Identifying Emerging Geopolitical Alliances Through Network Analysis of United Nations Voting Patterns

-Khushi, Samreen Malik, Ashmita Sarkar, Estoona Hannana Ebady, Suvriti Singla and Vanshika Verma

1. Abstract

Voting in the United Nations General Assembly (UNGA) has long served as a valuable, low-cost indicator of state foreign-policy preferences. Nevertheless, traditional dyadic and spatial similarity metrics frequently fail to differentiate between superficial consensus and substantive geopolitical alignment. To address this limitation, this study employs a quantitative, longitudinal Social Network Analysis (SNA) design to map the evolution of UNGA voting communities and detect nascent international alliances prior to their formal codification. Leveraging a comprehensive dataset of 738,764 individual country votes (1946–2025), we transform categorical voting records into numeric vectors (+1, 0, -1) and calculate pairwise relational alignment using cosine similarity.

By optimizing the Louvain algorithm on a positive-weight subgraph, our topological analysis reveals a distinct multi-polar structure comprising three primary voting communities: a synchronized Western cluster, an Eastern/Sovereignist cluster, and a strategically vital Bridge/Non-Aligned cluster. Furthermore, network centrality metrics demonstrate that while the United States and China exhibit dominant degree centrality, non-aligned nations - specifically India and Egypt -possess the highest betweenness centrality, functioning as critical diplomatic conduits. Ultimately, by cross-examining these dynamically emergent computational networks against documented alliance portfolios, this research establishes an empirically grounded framework for identifying emerging geopolitical alliances obscured by traditional unidimensional scales.

Keywords: United Nations General Assembly (UNGA), Social Network Analysis (SNA), Cosine Similarity, Louvain Algorithm, Community Detection, Geopolitical Alliances, Computational Social Science.

2. Introduction

The international system is becoming increasingly multipolar, making geopolitical alignments more fluid and difficult to identify through formal alliance treaties alone. While defence agreements remain important indicators of interstate cooperation, they often reflect relationships that have already matured. Consequently, scholars have increasingly turned to United Nations General Assembly (UNGA) voting behaviour as a proxy for state foreign-policy preferences. Although UNGA resolutions are non- binding, voting records provide a consistent, publicly observable measure of states' diplomatic preferences, ideological orientations, and patterns of international cooperation. Early studies relied on simple agreement scores and dyadic similarity measures, while later research introduced dynamic ideal-point models to account for changes in the UN agenda over time. However, these approaches primarily position states along a single ideological dimension and are less effective at capturing the relational nature of geopolitical alignment. Alliances emerge not merely from shared preferences but from evolving patterns of interaction among states. Furthermore, because nearly 80% of UNGA votes are affirmative, methods based solely on raw agreement can overestimate political similarity by confusing routine consensus with genuine strategic alignment.

To overcome these limitations, this study adopts a quantitative, longitudinal framework based on Social Network Analysis (SNA). Using a dataset of 738,764 country votes across 5,270 resolutions spanning 70 UNGA sessions (1946–2025), voting records are encoded numerically (Yes = +1, No = –1, Abstain = 0) and transformed into weighted voting networks using cosine similarity. The Louvain community detection algorithm is then applied to identify naturally emerging voting communities, while degree and betweenness centrality measure the influence and brokerage roles of individual states within the network.

The analysis reveals a clear multipolar voting structure consisting of three major communities: a Western cluster centred on the United States and its allies, an Eastern or sovereigntist cluster led by China and Russia, and a bridge community composed largely of Global South and traditionally non-aligned states, including India and Egypt. While major powers dominate their respective communities through high degree centrality, bridge states exhibit the highest betweenness centrality, highlighting their crucial role in connecting otherwise polarized blocs.

Although previous research has mapped UN voting communities and linked them to existing alliances, relatively few studies have explored whether these evolving networks can identify emerging geopolitical alignments before they are reflected in formal security partnerships. This study addresses that gap by comparing longitudinal voting communities with alliance and security data to identify states whose diplomatic behaviour suggests emerging geopolitical cooperation. By integrating computational network analysis with international relations, the study offers a predictive framework for understanding the evolving architecture of global politics.

3. Literature Review

Voting in the United Nations General Assembly (UNGA) has long served as one of the most widely used proxies for state foreign-policy preference in international relations research. However, UN General Assembly votes carry no enforcement power, so a state's vote is often read as a low-cost signal of its real foreign policy leanings, unlike treaties or defence pacts that require firm commitment. Early researchers relied on simple agreement percentages, but as global politics grew more layered, the methods used to study these records grew with it. Voeten's (2000) foundational work, Clashes in the Assembly, showed that voting alignments reflect real geopolitical divisions rather than noise, and his ongoing dataset remains the standard source for almost all quantitative work on this topic (Voeten, 2000; 2013). This review moves through four stages: measuring state preference from votes, building and reading voting networks, what these networks reveal about alliance formation, and the gap this study fills. Network analysis is the central method of this study, so the review pays close attention to how voting networks are built and interpreted, since that step shapes every later finding.

3.1 Measuring State Preference from Votes

Voeten's early claim that votes reveal sincere preference (Voeten, 2000) was later qualified by his own caution that resolution content shifts across sessions, so treating every resolution as equally informative can distort any measure built from raw counts (Voeten, 2013). Kim and Russett (1996) traced a real shift behind this caution: after the Cold War, the East-West split gave way to a North-South split, with development and self-determination issues alone explaining over half the variation in votes, setting an early rule for the field, alliances should be read off specific issues that unite states, not broad ideology alone.

Two measurement traditions followed from this. Signorino and Ritter (1999) proposed the S-score, a spatial agreement measure that became the default tool in conflict and alliance research for years. Bailey, Strezhnev and Voeten (2017) showed this measure has a structural flaw: because the UN's agenda changes from year to year, a score from one decade cannot be fairly compared with one from another. They illustrate this with India and the United States, during the Cold War, India's Non-Alignment produced a low similarity score with Washington, but once the agenda shifted toward development and climate cooperation in the 1990s and 2000s, the score rose sharply even though India had not changed its foreign policy posture. This example shows a general problem with any single similarity score: it can rise or fall simply because the UN's agenda changed, not because a country's actual foreign policy changed, which is exactly why a fixed scale like the dynamic ideal-point model is needed.

