# Spatiotemporal Adaptive Learning for Flow-Based Graph Neural Networks

**Fengjiao Li**[1]
School of Geographical and Earth Sciences, University of Glasgow, UK

**Meiliu Wu**
School of Geographical and Earth Sciences, University of Glasgow, UK

**Ana Basiri**
School of Geographical and Earth Sciences, University of Glasgow, UK

## Abstract

Spatial systems can often be represented in more than one relational way. For the same urban system, geographic adjacency, road connectivity, and mobility interaction provide different relational representations of how places are connected. However, most graph-based models either assume a single graph or combine multiple graphs without distinguishing their relative importance. This paper treats multi-graph spatiotemporal modelling as a problem of spatial perspective selection. We introduce adaptive spatial weights ($\alpha$) to learn the relative importance of predefined relational graphs, and temporal weights ($\beta$) in the loss function to distribute emphasis across forecasting horizons. Using COVID-19 dynamics in Glasgow and Edinburgh as a case study, we find that learned spatial weighting patterns differ across cities. Glasgow shows a more balanced configuration, whereas Edinburgh places stronger emphasis on road connectivity. Compared with single-graph and equal-weight baselines, the adaptive model achieves comparable performance with improved stability while making graph weighting explicit. These results suggest that multi-graph weighting is a context-dependent process of selecting among alternative relational representations.

**2012 ACM Subject Classification** Theory of computation → Network flows; Computing methodologies → Temporal reasoning; Computing methodologies → Spatial and physical reasoning; Computing methodologies → Artificial intelligence; Information systems → Geographic information systems

**Keywords and phrases** spatiotemporal adaptive learning, graph neural networks, multi-graph representation, GeoAI, flow-based networks

**Digital Object Identifier** 10.4230/LIPIcs.COSIT.2026.31

**Category** Short Paper

**Funding** This research was supported by the China Scholarship Council (CSC).

**Acknowledgements** The authors also acknowledge the use of publicly available data provided by Public Health Scotland (COVID-19 in Scotland dataset) and the Urban Big Data Centre (UBDC) at the University of Glasgow (mobility dataset).

**Declaration on GenAI/LLM use** ChatGPT was used to assist in the creation of Figure 1. The figure was generated for illustrative purposes, while the underlying scientific ideas, model design, and interpretation were developed by the authors. The generated visual content was reviewed, edited, and approved by the authors to ensure that it accurately reflects the content of the paper.

## 1 Introduction

Spatial systems are inherently relational and can be represented in multiple ways. The meaning of a location is shaped not only by its attributes, but also by the relations through which it is defined. In spatial information theory, space is increasingly viewed as a relational

[1] Corresponding author.


17th International Conference on Spatial Information Theory (COSIT 2026).
Editors: Sabine Timpf, Gabriele Filomena, Armand Kapaj, Rui Zhu, Nicholas Giudice, and Ed Manley; Article No. 31; pp. 31:1–31:9
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

construct rather than a fixed geometric container [2]. Similarly, research in spatial cognition shows that spatial representation depends on perspective, scale, and task [10]. As a result, different relational descriptions, such as geographic adjacency, infrastructural connectivity, and behavioural interaction, can provide valid but distinct views of the same system. This idea is also reflected in multilayer and multiplex network research, which formalises how multiple relational layers may jointly shape system dynamics [8, 3, 4]. In GIScience, related work has likewise emphasised that spatial representation is context-dependent rather than uniquely determined, and that different abstractions of the same system may coexist for different purposes. Together, these perspectives suggest that spatial systems admit multiple meaningful relational encodings rather than a single privileged structure.

Graph Neural Networks (GNNs) have become central tools for modelling relational structure in spatial and spatiotemporal systems, including applications such as traffic forecasting and epidemic modelling [9, 18, 16]. Spatiotemporal models such as STGCN and DCRNN have demonstrated strong forecasting performance [18, 9], while subsequent research has introduced attention mechanisms, graph learning, and multi-graph extensions to improve predictive flexibility [6, 17, 5, 15]. Most existing approaches either assume a single graph or combine multiple graphs using implicit mechanisms such as attention or graph learning [5, 15, 6]. However, these approaches primarily treat graph construction and graph combination as modelling or optimisation choices. Even when multiple relational graphs are used, fusion is usually implemented as weighted aggregation or attention-based combination, rather than being interpreted as explicit selection among coexisting spatial perspectives.

This issue is particularly important in real spatial systems, where different graphs encode different spatial logics. Geographic adjacency reflects spatial proximity, road networks capture infrastructural constraints, and mobility flows represent behavioural interaction across areas [1, 12]. These representations are not interchangeable, but instead provide complementary perspectives on the same system. From this perspective, multi-graph modelling can be understood not only as a problem of technical fusion, but also as a problem of representational selection. While previous studies have largely focused on improving predictive accuracy, comparatively less attention has been paid to understanding how alternative relational representations should be prioritised and what such prioritisation may reveal about the spatial processes being modelled.

