


ARTICLE

Generative AI and Topological Data Analysis of Longitudinal Panel Data

Badredine Arfi 

Department of Political Science, University of Florida, Gainesville, FL, USA; Email: barfi@ufl.edu

(Received 1 May 2024; revised 21 August 2024; accepted 22 November 2024)

Abstract

This article constructs an approach to analyzing longitudinal panel data which combines topological data analysis (TDA) and generative AI applied to graph neural networks (GNNs). TDA is deployed to identify and analyze unobserved topological heterogeneities of a dataset. TDA-extracted information is quantified into a set of measures, called functional principal components. These measures are used to analyze the data in four ways. First, the measures are construed as moderators of the data and their statistical effects are estimated through a Bayesian framework. Second, the measures are used as factors to classify the data into topological classes using generative AI applied to GNNs constructed by transforming the data into graphs. The classification uncovers patterns in the data which are otherwise not accessible through statistical approaches. Third, the measures are used as factors that condition the extraction of latent variables of the data through a deployment of a generative AI model. Fourth, the measures are used as labels for classifying the *graphs* into classes used to offer a GNN-based effective dimensionality reduction of the original data. The article uses a portion of the militarized international disputes (MIDs) dataset (from 1946 to 2010) as a running example to briefly illustrate its ideas and steps.

Keywords: generative AI; graph neural networks; topological data analysis; spatial auto-regressive model; Bayesian statistics

Edited by: Jeff Gill

1. Introduction

Analyzing longitudinal panel data in political and social sciences is essential for addressing critical research questions (Montfort and Oud 2010; Traunmuller, Murr, and Gill 2015; Reuning, Kenwick, and Fariss 2019; Zheng, Lv, and Lin 2021; Ye *et al.* 2023; Loaiza-Maya *et al.* 2022; Sterner *et al.* 2024; Mai, Zhang, and Wen 2018; Ferrari 2020). This type of data typically involves measurements of unit features grouped spatially and temporally, including features of dyads and structural properties of clusters or entire groups. When considering bilateral and multilateral relations cross-sectionally and temporally, along with the underlying structures, longitudinal panel data not only captures these dimensions but also incorporates relational dynamics and structural contexts, including dynamic blocks or clusters of entities that change over time (Lupu and Greenhill 2017; Olivella, Pratt, and Imai 2022).

Several methodologies, such as Hierarchical/Multilevel Modeling, Time Series Analysis, Causal Inference, Machine Learning Approaches, and Structural Equation Modeling (SEM), have been employed to analyze these aspects of the data. While dyadic analysis is predominant, cluster analysis, particularly through social network analysis, provides additional insights by accounting for correlations within and across panels, leading to more accurate estimations of standard errors (Arellano and Bonhomme 2023; Carlson, Incerti, and Aronow 2024). However, traditional dyadic and cluster analyses struggle with scalability as the data size grows. They are also limited by statistical constraints, such as

model assumptions, which become problematic with highly complex, non-linear models involving big datasets with a large number of features.

In contrast, deep learning approaches, particularly those utilizing Graph Neural Networks (GNNs), offer flexibility, adaptability, and scalability to handle big datasets with a huge number of parameters. GNNs excel at capturing complex relationships within graph-structured data by iteratively aggregating and updating information from neighboring nodes and subgraphs. GNNs are particularly effective for learning meaningful representations of nodes, edges, and more complex structures like clusters and blocks, capturing both local and global patterns in a data. They accommodate heterogeneous data types (numerical, images, video, texts, etc.) and manage complex interactions between them. GNNs, when combined with generative AI and deep learning, not only provide variables for Bayesian models but also uncover deep patterns, enhance data visualization, improve predictive accuracy, and strengthen interpretative power. Their scalability to large datasets with numerous features, without prior assumptions, sets GNNs when analyzed using generative AI apart from traditional statistical approaches. Furthermore, GNNs provide novel methods for analyzing deep-seated heterogeneities in graph data.

The analysis of the effects of unobserved heterogeneities has been flourishing in statistical approaches. Yet, the study of the effects of *topological heterogeneities* is rather lacking, especially in social sciences. This article proposes a framework for analyzing unobserved topological heterogeneities in longitudinal panel data by integrating generative AI, GNNs, and topological data analysis (TDA). Every dataset has an inherent topology, often containing heterogeneities that shape how the data explains empirical phenomena. This framework introduces a quantitative measure of topological heterogeneities, setting them apart from other forms of unobserved heterogeneities (Ferrari 2020). The article brings together Bayesian statistical analysis and TDA to compare the proposed measure with conventional statistical estimates of effects. Moreover, using generative AI applied to GNNs, the article proposes an approach to extract latent variables, in contrast with traditional methods that assume their existence and estimate their impact (Fong and Grimmer 2023; Loaiza-Maya *et al.* 2022; Mai *et al.* 2018; Montfort and Oud 2010; Reuning *et al.* 2019; Sterner *et al.* 2024; Trautmüller *et al.* 2015; Ye *et al.* 2023; Zheng *et al.* 2021). The combination of TDA, GNN, and generative AI for analyzing longitudinal panel data reconstructed as graphs offers new research opportunities, particularly in political and social sciences, where data complexity demands methods beyond traditional statistical models.

The article is structured as follows. It begins with a basic Bayesian spatial auto-regressive (SAR) model, incorporating a sociomatrix \mathbf{W} representing dyadic memberships in social networks (Kenny, Kashy, and Cook 2020; Wasserman and Faust 1994). This model establishes a framework for comparing the effects of measures obtained through TDA and generative AI with those from traditional parametric models. Second, the article introduces TDA of longitudinal data and converts its results into topological factors that are used in Bayesian statistics. Third, the article constructs a deep GNN machine-learning topology-based classification of longitudinal panel data after converting it to graphs, where nodes represent units of the panels and edges represent their dyadic relations, each with their respective features. Fourth, two generative AI models within the GNN framework are presented. The first model extracts latent factors of the data, constrained by topological heterogeneities, and shows their accuracy in representing underlying data features. Their statistical effects are then estimated through a Bayesian model. The second model performs data-dimensionality reduction by identifying key subgraphs that are used to generate a smaller longitudinal dataset. The article shows that such a reduced-dimensionality dataset achieves strong out-of-sample predictive performance in a Bayesian model. The article concludes by reflecting on the implications for social science research and a flowchart outlining the practical steps of the framework. An accompanying online supplement provides additional details and mathematical foundations. The article uses a dataset (1946–2010) on MIDAs for illustration (Hafner-Burton and Montgomery 2006; Kinne 2013; Lupu and Greenhill 2017; Pevehouse, Nordstrom, and Warnke 2004; Shannon, Morey, and Boehmke 2010, for details see the supplement).

Table 1. Bayesian estimation with sociomatrix expansion.

	Variable	mean	sd	hdi_3%	hdi_97%
Empirical variables	majpower	4.287	0.142	4.027	4.558
	allies	−0.676	0.121	−0.899	−0.445
	contgkf	1.386	0.138	1.118	1.634
	smldep	−8.060	0.244	−8.519	−7.599
	lcaprat	−5.234	0.333	−5.874	−4.622
	lgmidist	−8.137	0.272	−8.648	−7.628
	smlmat	−4.167	0.195	−4.540	−3.819
Sociomatrix effects	ρ	0.153	0.007	0.141	0.166
Zero-inflation	ψ	0.055	0.004	0.047	0.063

2. SAR Model

This section introduces a base Bayesian model for comparing later on the effects of the constructed latent variables to commonly used variables in the study of MIDs. I choose a spatially-auto-regressive (SAR) model to take into account the effects of the sociomatrix on the probability of MIDs. The SAR model is shown in Equation (2.1):

$$\mathbf{Y} = \rho \mathbf{W} \cdot \mathbf{Y} + \mathbf{X} \cdot \boldsymbol{\beta} + \boldsymbol{\varepsilon}, \quad (2.1)$$

where \mathbf{X} is a matrix of exogenous covariates, \mathbf{W} is a $N \times N$ sociomatrix representing the relations between various nodes (N dyads of states) of the network, ρ is a measure of the impact of social network effects on the dependent variable, and $\boldsymbol{\varepsilon}$ is an error term. To keep the computations feasible given the very large size of the sociomatrix, I expand \mathbf{W} in terms of powers of $\rho \mathbf{W}$. This makes sense by choosing $\rho \in [0, 1]$ in a typical SAR model and \mathbf{W} is normalized (with components in $[0, 1]$). As shown later, one need not go beyond ρ^p with $p \in (1, 2, 3, 4)$; higher powers do not in any meaningful way affect the results of the SAR model. Because the MID (outcome) variable is strongly zero-inflated, I use a zero-inflated binomial distribution for the likelihood function. All priors of the Bayesian model are chosen to be weakly informative. The vector of coefficients $\boldsymbol{\beta}$ includes the covariates of the empirical data.

Table 1 displays the results of the estimation. ρ has a mean of 0.153 with a 94% high density interval $[0.141, 0.166]$, which shows that the overall contribution of the main ρ term of the expansion has a small effect on the occurrence of MIDs in the dyads, with even smaller contributions of higher powers of the sociomatrix.

One can conclude that state membership in IOs considered as a network of relations shapes the likelihood of MIDs in relatively moderate ways (in conformity with earlier work (Kinne 2013)). I next construct a TDA of the data, the results of which—topological factors—will be added to the Bayesian SAR model as moderators. The introduction is meant to be illustrative for pedagogical reasons since TDA is not a topic that most social scientists are familiar with. The Supplementary Material provides more in-depth and mathematical introduction.

3. TDA

TDA begins with the premise that a data considered as a *whole* possesses certain invariant global properties—*topological features*—which constitute a signature of sorts for the data.¹ We first construe the data as a *point cloud*, that is, a (mathematical) *manifold*, wherein the data points are located. Each such manifold possesses properties of a qualitative nature which inform about its topologically connected

¹For much more, see Arfi (2024) and the references therein.