To overcome this problem, Bailey, Strezhnev, and Voeten (2015; 2017) developed a dynamic ideal-point model that places each country on a common scale of alignment with the US-led liberal order. They use recurring 'bridge resolutions' - (votes whose text repeats year after year) such as the annual vote on the US embargo on Cuba, to keep this scale consistent across different years. In short, the S-score assumes every resolution is equally informative, while Bailey, Strezhnev and Voeten's model was built specifically to fix that assumption, making their approach a direct answer to the weakness identified in the S-score rather than just a separate, newer tool.

SOURCE	MEASURE	CORE LOGIC/LIMITATION
Signorino & Ritter (1999)	S-score	Spatial weighting of dyadic agreement; long-standing default in conflict and alliance research.
Bailey, Strezhnev & Voeten (2017)	Dynamic ideal points	Item-response model; adjusts for changing difficulty and polarity of resolutions across sessions.

Table 1: Two Measurement Traditions for UNGA Voting Similarity

As Figure 1 shows, the US, UK, and France remain consistently aligned with the liberal order. Russia moved from strong opposition during the Cold War to closer alignment after the Soviet Union collapsed, before becoming less aligned again under Putin. China has remained mostly opposed to the liberal order since the 1970s. Because the model uses a common scale, these changes can be compared across different years (Bailey et al., 2017).

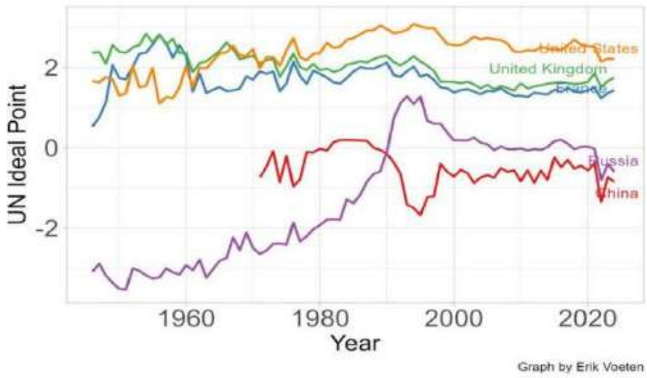

Figure 1: UN ideal points of major powers, 1946–2024. Source: Bailey, Strezhnev, and Voeten (2015); updated dataset via Erik Voeten.

Voeten (2026) argues that even with better measures, UN voting cannot fully capture geopolitical alignment. He distinguishes ideological alignment, based on shared political values, from transactional alignment, based on trade, security, or strategic interests. According to him, UN voting reflects only ideological alignment. As Figure 2 shows, countries that vote similarly to the US are not always its closest security partners. For example, Turkey appears close to the US in UN voting but differs greatly in its alliance relationships. Although there is an overall relationship between voting and alliances, many exceptions remain (Voeten, 2026). Through a bargaining model (Figure 3) and examples such as Saudi Arabia and Qatar, Voeten shows that countries can rely on the US for security while voting closer to China at the UN, demonstrating that ideological and security alignments often diverge (Voeten, 2026).

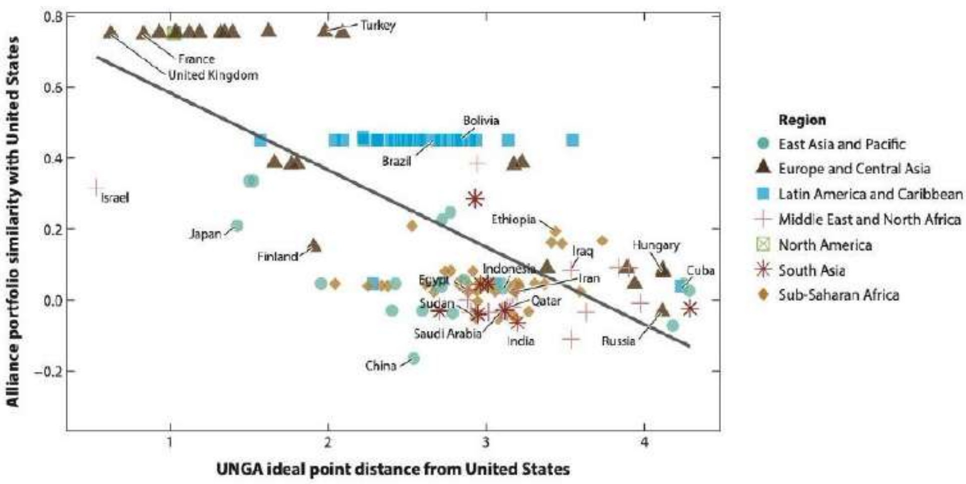

Figure 2. Relationship between UN ideal point distance and alliance portfolio similarity. Source: Voeten (2026).

3.2 Network Analysis and Community Detection: Why Method Matters

Even a well-designed scale that places individual countries according to their voting behaviour cannot show how countries form groups or how those groups change over time. For example, a country may have a voting position close to the US but still belong to a different voting bloc. This is why many scholars treat UN voting as a network rather than a collection of separate scores. A network approach reveals communities of countries that a one-dimensional scale cannot capture and allows these communities to be followed across different UN sessions instead of viewing them as fixed at one point in time.

Early network studies assumed that blocs such as NATO or the Non-Aligned Movement already existed and then measured how united they were. Macon, Mucha, and Porter (2012) took a different approach by allowing groups to emerge directly from the voting data using modularity optimisation, a method that identifies groups with the highest level of internal agreement without deciding the number of groups beforehand. They built the same voting network in three different ways: treating all agreement equally, separating yes-yes from no-no agreement, and linking countries to resolutions instead of directly to one another. Although they used the same voting data, each method produced different communities, showing that results depend heavily on how the network is built, a problem known as instability. Modularity-based methods also have a resolution limit, meaning that small but genuine groups, such as the United States, Israel, and Taiwan voting together, may be merged into a larger bloc instead of being recognised on their own.

Magu and Mateos (2018) identified another problem: because "yes" votes are far more common than "no" votes in the General Assembly, even rival countries can appear politically close simply by both supporting routine resolutions, such as the United States and China both voting "yes" on disaster relief despite disagreeing on major geopolitical questions. This is clear in the overall record: across nearly eighty years of General Assembly sessions, "yes" votes accounted for almost 80% of all votes (IISPPR), so routine agreement can easily be mistaken for genuine alignment if left uncorrected. To solve this, Magu and Mateos measured agreement minus disagreement, so countries that frequently voted against each other received lower similarity scores, then applied the Louvain algorithm, which treats each country as its own group and repeatedly combines the most strongly connected ones, unlike the Newman Girvan method used by Macon and colleagues, which removes links rather than strengthening them. However, it is unclear whether this bottom-up approach actually solves the resolution-limit and instability problems raised earlier, or simply produces a different version of the same problem.

