes from the following countries are coded from 01 January 2018 to present: Bahrain, Brunei, Jordan, Kuwait, Morocco, Oman, Qatar, and the United Arab Emirates.²⁵

²⁴ Jonathan D Moyer et al., 'When Heads of Government and State (HOGS) Fly: Introducing the Country and Organizational Leader Travel (COLT) Dataset Measuring Foreign Travel by HOGS', *International Studies Quarterly* 69, no. 2 (2025): sqaf013, <https://doi.org/10.1093/isq/sqaf013>.

²⁵ Jonathan D. Moyer et al., *Country and Organization Leader Travel (COLT) Codebook* (Frederick S. Pardee Institute for International Futures, Josef Korbel School of International Studies, University of Denver, 2024), <https://dataverse.harvard.edu/file.xhtml?fileId=11561869&version=2.0>.

Trip characteristics are recorded alongside the basic travel flow to provide context on the nature of the visit. For every trip, the dataset captures binary variables indicating specific activities performed by the leader. These include whether the leader held a bilateral meeting with the host country's leadership, attended a multilateral summit, or met with leaders from third-party countries. To capture the economic and cultural dimensions of diplomacy, the data further records whether the leader signed an official agreement, met with business leaders or private sector forums, gave a public address, or visited cultural and historical sites. This granular typology allows the GINA Diplomatic framework to differentiate between purely ceremonial visits, security-focused summits, and economic trade missions.

Table 1: Overview of Data Sources

Dataset	What It Measures	Source	Coverage	Level of Analysis	Use in GINA Military
UCDP Organized Violence Data	Records instances of organised violence resulting in at least one direct death.	Uppsala Conflict Data Program (UCDP), Georeferenced Event Dataset (GED).	1989–2024 (annual updates).	Event level (individual violent incidents) and actor level (Side A and Side B).	Events are filtered by user-selected years. Each event maps to an undirected edge between two actors. Edge weights represent the number of events recorded between actors. Actor set includes states, non-state armed groups, and civilians.
SIPRI Arms Transfers Data	Records international transfers of major conventional weapons.	Stockholm International Peace Research Institute (SIPRI), Arms Transfers Database.	1950–2024 (annual updates).	Transfer-level records (supplier → recipient).	Only delivered transfers with valid TIV values are used. Transfers map to directed edges. Edge weights represent total TIV delivered between states in the selected years. Actor set includes only states.

3.4. Countries Standardization

To ensure consistency across datasets and prevent mismatches in naming or country definitions, the GINA series uses an internal reference table as the master standard for all country and state entities. In GINA Diplomatic, all three datasets were parsed and all country entries were automatically matched to this reference; ambiguous or unmatched cases were resolved through manual verification. The table links a unique numeric identifier to a harmonised country name that is used across all pages, charts, and network components. This harmonisation step ensures that actors are represented consistently across time and across data sources, and establishes a coherent actor base before network construction begins.

4. Modelling

This section explains how GINA Diplomatic turns qualitative UN speeches and UNGA voting records into quantitative measures that can be used in network analysis. Section 4.1 describes how speeches from the UN General Debate are processed with Natural Language Processing (NLP): sentences mentioning other states are identified via geopolitically focused named entity recognition, their tone is classified with a transformer-based sentiment model, and these results are aggregated into a General Sentiment Index (GSI) that captures how positively or negatively one state speaks about another. Section 4.2 then explains how UN General Assembly votes are converted into a Voting Similarity Score (S-score), based on the extent to which pairs of states cast the same votes (yes, no, abstain) on the same resolutions within a given year.

4.1. Speech Sentiment Analysis

To transform the qualitative data of the UN General Debate Corpus into quantitative indicators of diplomatic relations, GINA Diplomatic employs a multi-stage Natural Language Processing (NLP) pipeline. The objective of this modelling phase is to isolate directed diplomatic rhetoric, determining not only which countries a speaker mentions but also the emotional tone and strategic context of those references. This process begins by segmenting each speech into individual sentences, which serve as the fundamental unit of analysis for both entity recognition and sentiment classification.