Figure 1. Examples of Cech and Rips–Vietoris simplicial complex.

and/or disconnected components. TDA uncovers this information. TDA proceeds by building high-dimensional equivalents of neighboring graphs using both connecting pairs and $(k+1)$ -uple of nearby data points. The objects so constructed are called *simplicial complexes*. They are used to identify key topological *defects* at various levels of dimensionality (Chazal and Michel 2021, 3). TDA therefore effectively projects the underlying topological properties of the data manifold into a space of simplicial complexes known as combinatorial graphs.² The nodes of these combinatorial graphs are the data points and the edges are whatever mutual proximity and relations that we decide to focus on. The goal of TDA is to find the invariants of this space, which are by extension a representation of the subterranean topological invariants of the data space (Ghrist 2014). Different prescriptions are used to build the simplicial complexes; examples of which are *Rips–Vietoris*, *Cech*, and *Alpha complexes* (Hatcher 2002). Figure 1 is a schematic illustration. The red dots are the original data points, and the colored geometric shapes are the identified simplicial complexes at different dimensions (Ghrist 2008, Figure 2). These prescriptions uncover the underlying (qualitative) topological properties of the data. A simplicial complex is in a sense a generalization of the notion of a graph (Ghrist 2017, 4). The study of these simplicial complexes is called *homology*, and the study of the most persistent features is called *persistence homology*.

I use TDA to probe how the persistent homological effects of the empirical data (taken as a whole) change over time. Following Arfi (2024), I focus on the first three topological effects, respectively, represented as H_0 , H_1 , and H_2 , each with two principal components. I start with the working proposition that these homological effects might be statistically relevant in explaining part of the variance of MID data. Due to the deep structural changes that have been occurring in the international system since World War II (for example, a rapid increase in the number of sovereign states and their joining existing and newly created IOs), one expects important changes in the topological features of the empirical data to be reflected in these changes.

The computation of persistent homology proceeds as follows. Using what is called a *filtration* process,³ we record the moments of *appearance* and *disappearance* of the topological features as we probe the data manifold (as in Figure 1). We thus record how discernible a topological feature (such as a loop, triangle, or tetra-pod) is as we are probing the data manifold through a process of filtration. The more persistent topological features are, the more important for the global topology of the data manifold they are. The information that is unearthed through filtration is saved in mathematical objects called *persistence diagrams* (PDs). These objects record the existence, nature, and concentration of the invariant topological features. They are usually visualized as a two-dimensional plot of the *disappearance or death*

²The theory of combinatorial graphs focuses on counting and enumerating various possible arrangements of vertices and edges, etc. (Harris, Hirst, and Mossinghoff 2008).

³Filtration is a method for building a nested sequence of topological spaces. Each space in this sequence captures the connections and disconnections between data points at a particular scale. These spaces are then converted into simplicial complexes that simplify computation. The persistence of features of these spaces across scales reveals the underlying topological structure of the data.

Persistent Diagrams

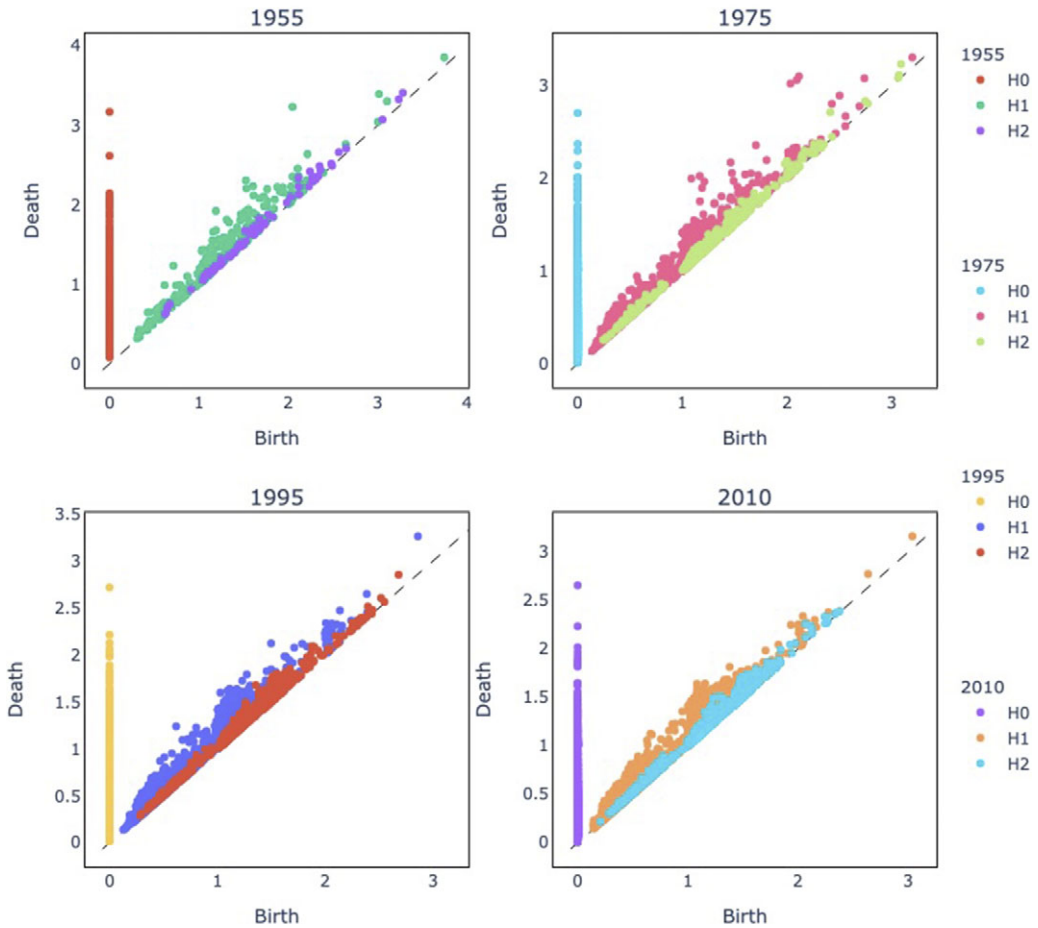


Figure 2. PDs for four years.

versus the *appearance or birth* of the invariant topological features of the data manifold as we go through the filtration.⁴ Four such PDs are displayed in Figure 2.

For any data manifold, we usually obtain a multitude of PDs which thus form a space. This space is then equipped with a generalized (non-Euclidean) notion of distance which makes it possible to analyze its properties. The PDs are complicated mathematical objects to deal with, which led students of algebraic topology to construct what is known as functional summaries of the persistent diagrams, which are then used to process the information stored in the PDs and use it in, for example, statistical analysis. One such functional summary is *persistence image functional* (Obayashi, Hiraoka, and Kimura 2018).⁵

PDs, such as those displayed in Figure 2, are essential tools for analyzing topological features in the data. Points that lie on or near the diagonal line represent features that persist across multiple spatial scales. These are considered significant and stable topological features of the data. The further a point is from the diagonal, the more persistent it is (the diagonal line represents death moment = birth moment).

⁴For a more in-depth discussion, see Arfi (2024), and references therein.

⁵See Arfi (2024) and its online supplementary material for many more details.

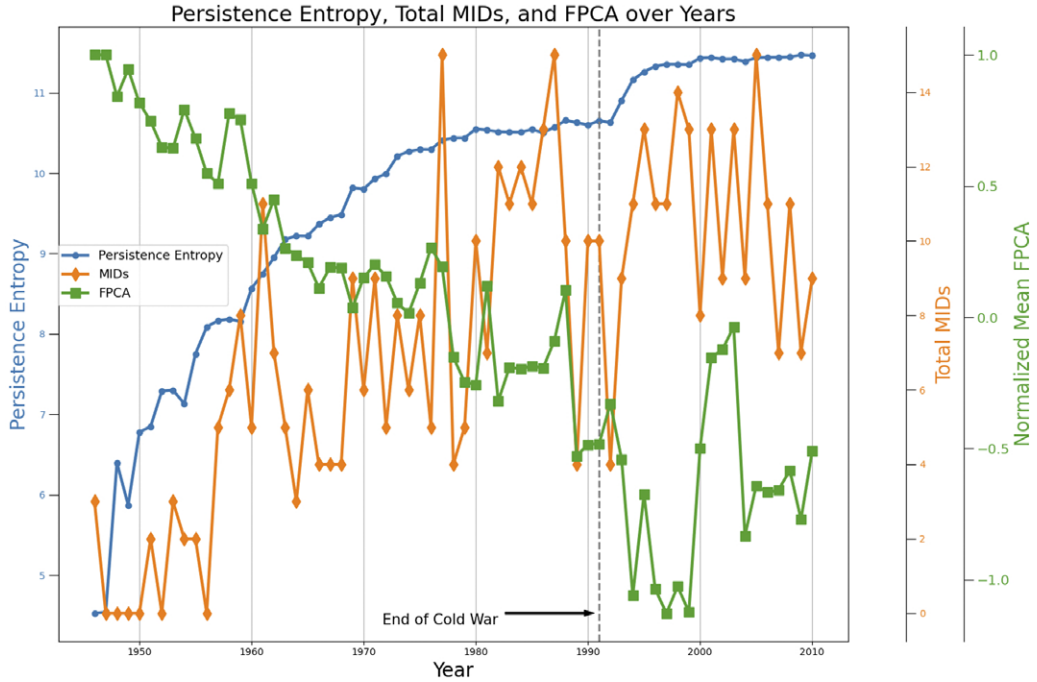


Figure 3. Persistence entropy, total MIDs and normalized mean-FPCA.

Features with a birth moment close to 0 and a death moment close to infinity correspond to connected components (0-dimensional features) labeled H_0 . Persistent topological holes in the data correspond to one-dimensional features, labeled H_1 , while voids such as triangles and tetra-pods correspond to higher-dimensional features, labeled H_2 .

The examination of the space of PDs is deepened by considering *persistence entropy*. In information theory and statistics, entropy is used as a measure of disorder and uncertainty. High entropy indicates strong disorder, while low entropy indicates weak disorder. Persistence entropy is used to quantify the complexity of and uncertainty in the topological features extracted from the data. Persistence entropy is obtained by discretizing the PDs into bins or intervals along the persistence axis (vertical axis). The entropy is then computed based on the distribution of points in these bins. High persistence entropy indicates that topological features are distributed more evenly across different persistence levels, suggesting greater topological complexity or randomness in the data. Low persistence entropy suggests that topological features are concentrated in specific persistence ranges, indicating a more structured or ordered dataset. Therefore, persistence entropy can be used to compare different datasets or to track changes in the topological structure of a dataset over time or across different conditions. The formula for persistence entropy is given in Equation (3.1):

$$E(PD) = - \sum_{i=1}^n \left(\frac{l_i}{L} \right) \log \left(\frac{l_i}{L} \right), \tag{3.1}$$

where PD is a persistence diagram consisting of n persistence pairs, where $l_i = d_i - b_i$ is the lifespan of the i -th feature, and $L = \sum_{i=1}^n l_i$ is the total persistence (sum of all lifespans) in the diagram. To illustrate the notion of persistence entropy, I plot in Figure 3 its variation for the illustrative data in conjunction with the total MIDs per year in the interstate system (as well as the normalized mean of the functional principal components (FPCA) discussed down below).

