

Identifying Resilient Communities in Road Networks: A Path-Based Embedding Approach

Christopher Wagner ✉️🏠^{ID}

Department of Statistics and Applied Probability, University of California, Santa Barbara, CA, USA

Somayeh Dodge ✉️🏠^{ID}

Department of Geography, University of California, Santa Barbara, CA, USA

Danial Alizadeh ✉️🏠^{ID}

Department of Geography, University of California, Santa Barbara, CA, USA

Abstract

Effective resilience analysis of road networks is fundamental to building sustainable and disaster prepared cities. Identifying which road segments share similar vulnerabilities is important for pinpointing high-risk areas within the network and implementing measures to safeguard them against future disruptions. Graph-based community detection can be applied to group together areas of the network sharing similar structural vulnerabilities. However, current graph-based community detection methods either struggle with integrating node features during partitioning or do not account for the path-based dependencies in road networks. This paper introduces the Path-based Community Embedding (PCE) model, an approach that leverages path-based embeddings to overcome these limitations. PCE combines the strengths of graph attention networks and Long Short-Term Memory models (LSTMs) to learn representations that incorporate both local neighborhood information and long-range path dependencies. Our results on the Santa Barbara road network show that PCE improves community detection performance for resilience analysis, thus offering a powerful tool for urban planners and transportation engineers to preemptively identify vulnerabilities in road networks.

2012 ACM Subject Classification Information systems → Geographic information systems

Keywords and phrases road networks, resilience analysis, machine learning, graph neural networks

Digital Object Identifier 10.4230/LIPIcs.GIScience.2025.9

Funding The authors gratefully acknowledge the financial support from the U.S. National Science Foundation through award #BCS-2043202.

1 Introduction

Road networks serve as the backbone of economic and social development by facilitating the flow of people, goods, and services. With the rapid increase in urbanization, continuous improvement of urban road systems is essential for maintaining the efficiency of transportation infrastructure [11]. A fundamental challenge in transportation engineering is evaluating the resilience of road networks, defined as their inherent capacity to recover performance when faced with disruptive events [16]. Resilience in road networks is crucial for ensuring stable mobility, minimizing economic losses, and enhancing emergency response capabilities in the face of disruptions caused by natural disasters, infrastructure failures, or congestion events [1].

The use of graph theory, where intersections can be represented as nodes and road segments can be represented as edges (or vice versa in a dual graph), has allowed researchers to ask and answer questions revolving around road networks and their resilience [2]. By modeling road networks as graphs, structural properties such as connectivity, centrality [7], and traffic flow can be quantitatively analyzed to assess a network's vulnerability to failures [1]. In particular, community detection in road networks has gained attention as an effective way



© Christopher Wagner, Somayeh Dodge, and Danial Alizadeh;
licensed under Creative Commons License CC-BY 4.0

13th International Conference on Geographic Information Science (GIScience 2025).

Editors: Katarzyna Sila-Nowicka, Antoni Moore, David O'Sullivan, Benjamin Adams, and Mark Gahegan;
Article No. 9; pp. 9:1–9:10



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

to identify areas that exhibit similar properties [6]. Community detection in graph theory is the process of identifying elements of the network that are closely connected to each other or share similar characteristics. This process reveals the underlying structure of a network by decomposing it into a set of subgroups, called communities or clusters [8]. For resilience analysis, partitioning roads into different communities can provide insights into which areas of the network share structural vulnerabilities or respond similarly to disruptions. During events such as flooding, roads in low-lying areas and near waterways are more susceptible to closures. If these areas can be identified beforehand as part of the same vulnerability community, transportation agencies can implement strategies to allocate resources efficiently to mitigate the effects of a disruption [17].

Traditional community detection methods leverage the structure of the network and partition the nodes into distinct groups by optimizing certain criteria. For instance, modularity maximization methods assign nodes to communities by maximizing the density of connections within the group compared to a random baseline [18]. Spectral clustering uses the eigenvalues of the graph's Laplacian to identify communities by minimizing the number of cuts between groups [20]. Hierarchical clustering recursively divides nodes according to their connectivity patterns to form a tree structure that reveals different levels of the community [6]. Although effective, these approaches often struggle to incorporate information from nodes such as traffic flow dynamics, road capacity, historical data of disturbance and spatial dependencies [1].