The identification of target states within these sentences relies on a hybrid Named Entity Recognition (NER) approach. The model utilises the spaCy Geopolitical Entity detection system alongside specialised country-recognition libraries to identify geographic references.²⁶ To ensure high precision, these detected entities are cross-referenced with custom dictionaries and standardised to ISO-3 country codes, which are then referenced to the harmonized system described in section 3.4. This allows the model to resolve variations in naming conventions, such as distinguishing between geographic locations and political state actors, and to attribute specific mentions to the correct node in the network.

Once a country mention is validated, the semantic tone of the containing sentence is evaluated using a transformer-based deep learning model. GINA Diplomatic utilises the Twitter-roBERTa-base sentiment model developed by Cardiff NLP, which is optimised for short-text sentiment classification.²⁷ For every sentence containing a state mention, the model assigns a probability score across three categories: positive, neutral and negative. This deep-learning approach allows the system to parse complex sentence structures and capture the general tone of the context in which a state is discussed, distinguishing between praise, factual citation and condemnation.

²⁶ All linguistic preprocessing is performed with spaCy, an open-source natural language processing library. Its pretrained pipelines provide tokenization, part-of-speech tagging, lemmatization and dependency parsing for English and other languages. See: spaCy documentation, "Linguistic Features: POS Tagging", available at: <https://spacy.io/usage/linguistic-features#pos-tagging>

²⁷ TweetNLP is an integrated platform for Natural Language Processing on social media text, offering task-specific transformer models for sentiment analysis, named entity recognition, emoji prediction and offensive language detection. See: Jose Camacho-Collados et al., "TweetNLP: Cutting-Edge Natural Language Processing for Social Media," EMNLP 2022 Demo Paper, arXiv:2206.14774v3 (2022), <https://doi.org/10.48550/arXiv.2206.14774>.

These sentence-level classifications are aggregated to compute the General Sentiment Index (GSI). This metric provides a normalised measure of the rhetorical warmth or hostility directed from the speaker to the referenced state. The GSI is calculated by subtracting the total count of negative mentions from the count of positive mentions and dividing the result by the total number of all mentions, including those classified as neutral. The resulting score ranges on a scale where values further from zero indicate a stronger polarity in diplomatic tone. It is important to note that while this method offers a robust proxy for public diplomatic stance, it captures the tone of the immediate textual context rather than a nuanced policy opinion and is subject to the inherent limitations of automated sentiment classification in complex political discourse.

4.2. Voting Similarity Scoring

To quantify the degree of political alignment between states, GINA Diplomatic computes a Voting Similarity Score (S-score) for every pair of member states within the United Nations General Assembly. This metric treats the voting record as an observable proxy for state preferences, allowing the model to measure how closely two countries align on global issues within a specific session. The approach operationalises agreement as instances where two states cast identical votes, categorising mutual Yes, mutual No, or mutual Abstention as converging preferences. Conversely, any scenario in which states cast differing votes is categorised as a disagreement. To ensure the measure reflects active diplomatic positioning rather than attendance capacity, instances where one or both states are absent or non-voting are excluded from the calculation.

The S-score is calculated by subtracting the number of disagreements from the number of agreements and normalising the result by the total number of shared votes (the sum of agreements and disagreements). This calculation produces a continuous index ranging from -1 to +1. A score of +1 indicates complete convergence where two states voted identically on every resolution in a given year, while a score of -1 indicates total opposition. A score of 0 suggests that agreements and disagreements are equally balanced.