Our work builds on these two lines of research. First, we recognise that multiple relational encodings of the same spatial system can be theoretically meaningful, with each graph representing a distinct spatial perspective and a corresponding conceptual flow mechanism. Second, we adopt a spatiotemporal graph learning framework in which these perspectives can be combined adaptively. Unlike most previous studies, which primarily view multi-graph learning as a problem of graph fusion for improving predictive accuracy, we interpret graph weighting as a process of spatial perspective selection. In this view, learned graph weights provide not only a mechanism for prediction, but also an interpretable indication of which relational representations may be more informative in a given context. Predictive performance therefore serves as supporting evidence for a broader representational question concerning the relative importance of alternative spatial perspectives.

In this paper, we frame multi-graph spatiotemporal modelling as a problem of *spatial perspective selection*. We make explicit how different relational perspectives are weighted within a unified model. To this end, we introduce an adaptive spatial weighting mechanism $\alpha$ to learn the relative contribution of different graphs. We also incorporate temporal weighting $\beta$ to account for differences across prediction horizons. Together, these provide a spatiotemporal weighting formulation, although our primary focus remains on the role of

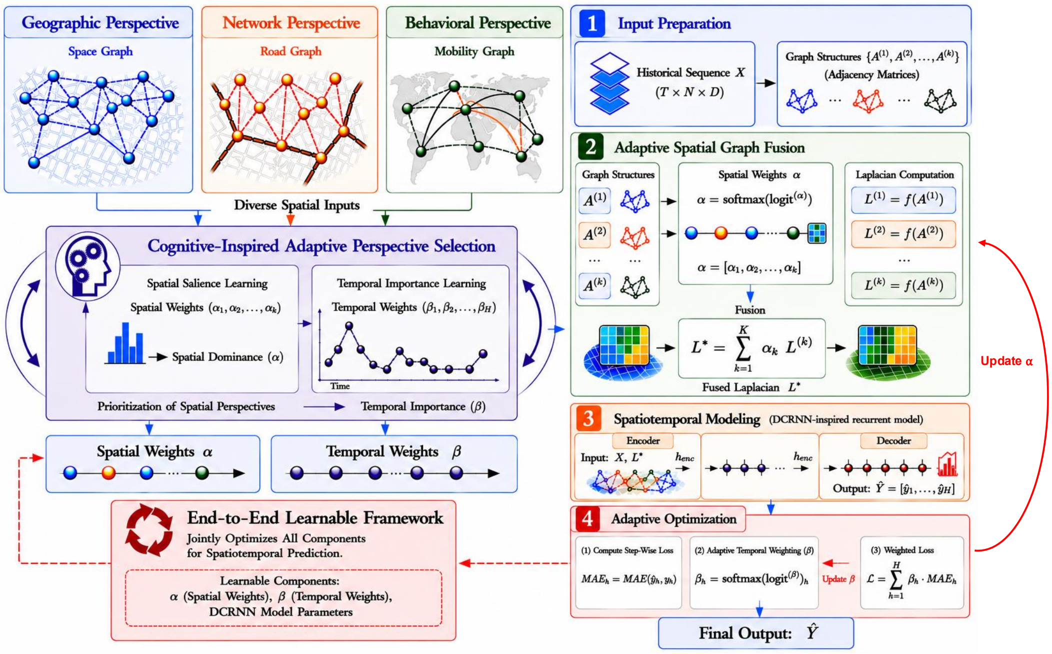


**Figure 1** Overview of the proposed spatiotemporal adaptive learning framework.

spatial weighting in representing different perspectives. To illustrate this idea, we implement the framework within a spatiotemporal graph neural model and evaluate it using epidemic dynamics in Glasgow and Edinburgh. Rather than focusing solely on predictive performance, we examine how learned spatial weights vary across contexts, providing insight into how different relational representations contribute in different urban settings.

This paper makes three contributions. First, it frames multi-graph modelling as a problem of spatial perspective selection. Second, it provides a practical formulation using $\alpha$ and $\beta$ to make relational weighting explicit. Third, it reveals empirical findings that spatial weighting patterns differ across cities, suggesting that relational importance is context-dependent and potentially linked to different underlying flow mechanisms represented by each graph.

## 2 A Novel Approach: Spatiotemporal Adaptive Learning

Existing spatiotemporal graph models often combine spatial fusion and temporal aggregation within a single adaptive process, making it difficult to distinguish weighting across relational graphs from weighting across prediction horizons. Here, we separate these two forms of weighting within a unified framework (Figure 1). Multiple spatial graphs are treated as distinct relational perspectives, including geographic adjacency, infrastructure connectivity, and mobility interaction. Spatial weighting is defined through $\alpha$, which determines the contribution of each graph and satisfies $\sum_{k=1}^{K} \alpha_k = 1$, producing a fused spatial operator $L^*$. Within the proposed framework, these weights can be interpreted as model-based proxies for the relative importance of different spatial perspectives.

Temporal weighting is introduced through $\beta$, which distributes emphasis across prediction horizons. While $\alpha$ governs spatial combination, $\beta$ controls the weighting of prediction errors over time. Taken together, they define a spatiotemporal weighting formulation that separates spatial perspective prioritisation from temporal focus allocation.