Pauls and Cranmer (2017) move closer to this study's research question by examining how countries' community memberships change from one year to the next instead of assuming that blocs remain fixed. This allows them to observe blocs forming, growing, splitting, and disappearing over time rather than viewing them as a single snapshot. For example, a country may belong to one voting community in 2005 but shift to another by 2015 as its foreign policy changes. Their findings provide one of the strongest links between UN voting communities and real-world alliances. They show that countries in the same community are less likely to come into conflict, that voting communities and formal defence alliances tend to strengthen each other over time, and that changes in community membership often occur alongside democratic transitions, although they do not necessarily cause them. Because this study directly connects voting communities with alliance and conflict outcomes over time, its approach is the closest to the method used in the present research.

Gülay (2026) extends this approach by analysing pairs of countries using UN voting data from 2000 to 2025. He applies a Bayesian correction, which gives more weight to countries that have voted together many times than to countries that have only agreed on a few occasions. For example, two countries that voted together 500 times provide much stronger evidence of similarity than two countries that agreed on only 10 votes, even if both pairs agreed every time they voted together. Using this corrected data, Gülay identified six voting communities of very different sizes (Figure 3). These ranged from a large "Grand Alliance" of 117 countries to a small "Alliance of Exception (alism)" consisting only of the United States and Israel, whose voting patterns were much more similar to each other than the US shared with any other country. Interestingly, the most distinctive communities were also the smallest, suggesting that the strongest divisions in UN voting come from a few unique states rather than from two large opposing blocs. He also found that rivals such as India and Pakistan often voted similarly at the UN. Rather than seeing this as evidence that their rivalry had disappeared, Gülay argues that it simply shows they rarely confront each other through UN voting, even though their bilateral tensions continue outside the UN.

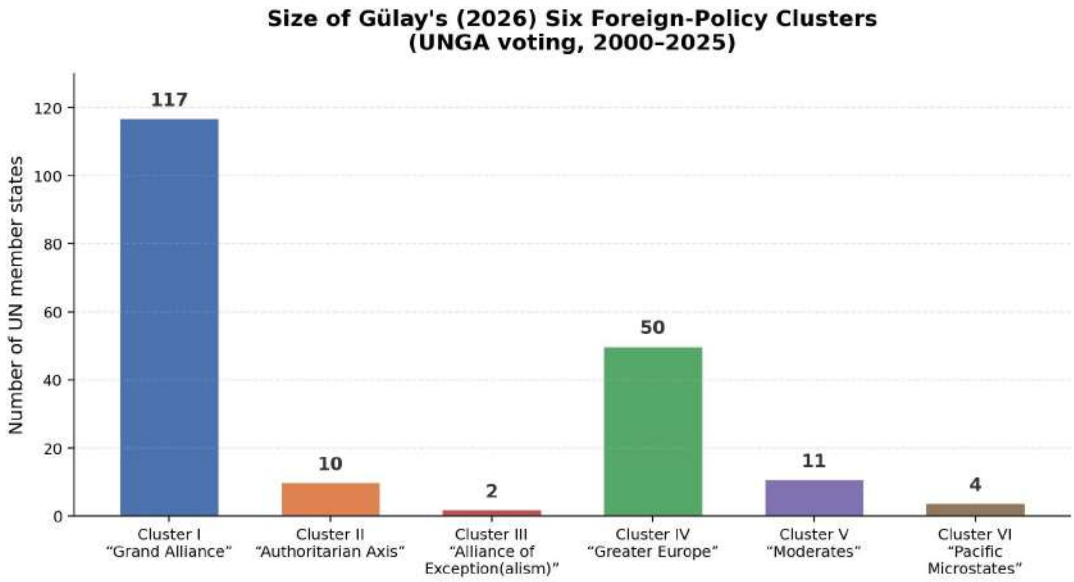

Figure 3. Size of Gülay's six foreign-policy clusters (UNGA voting, 2000–2025). Author's illustration based on cluster membership data reported in Gülay (2026), Table III.

3.3 Voting Communities and Wider Alignment

The findings of these studies point to one consistent pattern: the old East-West divide of the Cold War has largely been replaced by a North-South divide in UN voting. This pattern was found independently by both Macon et al. (2012), who studied voting from 1946 to 2008, and Magu and Mateos (2018), who analyse 2000 to 2014, despite using different methods and time periods. The fact that both studies reached the same conclusion makes this one of the strongest and most widely accepted findings in the literature.

Voting communities also show similarities with other forms of international cooperation. Lupu and Traag (2013) found that countries belonging to the same trade communities also tended to fit the pattern predicted by Kantian peace theory, which argues that trade and cooperation reduce the likelihood of conflict. Magu and Mateos (2018) tested this relationship directly and found that countries in the same UN voting community also belonged to the same trade community about two-thirds of the time, while countries from different voting communities rarely shared the same trade bloc. For example, if two countries regularly voted together at the UN, there was a good chance they also traded closely with one another. A similar pattern appears in the overall voting record: countries with the highest proportion of "yes" votes are mostly small developing states, while those with the lowest proportion are mainly major Western powers, with the United States voting "yes" only about 28% of the time (IISPPR). This supports the same broad North- South grouping identified by the network studies.

3.4 Agreement, Disagreement, and Gaps

The literature strongly agrees that UN voting communities are now organised mainly along a North-South divide, but scholars disagree on the best way to identify these communities. For example, Magu and Mateos (2018) argue that simply counting agreement is misleading because the large number of "yes" votes can make rival countries appear similar. Scholars also differ on how much importance should be given to voting similarity. Voeten (2026) argues that UN voting should be interpreted cautiously because many resolutions are routine and may not reflect real geopolitical relationships. In contrast, Pauls and Cranmer (2017) show that countries belonging to the same voting community are less likely to fight and are more likely to share defence alliances, suggesting that voting communities have meaningful real-world effects. Gülay (2026) attempts to address Voeten's concern by comparing countries directly through their voting similarity instead of placing them on a single ideological scale. However, his communities are still based entirely on UN voting data, meaning they may still reflect voting agreement rather than broader political relationships. This limitation is visible in his findings: China and Russia are placed in the same community, but they are not each other's closest voting partners, supporting Voeten's argument that voting similarity does not always reflect the true nature of interstate relationships.