5. Network Construction

This section describes how GINA Diplomatic translates speeches, voting, and high-level visit records into three relational network layers that capture distinct dimensions of interstate diplomatic interaction. Each layer is built directly from underlying event- or dyad-level data and follows clear rules for defining nodes, edges, weighting, filtering, and temporal aggregation. Section 5.1 details the construction of the speeches network based on country mentions and sentiment in UN General Assembly speeches. Section 5.2 outlines how bilateral voting similarity scores become an undirected voting alignment network. Section 5.3 explains how documented leader travel is transformed into a directed visits network. Section 5.4 describes the shared visual logic applied across all three layers. Together, these procedures ensure transparent, reproducible, and conceptually aligned representations of diplomatic discourse, behavioural alignment, and in-person engagement among states.

5.1. Speeches Network

The speeches network translates country mentions in UN General Assembly speeches into a relational structure that captures how states reference one another in diplomatic discourse. This operationalization builds directly on the UN General Debate Corpus described in Section 3.1 and the sentiment analysis described in section 5.1, which identifies which countries are mentioned in each speech and assigns positive, negative, or neutral sentiment to those mentions.

Node Set. Nodes represent all countries appearing in the speech mentions dataset as either actors (countries delivering speeches that mention others) or targets (countries mentioned in speeches). Coverage follows the dataset's country list with the naming standardisation described in Section 3.5, producing a state-level network consistent with the geographical definitions used across the GINA series. Only countries that made or received at least one mention during the selected temporal period are included in the node set.

Node Aggregation and Weighting. Node size reflects the total mention activity of each country, calculated as the sum of outgoing mentions (mentions made by the country in its speeches) and incoming mentions (mentions received from other countries' speeches). Countries that are highly active in diplomatic discourse, either as frequent speakers mentioning others or as frequently mentioned targets, appear larger in the visualisation.

Mention-to-Edge Translation. Each filtered mention record becomes a directed interaction between two states. Direction follows the discourse logic defined in the data structure: edges run from the actor country (the state delivering the speech) to the target country (the state mentioned in that speech). A mention of China in a speech by Brazil, for example, creates a directed edge from Brazil to China. Records lacking complete dyadic information or representing self-mentions (countries mentioning themselves) are excluded from network construction.

Edge Aggregation and Weighting. All mention records for the same ordered state pair within the selected time period are aggregated into a single directed edge. Edge weight equals the cumulative number of mentions between that pair across all years in the selection, with separate tracking of positive, negative, and neutral sentiment counts. Multiple mentions in different speeches or different years all contribute to the relationship weight.

Temporal Selection. Temporal filtering occurs before graph construction. Users may select a single year or a multi-year interval, and all mentions within the chosen period are aggregated into the network. When a single year is chosen, edge weights reflect that year's mention activity. When a multi-year period is selected, edge weights represent the total number of mentions across the entire interval, providing a cumulative view of sustained rhetorical relationships in UN General Assembly speeches.

Sentiment and Threshold Filtering. Users may apply two types of filters to focus on specific mention patterns. A sentiment filter allows visualization of all mentions, only positive mentions, or only negative mentions, enabling analysis of cooperative versus conflictual rhetoric. A minimum mention threshold filters out infrequent relationships, displaying only country pairs where the cumulative mention count meets or exceeds the specified value. When sentiment filtering is active, the threshold applies to the filtered sentiment category rather than total mentions. This dual filtering mechanism allows for cleaner visualizations that isolate substantive patterns of diplomatic attention and the affective dimensions of international discourse.

5.2. Voting Network

The voting network translates UN General Assembly voting records into a relational structure that captures how member states align or diverge in their voting behaviour. This operationalization builds directly on the voting data described in Section 3.2 and the S-scores computation described in section 4.2, where S-scores quantify the degree of similarity between each pair of countries across resolution votes within a session.

Node Set. Nodes represent all UN member states appearing in the voting records within the selected temporal period. Coverage follows the dataset's country list with the naming standardisation described in Section 3.4, producing a state-level network consistent with the geographical definitions used across the GINA series. Only states that participated in at least one recorded vote during the selected sessions are included in the node set.