With the rise of big data, modern approaches have leveraged the power of machine learning models on graphs to overcome these limitations by learning embeddings that encode both structural information and node (or edge) features. Models such as graph convolutional networks (GCNs) [13], graph attention networks (GATs) [21], and graph autoencoders (GAEs) [12] have enabled more adaptive community detection by integrating node and edge attributes along with temporal patterns. Additionally, machine learning models tend to be more flexible than traditional models, which rely on fixed assumptions about network topology. For example, spectral clustering implicitly assumes that the cluster structure is encoded in the leading eigenvectors of the graph's Laplacian, which holds when the Laplacian has a few small eigenvalues corresponding to well-separated communities [20]. Unlike traditional approaches, machine learning models can identify patterns in how disruptions affect different segments of a road network without many assumptions. Leveraging these methods can help identify communities that not only reflect connectivity patterns, but also resilience related properties (e.g. vulnerability, centrality, proximity to fire, distance to flood risk areas, etc.) in a data-adaptive manner. While these advancements have improved adaptive community detection, they mainly focus on local connectivity patterns and typically overlook the broader structural dependencies that influence network behavior during disruptions. Incorporating path-based embeddings addresses this limitation by capturing sequences of interconnected nodes rather than neighborhoods to create a more comprehensive representation of road network topology.

In this paper, we highlight the role of machine learning, particularly graph-based neural architectures, in road network community detection to quantify disruption resilience. Additionally, we introduce a model that accounts for path-based dependencies in order to reveal structural patterns linking roads with similar resilience characteristics. We evaluate our approach on a real world road network in Santa Barbara, California to assess its effectiveness in community detection compared to the baseline methods.

The rest of this paper is structured as follows: section 2 describes our methodology, including data preprocessing, the embedding generation, and the proposed clustering approach; section 4 presents experimental results, the analysis, and the discussion; and section 5 discusses conclusions and future directions.

2 Methods

This section introduces our proposed Path-based Community Embedding (PCE) model and formalizes the problem of generating embeddings for community detection.

Given a road network represented as a graph $G = (V, E)$, where $V = \{v_1, v_2, \dots, v_n\}$ is the set of nodes (i.e., intersections) and $E \subseteq V \times V$ is the set of edges (i.e., road segments), the goal is to learn node embeddings that facilitate community detection. A community is defined as a set of nodes that exhibit strong structural and functional similarity, which we infer from the graph-based embeddings and agglomerative clustering. Each node v_i has a feature vector $\mathbf{x}_i \in \mathbb{R}^d$, and the network structure is captured by an adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, where $A_{ij} = 1$, if $(v_i, v_j) \in E$, otherwise $A_{ij} = 0$.

The proposed model consists of three main components: a Graph Attention Network (GAT) encoder [21], a path-based LSTM embedding module [10], and a reconstruction-based decoder. The GAT encoder generates node embeddings by aggregating features from neighbors with learned attention weights:

$$\mathbf{z}_i = \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij} \mathbf{W} \mathbf{x}_j \right), \quad i \in V \quad (1)$$

where $\mathbf{z}_i$ represents the learned embedding for node i , which is computed by aggregating information from its neighbors in $\mathcal{N}(i)$, the term $\mathbf{x}_j$ is the input feature vector of node j , and $\mathbf{W}$ is a trainable weight matrix that transforms the input features into a new feature space. The function $\sigma(\cdot)$, the ReLU activation function, is applied element-wise to introduce non-linearity in the learned embeddings. The attention coefficient α_{ij} represents the importance weight assigned to the feature vector of the neighbor node j when aggregating information for the node i , which is defined as [21]:

$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(\mathbf{a}^T [\mathbf{W} \mathbf{x}_i \parallel \mathbf{W} \mathbf{x}_j]))}{\sum_{k \in \mathcal{N}(i)} \exp(\text{LeakyReLU}(\mathbf{a}^T [\mathbf{W} \mathbf{x}_i \parallel \mathbf{W} \mathbf{x}_k]))} \quad (2)$$