Persistence entropy steadily increases over the whole time period. The topological heterogeneities are, therefore, becoming increasingly more diverse and their numbers are increasing over the years. The

deep structure of the international system is. therefore, increasingly more complex and deeply more heterogeneous in nature. As displayed in Figure 3, this concords with more MID's in the international system over the years; the two trends are remarkably concordant. The vertical line points to the end of the Cold War and collapse of the Soviet Union. This corresponds to a jump in both the persistence entropy and the total MID's, with both more or less reaching a plateau between the end of the Cold War and 2010. These plateaus are preceded by prior plateaus in the 1980's decade. This decade corresponds to tumultuous changes in the international scene during the Reagan era as the architecture of the world order once again was at stake, much like at the end of World War II. The biggest increases in both persistence entropy and total MID's occur (in tandem) between the end of World War II and 1980. These decades did indeed witness an overhaul of the international system with the start of the Cold War and its rivalry, but also with increases in the levels of cooperation through multilateral organizations and institutions.

Another way to probe the extent to which topological heterogeneities affect the evolution of the international system is to transform the information stored in the PD's into quantities that can be manipulated using usual algebraic frameworks (PD's are multi-sets and hence cannot be manipulated using usual algebra). Following closely Arfi (2024) and Wang (2016), I transform the information in the PD's into persistence images which are then vectorized and transformed into FPCA (a generalization of principal components analysis to functions). A persistence image represents complex information stored in the PD's by capturing its significant features, and turning this information into a simple image (see Supplementary Material for mathematical details). We then use PCA to reduce the dimensionality of the vectorized persistence images and capture the main features using the top principal components, called topological components. I choose to represent each topological component by two FPCA; these are

$$H_0 \Rightarrow (H_{0_1}, H_{0_2}), H_1 \Rightarrow (H_{1_1}, H_{1_2}), H_2 \Rightarrow (H_{2_1}, H_{2_2}).$$

Figure 3 displays, in addition to entropy and MID's, the mean FPCA which shows a decreasing trend over the years if with some irregularities in this respect, increasing sharply at around 2003 (which, historically, corresponds to the US invasion of Iraq) and then decreasing by the mid-decade to what roughly is a plateau. This is roughly in accordance with the observations made on the evolution of persistence entropy, which is not surprising since both are different functional transformations of the topological information stored in the PD's.

In the next section, I augment the list of covariates (representing the empirical data) with two homological factors— H_{1_1} and H_{1_2} —which are the two first FPCA corresponding to the one-dimensional topological effects (one-dimensional loops). The question is whether these topological effects are statistically relevant for explaining part of the variance in the posterior distribution of the dependent variable, MID.

4. SAR Model with Topological Factors as Moderators

Because topology reflects global properties of the data, it makes sense to consider topological effects as moderators used to probe how micro-level features (covariates) are conditioned by macro-level topological properties. To this end, the SAR model Equation (2.1) is modified as follows:

$$\begin{aligned} \gamma_i &= \sum_{j=1}^{n_{\text{vars}}} X_{ij} \left(\sum_{k=1}^{n_H} H_{ik} \delta_k \right) + \sum_{k=1}^{n_H} H_{ik} \zeta_k, \text{ moderation effects} \\ \eta_i &= \gamma_i + (1 + \rho \mathbf{W} + \rho^2 \mathbf{W}^2 + \rho^3 \mathbf{W}^3 + \rho^4 \mathbf{W}^4) \cdot \mathbf{X}_i \cdot \beta \\ y_i &\sim \text{ZeroInflatedBinomial}(\psi, N_{\text{obs}}, \text{sigmoid}(\eta_i)), \end{aligned} \quad (4.1)$$

where ζ are the coefficients for topological effects; δ coefficients for moderation terms; \mathbf{H} is the matrix of moderators/topological factors (dimension: $N \times n_H$), with the rest of the quantities defined as before. A Bayesian estimation leads to the results in Table 2 (just showing the topological effects).

Table 2. Model with sociomatrix and topological moderator effects.

	Variable	mean	sd	hdi_3%	hdi_97%
Topological effects	H_1_1	−5.124	0.357	−5.760	−4.421
	H_1_2	−10.284	0.585	−11.378	−9.198
Moderator effects	δ_{H11}	2.388	0.184	2.056	2.748
	δ_{H12}	5.022	0.344	4.346	5.635

H_1_1 and H_1_2 have strong negative effects with high credible intervals not including zero, indicating important negative relationships. The moderator effects δ_{H11} and δ_{H12} show strong positive effects, suggesting important positive relationships between these moderators and the dependent variable.

I compare the results of this model to a *bare* model (with no sociomatrix and no topological effects) using the *Expected Log Pointwise Predictive Density* (ELPD) measure (Vehtari, Gelman, and Gabry 2017) (see Supplementary Material). The model containing the sociomatrix and topological effects as moderators is ranked highest. This is not surprising since topological effects are global characteristics of the data considered as a whole and are thus expected to shape the analysis of the data (given the rough comparison of the evolution of total MIDs and persistence entropy in the previous chapter). This suggests that specifying topological effects as moderating linear predictors is an adequate way to put to work the topological information unearthed through TDA. However, the topological information extracted via TDA and stored in the topological measures can also be put to work in other more versatile ways. In the remaining parts of the article, I show that specifying the role of topological effects differently does indeed demonstrate their importance to the analysis—that is, using them as classifying and conditioning factors of the data to uncover patterns within the data and to probe its deep latent features. This is done using deep machine learning and generative AI analysis.

5. Deep GNN Analysis

As shown in many studies social networks and other structures such as blocks and clusters play important roles in the analysis of longitudinal panel data. I propose that GNNs take this to new levels. To this end, I restructure the data as a set of graphs, with one graph per each yearly panel, (in the empirical illustration, states are nodes and their mutual relations as edges). The nodes can have many attributes (just two in the illustration: being a major power and being involved in a number of MIDs). The dyadic relations between nodes are taken as multiple attributes of the edges (such as joint memberships in IOs and other features in the illustration). This approach makes it possible to consider simultaneously nodes and their dyadic relations (edges) in the same framework, each with its *multiple attributes*. In conventional graph analysis (such as social network analysis), edges can have weights which reflect the strength of the relations between nodes. Using GNNs, allows us to include neighborhood and cluster relations between nodes and how these entities communicate in a graph to influence one another. This implies that not only are we including dyadic relations, but also more complicated relations. This is captured through the process of *attention* (called graph attention network [GAT]) as explained later (and more so in the Supplementary Material).⁶

I carry out the analysis by constructing three GNN models, one for node classification and reconstruction, one for graph classification, and one for graph generation. Node classification is within each graph with the values of the dependent variable (MID) taken as classification labels. This *node classification* paves the way for identifying latent factors, which are then used as covariates in a Bayesian model. The construction of these latent variables is done under a regularization (conditioning)

⁶Attention mechanisms assign different weights to different nodes in a neighborhood based on their relations with other nodes and their attributes. The attention mechanism are modelled as learnable weights.




Figure 4. GCN model (Kipf and Welling 2017).

scheme using a discretized version of the topological factors obtained through TDA in Section 2. The regularization process is the GNN generalization of the moderation conditioning process considered earlier in the article. *Classification of graphs* uses the topological factors as labels for graphs which paves the way for a topology-based clustering of graphs. The third is a GNN model which offers a process of graph generation that allows us to zero-in on the most important subgraph of each graph (as explained down below), the effects of all such subgraphs on the occurrence of MIDs are then estimated through a Bayesian model.

5.1. GNN Models

GNNs are a generalization of neural networks to graph data (Labonne 2023). GNNs encode representations of nodes and edges that depend on the graph structural properties, as well as node and edge features. A GNN functions through *neural message passing* which passes information between nodes and is updated through the neural networks. One of the most versatile forms of GNNs is the Graph Convolutional Network (GCN) architecture (Kipf and Welling 2017). This is an adaptation of Convolutional Neural Networks (CNNs) usually used for deep learning about images and texts to graphs (Zhou *et al.* 2020). A GCN is specifically defined with a layer-wise propagation rule, which usually is a non-linear function that updates node and edge features by aggregating the features of their neighbors. GCNs efficiently leverage the graph structure by performing convolutions, thereby enabling effective feature extraction for tasks like node classification and embedding, edge embedding and prediction, and graph embedding and classification. In this article, I first use a deep GNN model which includes multiple layers in the model propagation as shown in the generic example of Figure 4.

Generally, there are three types of layers: input, hidden, and output. These layers are chosen depending on the task at hand. In this article, I choose as the basic GNN layer for purpose of graph classification what is known as GAT to incorporate not only node features but also (as explained before) various edge features. GATs make use of *attention processes* to give neighboring nodes varying weights depending on their relative relevance in the graph (Zhou *et al.* 2020), with the weights learned during the training phase. This enables the model to zero in on the most relevant information in the graph. GAT thus dynamically enhances the ability of the model to capture highly complex interactions in graph-structured data. I specifically use the *EdgeGATConv* layer (see Supplementary Material) as the base for constructing three different architectures geared, respectively, toward (a) graph classification using topological labels, (b) conditioned extraction of latent features, and (c) graph generation.

5.2. Graph Classification

The topological information stored in the FPCA is discretized into four categories which are then used to classify the graphs into four topological classes through supervised deep machine learning.⁷ The list of

⁷From now and onward, I only consider the H1 component of topological effects, which corresponds to topological loops.