Node Aggregation and Weighting. Node size reflects the number of edges connected to each country, indicating the breadth of its voting relationships that meet the threshold criterion. Countries with many significant voting alignments or divergences (as defined by the edge threshold filter) appear larger in the visualisation, while those with fewer qualifying relationships appear smaller. This sizing approach highlights states that maintain meaningful voting patterns with a wide range of partners, regardless of whether those patterns reflect agreement or disagreement.

S-score-to-Edge Translation. Each filtered dyadic voting record becomes an undirected interaction between two member states. Unlike the directional logic of the speeches and visits networks, voting similarity is inherently symmetric: if State A votes similarly to State B, then State B votes similarly to State A by the same magnitude. The network therefore uses undirected edges, with each S-score record between two countries represented as a single edge connecting both nodes.

Edge Aggregation and Weighting. All voting records for the same unordered state pair are aggregated into a single undirected edge. When multiple sessions fall within the selected period, the S-score for each country pair is calculated by taking the simple average of that pair's S-scores across all sessions. This averaging treats each session equally, regardless of the number

of resolutions voted upon. Edge weight in the visualisation is determined by the absolute value of the averaged S-score, ensuring that both strong positive alignment (high positive average S-score) and strong negative alignment (high negative average S-score) appear as prominent relationships in the network structure.

Temporal Selection. Temporal filtering occurs before graph construction. Users select a year range, and the system automatically matches these years to their corresponding UN General Assembly sessions, including both regular and special sessions. All S-scores from sessions within the chosen period are aggregated into the network. When a single year is chosen, edge weights reflect that year's voting patterns across all sessions. When a multi-year period is selected, edge weights represent the averaged S-scores across all sessions in the entire interval, providing a cumulative view of sustained voting alignment or divergence.

Threshold Filtering. Users may apply a minimum absolute S-score threshold to focus on politically significant relationships. The threshold operates on the absolute value of the averaged S-score, meaning it captures both strong positive alignments (countries voting together) and strong negative alignments (countries voting oppositely). Only country pairs with absolute averaged S-scores meeting or exceeding the threshold appear in the visualization. Edge colors distinguish alignment type: blue for positive alignment (S-score above 0.5), red for negative alignment (S-score below negative 0.5), and gray for neutral relationships. This allows for cleaner visualizations that isolate patterns of consistent voting behaviour, whether aligned or opposed.

5.3. Visits Network

The diplomatic visits layer translates bilateral travel data into a relational structure that captures how states engage through high-level diplomatic encounters. This operationalization builds directly on the travel records described in Section 3.3, where visits by heads of state and government create documented interactions between countries.

Node Set. Nodes represent all states appearing in the travel dataset as either sending countries (leader's home state) or receiving countries (visited state). Coverage follows the dataset's state list with the naming standardization described in Section 3.4, producing a state-level network consistent with the geographical definitions used across GINA Diplomatic.

Node Aggregation and Weighting. Node size reflects the total diplomatic activity of each country, calculated as the sum of outgoing visits (visits made by the country's leaders) and incoming visits (visits received from other countries' leaders). Countries that engage in high levels of diplomatic travel, either as hosts or as visitors, appear larger in the visualization. This sizing approach captures the overall diplomatic activity level of each state, regardless of whether that activity is primarily outbound, inbound, or balanced. States with minimal travel activity in either direction appear smaller, while those actively participating in bilateral diplomacy through frequent visits appear more prominent.

Visit-to-Edge Translation. Each filtered travel record becomes a directed interaction between two states. Direction follows the travel logic defined in the data structure: edges run from the leader's home country to the visited country. A visit by Bulgaria's head of government to Austria, for example, creates a directed edge from Bulgaria to Austria. Records lacking complete dyadic information (missing leader or visited country identifiers) are excluded from network construction.

Edge Aggregation and Weighting. All visit records for the same ordered state pair within the selected time period are aggregated into a single directed edge. Edge weight equals the cumulative number of visits between that pair across all years in the selection. Multiple visits by different leaders or the same leader on separate occasions all contribute to the relationship weight. This aggregation reveals the sustained intensity of bilateral diplomatic engagement, distinguishing between occasional encounters and patterns of regular interaction over time.

