United Nations General Assembly Vote Similarity Networks

Rijul Magu(✉) and Gonzalo Mateos

University of Rochester, Rochester, NY, USA
rmagu2@cs.rochester.edu, gmateosb@ur.rochester.edu

Abstract. We construct a network of member states of the United Nations General Assembly based on how similarly they vote on resolutions. We describe a similarity metric that we feel better describes the inter-nation relationships than previously proposed models. Next, we introduce a mechanism to infer the best diplomatic path between countries that do not have high similarity in voting. Lastly, we create a bilateral commodity trade network between countries and evaluate the overlap between the trade and voting networks by applying community detection analysis. Our findings show that generated communities mimic real-world groupings and that there indeed is an alignment between voting and trade networks, paving the way for further studies on the connection between economic dependence and voting behavior.

1 Introduction

The United Nations, which started towards the end of the second world war, is an international body that comprises of 193 sovereign nations. The main policy-making organ of the United Nations is the General Assembly which includes all the member states of the United Nations. The General Assembly meets regularly where it considers current issues of critical importance to the international community in the form of high-level thematic debates organized by the President of the General Assembly in consultation with the membership.

The aim of these meets is to resolve global issues by incorporating the points of view of as many stakeholders as possible, followed by passing recommendations and directives, which it attempts to do through a voting mechanism. The official website of the United Nations states Each of the 193 Member States in the Assembly has one vote. Votes taken on designated important issues such as recommendations on peace and security, the election of Security Council and Economic and Social Council members, and budgetary questions require a two-thirds majority of Member States, but other questions are decided by a simple majority. In recent years, an effort has been made to achieve consensus on issues, rather than deciding by a formal vote, thus strengthening support for the General Assemblys decisions. The President, after having consulted and reached agreement with delegations, can propose that a resolution be adopted without a vote.

C. Cherifi et al. (eds.), *Complex Networks & Their Applications VI*,
Studies in Computational Intelligence 689, https://doi.org/10.1007/978-3-319-72150-7_95

That said, as would be expected noting the current voting model, General Assembly resolution voting processes serve as active lobbying grounds where countries often attempt to sway the opinion of fellow member states in their favor. Thus, it becomes critical for countries to know which other countries are important for which vote and how to approach them in the event relationships are not friendly.

In this paper, we attempt to model the voting relationships between countries through a network science perspective. We analyze existing metrics to define similarity in voting and provide our own measure that more accurately captures the underlying level of agreement. We then run the Louvain method [3] on the constructed graph of country voting similarities to unveil communities from the network and evaluate whether we can recover expected groupings that represent voting blocs. The communities that we uncover align well with the real world clusters.

We also present a unique methodology to identify the optimal diplomatic path between two countries, i.e. an automated way to suggest a sequence of countries to approach if one nation were to try and influence another nation with which it shares low voting similarity and therefore, likely weak ties.

In our last endeavor, we attempt to show the connection between voting and bilateral trade by comparing the communities churned through applying the Louvain method on each network. We observe how the Jaccard similarities for the corresponding communities of each network turn out to be high, indicating a relationship between the networks.

2 Related Work

There have been previous works that attempted to model either United Nations voting or bilateral trade. The relevant work involving General Assembly voting similarity networks often look at only the variations of agreements (for instance, both countries vote yes, both vote no, both abstain or any combination of these) between countries to form their measure [6,9]. They do not take into account the vital role of disagreements, particularly in a dataset where the number of 'yes' votes greatly outnumber 'no' votes. Hence, in Sect. 4, we propose a novel country similarity metric and describe how we incorporate disagreements to construct the voting graph. Co-voting agreements have also been studied within other political contexts, for instance, by modelling the European Parliament roll-call votes data [4].

It is worth noting that not much work beyond applying community detection to the generated voting network has been reported to date. In this paper, we describe a unique application of the constructed adjacency matrix, namely that of finding the optimal diplomatic path between any two countries.

On the other hand, network representations of international bilateral trade has well documented merits, where several techniques spanning numerous applications have been put to use [1,7,11]. For instance, the authors in [2] use a normalized mutual information index to study product network structures across time. Furthermore, earlier efforts have also involved incorporating community detection analysis for international trade [13]. In prior work, one interesting inference is that of bilateral trade depending more upon geographic proximity than number of trade agreements. Other studies focus on the effects of political alliances on trade but do not attempt to measure the overlap between the two independent networks of voting similarity and bilateral trade.