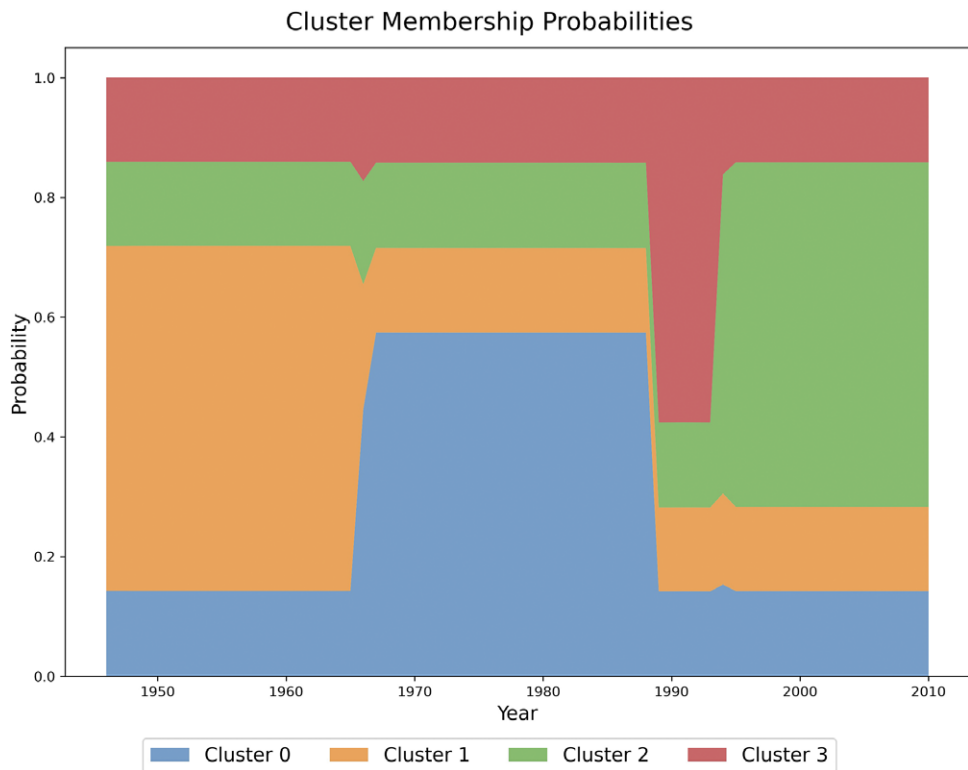


Figure 5. Proportions of clusters over the years.

graphs is split into training and testing sets of graphs.⁸ The goal is to calculate the accuracy of predicting graph distribution after training the GNN classification model.

Having trained and saved the model, we can use the saved model to examine, for example, how the membership probabilities in topological classes change from graph to graph, that is, how the topological structure of the international system of states evolves from year to year.

Figure 5 displays the yearly proportions of topological classes. The colors indicate the different topological classes from the data. The boundaries between the class probabilities change over the years, which means that the topological structure of the international system varies over the years, with most variation occurring between the mid-1960s and the late 1980s, the latter corresponding to the end of the cold war, with Class 0 being predominant by far. During these twenty years or so the international system saw many changes in addition to the re-ignition of the cold war. We then see important shifts in the boundaries, with Class 3 gaining in membership whereas the others shrinking, if moving toward stabilization, during the 1995–2010 period, with Class 2 being the dominant. The few years in the aftermath of the cold war, which witnessed much instability in world order, are manifested in the sharp variation of the probabilities for the four classes, with Class 3 being dominant during that period.

So far, these plots have been done using the predictive power of a graph classification model. However, the basic layers of the model utilize, as explained before, the attention mechanisms, the weights of which can be learned. These weights are edges weights (not nodes) and hence we can use them to extract some useful information concerning how edge importance is propagated in the neighborhoods (which is what the attention mechanism focuses on). I first look at the global picture by focusing on the whole set of graphs (that is, from 1946 to 2010). Figure 6 shows how the maximum attention weights

⁸See Supplementary Material for the distribution of the *true* values of the graph labels.

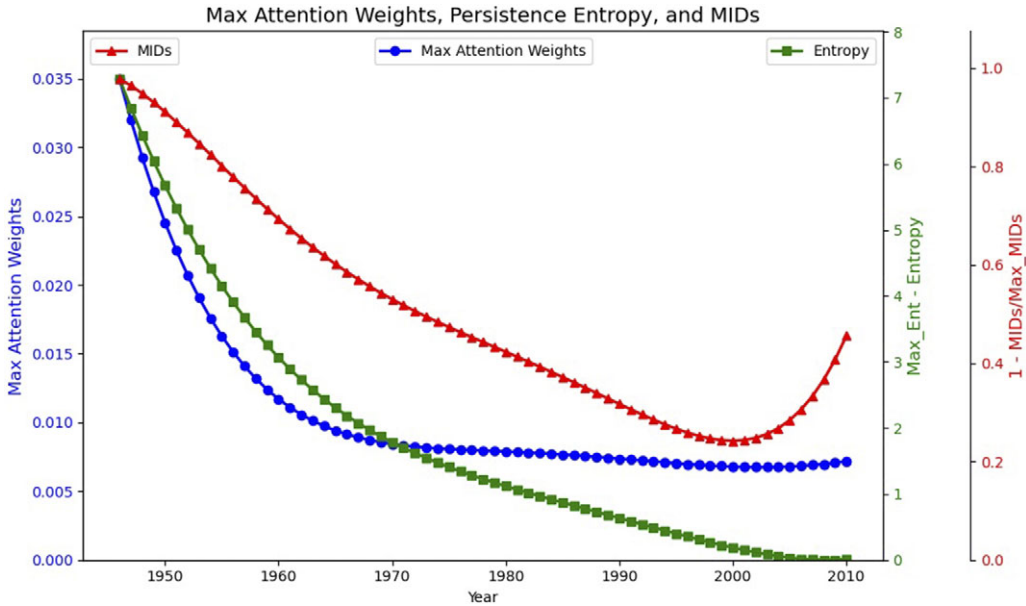


Figure 6. Max attention weights, persistence entropy, and MIDs.

vary in conjunction with persistence entropy and total MIDs in the system (all three have been spline-smoothed for comparison clarity purposes).

The time trends of the three quantities are quite similar in general. This signifies that the attention mechanisms are also able to capture the dynamics in terms of MIDs, much like persistence entropy does. Persistence entropy is a measure of topological disorder, total MIDs per year can be understood as a measure of political disorder in world politics, the evolution of both is roughly speaking picked up by the maximum attention weights of the edges (relations between states). The upward behavior of the total MIDs curve is not captured by either the persistence entropy or the attention weights.⁹

Second, we can also zoom in by finding which edges have the top k attention weights in any single year (graph). Plotting these, we obtain Figure 7 which displays, at the top, the subgraphs with the eight highest attention weights for 1985 and 2010. The number displayed in the edges are the normalized attention weights between the respective nodes (states). Not only does the structure of the subgraph change but also the nodes (states) composing the subgraphs. The lack of an edge between any two nodes means that there isn't any meaningful attention propagation between them. Remember that these attention mechanisms provide information on which edges are important in their neighborhoods, and how information is passed through the neighborhood. Edge-level attention focuses on assigning weights to edges rather than nodes. This is particularly useful in scenarios where the relationships between nodes (represented by edges) carry different levels of importance. In the empirical illustration, much information consists of features of the edges in a graph, and hence attention weights allow us to learn the importance of these features in the model, both locally and globally. The bottom two plots represent the subgraphs for the same eight states with the normalized number of shared IOs as a weight (red) in 1985 and 2010. These IO-sharing membership weights are not learned; they are computed from the sociomatrix. The attention weights are optimally learned, which means that they capture the dynamic relations as the neural networks are fired throughout training, that is, they reflect the dynamic and highly complex short and long "range" correlations in the data across the neural networks. While not done in here, attention weights can be used for classification as well as prediction purposes. They can

⁹Perhaps including other topological components such as (H_1_2 and H_2_1) might be able to rectify this.

IO-Weighted and Attention-Weighted SubGraphs

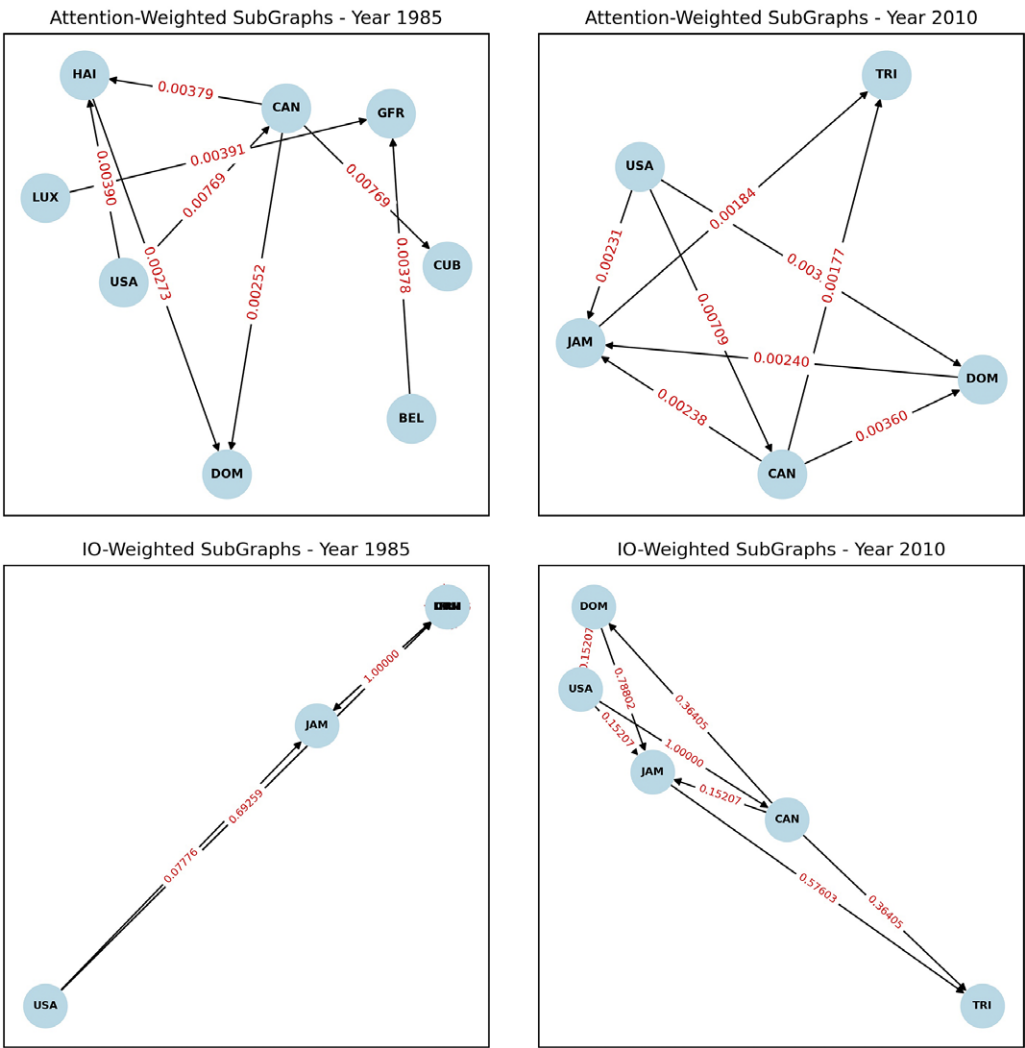


Figure 7. Subgraphs with top eight attention and IOs-sharing weights.

also be used to draw heatmaps of the various topological classes as a function of time, thereby providing useful information on the evolution of the system of graphs.