To capture long-range dependencies, paths $P = \{p_1, p_2, \dots, p_m\}$ are sampled from the network, where each p_l is a sequence of nodes. The embeddings along a path are aggregated as:

$$\mathbf{z}_{p_j} = \frac{1}{|p_j|} \sum_{v \in p_j} \mathbf{z}_v, \quad j = 1, \dots, m \quad (3)$$

These path embeddings are processed using a bidirectional LSTM to extract sequential dependencies, producing refined path embeddings $\mathbf{h}_{p_l}$. The final node embeddings are then obtained by aggregating the node embeddings and the path embeddings that the node participates in:

$$\mathbf{z}_i^{\text{final}} = \mathbf{z}_i + \frac{1}{|\mathcal{P}(i)|} \sum_{p_l \in \mathcal{P}(i)} \mathbf{h}_{p_l} \mathbb{I}(i \in p_l) \quad (4)$$

where $\mathcal{P}(i)$ is the set of paths that node i belongs to and $\mathbb{I}(i \in p_l)$ is the indicator function that equals 1 if node i is in path p_l and 0 otherwise.

To ensure the embeddings capture meaningful structure, a decoder reconstructs the original node features using a linear transformation followed by a non-linear activation function. This maps the learned embeddings back to the input feature space. To ensure the community structure is learned, the model is trained in an unsupervised manner using a

joint loss function that combines reconstruction and disruption risk contrastive components. Let $\mathbf{x}_i$ denote the input features for node i and $\hat{\mathbf{x}}_i$ the reconstructed features. Next, we compute a composite risk score r_i for each node using normalized signals: betweenness centrality $\tilde{b}_i$, distance to fire boundaries $\tilde{d}_{\text{fire},i}$, and distance to flood zones $\tilde{d}_{\text{flood},i}$. Specifically, $r_i = w_1 \cdot \tilde{b}_i + w_2 \cdot (1 - \tilde{d}_{\text{fire},i}) + w_3 \cdot (1 - \tilde{d}_{\text{flood},i})$, where w_1 , w_2 , and w_3 are scalar weights. We define a contrastive loss over pairs of nodes (i, j) where $|r_i - r_j| > \delta$ which is given by:

$$\mathcal{L}_{\text{total}} = \alpha \cdot \frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{x}}_i - \mathbf{x}_i\|_2^2 + \beta \cdot \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} \max(0, \cos(\mathbf{z}_i, \mathbf{z}_j) - \tau) \quad (5)$$

where $\mathcal{P}$ is the set of high risk difference node pairs and α , β , and τ are hyperparameters controlling the tradeoff between reconstruction and risk aware separation.

After training, agglomerative clustering is applied to $\mathbf{z}^{i^{\text{final}}}_i \in V$ to detect communities. This is a hierarchical clustering method which iteratively merges the most similar node embeddings based on the Euclidean distance and Ward’s linkage criterion [22], which minimizes the variance within each cluster at every merging step. This clustering method helps preserve spatial continuity while allowing for flexible determination of the number of communities based on the road network topology.

3 Experiments

To evaluate the proposed model, we use a road network representing the downtown Santa Barbara area in California. The city of Santa Barbara is a south-facing coastal town, situated between the Santa Ynez Mountains and the Pacific Ocean. The downtown area, depicted in Figure 1 is located in the lower coastal plain where there is a historic flood risk (shown in blue), while the higher grounds in the foothills are closer to chaparral and forest areas with heightened wildfire risk based on historical fires (shown in orange). The road network is typical of a coastal city in California with restricted geography, featuring Interstate Highway 101 running through the city and serving as a major transportation corridor. The highway closely parallels the coastline, connecting the downtown area to nearby regions while navigating the narrow space between the mountains and the ocean. The road network data, derived from OpenStreetMap (OSM) using OSMnx package [4], consists of 2105 nodes (intersections) and 4234 edges (roads), includes all road types, from major highways to residential streets. For each road segment (edge), we incorporated several features relevant to community resilience and vulnerability, including: proximity to previous fire perimeter boundaries, proximity to historical flood risk zones, betweenness centrality (measuring the importance of a road segment for network flow), closeness centrality (measuring the accessibility of a road segment to all other segments), degree centrality (the number of connections a road segment has). These features were chosen to capture both the topological characteristics of the road network and its exposure to various hazards. The betweenness, closeness, and degree centrality measure were each calculated using the formulas as specified in [19] and are implemented using the NetworkX Python package [9].