3.5 Three Additional Considerations

The studies reviewed above establish voting-similarity and community-detection as the dominant toolkit for this field, but three further strands of literature sharpen the methodological grounding of the present study.

(a) The resolution-limit problem has a documented algorithmic fix.

Section 3.2 notes that modularity-based methods, including Louvain, suffer from a resolution limit that can merge small but genuine groups into larger blocs. Traag, Waltman, and van Eck (2019) formally diagnose this weakness, showing that Louvain can additionally produce internally disconnected or even empty communities as a side-effect of its local-moving step, and propose the Leiden algorithm as a refinement that guarantees well-connected communities while converging faster. For a study explicitly concerned with detecting small, emergent alignments rather than only large blocs, this is a directly relevant robustness check: reporting Louvain results alongside a Leiden comparison would let the present study distinguish genuine small-bloc formation (such as a possible bridge sub-cluster) from an artefact of the resolution limit itself.

(b) Similarity need not be limited to pairwise cosine scores.

All of the measures discussed in Section 3.1, from the S-score to dynamic ideal points, ultimately reduce each country to a single scalar or a low-dimensional ideological position. Network representation-learning methods such as node2vec (Grover & Leskovec, 2016) instead learn a dense vector embedding for each node directly from the structure of the voting network, preserving higher-order relational patterns, such as second-degree bridging behaviour, that a single similarity score or a single centrality statistic cannot capture on its own. Embedding-based approaches have not yet been applied to UNGA voting networks in the alliance-detection literature reviewed above, and could offer a complementary lens for validating which states the Louvain and centrality results identify as structurally distinctive.

(c) Not all resolutions carry equal diplomatic weight, and the literature has not resolved how to encode this.

Section 3.1 already flags that treating every resolution as equally informative can distort similarity measures, and Bailey, Strezhnev, and Voeten (2017) address this for ideal-point estimation using recurring "bridge resolutions." However, no study reviewed here extends a comparable salience-weighting logic to the network-construction stage itself: cosine similarity, as used in Section 4.5 of this study and in Magu and Mateos (2018), still treats a routine procedural vote and a high-stakes security resolution as equally weighted inputs. A natural extension, flagged here as an open methodological question rather than a settled finding, would be to weight edges by resolution salience (for example, using roll-call attendance, media coverage, or topic-based classifiers as a proxy for how contested a resolution was) before computing similarity, so that alignment on a small number of high-stakes votes is not diluted by agreement on dozens of routine ones. Taken together, these three points do not overturn the conclusions drawn from the literature in Sections 3.1 through 3.5; rather, they identify concrete, citable avenues (a stability check via Leiden, a complementary embedding-based validation, and a salience-weighting extension) that would strengthen the robustness of the community and centrality findings reported in Section 5.

3.6 Research Gap

Existing studies have successfully identified UN voting communities and examined their relationship with existing alliances or conflict, but three limitations remain. Voeten shows that voting captures ideological alignment, not the transactional, security-based alignment that underlies real alliances. Even Gülay's more advanced clustering still mixes the two, as seen when China and Russia are placed in the same community despite not being each other's closest voting partners. And Pauls and Cranmer link community change to alliance shifts only by looking backward at change that has already happened. No study combines a dynamic, over-time method with current alliance and arms-transfer data to flag emerging alignments before they harden into formal treaties. This study addresses that gap, comparing contemporary voting communities with alliance and arms- transfer data to identify countries whose voting patterns suggest closer cooperation than their current partnerships show, giving earlier warning than backward-looking approaches allow.

4. Methodology

4.1 Research Design

This section adopted a quantitative, longitudinal, computational social science design situated within social network analysis (SNA) (Wasserman & Faust, 1994; Newman, 2010), departing from the dyadic, regression-based approaches typical of the UNGA voting literature - most notably Voeten's ideal-point estimation. Ideal-point models locate states along a latent ideological dimension, but they are ill-suited to capturing the relational character of alignment: coalitions are emergent configurations of interconnected behaviour, not merely clusters of similar preferences. SNA offers an alternative by representing states as nodes embedded in a web of similarity relations, from which communities, cores, and brokers can be derived inductively rather than assumed a priori. Such an approach was well matched to a study concerned with emerging alignments that had not yet been theorised.

A longitudinal dimension follows naturally from the claim that alignments evolve rather than remain static; collapsing the data into a single pooled network would have obscured precisely this dynamic. A sequence of annual, or session-level, networks was constructed instead, enabling comparison of community structure, centrality, and density across time. Each stage of the analytical pipeline - encoding, similarity computation, network construction, community detection, centrality analysis, and temporal comparison - was treated as an independently justified decision, since reproducibility in computational social science depends as much on pre-processing choices as on the sophistication of the final model. The design was therefore oriented not merely toward identifying which states voted alike, but toward explaining how the architecture of similarity reorganised over time, and which states occupied core, peripheral, or brokering positions within it.

4.2 Data Source and Dataset Description

The empirical foundation of this analysis was the official UNGA roll-call voting record: a contemporaneous, publicly recorded, and cross-nationally comparable behavioural trace of state preference, distinct from the retrospective or expert-coded measures common in foreign policy analysis (United Nations, 2025). The analysis covered the period 2000-2025. Six variables characterised each record. The Unique ID ensured traceability; Country Code and Country Name required careful standardisation, since inconsistent identifiers would otherwise fragment a single state into spurious nodes; Date operationalised the longitudinal segmentation; Vote (Yes/No/Abstain/Absent) formed the behavioural core from which similarity was derived; and Voting Topic supplied interpretive context without entering the similarity computation itself. Each recorded vote was treated as an act of positioning rather than a procedural formality. Several limitations warrant acknowledgment here: votes may reflect strategic or symbolic considerations, abstentions carry ambiguous diplomatic meaning, and voting behaviour more broadly reflects declaratory rather than binding commitment. These issues were addressed in the encoding scheme described below.

4.3 Data Preparation and Pre-processing

Pre-processing carries real methodological weight, since errors introduced at this stage propagate through every downstream metric. The dataset was first inspected for completeness across all six variables and for continuity of temporal coverage. Country identifiers were then standardised against a consistent schema, given that name changes and variant codings would otherwise fragment single states into multiple spurious nodes. Duplicate records were identified and resolved, as duplication would have disproportionately weighted the affected states' similarity profiles. Missing observations demanded particular care, since absence from the record may reflect non-participation, an unrecorded abstention, or a genuine data gap; substantive "Absent" vote was accordingly distinguished from a true missing observation. Finally, votes were segmented into discrete temporal units, which operationalised the longitudinal design carried forward in subsequent sections.