All results obtained in this section, used the TDA-generated topological measures as “classification labels”. In the next section, I construct a model which uses the topological effects as a constraint that conditions the space of latent factors. This is an extension of the idea of using topology effects as moderators in the Bayesian model as explained earlier.

5.3. GNN-Extracted Latent Features

Analysis of latent factors is used in statistical analysis to reduce dimensionality, improve model accuracy, and uncover relationships not immediately observed, using tools such as Confirmatory Factor Analysis (CFA) and SEM. Latent factors are important for understanding causality by controlling for unobserved con-founders, thereby leading to more precise conclusions. Because GNNs handle




Figure 8. Conditioned ARGVA model with the topological conditioning done through embedding: $Z \times \zeta \rightarrow Z$.

well high-dimensional data and capture intricate, non-linear, and complex relationships by learning representations that fall within the graph's topology, they are particularly, suited for detecting latent factors, thereby providing insights into the underlying structure that traditional methods might miss. This is due to the internal iterative and dynamic training and validation processes built into the GNN models. This section puts into play these strengths of GNN models to unearth latent factors of longitudinal panel data.

This section extends a model called *Adversarially Regularized Graph Variational Auto-Encoder* (ARGVA) architecture (Pan *et al.* 2018), to extract latent space features *conditioned* with the TDA features. The two key aspects of this model are for being *adversarial* and a *variational auto-encoder*, plus the important role that *regularization* plays overall. A generative GNN model can additionally be specified by imposing a topological constraint on extracting latent variables by adding a regularizer to the process of propagation through the various neural networks. This is based on the argument that a latent variable must be conditioned by the deep seated topological properties of the observed data since it is generated in the same manifold as the observed data, the deep properties of which are reflected in the TDA-unearthed topological factors. The model is schematically represented in Figure 8, where ζ represents the topological effects constraint and Z the latent variables.

The model effectively encodes the structure, node, and edge contents of a graph into a compact representation-embedding. The embedding is *moderated* by the topological effects (ζ). The embedding is then submitted to a decoder which is trained to reconstruct the input graph structure. The intermediate latent representation is internally forced to match a learnable probability distribution through an adversarial training (discriminator) module. We then jointly optimize the graph encoder learning and adversarial regularization to obtain the best graph embedding in a lower dimensionality compact and continuous feature space (from a latent space with, for example, 64 dimensions to a two-dimensional space in this article because the illustration has two features per node/state), while preserving the information about the graph structure, topological constraint, and node and edge features into the embeddings (Pan *et al.* 2018; Zhang *et al.* 2017).¹⁰ I use *EdgeGATConv* as the base layer in the encoding, decoding, and discriminating phases, which includes both node and edge features in the computation as well as attention mechanisms. The Encoder, Decoder, and Discriminator *qua* neural networks are constituted of large numbers (100) of these layers stacked together, with thousands of learnable parameters. Formally, the neural networks are instantiated through so-called weights, which

¹⁰See Supplementary Material for a short mathematical summary of this architecture, (Pan *et al.* 2018).

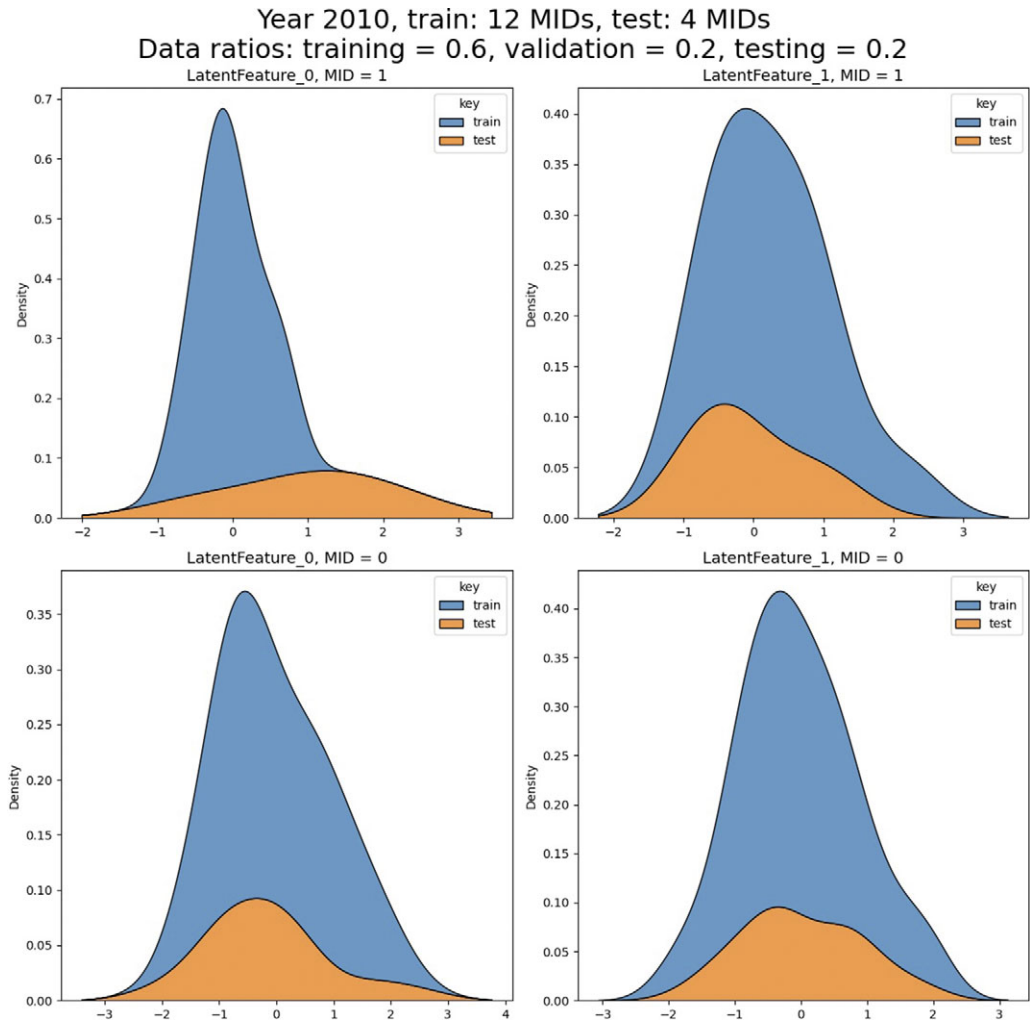


Figure 9. KDE of latent space PCA components in year 2010.

is what is learned in the training (see Supplementary Material for mathematical details) and then saved. Once the latent space is learned, I use PCA to reduce the high-dimensionality space into two principal components. To preserve the structure of the original data, the model learns separately the latent space features for each graph (year). We are working not with whole graphs (as in the classification model of the previous section) but rather with data within each graph as the task is to reconstruct each graph data separately to generate latent variables that are faithful for every panel of the data. I, thus, split the nodes and corresponding edges into training, validation, and testing subsets in each graph.

Figure 9 displays the KDE (Kernel Density Estimation) for year 2010 (as an illustration) for each for the two possible values (0,1) of the MID variable *disp* for training and testing data. The overall shapes for both latent space features are not that different going from training to testing data for both values (0,1) of the MIDs. The variability in the latent factors between the two values (0,1) of the MIDs suggests a possible statistical correlation between MID and the latent factors. This will be probed using a Bayesian model as done in previous sections.

As a second illustration, consider the question of nodes (states) memberships in the topological classes (clusters), and how it changes over the years. We can plot the time variation of the cohesion of all




Figure 10. Cluster cohesion over time.

clusters, where cohesion indicates how close the nodes within the same cluster are to each other. We can then compare its trending behavior, for example, to the variation of MIDs, as shown in Figure 10, which displays spline-smoothed plots of both the cohesion measures of the clusters and the MIDs measure over time. The plot provides a visualization of the dynamic relationship between international system cohesion and military disputes. It suggests that while global disputes increased with rising cohesion during the mid-20th century, the relationship became more complex towards the end of the century, with some clusters maintaining or increasing cohesion even as overall disputes declined. During the 1946–1965 period, three clusters (0,2,3) show increasing cohesion, which coincides with an increase in the number of MIDs. Cluster 1 remains almost constant until about 1980. This suggests that as the international system became more cohesive, possibly due to alliances or blocs at the beginning of the Cold War and as many formerly colonized states achieved independence and leaned toward either side of the Cold War. It is also the era when conflicts intensified such as in the Korean war, Vietnam war, and a number of wars of independence. From 1985 to 2010, after the peak of the Cold War, the number of MIDs declines, and we observe varied cohesion trends across clusters. Three clusters (1,2,3) display a maximum and then decline between 2000 and 2010, whereas Cluster 0 reaches a local minimum and then bounces back upward. The decline in MIDs matches the behavior of clusters (1,2,3). This corresponds to the end of the Cold War, followed with a decrease in large-scale international disputes. The differing cohesion trends seem to indicate different regional dynamics or new forms of international cooperation, or conflict.