The proximity values are obtained by measuring the shortest Euclidean distance from each intersection (node) in the road network to the nearest boundary of the fire perimeter or flood risk zone. Figure 1 presents a map of these boundaries in the city of Santa Barbara. Each intersection is treated as a single point and the distance was computed to the closest edge of the respective hazard polygon. The fire and flood data were obtained from the Santa Barbara County historical fire database and the Federal Emergency Management Agency (FEMA) 100 and 500 year flood risk zones for Santa Barbara, which were both accessed from DataBasin [5].




Figure 1 Downtown Santa Barbara with historic fire and flood risk zones, generated from DataBasin.

To evaluate the performance of our model, we compare it against several community detection models. First, the K-means clustering, which partitions nodes into different clusters by minimizing intra-cluster variance in the feature space [15]. It is noted that the K-means clustering method does not consider the graph structure nor does it produce embeddings. The Louvain algorithm is a hierarchical method that iteratively merges nodes into communities to maximize modularity [3]. This method does not require specifying the number of clusters in advance, thus we tune the resolution parameter which controls the granularity of the detected communities. Region2Vec is a spatially-aware graph embedding method that incorporates both node attributes and spatial interactions to detect communities in spatial networks [14]. It uses Graph Convolutional Networks (GCNs) to generate embeddings that balance structural connectivity and spatial proximity before applying clustering to detect regions with similar properties. Graph Autoencoder (GAE) is an unsupervised model that learns node embeddings by reconstructing the graph's adjacency matrix using a GCN [12]. The Spatial Graph Autoencoder (SGAE) extends GAE by simply weighting neighboring nodes based on their spatial proximity, which helps enforce spatial contiguity in the learned embeddings. For each model, we chose the number of clusters to be four for all models except Louvain, since pre-defining this hyperparameter is not supported. This number of clusters provides a good balance between interpretability and meaningful distinctions in the road network's structure.

We compare the models using three key metrics following the convention used in [14]. The first is cosine similarity, which measures the angular similarity between two node embeddings. The resulting value ranges from -1 to 1 , where 1 indicates identical vectors (perfect similarity), 0 indicates orthogonality (no similarity), and -1 indicates completely opposite vectors. Second, Join Count Ratio (JCR) [14] measures the proportion of edges within the road network that connect nodes belonging to the same community. Given a graph $G = (V, E)$ and a community assignment function $C : V \rightarrow \{1, 2, \dots, K\}$, we define:

$$\text{JCR} = \frac{J_{\text{same}}}{J_{\text{same}} + J_{\text{diff}}}, \quad (6)$$

where J_{same} is the number of edges $(u, v) \in E$ where $C(u) = C(v)$, and J_{diff} is the number of edges where $C(u) \neq C(v)$. A higher JCR indicates that the detected communities are more spatially contiguous. Third, Modularity (Q) [18] evaluates the strength of the community structure by comparing the fraction of edges within detected communities to the expected fraction in a random graph with the same degree distribution. It is defined as:

$$Q = \frac{1}{2m} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2m} \right] \delta(C_i, C_j), \quad (7)$$

where A_{ij} is the adjacency matrix of the graph, k_i and k_j are the degrees of nodes i and j , m is the total number of edges in the graph, and $\delta(C_i, C_j)$ is an indicator function that equals 1 if nodes i and j belong to the same community and 0 otherwise. Higher modularity values correspond to structurally cohesive communities, since they show a greater presence of intra-community edges compared to inter-community edges. These metrics allow us to assess the quality of the embeddings along with the spatial contiguity and structural strength of the predicted communities.

4 Results and Discussion

Figure 2 presents the detected communities using our proposed Path-based (PCE) Community Embedding model, compared to other baseline community detection models: K-Means, Louvain, Region2Vec, Graph Autoencoder (GAE). Table 1 presents a quantitative comparison across the three metrics described in Section 3. The results indicate that the PCE model outperforms most baselines across the three evaluation metrics which can highlight its usefulness in capturing community structures in the road network. It exhibits the highest cosine similarity, which means the embeddings effectively capture the similarity structure of nodes. Additionally, it achieves the highest modularity score and the second highest Join Count Ratio, hence the detected communities form strongly connected groups while maintaining spatial contiguity. The overall results from the community detection across different methods is shown in Figure 2.