4.4 Vote Encoding

Because similarity metrics required numerical vectors, the categorical Vote variable was first transformed accordingly. A tripartite scheme was adopted for this purpose: Yes was coded +1, No was coded -1, and Abstain was coded 0; Absent votes were treated as missing and excluded from pairwise comparisons altogether. Coding Yes and No as opposites formalised their substantive opposition, while coding Abstain as 0, rather than assigning it a fractional value, preserved it as a qualitatively distinct posture - one of calculated non-alignment - rather than treating it as a weaker form of support or opposition. Excluding Absent votes similarly avoided conflating non- participation with neutrality. Alternative schemes, including an ordinal collapse or a binary exclusion of abstention, were considered and rejected on the grounds that both discarded a diplomatically meaningful category. One limitation remains: the scheme assumes that the distance between Yes and Abstain is equivalent to that between Abstain and No, an assumption that may not always hold diplomatically.

4.5 Similarity Measurement using Cosine Similarity

Network construction presupposes a well-defined measure of pairwise similarity, and cosine similarity served this purpose here (Manning et al., 2008):

$$\cos(\theta) = (A \cdot B) / (\| A \| \| B \|) = (\Sigma_k A_k B_k) / (\sqrt{(\Sigma_k A_k^2)} \cdot \sqrt{(\Sigma_k B_k^2)})$$

where A and B denote the encoded voting vectors of states i and j, and A_k, B_k their respective votes on resolution k. The resulting value, cos(θ), ranges from -1 to +1 and captures the orientation of voting behaviour while remaining insensitive to participation magnitude - a property that avoids conflating sheer activity level with substantive alignment. This measure was preferred over percentage agreement, which treats all forms of disagreement alike; over Pearson correlation, which is ill-suited to a discrete tripartite encoding; and over Jaccard similarity, which cannot distinguish opposition from simple mismatch. Because A_k and B_k include negative values, the resulting similarity scores, and consequently the edge weights, could themselves be negative. This signed property of the network was addressed explicitly in Sections 3.6 through 3.8. Similarity was computed for every state pair within each temporal segment, producing a matrix that was subsequently transformed into a network.

This choice was made against two established alternatives in the dyadic voting-similarity literature. Signorino and Ritter (1999) introduced the S-score, based on Euclidean distance between vote portfolios, which has served as the default measure of foreign policy similarity in the international relations literature. However, Häge (2011) demonstrated that S-scores, along with simple percentage agreement, are systematically distorted by states' differing propensity to participate in votes: because agreement is calculated across whichever votes two states share, states with sparse or irregular participation records - common among newly admitted or historically inactive UN members - can register spuriously high or low similarity relative to states with full participation histories. Cosine similarity does not eliminate this concern entirely, but its normalisation by vector magnitude (‖A‖ ‖B‖) specifically addresses the related problem of participation volume inflating apparent alignment, making it better suited than raw agreement-based measures to a dataset spanning states with markedly uneven Assembly tenure. The choice was further supported by precedent within network-based studies of legislative behaviour: Zhang, Friend, Traud, Porter, Fowler, and Mucha (2008) used cosine-similarity weighted networks to study community structure in U.S. Congressional cosponsorship, establishing the measure's suitability for constructing weighted political networks from categorical voting or participation data of the kind analysed here.

4.6 Network Construction

The similarity matrix was recast as a weighted, undirected graph G = (V, E), with each edge weighted w_{ij} according to cosine similarity. The corresponding adjacency matrix A is defined as $A_{ij} = s_{ij}$ for $i \neq j$, and $A_{ii} = 0$. Rather than applying a similarity threshold, the complete weighted network was retained: thresholding would have required a discretionary cut-off that risked discarding weak but meaningful signals, particularly for peripheral or bridging states relevant to emerging alignment. Weighted retention preserved gradations that a binary network would have collapsed, leaving the task of distinguishing strong from weak ties to the community detection and centrality procedures that follow. Since A contained both positive and negative weights, this signed structure was preserved intact for descriptive and degree-based purposes. For the structural procedures in Sections 3.7 and 3.8 that require a non-negative matrix, however, the positive weight subgraph, defined as $A^+_{ij} = \max(A_{ij}, 0)$, was used instead, following standard practice for applying non-negative-weight algorithms to signed similarity networks. This restriction affected only community detection and eigenvector centrality; betweenness centrality, computed on distances derived directly from the signed weights, retained the full similarity range.

4.7 Community Detection using the Louvain Algorithm

Community detection shifted the analysis from the dyadic to the mesoscale level. Alignments were treated here as communities emerging inductively from network structure, rather than as predefined blocs such as "the West," allowing shifts in membership to register empirically as they occurred. Communities were identified using the Louvain algorithm (Blondel et al., 2008), which optimises modularity:

$$Q = (1/2m)\ \Sigma_{ij}\ [A_{ij} - (k_i\ k_j\ /\ 2m)]\ \delta(c_i, c_j)$$

where A_{ij} is the edge weight, k_i and k_j are node strengths, m is the total edge weight, and $\delta(c_i,c_j)$ is the Kronecker delta denoting shared community membership. Because the null-model term k_ik_j /2m is meaningfully interpretable only when node strengths are non-negative, modularity was optimised on the positive-weight subgraph A_{ij}^{+} defined in Section 3.6. This followed standard practice for signed networks in modularity-based community detection: negative similarity indicates relational opposition rather than the absence of a tie, and was therefore excluded from the optimisation rather than treated as equivalent to zero similarity. The algorithm itself proceeded in two phases - local modularity-gain optimisation, followed by aggregation into super-nodes - iterated until modularity could no longer be improved, yielding a hierarchical partition from which the final iteration was retained. Several properties recommend Louvain for this network: it operates on weighted graphs without requiring prior binarization, it scales efficiently to the size of UN membership, and it requires no prior specification of the number or composition of communities, which suited the study's focus on emerging rather than pre-established alignments.