Figure 11 displays the cluster membership probabilities for six states around the world. The USA and RUS (Russia) plots show relatively stable membership probabilities across different clusters, indicating their central roles in global systems that persist over time. For instance, after 1960 the USA remains consistently likely to belong to all four clusters, with some variation. Similarly, Russia has high membership probabilities in Clusters 0 and 2, with a noticeable transition period in the early 1990s, likely reflecting the post-Soviet domestic and geopolitical changes. The China plot displays a more dynamic shift between clusters, especially after the early 1970s, which correspond to the period around the People's

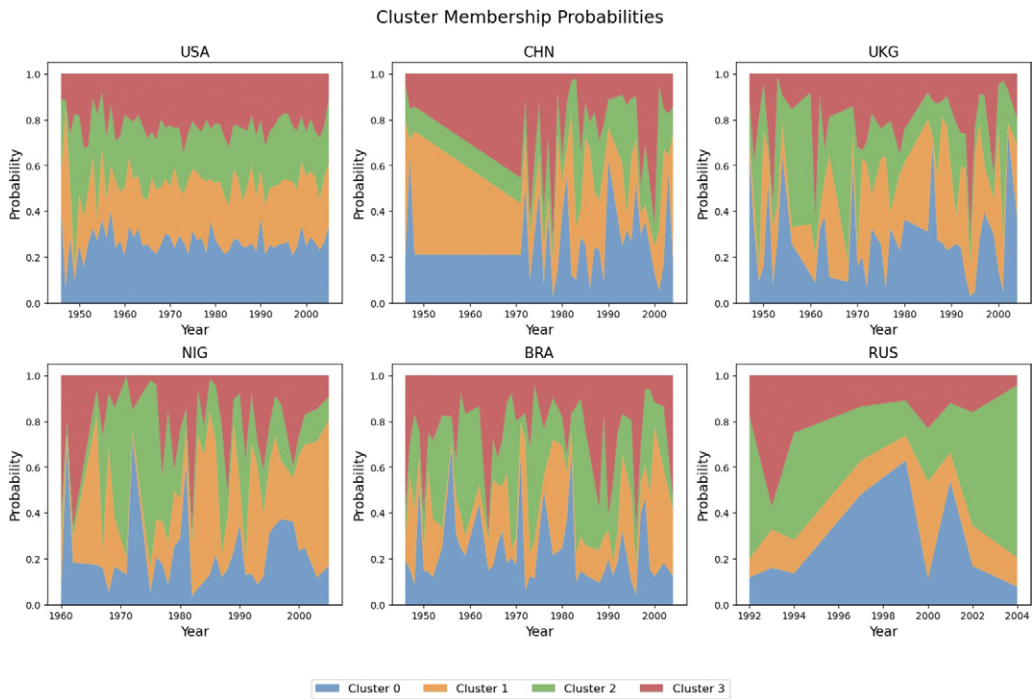


Figure 11. Probabilities of state membership in four clusters.

Republic of China became more active internationally during the “rapprochement” with the US. The UK, NIG (Nigeria), and BRA (Brazil) plots exhibit more fluctuation in cluster membership, indicating that their roles or the nature of their international interactions are more variable or influenced by regional and global changes. Nigeria and Brazil, for example, have periods where they shift significantly rapidly between different clusters, likely reflecting changes in their respective regional influences and external relations.

Building on these results, I now consider a Bayesian model that includes all empirical variables, the expanded SAR variables, the topological effects, and the latent variables. The latter are construed as moderators much like the topological effects (with a similar specification in the model equation). The modified model specification of the augmented SAR model Equation (5.1) is as follows:

$$\begin{aligned}
 \gamma_i &= \sum_{j=1}^{n_{\text{vars}}} X_{ij} \left(\sum_{k=1}^{n_H} H_{ik} \delta_k \right) + \sum_{k=1}^{n_H} H_{ik} \zeta_k, \text{ topological moderation effects} \\
 \omega_i &= \sum_{j=1}^{n_{\text{vars}}} X_{ij} \left(\sum_{k=1}^{n_L} L_{ik} \theta_k \right) + \sum_{k=1}^{n_L} L_{ik} \xi_k, \text{ latent-space moderation effects} \\
 \eta_i &= \gamma_i + \omega_i + (1 + \rho \mathbf{W} + \rho^2 \mathbf{W}^2 + \rho^3 \mathbf{W}^3 + \rho^4 \mathbf{W}^4) \cdot \mathbf{X}_i \cdot \boldsymbol{\beta} \\
 y_i &\sim \text{ZeroInflatedBinomial}(\psi, N_{\text{obs}}, \text{sigmoid}(\eta_i)),
 \end{aligned} \tag{5.1}$$

where $\boldsymbol{\xi}$ are the coefficients for latent-space effects; $\boldsymbol{\theta}$ coefficients for moderation terms; \mathbf{L} is the matrix of moderators/latent factors (dimension: $N \times n_L$), with the rest of the quantities defined as before in Equation (4.1). A Bayesian estimation leads to the results in Table 3.

I compare these results and all other estimations using the ELPD measure (Vehtari *et al.* 2017). The all-inclusive model and the model including the sociomatrix and topological effects are very close to one another with a 0.3% difference in the elpd_loo measure. This is not surprising as we see on Table 3 that the mean for the latent variables and moderators are very small and the HDI intervals both include

Table 3. All encompassing model.

	Variables	mean	sd	hdi_3%	hdi_97%
Empirical variables	majpower	1.077	0.191	0.721	1.439
	allies	-2.808	0.204	-3.179	-2.413
	contgkf	0.866	0.163	0.553	1.161
	smldep	-7.297	0.247	-7.762	-6.837
	lcaprat	-3.072	0.341	-3.707	-2.450
	lgmidist	-7.138	0.304	-7.723	-6.585
	smlmat	-4.001	0.220	-4.410	-3.593
Topological effects	H_1_1	-5.125	0.359	-5.785	-4.447
	H_2_1	-10.269	0.586	-11.319	-9.156
	δ_{H11}	2.389	0.185	2.051	2.740
	δ_{H21}	5.013	0.347	4.350	5.641
Latent variables	L0	-0.006	0.985	-1.848	1.855
	L1	0.011	1.017	-1.790	2.023
	θ_{L0}	0.011	1.004	-1.931	1.851
	θ_{L1}	-0.005	0.989	-1.831	1.853
Sociomatrix	ρ	0.080	0.009	0.063	0.097
Zero-inflation	ψ	0.085	0.008	0.071	0.099

zero value.¹¹ However, they both perform much better than the bare and SAR models. Continuing the task of probing how to make use of the information about topological heterogeneities, I introduce next another GNN model used to optimally reduce the dimensionality of data.

5.4. GNN-Based Dimensionality Reduction of Data

This section proposes an approach that radically reduces the dimensionality of data, while still preserving essential properties of the original data. It is based on a GNN classification model with topological factors taken as graph labels for graphs and MIDs values as node labels.

The entry point to building this approach is an off-shoot to the question of explaining the decisions underpinning the learning process in GNN models, a very active area of research (Agarwal *et al.* 2023; Zhou *et al.* 2020). SubGraphX model addresses this issue by constructing an algorithm that finds subgraphs of a graph that are most influential in a model's predictions using a Monte Carlo Tree Search (MCTS) based strategy; that is, finding subgraphs, when removed, most significantly affect the output of the GNN model. This clearly makes it a good candidate for dimensionality reduction of complex data. This idea is formalized in Equation (5.2).

$$\text{Score}(S) = \frac{1}{N} \sum_{i=1}^N L(y_i, f(G \setminus S; \theta)). \quad (5.2)$$

S represents a candidate subgraph, $G \setminus S$ denotes the graph G excluding the subgraph S , y is the ground truth label (which in the illustration corresponds to MID values, (0,1)), f denotes the GNN with parameters θ , and L is a loss function measuring the discrepancy between the GNN output and the

¹¹This suggests that the Bayesian approach is not best suited for putting to work the unearthed latent space variables. They do indeed provide other kinds of insights on the properties of the data as discussed earlier in the section.

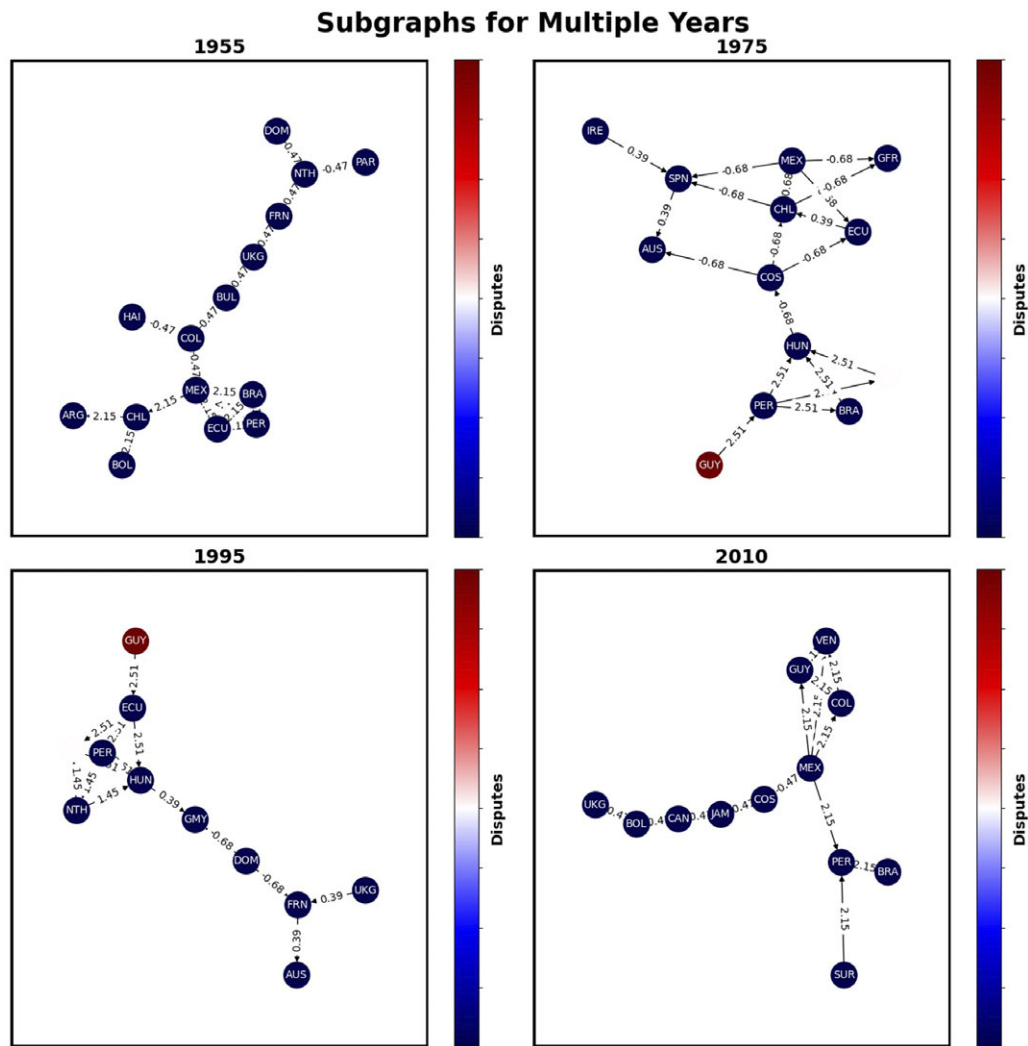


Figure 12. Most important subgraphs for four years.

ground truth. The importance of each subgraph is evaluated using Shapley values as scores.¹² Subgraphs with higher scores are considered more influential on the output of the GNN, and, therefore, good explainers of the behavior of the model performance.