3 Dataset

Our work relies on two datasets. Primarily, we work with the most recent version (as of the time of writing this paper) of the United Nations General Assembly Voting Data first used in [12]. The dataset is composed of votes of every member nation for every General Assembly resolution between the years 1946 and 2015. Each vote cast can belong to either one of the following categories (and their respective representations in the dataset): yes (1), abstain (2), no (3), absent (8) and not a member(9). Additionally, the data indicates under which thematic categories the resolution lies. These are: Votes relating to the Palestinian conflict (me), votes relating to nuclear weapons and nuclear material (nu), votes relating to arms control and disarmament (di), votes relating to human rights (hr), votes relating to colonialism (co) and votes relating to (economic) development (ec). Finally, starting from session 39, for each vote there is an additional attribute indicating whether it was identified as important by the U.S. State Department report on Voting Practices in the United Nations.

Our second dataset is a combination of international bilateral product trade data released by The Center for International Data from Robert Feenstra [8] (for the years 1962 to 2000) and UN COMTRADE (for the years 2001–2014)[5]. The full dataset is comprehensive, consisting of an exhaustive list of individual products imported and exported between each country pair for each year.

4 Voting Network

Using the voting data described above, we propose a similarity metric to capture pairwise relationships between different voting countries. The formula for the similarity score between countries i and j is represented mathematically as:

$$VoteSimilarity(i,j) = \frac{\#agreements(i,j) - \#disagreements(i,j)}{TotalMutualVotes(i,j)} \tag{1}$$

An agreement occurs when two countries both vote "yes" or both vote "no" on a specific resolution. A disagreement takes place when one country votes yes and the other votes no or vice-versa. The terms *#agreements(i,j)* and *#disagreements(i,j)* indicate the total count of agreements and disagreements, for a given range of years. *TotalMutualVotes(i,j)* represent the total number of resolutions that the two countries have voted upon, for the same years. Since the difference between *#agreements(i,j)* and *#disagreements(i,j)* always turns out

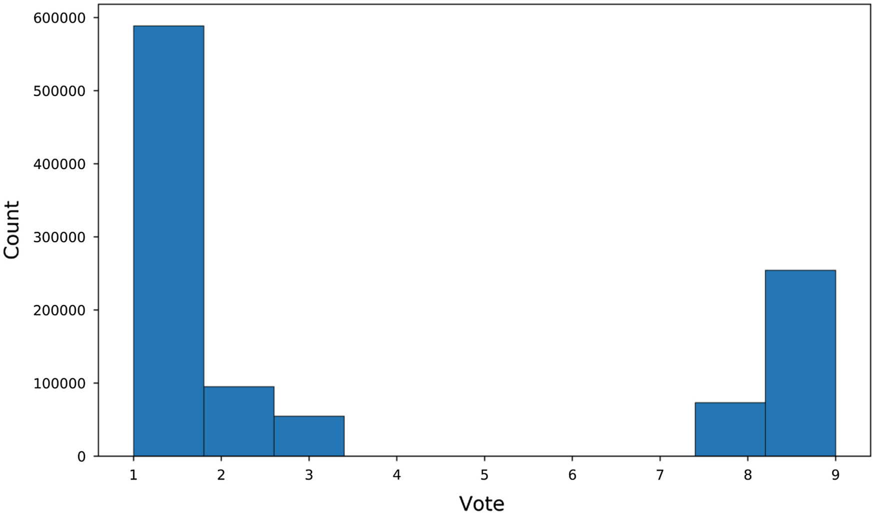

Fig. 1. Vote distribution in General Assembly voting data. 1: "Yes", 2: "Abstain", 3 : "No", 8: "Absent", 9: "Not a Member". Note the large number of instances of "Yes" votes, in stark contrast to "No", which have the lowest counts.

to be a positive quantity and that *TotalMutualVotes(i,j)* is a value larger than this difference, *VoteSimilarity(i,j)* is then a symmetric score between 0 and 1. Unlike in previous work [9] where the focus was on accounting for the agreements between country votes, we decided to penalize our similarity scores by subtracting the disagreements. One major motivating factor behind this decision was that an extremely large number of votes in the dataset are yes votes. If solely "yes" votes are considered, they positively bias country pairwise similarity. This is because in the dataset, agreements are almost always agreements on the yes votes. Predictably, this leads to unusual behavior such as countries that are normally expected to be adversaries end up having a similarity score approaching one. Figure 1 displays the number of votes for the choices- yes, no and abstain. Note that absent and not a member categories were not considered for our analysis because they do not provide any information about the behavior of the respective countries. Secondly, we remove resolutions where all countries voted "yes" because these resolutions do not help in differentiating between countries.