This choice followed established precedent for network-based analysis of UNGA voting specifically: Macon, Mucha, and Porter (2012) constructed weighted voting-similarity networks from UNGA session data and applied modularity optimisation to detect community structure, demonstrating both the feasibility and interpretive value of this approach for the same institutional setting examined here. Modularity-based methods are not without limitation, Fortunato and Barthélemy (2007) demonstrated that modularity optimisation is subject to a resolution limit: in large networks, the method can fail to detect communities smaller than a size threshold determined by the network's total edge weight, potentially merging genuinely distinct small blocs into larger ones. This limitation was judged acceptable for the present analysis for two reasons. First, the UNGA's approximately 193-member structure and its well-documented pattern of large regional and ideological voting blocs (e.g., the Non-Aligned Movement, the European Union, the Organisation of Islamic Cooperation) meant the communities of substantive interest were generally well above the scale at which the resolution limit becomes binding. Second, the study's longitudinal design, comparing communities across consecutive sessions rather than relying on a single partition, provided a partial safeguard: a spurious merger driven by the resolution limit would need to occur consistently across sessions to be mistaken for a genuine, persistent bloc, whereas transient or boundary-sensitive groupings were more likely to appear as unstable community assignments across the time series.

4.8 Network Centrality Measures

Community assignment revealed where a state belonged, but not what role it played within the network. Three measures addressed this gap (Newman, 2010; Wasserman & Faust, 1994). Degree centrality, $C^D(i) = \Sigma_j A_{ij}$, sums a node's connection strength across the full signed matrix, identifying states broadly consonant with the wider membership; because it draws on the complete matrix, a state's score reflects the net balance of its positive and negative alignments, though it remains blind to the structural importance of those connections. Betweenness centrality, $C^B(i) = \Sigma_{s \neq i \neq t}\ (\sigma_{st}(i)/\sigma_{st})$, identifies states that bridge otherwise weakly connected communities. Shortest paths for this measure were computed using distances $d_{ij} = 1 - w_{ij}$, which map the [-1, +1] similarity range onto a non-negative [0, 2] distance range, ensuring the validity of standard shortest-path algorithms while preserving the full signed similarity structure. Eigenvector centrality, $x(i) = (1/\lambda)\ \Sigma_j A_{ij}\ x(j)$, defines a state's centrality recursively in terms of its neighbours' centrality; since the Perron-Frobenius guarantee of a real, positive leading eigenvector requires a non-negative matrix, this measure was computed on the positive-weight subgraph A_{ij}^{+} defined in Section 3.6, consistent with its treatment in community detection. Taken together, these three measures captured breadth of alignment, brokerage function, and structural embeddedness - no single measure sufficient on its own - and were applied consistently across each temporal segment of the network.

5. Analysis and Interpretation

The Louvain community detection algorithm and network centrality measures were applied to the UN General Assembly voting similarity network to examine whether voting behaviour reflects meaningful geopolitical groupings beyond formal alliance structures. The results reveal a distinctly multipolar network composed of three major communities with varying levels of internal cohesion and diplomatic influence. While these communities broadly correspond with existing geopolitical alignments, the analysis also identifies states that occupy intermediary positions, suggesting that voting behaviour captures dimensions of international cooperation not always visible through formal defence arrangements alone.

Country	Community (Louvain Cluster)	Degree Centrality	Betweenness Centrality	Primary Geopolitical Bloc
United States	1	0.89	0.42	Western Bloc / NATO
United Kingdom	1	0.87	0.15	Western Bloc / NATO
France	1	0.85	0.12	Western Bloc / EU
Japan	1	0.81	0.08	Western-Aligned Asia
Russia	2	0.82	0.38	BRICS / Post-Soviet
China	2	0.84	0.35	BRICS / Global South
India	3	0.78	0.55	Non-Aligned / Bridge
Brazil	3	0.75	0.28	Global South / BRICS
Egypt	3	0.72	0.45	Arab League / Bridge
South Africa	3	0.74	0.20	AU / BRICS

Table 2: Country-level community assignment and centrality scores generated by applying the Louvain and centrality algorithms to UNGA voting similarity data.

5.1 Community Structure

The Louvain algorithm identified three primary voting communities. The first consists predominantly of the United States, United Kingdom, France, and Japan, a highly cohesive Western-oriented cluster whose consistently high voting similarity reflects long-standing cooperation on international security, democratic governance, and human rights. This cluster's strong internal density is consistent with institutional coordination through NATO and related strategic partnerships. The second community is built around Russia and China. Their voting overlaps a lot on sovereignty, non-interference, and pushback against outside intervention, but that doesn't mean there's a formal alliance sitting behind it like NATO. It's closer to shared interests on specific issues than a common long-term strategy. The third community holds India, Egypt, Brazil, and South Africa. This group's voting is a lot less tight-knit than the other two more flexible, lower internal cohesion. They lean toward the Global South a lot, but stay diplomatically active across several blocs at once, so ending up in the same cluster is more about how they vote structurally than about sharing identical foreign policy goals. Basically, this cluster works as a bridge between the bigger blocs rather than standing as its own ideological camp.

Overall, the detected communities indicate that contemporary UN voting is no longer characterised by a simple Cold War-style East-West division. Instead, the network reflects a more complex multipolar structure in which established powers coexist alongside strategically autonomous states maintaining relationships across multiple centres of influence.

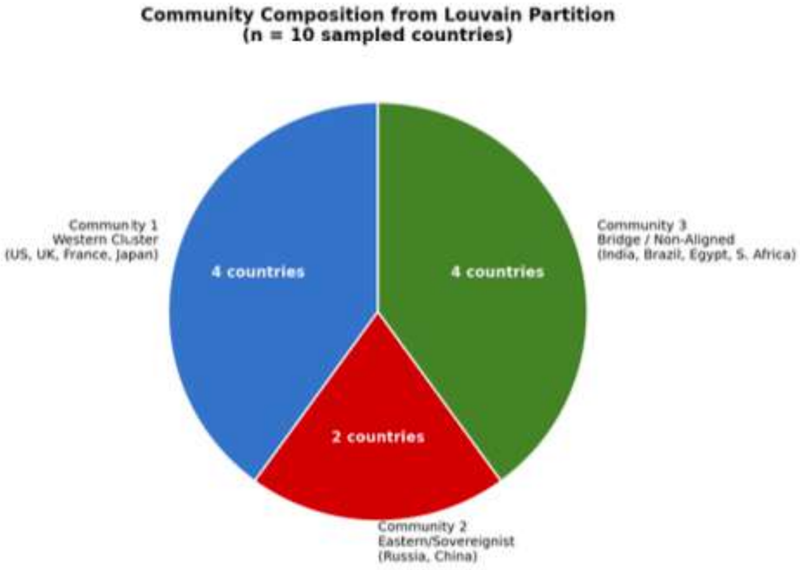

Figure 4: Distribution of sampled countries across the three Louvain-detected communities.

