Emergence of Consensus Proceeds in Phases: Tracking Themes Across the Negotiations of the WHO Pandemic Agreement

Keywords: NLP; networks; consensus; pandemic, diplomacy

Extended Abstract

Motivation. Traditional analysis of international negotiations, while indispensable, are inherently limited in their ability to quantify subtle, high-dimensional shifts in position, the latent semantic structures of disagreement, and the complex dynamics of coalition formation that occur across thousands of pages of documentation and hours of debate. In this work, we examine the negotiations that led to the adoption of the WHO Pandemic Agreement, one of the most important and recent international agreements, using tools of computational diplomacy (an emerging discipline which applies the tools of natural language processing, network analysis and machine learning to the unstructured data generated by diplomatic processes). We aim to answer the following questions: (1) How consensus evolves across the discussions, (2) what are the most prominent themes identifiable in the discussions leading to the agreement, and (3) whether these themes are persistent across negotiations.

Approach and Methodology. The dataset we used in this study is a collection of documents published as part of the negotiations leading to the Pandemic Agreement. We performed an automated download of the documents from the agreement's official website [2], obtaining a total of 214 documents. From each downloaded document, we extract the main text and chunk it hierarchically (i.e. according to sections, paragraphs, and sentences), obtaining a collection of text chunks which we then subject to topic modeling using `BERTopic` [1]. We then construct a network of the identified topics across the negotiation rounds ($N = 259, E = 492$), with the topics as the nodes and the Jaccard similarity between keyword sets associated with each topic as edges.

Results. We find a repeated pattern of the number of topics identified in the meeting documents 1 This pattern allows us to segment the entirety of the negotiation process into three phases, or regimes, each characterized by this "wave" pattern: an initial phase (Rounds 1 to 5), a middle phase (Rounds 6 to 8), and an end phase (Rounds 9 to 13).

The most prominent themes extracted from the documents pertain to the importance of international cooperation among countries in response to the pandemic. In contrast, the most persistent topics across negotiation rounds are none of these; the most persistent topic is about concerns that stakeholders should have their due place and engagement in the negotiations 2.

Conclusions and Outlook. We have examined the texts of negotiations for the WHO Pandemic Agreement using the tools of Computational Diplomacy, and identified the existence of phases as negotiations unfolded across time; that is, the emergence of consensus did not occur linearly over the rounds of negotiation, but in a succession of well-bounded phases.

Interestingly, we also found out that the most prominent themes (based on support from the documents) are not necessarily those which are persistent across time (based on semantic similarity). Studies to further elucidate these results are currently ongoing.

References

[1] Maarten Grootendorst. Bertopic: Neural topic modeling with a class-based tf-idf procedure, 2022.

[2] World Health Organization. Intergovernmental Negotiating Body. `https://inb.who.int`, 2021. Accessed: 2026-01-26.

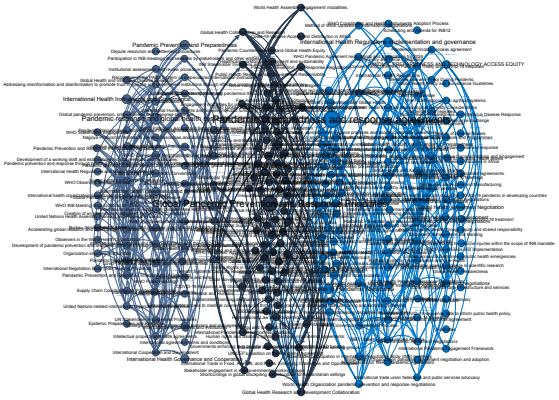

Figure 1: **Network of topics identified in discussion documents leading to the WHO Pandemic Agreement.** Topics were identified using the `BERTopic` framework, and labels were generated using a large language model from keywords extracted as part of the topic modeling; they are otherwise not used in the analysis. Nodes are topics; sizes are proportional to the average number of text chunks supporting that node; edge weights are Jaccard similarities between the paired nodes' respective keyword sets.

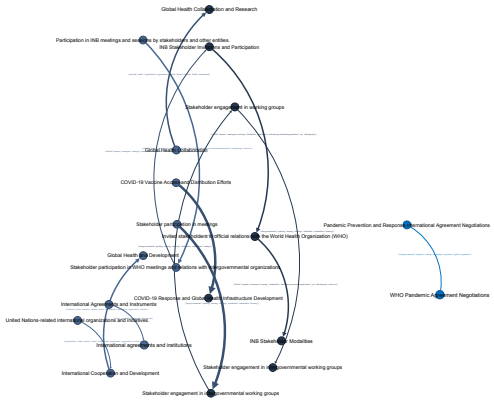

Figure 2: **Most persistent topics across negotiation rounds.** A topic is considered persistent across multiple negotiation rounds if the computed Jaccard similarites are 0.5 or higher.