Upon cleaning the data, we use *VoteSimilarity(i,j)* to define the edge weights between every pair of countries (nodes) to create a complete voting network. To capture the strongest interactions among countries, we introduce some level of edge sparsity in the resulting graph by pruning all edges below a threshold. As a heuristic, we note that the expected similarity for any given pair of countries in the general incarnation of the network would be 0.5 (for a similarity score ranging from 0 to 1) which suggests a fair, plausible choice of threshold. In practice, we select the threshold to be 0.55, including a positive bias as a result of observing that the empirical distribution of computed scores is right-skewed towards 1.

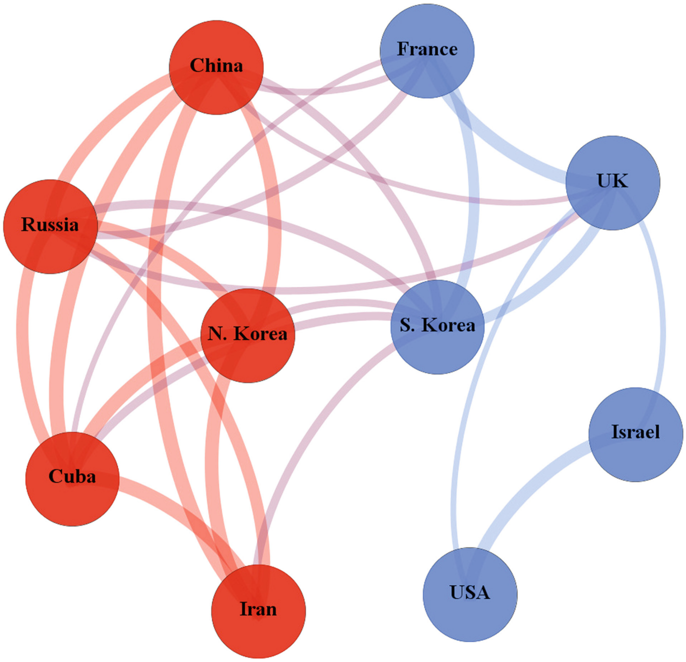

Fig. 2. Subgraph of the voting network containing a few core representative members from each community (differentiated by color). Thicker edges indicate higher similarity.

The value 0.55 is also low enough to ensure that the graph remain connected, since all nodes have a comfortably large number of edges with weights greater than that value. That said, we fix this threshold value and use it for pruning each voting networks we create.

Next, we run the locally greedy Louvain method on the network to unveil communities within the constructed graph. We select a range of years between 2000 and 2014 to capture the modern political landscape. For this version of the network, we use 180 countries after removing certain nations for which relevant data was missing. As a result, we observe the formation of 2 communities that closely align with real world international dynamics; see Table 1. For example, we see that the first community largely comprises developed western countries such as the United States of America, United Kingdom and France that are known to have common diplomatic interests over a large number of issues. The second community consists mainly of Asian and African nations, many of which are not as prosperous as their European counterparts and therefore vote very differently, particularly on resolutions involving development. Figure 2 visually depicts vote similarity relationships of some of the interesting nodes within the network.

Table 1. Representative countries of communities identified in the constructed voting network for the years 2000–2014

Community A	Community B
Canada, Spain, Australia, Netherlands, New Zealand, Albania, South Korea, France, Norway,United States of America, Ukraine, Finland, Sweden, Bulgaria, Romania, Portugal, Cyprus, Austria, Japan, Italy, Marshall Islands, Belgium, Georgia, Denmark, Poland, Israel, Iceland, Ireland, Hungary, United Kingdom, Greece	Ethiopia, Sri Lanka, Saudi Arabia, Kuwait, Liberia, Pakistan, Oman, India, Kenya, Afghanistan, Eritrea, Somalia, Peru, Cuba, China, Dominican Republic, Bahrain, Tonga, Libya, Indonesia, Vietnam, Russia, South Africa, Malaysia, Mozambique, Uganda, Brazil, Ivory Coast, Nigeria, Bangladesh, Iran, Algeria, Morocco, Uruguay, Lebanon, Egypt, Colombia, Sudan, Nepal, Philippines, Iraq, North Korea, Syria, Mexico, Congo