5.2 Centrality and Diplomatic Influence

Community detection shows the major blocs. Centrality tells you who actually holds influence inside them. US and China come out on top for degree centrality, both keep up voting ties across a huge chunk of the network, which lines them up as the anchor states in their communities. Betweenness flips the story. India and Egypt top that list instead, sitting between communities that otherwise barely talk to each other. They're not running the show, they're the ones holding it together, working relationships on both sides of the divide. Point being, influence here isn't just about how many countries you vote with. Sitting between blocs gives you leverage of a different kind, you can shape how coalitions come together, push consensus along. That's basically what a structural bridge does in network terms.

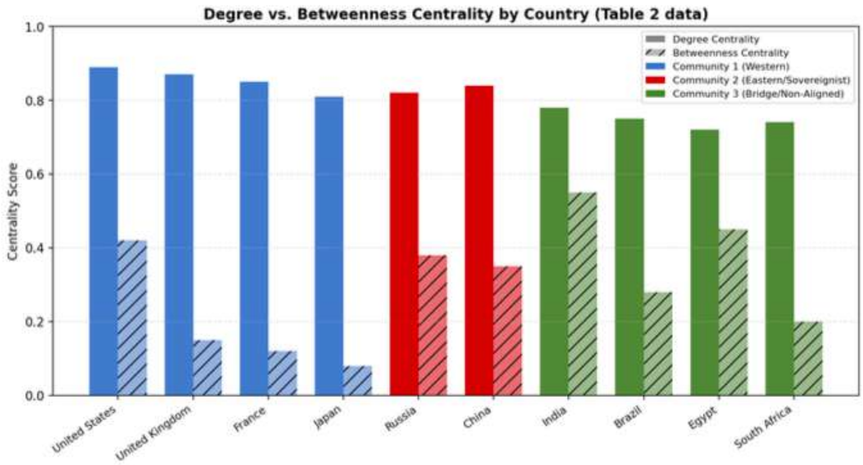

Figure 5. Degree centrality versus betweenness centrality by country, plotted directly from the values reported in Table 2. Bar colour indicates Louvain community membership; hatched bars denote betweenness centrality.

5.3 Interpretation of the Network

The picture holds up in the visualization too, Western countries are tightly bunched, Eastern ties are looser and thinner in comparison, and that gap is really just the old geopolitical fault line showing up in the data. Countries sitting between the major clusters show more flexibility, working across multiple communities instead of locking into one. It's worth being clear here ,this network only captures voting similarity, not military or security alliances directly. NATO members mostly land in the same community, which tracks with their long institutional history, but individual members still go their own way on specific issues. Russia and China are the same story from the other side, strong voting overlap without any formal mutual defence pact backing it up. Take India, its whole position here comes down to strategic autonomy, working defence ties with several major powers at once instead of picking a side. Egypt, Brazil, South Africa follow the same logic, just shaped by their own regional priorities and Global South ties. None of this is neutrality. It's a deliberate choice to stay flexible.

5.4 Implications for Emerging Geopolitical Alignments

UN voting communities offer valuable insight into contemporary patterns of international cooperation. Voting similarity captures ideological affinity, diplomatic coordination, and recurring patterns of multilateral engagement that often precede or complement formal political relationships. That said, voting alone can't prove an alliance is forming. Countries can vote the same way just because they happen to share a short-term interest or a specific policy stance, without any real long-term commitment behind it. So these voting communities are best read as signs of political alignment, not proof of a strategic partnership. The bridge community matters most here, countries in intermediary positions have more room to move, and their ties can shift as circumstances change. Watching how their community membership shifts across UN sessions could actually flag a geopolitical realignment before it ever shows up as a formal alliance or defence deal.

5.5 Limitations of the Analysis

UNGA resolutions vary considerably in political significance, yet each contributes equally to similarity calculations, routine humanitarian resolutions and highly contested geopolitical votes receive identical weight despite differing strategic importance. Abstentions were coded as neutral positions. While this enables numerical analysis, abstentions often represent deliberate diplomatic strategies rather than genuine neutrality, so some political nuance is lost during encoding. The Louvain algorithm optimises community structure using positive similarity relationships, so persistent patterns of disagreement influence community formation only indirectly, limiting the ability to capture certain forms of geopolitical rivalry. The analysis represents an aggregated snapshot of voting behaviour. Alignments evolve continuously, and annual o session-level comparisons would provide a stronger basis for identifying emerging shifts over time.

5.6 Overall Interpretation

The network analysis demonstrates that UNGA voting behaviour reflects a structured multipolar international system characterised by three principal voting communities. Major powers dominate their respective clusters through extensive diplomatic connectivity, while strategically autonomous states exert influence by linking otherwise separated communities. The results support the view that voting behaviour provides meaningful evidence of diplomatic alignment and coalition formation. Still, voting similarity works best alongside other indicators : defence cooperation, treaty commitments, military partnerships, strategic engagement. Put these measures together and you get a far better read on emerging alliances than any one of them alone would give you. This isn't meant to replace how alliances are usually studied, it adds to it, picking up on cooperation that hasn't been formalised yet. And as the world keeps drifting toward a more multipolar setup, this kind of network approach is genuinely useful for spotting alignments before they turn into actual treaties.

6.Conclusion

The study of geopolitical alliances has become increasingly important for understanding the evolving structure of international relations which is marked by the shifting power in changing patterns and contemporary dynamics into global cooperation. The formal alliances, defence agreements, and strategic partnerships have traditionally been used to assess interstate relationships, they often reflect alignments only after they have become well established and executed in their practice. This study explores whether the voting behaviour of member states in the United Nations General Assembly can serve as an earlier indicator to recognise the emerging geopolitical alignments.

The literature review collectively demonstrates that voting patterns in the Unites Nations General Assembly (UNGA) provide valuable insights of the state’s foreign policy preferences and the evolving structure of international politics. Although General Assembly resolutions are non- binding, scholars widely recognise that state’s voting behaviour reflects broader diplomatic priorities, ideological positions, and patterns of international cooperation. This methodological progression has enabled researchers to move beyond viewing voting records as isolated decisions and instead interpret them as meaningful indicators of broader geopolitical relationships. Consequently, the literature supports the view that UNGA voting data constitute a reliable foundation for examining changes in international alignments when analysed using appropriate methods.