The Louvain method, a modularity-based community detection algorithm, achieves the highest JCR. Unlike other models, it relies solely on network topology and excels at preserving connectivity, which may cause it to miss higher-level feature similarities. Thus, it is highly effective at grouping nodes into spatially contiguous regions. Interestingly, although the Louvain method is designed to maximize modularity, the GAE, SGAE, and PCE models achieve higher modularity. Region2Vec and the Graph Autoencoder (GAE) have strong performance across all metrics. Region2Vec, trained to capture spatial proximity and network structure, performs particularly well in cosine similarity and JCR. The K-Means algorithm has the lowest performance across all metrics, particularly modularity and join count ratio.




Figure 2 Results from all community detection methods for the Santa Barbara road network.

Table 1 Comparison of community detection methods. Highest values are bolded.

Method	Cosine Similarity	Join Count Ratio	Modularity
K-Means	—	0.7614	0.3825
Louvain	—	0.9468	0.6060
Region2Vec	0.9314	0.8559	0.5972
GAE	0.8908	0.8669	0.6152
SGAE	0.9820	0.9374	0.6478
PCE	0.9883	0.9432	0.6641

This highlights an expected limitation of feature-space clustering methods applied to road networks: they often disregard the underlying graph structure. Because K-Means ignores connectivity constraints, it may assign distant nodes to the same cluster based on feature similarity, leading to fragmented and spatially disjoint communities. This is reflected in its results, which highlight that clustering methods that overlook graph topology are not well-suited for road network community detection.

PCE’s strong performance can be attributed to its hybrid approach. Its GAT encoder allows nodes to selectively aggregate information from relevant neighbors, while the LSTM-based path encoder captures long-range dependencies in order to ensure that detected communities align with observed connectivity patterns. The integration of neighborhood-based and path-based embeddings ensures that communities are both spatially contiguous and structurally meaningful (respecting graph topology). PCE’s high modularity suggest that it detects highly interconnected subgraphs, which are less vulnerable to isolated disruptions. This is ideal in the context of disruption resilience where communities should represent network regions that can maintain connectivity and functionality during disruptions. In applications like traffic flow management where it’s important to ensure that disruptions in one region don’t severely impact connected areas, this model and other machine learning approaches can potentially help identify regions needing reinforcement to maintain network stability.

Furthermore, the physical and spatial characteristics of the detected communities (Figure 2) offer valuable insights into their resilience profiles. In the visualization of the clusters obtained from PCE in Figure 2f, the densely connected Cluster 3 (yellow) is located in a suburban area with inherent capacity to absorb localized disruptions due to its strong internal connectivity. Conversely, the more dispersed Cluster 0 (blue) shows a community potentially more vulnerable to fire disruptions along critical connecting routes, despite its geographical spread. The core of the network represented by Cluster 1 (red) has high internal connectivity since it contains the roads with the highest centrality, which can provide alternative routing strategies to prevent cascading failures in the event of major disruptions. Finally, Cluster 2 (green) is the one situated closest to the water and the flood risk zones which may mean increased flood risk vulnerability. These clusters demonstrate PCE’s ability to not only identify community structures but also reveal information about their strengths and weaknesses under disruptive events. Our findings suggest that our model can effectively capture the underlying structural characteristics that contribute to resilience in road networks which is valuable for targeted interventions and resilience planning.

5 Conclusion

In this paper, we tackled the challenging problem of community detection in road networks, focusing specifically on identifying communities with shared resilience characteristics. We introduced a novel graph-based embedding model that effectively captures both local and global structural information within the network. Our model PCE leverages Graph Attention Networks to learn local patterns and path-based LSTM to learn long-range global dependencies. This allows the model to understand how nodes are connected across the network, which captures broader structural relationships that are crucial for identifying resilient communities. The combined embeddings that incorporate both local and global perspectives are then used to reconstruct the original node features in a self-supervised manner. Finally, we employ agglomerative clustering on these learned embeddings to reveal the community structure.