5 Diplomatic Path

As part of understanding the underlying relationships of different member nations, we implemented a mechanism to approximate the real-world diplomatic channels that result from these relationships. For example, consider the case of two countries i and j, where the edge weight between i and j is small, which in turn signifies low voting similarity. Within context, this is generally indicative of a sour relationship between the countries. Now, in a hypothetical situation (that occurs frequently in the real world) where it is important for i to communicate with j to lobby for a desired vote on a resolution, it becomes imperative that i engages with a diplomatic chain of mediators to ultimately achieve its goal. Assuming we restrict ourselves to a node cardinality of 4 (including i and j), the sequence may look something like i - k - l - j. The 3 edges ik, kl, and lj have an average edge weight value higher than the average edge weights of all possible paths between i and j containing 4 or lesser countries.

In general, for any two given countries we determine the path with the highest mean value. That is, for each path *P(i,j)* (a sequence of edge values) between two nodes i and j, with path length *L(i,j)*, we evaluate the value of,

$$AvgPathValue = \frac{\sum P(i,j)}{L(i,j)} \tag{2}$$

The path with the highest average value is selected as the diplomatic path between the nodes. It is important to note that most countries will not pursue a list of countries beyond a certain reasonable limit as the minor incremental benefit would not be worth the time and effort. Therefore, we limit the maximum number of edges the algorithm is allowed to traverse.

To illustrate our findings, let us take a bilateral relationship that is known to be volatile- that of the United States of America and Democratic Peoples

Republic of Korea (North Korea). We run our process onto the nodes, upper bounding the maximum path length by 3. The chain we obtain is United States of America - Japan - Singapore - North Korea. Ties between each successive pair are known to be warm, and we therefore have a diplomatic path that can very well be put to use. We have listed the results of some of the other simulations in Table 2.

Table 2. Some sample paths

Source	Target	Path
India	Pakistan	India - Egypt - Oman - Pakistan
USA	Israel	USA - Israel
Japan	China	Japan - Kuwait - Oman - China
UK	Argentina	UK - Cyprus - Chile - Argentina

Secondly, it is also interesting to analyze countries that occur most often within paths. It can be inferred that these countries try to maintain a balancing act between opposing factions of ideologies that determine voting blocs. The top 10 countries, along with the number of paths they occur in is given in Table 3.

Table 3. Top 10 countries occurring in most diplomatic paths (in descending order)

Country	Number of occurrences
Chile	5535
Oman	5156
Sri Lanka	2784
Indonesia	2027
Qatar	1913
Mexico	1817
Guyana	1730
Malaysia	1353
Kuwait	1351
Algeria	1138

6 Trade Network

We move on to analyzing how voting relationships align with other international pairwise associations. Consequently, we use the international bilateral product trade data described earlier, to create a network for all trade activity between 1990 and 2014.

Table 4. Community representatives found for trade network for the years 1990–2014

Community A	Community B
United States of America, Japan, United Kingdom, Canada, China, Australia	Syria, Russia, Iraq, Iran, Turkey, Libya, Egypt, Morocco, Algeria

The difference between the voting and trade networks is that the trade network is a multigraph with two edges between every possible pair of nodes, instead of one. The edges weights denote the relative importance of one country to the other in terms of trade. For example, for a country i, if the sum total of exports and imports with country j is 10% of the sum total of exports and imports with all countries i is in a trade relationship with, then the one edge from i to j would be a value of 0.1. Similarly, from the perspective of j, if the trade of j with i is 85% of all its trade, then the second edge from i to j is 0.85. The reason absolute trade volumes (or some normalized version of it) were not used is so that the dominant high volume trade flows do not dominate the low volume (but perhaps equally important in relative terms) relationships. For example, the trade between two G20 countries (which might not even be allies in the political sense) is likely to be many orders in magnitude higher than the trade between two ally developing countries. Therefore, it becomes imperative to use a relative measure that better represents their association. In a similar manner to the voting network, we remove some of the edges that are too low. In this case, we remove edges that have a value below 0.005.