Temporal Selection. Temporal filtering occurs before graph construction. Users may select a single year or a multi-year interval, and all visits within the chosen period are aggregated into the network. When a single year is chosen, edge weights reflect that year's activity. When a multi-year period is selected, edge weights represent the total number of visits across the entire interval, providing a cumulative view of sustained diplomatic relationships.

Threshold Filtering. Users may apply a minimum visit threshold to focus on substantive bilateral relationships. Only country pairs with cumulative visit counts meeting or exceeding the threshold appear in the visualization. This allows for cleaner visualizations, reducing noise and isolating patterns of repeated engagement.

5.4. Visual Representation

All three network layers (Speeches, Voting, and Visits) share a common visual logic to ensure that structural differences between networks reflect the underlying data rather than presentation choices. The layout algorithm and node colouring mechanisms are consistent across layers, providing a basis for comparison.

Layout Algorithm. Nodes are positioned using a ForceAtlas2-based force-directed layout. This algorithm places actors closer together when they share strong or frequent connections and farther apart when their ties are weaker or absent²⁸. This creates a map of the network where clusters emerge naturally from the data: actors involved in many interactions tend to appear near the centre, while those with fewer links move toward the edges. Because all three layers use the same layout rules, differences in how the networks look stem from the underlying relationships rather than from the visual method itself.

Node Colouring. Users can optionally colour nodes by communities or geographic regions. Community detection uses a greedy modularity algorithm, which groups actors that are more closely connected to one another than to the rest of the network.²⁹ Geographic region colouring is based on world regions classified by GINA's country harmonization described in section 3.4. Using the same colouring approaches across all three layers ensures that visual patterns reflect the structure of each network rather than predefined presentation choices.

Edge Colouring. While nodes and layout follow consistent rules, edge colouring varies by layer to reflect the nature of each relationship. In the speeches layer, edge colours indicate overall sentiment: blue for predominantly positive mentions, red for negative, and grey for neutral. In the

²⁸ Mathieu Jacomy et al., 'ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software', *PLOS ONE* 9, no. 6 (2014): e98679, <https://doi.org/10.1371/journal.pone.0098679>.

²⁹ Zhong, J., He, B. and Zhang, X. "Multilevel *k*-means clustering algorithm." *Scientific Reports* 6, Article 30750 (2016). Available at: <https://www.nature.com/articles/srep30750>

voting layer, edge colours show a gradient of the average S-score (voting agreement) across sessions. In the visits layer, edges have no colour variation. This differentiation allows each layer to communicate the qualitative dimension most relevant to its data while maintaining consistency in structural representation.

Table 2: Summary of Network Construction Choices in GINA Diplomatic

Feature	Speech Layer	Voting Layer	Visits Layer
Underlying Data	UN General Debate Corpus (UNGDC)	UN General Assembly Voting Data	COLT Leader Travel Data
Directionality	Directed	Undirected	Directed
Layout	ForceAtlas2-based force-directed layout (shared across layers)	ForceAtlas2-based force-directed layout (shared across layers)	ForceAtlas2-based force-directed layout (shared across layers)

Nodes	Countries that made or received (identified) mentions in UN General Assembly speeches during the selected period	Countries participating to at least one vote in the selected period	Recorded leaders' countries and visited countries in the selected period
Node size	Outgoing mentions + Incoming mentions	Number of edges	Outgoing visits + Ingoing visits
Node colour	Community detection (Greedy modularity algorithm) or World region	Community detection (Greedy modularity algorithm) or World region	Community detection (Greedy modularity algorithm) or World region