To illustrate, Figure 12 displays four such subgraphs obtained for different years.¹³ The basic layer-architecture of the GNN model that I use to generate the subgraphs is similar to the one used in the previous section. The saved trained model is deployed to find the most important subgraphs for each year. I then convert these sugraphs into a longitudinal panel data the size of which will be much smaller (1118 instead of 240,273 for the original data). The generated data is then used to estimate the basic Bayesian model defined before. The results are displayed in Table 4.

¹²Wikipedia contributors 2024.

¹³The weights on the edges of the subgraphs indicate the number of shared international organizations. The color of the nodes (countries shown with abbreviated names) indicate the number of disputes that the node is involved in divided by the total number of disputes in the set of dyads at a specific year. The colorbar on the right side of each subplot is the corresponding scale for this.

Table 4. Bayesian estimates with SubGraphX-generated data.

	Variables	mean	sd	hdi_3%	hdi_97%
Generated variables	majpower	9.362	5.900	-20.390	0.020
	allies	-8.082	1.931	-11.801	-5.053
	contgkf	-2.620	4.111	-8.849	6.192
	smldep	-5.861	2.617	-10.642	-1.508
	lcaprat	1.431	3.512	-2.812	7.376
	lgmidist	-4.151	2.085	-8.110	-0.976
	smlmat	-15.948	4.614	-24.708	-7.707
Zero-inflation	ψ	0.205	0.060	0.098	0.320

**Figure 13.** ROC curve using empirical data as testing data.

To test the robustness of this model, I run an out-of-sample predictive check using the original empirical data (minus the data for those dyads that are included in the SubgraphX-generated data to avoid overfitting). I use the ROC AUC curve to display the results shown in Figure 13. The AUC value indicates a 75% chance that the model will be able to distinguish between a random positive and a random negative example. The accuracy level is 99.4%. Given that the SubgraphX-generated longitudinal panel data has just 1118 observations as compared to 242,000 observations of testing data, the predictive power is quite impressive.¹⁴

¹⁴These results can be improved by running the classification code for a much longer number of epochs and increasing the Shapley threshold value of the SubGraphX code from 100 to much higher values—see the docstring in the corresponding code file for more details.

6. Conclusion

This article introduces a pioneering approach to analyzing longitudinal panel data by integrating TDA, generative AI, and GNNs, offering a glimpse into the vast potential of these methods in social sciences. The framework presented here, while illustrated with a subset of the MIDs data, is highly transportable and ready for application to the full dataset, including a much broader range of variables.

The versatility of this approach extends beyond the binary nature of the dependent variable used here. The models can be easily adapted for multi-categorical or continuous dependent variables, making the framework applicable to various types of longitudinal data. The transformation of *dyadic* data into graphs, while advantageous, is not essential; non-dyadic data can also be effectively analyzed using GNNs.

While this article focused on numerical, structured data, the GNN models are equally capable of handling heterogeneous data types, including text, images, audios, videos, and more. This adaptability is supported by Python packages like PyG and DGL, which seamlessly integrate different data types within what is known as heterogeneous graphs.

A key advantage of GNNs and deep learning models is their ability to be trained, saved, and deployed for predictions on new data as well as used to generate synthetic data which very tightly resembles the original data; all of which can be done by constructing an automation code of the entire analysis pipeline. This makes it possible to develop user-friendly applications for broader use while maintaining access to the underlying code for customization and extension.

TDA was used to examine the effects of topological defects in the illustrative MIDs data by converting TDA-extracted topological information into numerical data through persistence images, which in turn were transformed via functional principal analysis (FPCA). However, persistence images can also be directly used in generative AI and GNN models *qua* images (this is currently a work of mine in progress) without the need for FPCA. The GNN models proposed here can straightforwardly analyze data composed of numerical structured data, persistence images, and other types of data such as symbolic pictures and mimics.

The methods outlined here are not limited to graph and node classifications; they also extend to edge (link) classification and prediction, which are crucial for developing recommendation systems.

More generally, the new methodology, integrating TDA, generative AI, and GNNs, can be applied to various research-areas of political science to build policy recommendations. Below are some, if very brief, suggestions on how one might leverage this approach.

- *Voter Behavior Analysis*: GNNs can be used to analyze voter behavior by integrating data from various sources, including social media, polling data, and historical voting records. TDA can identify underlying topological structures in voter networks, uncovering deep-seated political affiliations and trends. This approach could inform targeted campaign strategies or identify emerging voter concerns.
- *Partisan Dynamics*: The methodology can examine the evolving structure of partisan alliances in Congress, identifying key nodes (influential legislators) and edges (coalition patterns) that could be targeted to foster bipartisan cooperation. Democrats, Republicans, and Independents can be construed as types of nodes, and caucuses can be construed as different types of edges; all of which with a variety of features.
- *Regime Stability Analysis*: TDA and GNNs can be used to compare political regimes across countries by analyzing longitudinal data on governance, economic indicators, and civil unrest. This can help identify structural (topological) factors that shape regime topological cohesion and stability (or instability), providing insights for understanding the logic of interventions in fragile states.
- *Social Movements*: The topological structures of social movements across different countries can be analyzed using TDA, combined with GNNs to model how these movements evolve and interact with state structures. Social movements and states will be construed as different types of nodes with a variety of idiosyncratic features. Social unrest or promoting democratic engagement can then be analyzed via TDA, and generative AI can be deployed to predict new outcomes.

Work Flowchart

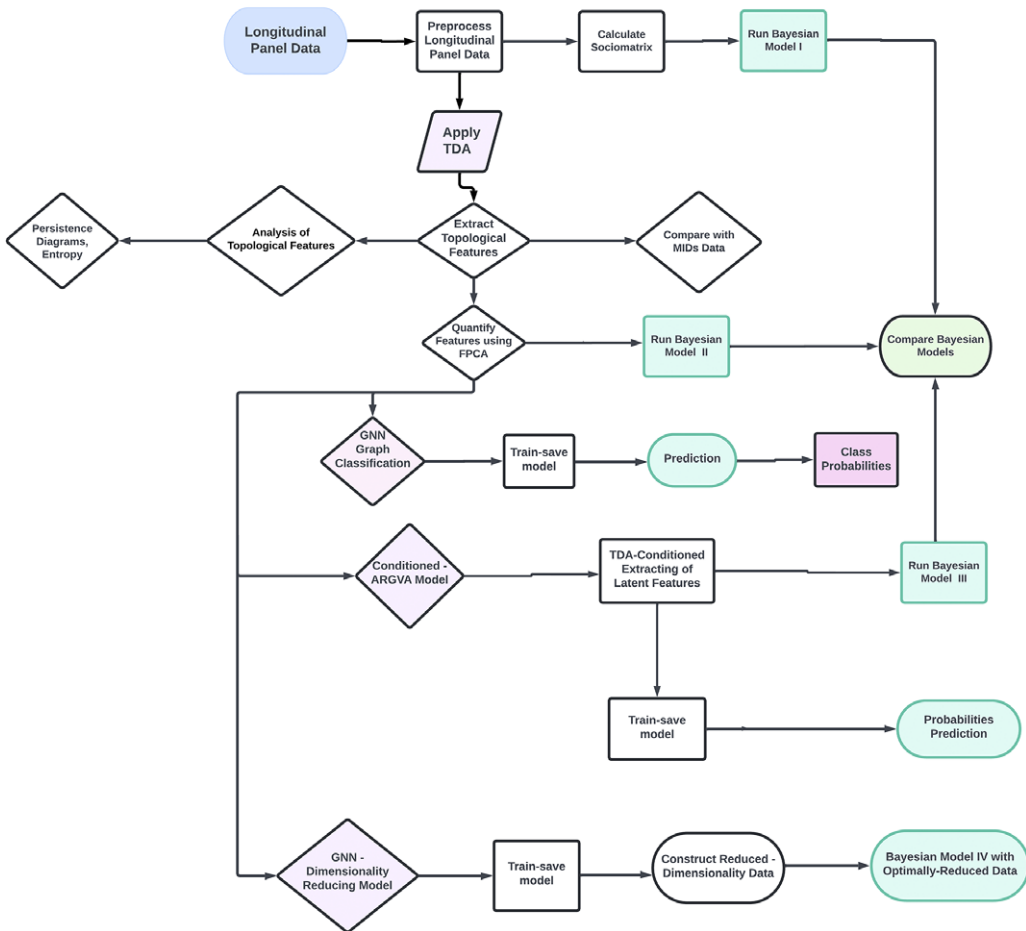


Figure 14. Flowchart steps.

- *Conflict Prediction and Prevention:* GNNs can model the network of international alliances and conflicts by integrating historical data, economic indicators, and diplomatic ties. TDA can identify latent topological features that signal potential conflicts, allowing policymakers to model forward-looking strategies through generative AI models.
- *Diplomatic Strategy:* The methodology can develop sophisticated diplomatic strategies by analyzing the relational dynamics between states in different regional contexts. Different types of nodes such as states, regional organizations, informal partnerships, alliances, and economic actors with a variety of interactions would make it possible, using generative AI, to model the impact of diplomatic actions on bilateral relations, alliance formations, and regional stability, etc.