Next, just like for the voting network we run the Louvain method to detect trading communities or blocks, which in this case turn out to be 2 as well. Again, this results in an intuitive set of country groupings. For instance, we notice United States of America, United Kingdom, Japan and others grouped together in one community, with Russia, Iraq, Syria, Iran and others in the second. Representative countries of each community are listed in Table 4.

Finally, we try to estimate the overlap between the voting and trade networks by comparing the members of the resulting communities. Just like in the first section, we create a voting network, but this time for the years 1990 to 2014, to be able to compare against the trade network. While we do this, we also ensure we are using the same set of countries for both networks. Next, we run the Louvain method on this new voting network and output the communities. Finally, we use Jaccard index [10] as a set matching metric to compare the groupings across networks. For any sets X and Y, the Jaccard index is defined as

$$J(X,Y) = \frac{|X \cap Y|}{|X \cup Y|} \tag{3}$$

We note significant similarity between the two sets of communities in Table 5. Voting community B has a Jaccard index of 0.718 with trade community B. If we remove voting community B and trade community B from the set, we observe

Table 5. Jaccard similarities for different community combinations

Vote community	Trade community B	J(X,Y)
A	A	0.487
A	B	0.666
B	A	0.176
B	B	0.718

that the remaining pair, that of voting community A and trade community A have a Jaccard similarity of 0.487, which is also considerable. This strongly suggests an overlap of the two networks, further indicating a relationship between bilateral trade and voting in the United Nations General Assembly.

7 Conclusion

We studied the interactions of countries through their voting behavior in the United Nations General Assembly and through their economic transactions with each other. We were able to derive an optimal diplomatic path. The diplomatic path is a series of countries to sequentially communicate with in the event a given nation wants to engage with another country with which it shares low voting similarity. We also showed the structural similarity of the voting network with the trade network by applying community detection on both networks and analyzing the result. The findings indicated that there was considerable overlap between the internal structures of the two networks because of the high Jaccard similarities between corresponding communities.

A few areas we wish to continue working on primarily involve exploring alternative metrics to better represent the inter-country relationships. More work can also be done on using specific resolutions and consequently specific commodity trade for analysis. For example, it would be interesting to observe the results of only using resolutions that deal with weapons juxtaposed with the trade of arms and weaponry to see if we get a stronger alignment. Finally, dynamic networks can be studied to understand how relationships evolve over time.

References

1. Barigozzi, M., Giorgio F., Diego, G.: Multinetwork of international trade: a commodity-specific analysis. Phys. Rev. E **81**(4), 046104 (2010)
2. Barigozzi, M., Fagiolo, G., Mangioni, G.: Identifying the community structure of the international-trade multi-network. Phys. A **390**(11), 2051–2066 (2011)
3. Blondel, V.D., et al.: Fast unfolding of communities in large networks. J. Stat. Mech. Theory Exp. **2008**(10), P10008 (2008)
4. Cherepnalkoski, D., et al.: Cohesion and coalition formation in the European parliament: roll-call votes and twitter activities. PloS one **11**(11), e0166586 (2016)
5. Comtrade, U.N.: UN Comtrade database. UN Comtrade Online (2015)

6. Crooks, A., et al.: International relations: state-driven and citizen-driven networks. Soc. Sci. Comput. Rev. **32**(2), 205–220 (2014)
7. Fan, Y., et al.: The state's role and position in international trade: a complex network perspective. Econ. Modell. **39**, 71–81 (2014)
8. Feenstra, R.C., et al.: World trade flows: 1962–2000. National Bureau of Economic Research No. w11040 (2005)
9. Ha, H., et al.: Proximity based circular visualization for similarity analysis of unga voting patterns. In: Workshop on Big Data Visual Analytics (BDVA), 2015. IEEE (2015)
10. Jaccard, P.: tude comparative de la distribution florale dans une portion des Alpes et des Jura. Bull Soc Vaudoise Sci Nat **37**, 547–579 (1901)
11. Lupu, Y., Traag, V.A.: Trading communities, the networked structure of international relations, and the Kantian peace. J. Conflict Resolut. **57**(6), 1011–1042 (2013)
12. Voeten, E.: Data and analyses of voting in the UN General Assembly (2012)
13. Zhu, Z., et al.: The rise of China in the international trade network: a community core detection approach. PloS one **9**(8) (2014)