Edges	Speech mention records from the actor country (speaker) to the target country (mentioned) nodes	Similarity in voting behaviour across the selected period	Travel records from the leader's country to the visited country nodes
Edge size	Number of mentions from country A to country B in the selected period	Absolute value of the average S-score	Number of visits from country A to country B in the selected time period
Edge colour	Gradient based on dominant sentiment: Blue for positive mentions, Red for negative mentions, Gray for neutral mentions	Gradient of the average (over sessions) S-score in the selected time period	None

6. Analysis and Visualization

This chapter outlines how GINA Diplomatic transforms its processed voting, speech and travel data into substantive analytical insights and visual patterns. Section 6.1 explains how UN General Assembly voting is examined at global, state and dyadic levels, using modularity, centrality, ideal points and voting similarity to assess the evolution of voting blocs, policy alignment and systemic fragmentation. Section 6.2 describes how sentiment extracted from UN General Debate speeches is analysed to track changes in global diplomatic tone, bilateral rhetorical dynamics and the behaviour of key actor groupings such as EU candidates or BRICS members. Section 6.3 focuses on physical diplomacy, showing how patterns of high-level state visits reveal shifting alliances, strategic courtship and the intensity of bilateral engagement.

6.1. Voting behaviour

Voting on a global scale enables analysis of systemic patterns in UN General Assembly voting. The key analytical possibilities include tracking the formation and evolution of voting blocs through the Community Score, which identifies clusters of states that vote together frequently. The Centrality Score reveals how interconnected the overall voting network is at any given time, indicating whether international consensus is fragmenting or consolidating.

The post-Cold War era was initially characterized by unprecedented convergence in UNGA voting, as the ideological cleavages of bipolarity dissolved and a liberal international order appeared. Three decades later, analysts increasingly speak of renewed great power competition and a fracturing multilateral system. The GINA Diplomatic dashboard enables systematic testing of this narrative through its community modularity. By tracking modularity over time, researchers can assess whether distinct voting blocs have indeed re-emerged and hardened since the early 2000s. Periodising the analysis into the unipolar moment (1991-2008), emerging multipolarity (2008-2014), and intensified great power competition (2014-present) allows for structured comparison. Such an analysis would provide empirical grounding for claims about multilateral fragmentation.

At the state level, the Ideal Points Index allows you to track how an individual state's voting position evolves over time, whether it drifts toward or away from particular blocs. The Voting Similarity Score enables dyadic analysis: how closely State A votes with State B across sessions, and how this alignment has changed. One can also examine a state's centrality within the network (whether it is a peripheral outlier or a well-connected node) and its community membership (which voting bloc it cluster with, and whether this has changed).

A central assumption underlying EU enlargement policy is that candidate countries progressively adopt Union positions as part of their accession trajectory, yet the timing and depth of this alignment remains underexplored. Using the voting similarity score, researchers can track how closely current candidates such as Ukraine, Moldova, Turkey, and the Western Balkan states align with EU member voting patterns over time. This enables testing whether diplomatic convergence precedes formal accession negotiations, accelerates during candidacy, or only consolidates after membership. The ideal points index adds further precision by positioning candidates on a continuous scale relative to EU members and other reference groups. For contemporary policy relevance, comparing Ukraine's post-2014 voting trajectory with earlier

candidate pathways (such as Poland or the Baltic states in the 1990s) could reveal whether its alignment is historically typical or exceptionally rapid. Such analysis directly informs debates about enlargement readiness and the substantive meaning of candidate status.

6.2. Speech Analysis

The General Sentiment Index captures the aggregate diplomatic tone across all UNGA General Debate speeches. This allows the tracking of whether global diplomatic discourse is becoming more positive/cooperative or negative/confrontational over time. One can identify years where sentiment shifted notably and correlate these with international crises or breakthroughs. By comparing the General Sentiment Index across Sessions 76–78 (2021–2023), researchers can examine whether the overall tone of multilateral discourse became measurably more negative before the invasion materialized, or whether the shift occurred only in its aftermath. This offers an empirical test of whether diplomatic rhetoric at the UNGA functions as a leading indicator of geopolitical rupture or merely reflects events post hoc.