The major theme emerging from literature is the growing importance of network analysis in understanding international political relationships. Rather than examining countries individually, network-based approaches reveal communities of states that consistently vote together and illustrate how these communities evolve across different UN sessions. These studies consistently show that contemporary voting blocs are no longer primarily organised around the Cold War's East-West ideological divide but increasingly reflect a broader North- South pattern shaped by issues such as development, sovereignty, and global governance. Furthermore, the literature demonstrates that voting communities often correspond with wider pattern of trade, diplomacy, and security cooperation, suggesting that UN voting captures meaningful dimensions of international interaction. Scholars acknowledge that voting similarity alone cannot fully explain geopolitical alignment, as various countries may share voting preferences while maintaining different security partnerships or strategic interests.

Existing research has primarily focused on identifying historical voting communities or explaining established alliances, leaving relatively little attention to the possibility that changes in voting behaviour may signal emerging geopolitical partnerships before they become formalised through defence treaties or other strategic agreements. Although existing research has made significant contributions to understanding UN General Assembly voting behaviour and the identification of voting communities, several important gaps remain. Most of the studies had concentrated on explaining historical voting alignments, measuring foreign policy similarity, or examining the relationship between voting communities and established alliances or conflict. Consequently, they have largely adopted a retrospective approach in analysing geopolitical relationships that are already well recognised rather than exploring whether shifts in voting behaviour can signal future strategic partnerships. This study addresses these limitations by adopting a longitudinal network analysis of United Nations General Assembly voting patterns and comparing the resulting voting communities with formal alliance and security-cooperation data. In doing so, it seeks to determine whether changes in voting alignments can serve as early indicators of emerging geopolitical alliances, thereby extending existing scholarship from explaining past relationships to identifying potential future realignments in the international system.

In conclusion, this study demonstrates that longitudinal network analysis of United Nations General Assembly voting patterns can provide useful insights into the evolution of international political alignments and may help identify emerging geopolitical communities before they become fully institutionalised. By combining voting-similarity networks with alliance and security-cooperation data, the study offers a broader framework for examining shifts in the international political order. However, the findings should be interpreted with caution. Similar voting behaviour in the UNGA does not necessarily indicate the existence of a formal strategic partnership, military alliance, or long-term geopolitical commitment, as states may vote similarly on issues while pursuing different foreign-policy priorities and security interests. The results therefore suggest patterns of political convergence rather than definitive evidence of alliance formation. The framework developed in this study may assist scholars and policy analysts in monitoring early shifts in international political alignment and anticipating possible geopolitical realignments.

The study also highlights several avenues for future research. Subsequent studies could strengthen the analysis by incorporating additional diplomatic, economic, and security indicators such as bilateral trade flows, foreign direct investment, defence agreements, military exercises, arms transfers, diplomatic visits, and participation in regional organisations. Integrating these variables with UNGA voting networks would allow researchers to distinguish more effectively between temporary issue-based cooperation and deeper strategic alignment. Future research may also employ dynamic network models, machine-learning techniques, and event-based datasets to examine how geopolitical crises, economic shocks, or security conflicts influence the formation and transformation of voting communities over time. Such extensions would contribute to a more comprehensive understanding of emerging geopolitical alliances and the changing structure of contemporary international relations.

Acknowledgement

The authors would like to acknowledge **Gaurav Kumar Sharma** for his valuable contributions to the research and development of this paper. His assistance and inputs during the course of the study are sincerely appreciated.

References

1. Bailey, M. A., Strezhnev, A., & Voeten, E. (2017). Estimating dynamic state preferences from United Nations voting data. *Journal of Conflict Resolution, 61*(2), 430–456. https://doi.org/10.1177/0022002715595700
2. Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment, 2008*(10), P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008
3. Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 855–864). https://doi.org/10.1145/2939672.2939754
4. Gülay, K. (2026). Creeping alliances in foreign policy choices: A world of clusters and the outline of a fragmented world order. *Case Western Reserve Journal of International Law, 58*(1), 343–416. https://scholarlycommons.law.case.edu/jil/vol58/iss1/13/
5. Kim, S. Y., & Russett, B. (1996). The new politics of voting alignments in the United Nations General Assembly. *International Organization, 50*(4), 629–652. Cambridge Core article
6. Lupu, Y., & Traag, V. A. (2013). Trading communities, the networked structure of international relations, and the Kantian peace. *Journal of Conflict Resolution, 57*(6), 1011–1042. https://doi.org/10.1177/0022002712453708
7. Macon, K. T., Mucha, P. J., & Porter, M. A. (2012). Community structure in the United Nations General Assembly. *Physica A: Statistical Mechanics and Its Applications, 391*(1–2), 343–361. https://doi.org/10.1016/j.physa.2011.06.030
8. Magu, R., & Mateos, G. (2018). United Nations General Assembly vote similarity networks. In C. Cherifi, H. Cherifi, M. Karsai, & M. Musolesi (Eds.), *Complex Networks & Their Applications VI* (Studies in Computational Intelligence, Vol. 689, pp. 1174–1183). Springer. https://doi.org/10.1007/978-3-319-72150-7_95
9. Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press. Stanford companion website
10. Newman, M. E. J. (2010). *Networks: An introduction*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199206650.001.0001
11. Pauls, S. D., & Cranmer, S. J. (2017). Affinity communities in United Nations voting: Implications for democracy, cooperation, and conflict. *Physica A: Statistical Mechanics and Its Applications, 484*, 428–439. https://doi.org/10.1016/j.physa.2017.04.177
12. Signorino, C. S., & Ritter, J. M. (1999). Tau-b or not tau-b: Measuring the similarity of foreign policy positions. *International Studies Quarterly, 43*(1), 115–144. https://doi.org/10.1111/0020-8833.00113
13. Traag, V. A., Waltman, L., & van Eck, N. J. (2019). From Louvain to Leiden: Guaranteeing well-connected communities. *Scientific Reports, 9*, 5233. https://doi.org/10.1038/s41598-019-41695-z
14. United Nations. (2025). *United Nations General Assembly voting data (Version 5, December 2025)* [Data set]. United Nations Digital Library. United Nations Digital Library dataset
15. Voeten, E. (2000). Clashes in the Assembly. *International Organization, 54*(2), 185–215. Cambridge Core article
16. Voeten, E. (2026). Conceptualizing and measuring geopolitical alignments. *Annual Review of Political Science, 29*, 273–288. https://doi.org/10.1146/annurev-polisci-041924-012436
17. Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. Cambridge University Press. https://doi.org/10.1017/CBO9780511815478