Our key contribution lies in the unique combination of local and global graph information, enabling the identification of communities that are not only spatially contiguous but also share resilience properties due to their structural organization. We demonstrated the effectiveness of our approach on the Santa Barbara road network, where we were able to identify distinct communities. These findings suggest that our model can effectively capture the underlying structural characteristics that contribute to resilience in road networks. Furthermore, the unsupervised nature of our approach makes it applicable to a wide range of road network analysis tasks, even when labeled data is scarce or unavailable.

There are several limitations to this study, the main one being the computational complexity of the model and the calculated features, namely betweenness and closeness centrality. As the size of the network increases (e.g. over 100,000 nodes), the computation time will skyrocket and thus may be infeasible. Future work could explore ways to decrease the time complexity of computing these metrics and the model itself to enhance the scalability and assess the generalizability of the model. The model can also be further extended to incorporate dynamic network information (e.g, real-time movement flows) to further enhance its ability to identify resilient communities.

References

- 1 Danial Alizadeh and Somayeh Dodge. Disaster vulnerability in road networks: a data-driven approach through analyzing network topology and movement activity. *International Journal of Geographical Information Science*, pages 1–22, 2024. doi:10.1080/13658816.2024.2411001.
- 2 Chandra Balijepalli and Raphael Oppong. Measuring vulnerability of road network considering the extent of serviceability of critical road links in urban areas. *Journal of Transport Geography*, 46:65–75, 2015.
- 3 Vincent D. Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008, 2008. doi:10.1088/1742-5468/2008/10/P10008.
- 4 Geoff Boeing. Modeling and analyzing urban networks and amenities with osmnx. *Geographical Analysis*, 2025.
- 5 Santa Barbara County Fire Department and CAL FIRE. Fire history, santa barbara county, 1990-2020, 2021. Accessed: 2025-02-19. URL: <https://databasin.org/datasets/322c741a26d24400a1ab948d40887fad/>.
- 6 Santo Fortunato and Darko Hric. Community detection in networks: A user guide. *Physics Reports*, 659:1–44, 2016.
- 7 Linton C Freeman et al. Centrality in social networks: Conceptual clarification. *Social network: critical concepts in sociology*. Londres: Routledge, 1(3):238–263, 2002.
- 8 Michelle Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002. doi:10.1073/pnas.122653799.
- 9 Aric Hagberg, Daniel Schult, and Pieter Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference (SciPy)*, pages 11–15, 2008.
- 10 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. doi:10.1162/neco.1997.9.8.1735.
- 11 Yang Hong and Yao Yao. Hierarchical community detection and functional area identification with osm roads and complex graph theory. *International Journal of Geographical Information Science*, 33(8):1569–1587, 2019. doi:10.1080/13658816.2019.1584806.
- 12 Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016. arXiv:1611.07308.
- 13 Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations (ICLR)*, 2017. arXiv:1609.02907.
- 14 Yunlei Liang, Jiawei Zhu, Wen Ye, and Song Gao. Geoai-enhanced community detection on spatial networks with graph deep learning. *Computers, Environment and Urban Systems*, 117:102228, 2025. doi:10.1016/j.compenvurbsys.2024.102228.
- 15 J. MacQueen. Some methods for classification and analysis of multivariate observations. *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.

- 16 B. Martín, E. Ortega, R. Cuevas-Wizner, A. Ledda, and A. De Montis. Assessing road network resilience: an accessibility comparative analysis. *Transportation Research Part D: Transport and Environment*, 95:102851, 2021. doi:10.1016/j.trd.2021.102851.
- 17 Lars-Göran Mattsson and Erik Jenelius. Vulnerability and resilience of transport systems—a discussion of recent research. *Transportation Research Part A: Policy and Practice*, 81:16–34, 2015.
- 18 M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
- 19 M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- 20 Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 849–856, 2001. URL: <https://proceedings.neurips.cc/paper/2001/hash/801272ee79cfde7fa5960571fee36b9b-Abstract.html>.
- 21 Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *International Conference on Learning Representations (ICLR)*, 2018.
- 22 Joe H. Ward. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. doi:10.1080/01621459.1963.10500845.