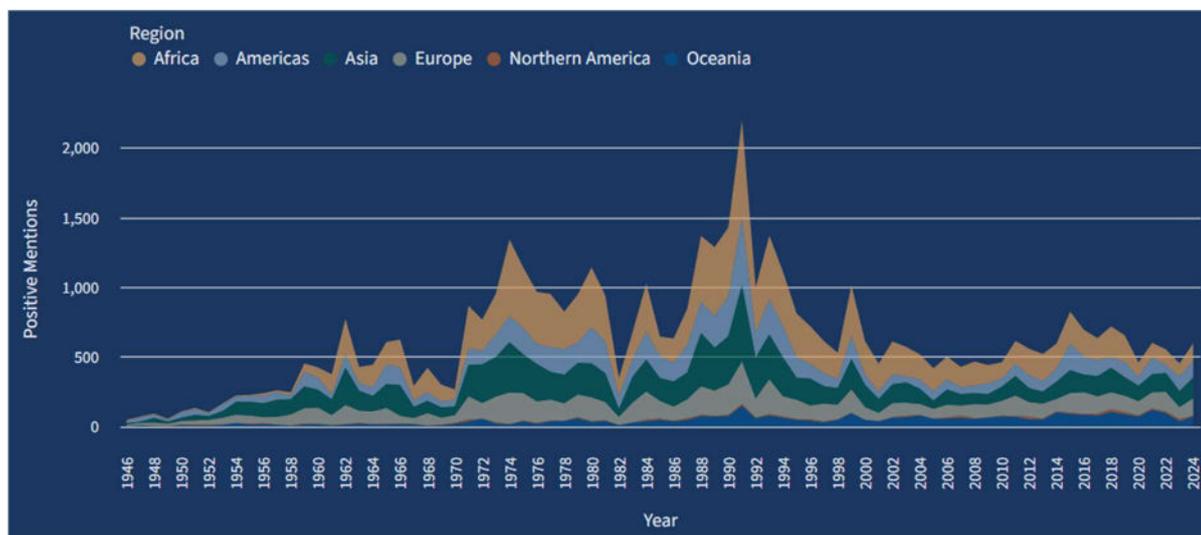


Figure 1 - Positive mentions in UNGD speeches (1946-2023)

The State Sentiment Index enables analysis of how individual states are discussed by others (sentiment toward them) and how they discuss other states in their own speeches. This creates possibilities for examining bilateral sentiment dynamics, is State A's rhetoric toward State B reciprocated? One can also track how a state's diplomatic tone evolves over time or identify which states receive predominantly positive versus negative mentions in the corpus.

The BRICS grouping presents itself as a coherent voice for the Global South, yet the extent to which this institutional partnership translates into genuine diplomatic solidarity, particularly under crisis conditions, remains contested.

The GINA Diplomatic speech analysis tools allow researchers to examine whether China, India, Brazil, and South Africa have maintained, softened, or distanced their rhetorical positioning toward Russia following the 2014 Annexation of Crimea and the February 2022 invasion of Ukraine. By tracking the State Sentiment Index for mentions of Russia and Ukraine in each country's General Debate speeches from 2014 through 2023 (2024, new model), analysts can detect shifts in tone that voting records alone might obscure. A state may vote alongside Russia while rhetorically hedging through neutral or even mildly critical language toward Moscow, or

conversely, express continued partnership with Russia while offering sympathetic framing of Ukraine's situation. The contrast between how each BRICS member discusses Russia versus Ukraine reveals the extent of diplomatic hedging: does Ukraine receive dismissive, neutral, or implicitly supportive treatment, and does sentiment toward Russia remain warm or cool noticeably? This analysis tests whether BRICS functions as a genuine alignment with shared diplomatic commitments or remains a looser coordination mechanism where members preserve independent rhetorical positioning when core interests diverge.