The ideas proposed in this article are just the tip of a huge iceberg. TDA and generative-AI GNN models open new avenues of research not charted before in social sciences. Going beyond political science, much has been done with TDA and GNNs. However, very few have considered the framework suggested in this article. Having said this, the methodological framework has broad applications across diverse scientific domains. In sociology, TDA could capture persistent patterns in social stratification networks, GNNs could classify community structures, and generative AI could model social mobility pathways. In neuroscience, TDA could identify topological features in brain connectivity patterns,

while GNNs could classify different cognitive states, and generative AI could infer missing neural connections. In financial markets, TDA could detect market regime changes through topological shifts in asset correlations, GNNs could classify trading patterns, and generative AI could help predict systemic risk propagation. In epidemiology, TDA could uncover hidden structures in disease transmission networks, GNNs could classify outbreak patterns, and generative AI could generate possible transmission scenarios. In climate science, TDA could identify persistent topological features in climate networks, GNNs could classify weather patterns, and generative AI could help fill gaps in spatial-temporal data. The framework's unique combination of topology-aware feature extraction, graph-based learning, and generative modeling makes it particularly valuable for complex systems, where both structural patterns and their evolution need to be understood and predicted.

Figure 14 details all steps undertaken in this article to make its case.

Acknowledgements I thank Jeff Gill and four peer reviewers for great comments and suggestions which have made the article a much better product.

Funding Statement. The author did not receive any financial support or funding for this project.

Data Availability Statement. Replication code for this article has been published in the Political Analysis Harvard Dataverse at <https://doi.org/10.7910/DVN/F4UHHW> (Arfi 2025).

Competing Interests. The author declares no competing interests.

Supplementary Material. For supplementary material accompanying this paper, please visit <https://doi.org/10.1017/pan.2025.10019>.

References

- Agarwal, C., O. Queen, H. Lakkaraju, and M. Zitnik. 2023. "Evaluating Explainability for Graph Neural Networks." *Scientific Data* 10 (1): 144.
- Arellano, M., and S. Bonhomme. 2023. "Recovering Latent Variables by Matching." *Journal of the American Statistical Association* 118 (541): 693–706. <https://doi.org/10.1080/01621459.2021.1952877>
- Arfi, B. 2024. "The Promises of Persistent Homology, Machine Learning, and Deep Neural Networks in Topological Data Analysis of Democracy Survival." *Quality & Quantity* 58 (2): 1685–1727. <https://doi.org/10.1007/s11135-023-01708-6>
- Arfi, B. 2025. "Replication Data for: Generative AI and Topological Data Analysis of Longitudinal Panel Data." <https://doi.org/10.7910/DVN/F4UHHW>
- Carlson, J., T. Incerti, and P. M. Aronow. 2024. "Dyadic Clustering in International Relations." *Political Analysis* 32 (2): 186–198. <https://www.cambridge.org/core/product/521F1378ED4F8411B2A4AB0E8FA172A0>
- Chazal, F., and B. Michel. 2021. "An Introduction to Topological Data Analysis: Fundamental and Practical Aspects for Data Scientists." *Frontiers in Artificial Intelligence* 4: 108. <https://www.frontiersin.org/article/10.3389/frai.2021.667963>
- Ferrari, D. 2020. "Modeling Context-Dependent Latent Effect Heterogeneity." *Political Analysis* 28 (1): 20–46. <https://www.cambridge.org/core/product/B7B0AF067DF97A1A8F0B50646EF64F24>
- Fong, C., and J. Grimmer. 2023. "Causal Inference with Latent Treatments." *American Journal of Political Science* 67 (2): 374–389. <https://doi.org/10.1111/ajps.12649>
- Ghrist, R. 2008. "Barcodes—The Persistent Topology of Data." *Bulletin of the American Mathematical Society* 45: 61–75.
- Ghrist, R. 2014. *Elementary Applied Topology*, ed. 1.0. Amazon: Createspace.
- Ghrist, R. 2017. "Homological Algebra and Data." *The Mathematics of Data, IAS/Park City Mathematics* 25: 273–325.
- Hafner-Burton, E. M., and A. H. Montgomery. 2006. "Power Positions: International Organizations, Social Networks, and Conflict." *Journal of Conflict Resolution* 50 (1): 3–27. <https://doi.org/10.1177/0022002705281669>
- Harris, J. M., J. L. Hirst, and M. Mossinghoff. 2008. *Combinatorics and Graph Theory*, 2nd edition. New York: Springer.
- Hatcher, A. 2002. *Algebraic Topology*. New York: Cambridge University Press.
- Kenny, D. A., D. A. Kashy, and W. L. Cook. 2020. *Dyadic Data Analysis*. 2020th edition. New York: The Guilford Press.
- Kinne, B. J. 2013. "IGO Membership, Network Convergence, and Credible Signaling in Militarized Disputes." *Journal of Peace Research* 50 (6): 659–676. <https://doi.org/10.1177/0022343313498615>
- Kipf, T. N., and M. Welling. 2017. "Semi-Supervised Classification with Graph Convolutional Networks." In *International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=SJU4ayYgl>
- Labonne, M. 2023. *Hands-On Graph Neural Networks Using Python: Practical Techniques and Architectures for Building Powerful Graph and Deep Learning Apps with PyTorch*. Birmingham: Packt Publishing Ltd.
- Loaiza-Maya, R., M. S. Smith, D. J. Nott, and P. J. Danaher. 2022. "Fast and Accurate Variational Inference for Models With Many Latent Variables." *Journal of Econometrics* 230 (2): 339–362. <https://www.sciencedirect.com/science/article/pii/S0304407621001330>

- Lupu, Y., and B. Greenhill. 2017. "The Networked Peace: Intergovernmental Organizations and International Conflict." *Journal of Peace Research* 54 (6): 833–848. <https://doi.org/10.1177/0022343317711242>
- Mai, Y., Z. Zhang, and Z. Wen. 2018. "Comparing Exploratory Structural Equation Modeling and Existing Approaches for Multiple Regression with Latent Variables." *Structural Equation Modeling: A Multidisciplinary Journal* 25 (5): 737–749. <https://doi.org/10.1080/10705511.2018.1444993>
- Montfort, K., J. H. L. Oud, and A. Satorra. 2010. *Longitudinal Research with Latent Variables*. Springer Verlag.
- Obayashi, I., Y. Hiraoka, and M. Kimura. 2018. "Persistence Diagrams with Linear Machine Learning Models." *Journal of Applied and Computational Topology* 1 (3): 421–449. <https://doi.org/10.1007/s41468-018-0013-5>
- Olivella, S., T. Pratt, and K. Imai. 2022. "Dynamic Stochastic Blockmodel Regression for Network Data: Application to International Militarized Conflicts." *Journal of the American Statistical Association* 117 (539): 1068–1081. <https://doi.org/10.1080/01621459.2021.2024436>
- Pan, S., R. Hu, G. Long, J. Jiang, L. Yao, and C. Zhang. 2018. "Adversarially Regularized Graph Autoencoder for Graph Embedding." In *IJCAI*, 2609–2615.
- Pevehouse, J., T. Nordstrom, and K. Warnke. 2004. "The Correlates of War 2 International Governmental Organizations Data Version 2.0." *Conflict Management and Peace Science* 21 (2): 101–119. <http://www.jstor.org/stable/26273548>
- Reuning, K., M. R. Kenwick, and C. J. Fariss. 2019. "Exploring the Dynamics of Latent Variable Models." *Political Analysis* 27 (4): 503–517. <https://www.cambridge.org/core/product/CBE116F37900DAE957B2D7EB53DB0907>
- Shannon, M., D. Morey, and F. J. Boehmke. 2010. "The Influence of International Organizations on Militarized Dispute Initiation and Duration." *International Studies Quarterly* 54 (4): 1123–1141. <https://doi.org/10.1111/j.1468-2478.2010.00629.x>
- Sterner, P., F. Pargent, D. Deffner, and D. Goretzko. 2024. "A Causal Framework for the Comparability of Latent Variables." *Structural Equation Modeling: A Multidisciplinary Journal* 31 (5): 747–758. <https://doi.org/10.1080/10705511.2024.2339396>
- Trautmüller, R., A. Murr, and J. Gill. 2015. "Modeling Latent Information in Voting Data with Dirichlet Process Priors." *Political Analysis* 23 (1): 1–20. <https://www.cambridge.org/core/product/338F8FC146746CB76EAA19D1F8517A26>
- Vehtari, A., A. Gelman, and J. Gabry. 2017. "Practical Bayesian Model Evaluation Using Leave-One-Out Cross-Validation and WAIC." *Statistics and Computing* 27 (5): 1413–1432. <https://doi.org/10.1007/s11222-016-9696-4>
- Wang, J., J. M. Chiou, and H. G. Müller. 2016. "Functional Data Analysis." *Annual Review of Statistics and Its Application* 3 (1): 257–295. <https://doi.org/10.1146/annurev-statistics-041715-033624>
- Wasserman, S., and K. Faust. 1994. *Social Network Analysis: Methods and Applications*, Structural Analysis in the Social Sciences. Cambridge: Cambridge University Press. <https://www.cambridge.org/core/books/social-network-analysis/90030086891EB3491D096034684EFFB8>
- Wikipedia contributors. 2024. "Shapley Value—Wikipedia, The Free Encyclopedia." [Online; Accessed 18-March-2024]. https://en.wikipedia.org/w/index.php?title=Shapley_value&oldid=1198571674
- Ye, Y., Z. Liu, D. Pan, and Y. Wu. 2023. "Regression Analysis of Logistic Model with Latent Variables." *Statistics in Medicine* 42 (6): 860–877. <https://doi.org/10.1002/sim.9647>
- Zhang, D., J. Yin, X. Zhu, and C. Zhang. 2017. "Network Representation Learning: A Survey." *IEEE Transactions on Big Data* 6: 3–28. <https://api.semanticscholar.org/CorpusID:1479507>
- Zheng, Z., J. Lv, and W. Lin. 2021. "Nonsparse Learning with Latent Variables." *Operations Research* 69 (1): 346–359. <https://doi.org/10.1287/opre.2020.2005>
- Zhou, J., et al. 2020. "Graph Neural Networks: A Review of Methods and Applications." *AI Open* 1: 57–81. <https://www.sciencedirect.com/science/article/pii/S2666651021000012s>

Author Biography. Dr. Badredine Arfi is currently a professor of political science and international relations at the University of Florida. He holds dual doctorates from the University of Illinois at Urbana-Champaign: a Ph.D. in Theoretical Physics (1988) and a Ph.D. in Political Science/International Relations (1996). His current teaching and research span international politics and security, ethnic conflict, and the methodological and epistemological foundations of artificial intelligence. His work integrates advanced tools such as deep machine learning, generative AI, TDA, fuzzy logic, quantum-theoretic approaches to social theory, and statistical modeling to explore complex social and political phenomena.