6.3. Visiting States

The State Visit Frequency metric enables analysis of bilateral diplomatic engagement intensity and its evolution over time. By tracking the frequency of visits between heads of state and government representatives, researchers can identify which relationships are being actively cultivated versus maintained at routine levels. This allows testing whether deepening partnerships manifest through increased visit frequency before materializing in formal agreements or institutional frameworks.

Middle powers such as Turkey, Saudi Arabia, and the UAE have pursued increasingly autonomous foreign policies since 2015, yet the extent to which this translates into concrete relationship-building with non-Western powers remains underexplored. By comparing visit frequency patterns between these states and China versus their traditional Western partners (US, major EU states) from 2010 through 2024, analysts can assess whether diplomatic autonomy is rhetorical or substantive. A marked increase in visits with Beijing, while Western engagement declines, would indicate strategic shifts in diplomatic relationships.

7. Limitations

This chapter highlights key constraints in the data and methodological choices underpinning GINA Diplomatic. Section 7.1 outlines limitations in the underlying datasets, including imperfect detection of historical states in speech transcripts, the exclusion of consensus resolutions in UN voting data, and the omission of non-public diplomatic travel. Section 7.2 explains how modelling choices – such as sentence-level sentiment analysis, equal weighting of all UNGA resolutions and annual aggregation – introduce analytical simplifications. Together, these limitations clarify the caution needed when interpreting network patterns and derived indicators.

7.1. Data Limitations

The reliability of the GINA Diplomatic dashboard is intrinsically linked to the nature of the open-source datasets employed. A primary limitation within the UN General Debate Corpus (UNGDC) is the detection of historical state entities. The geopolitical entity detection model is less accurate at identifying states that no longer exist or have undergone significant name changes, particularly in sessions predating the Cold War. Furthermore, the dataset relies on the official transcripts of speeches; if multiple states are mentioned within a single sentence, the current visualization logic may only highlight one actor, potentially underrepresenting complex multilateral references.

Regarding the UN General Assembly Voting Data, the dataset is restricted to "recorded votes." It excludes resolutions adopted by consensus or without a vote – a common practice for less controversial or procedural matters. Consequently, the data may overrepresent moments of division while underrepresenting areas of unanimous international cooperation. Finally, the COLT Leader Travel Data, while comprehensive, relies on public records, media reports, and official press releases. As such, it captures only public diplomacy; secret negotiations, unannounced visits, or travel by lower-level officials not tracked by the dataset are excluded, potentially omitting discreet but significant diplomatic channels.

7.2. Methodological Limitations

The methodological choices made to translate qualitative and quantitative data into network structures introduce specific analytical constraints. In the Speech Sentiment Analysis, the NLP model operates on a sentence-by-sentence basis. This approach ensures interpretability but lacks broader contextual awareness; for instance, the model may fail to link a pronoun in one sentence to a country mentioned in the previous one, leading to missed connections. Additionally, sentiment is classified based on the highest probability category (positive, negative, or neutral). This "winner-takes-all" classification strips away nuance, as a statement with a 51% probability of being negative is treated identically to one with 99% probability, potentially flattening the spectrum of diplomatic signalling.

In the Voting Network, the calculation of the Voting Similarity Score (S-score) treats all resolutions equally. A procedural vote on meeting scheduling contributes the same weight to the relationship score as a vote on a major security crisis. This lack of thematic weighting means that high alignment scores could theoretically be driven by trivial administrative agreements rather than deep strategic convergence. Furthermore, the network visualization relies on annual aggregation. While necessary for longitudinal analysis, this yearly resolution can obscure rapid intra-year shifts in alliances caused by sudden geopolitical shocks, presenting a smoothed average of interactions rather than a real-time diplomatic pulse.



The Hague Centre
for Strategic Studies

HCSS
Lange Voorhout 1
2514 EA The Hague

Follow us on social media:
[@hcssnl](#)

The Hague Centre for Strategic Studies
Email: info@hcss.nl
Website: www.hcss.nl