**Table 1** Conceptual correspondence between relational graphs and spatial perspectives.

| Graph | Spatial Perspective | Flow Mechanism |
|---|---|---|
| Geographic adjacency | Metric continuity | Local spillover |
| Road connectivity | Infrastructural constraint | Transport-mediated diffusion |
| Mobility interaction | Behavioural interaction | Movement-driven transmission |

Conceptually, the proposed framework distinguishes two related but different questions. The first concerns which spatial perspectives should receive greater emphasis when multiple relational representations coexist. The second concerns which forecasting horizons should receive greater emphasis during model optimisation. By separating adaptive spatial weighting ($\alpha$) from adaptive temporal weighting ($\beta$), the framework allows learned graph weights to be interpreted as model-based indicators of relative spatial importance rather than as components of a generic attention mechanism. This distinction helps clarify how spatial and temporal information contribute to the overall learning process.

# 3 Data and Methods

## 3.1 Data

The empirical analysis focuses on Glasgow and Edinburgh. Daily COVID-19 confirmed case counts at the InterZone level (March 2020–February 2023) were obtained from Public Health Scotland [13]. Origin–destination mobility matrices were obtained from the Urban Big Data Centre (UBDC) [14], and road and rail networks were downloaded from OpenStreetMap via Geofabrik [11]. These datasets support the construction of geographic, infrastructural, and mobility-based graphs.

Let $X_t \in \mathbb{R}^N$ denote infection rates at time $t$, and $Y \in \mathbb{R}^{H\times N}$ denote future values over $H$ horizons. A history window $T = 7$ is used to predict $H = 7$ days. Inputs are standardised using training statistics and evaluated on the original scale.

## 3.2 Methods

We construct three relational graphs representing geographic adjacency, road connectivity, and mobility interaction (Table 1), all defined on the same node set.

**Geographic adjacency.**

$$A_{ij}^{\text{space}} = \begin{cases} 1, & \text{if zones } i \text{ and } j \text{ share a boundary,} \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

**Road connectivity.**

$$\begin{aligned} D_{ij}^{\text{road}} &= \min\left\{D_{ij}^{\text{road}}, D_{ij}^{\text{bike}}, D_{ij}^{\text{walk}}, D_{ij}^{\text{road-rail}}, D_{ij}^{\text{walk-rail}}\right\}, \\ A_{ij}^{\text{road}} &= \begin{cases} \dfrac{1}{d_{ij}^{\text{road}} + \epsilon} & d_{ij}^{\text{road}} < \infty, \\ 0 & d_{ij}^{\text{road}} = \infty, \end{cases} \quad i \neq j. \end{aligned} \tag{2}$$

**Mobility interaction.**

$$A_{ij}^{\mathrm{mob}} = \frac{f_{ij}}{\sum_j f_{ij}}, \tag{3}$$

**Graph normalisation and diffusion operators.**

$$A^{(k)} \leftarrow \frac{A^{(k)} + (A^{(k)})^\top}{2},$$
$$L_1^{(k)} = (D^{(k)})^{-1}A^{(k)}, \quad L_2^{(k)} = (D^{(k)})^{-1}(A^{(k)})^\top, \quad L^{(k)} = \frac{L_1^{(k)} + L_2^{(k)}}{2}.$$

**Adaptive diffusion fusion.**

$$\alpha_k = \mathrm{Softmax}\Big(\mathrm{logit}^{(\alpha)}\Big)_k, \quad \sum_{k=1}^{K}\alpha_k = 1, \quad L^* = \sum_{k=1}^{K}\alpha_k L^{(k)}.$$

where $\mathrm{logit}^{(\alpha)}$ is a learnable graph scoring vector, $\alpha_k$ denotes the weight assigned to graph view $k$, $K$ is the number of graph views, $L^{(k)}$ is the graph-specific diffusion operator, and $L^*$ is the fused diffusion operator.

**Spatiotemporal modelling.**

$$\mathbf{g}_\theta * \mathbf{x} = \sum_{k=0}^{K}\theta_k T_k(L^*)\mathbf{x}, \quad \beta_h = \mathrm{Softmax}\Big(\mathrm{logit}^{(\beta)}\Big)_h, \quad \sum_{h=1}^{H}\beta_h = 1,$$
$$\mathcal{L} = \sum_{h=1}^{H}\beta_h \frac{1}{N}\sum_{i=1}^{N}\left|y_i^{(h)} - \hat{y}_i^{(h)}\right|.$$

where $T_k(\cdot)$ denotes the $k$-th Chebyshev polynomial, $\theta_k$ is a learnable filter coefficient, $\beta_h$ is the learned weight assigned to forecast horizon $h$, and $\mathcal{L}$ is the horizon-weighted MAE loss.

**Training and evaluation.** Data are split chronologically (70/15/15). Models are trained with Adam [7], early stopping, and gradient clipping, and repeated five times. Performance is evaluated using

$$\mathrm{MAE} = \frac{1}{N}\sum|y-\hat{y}|, \quad \mathrm{RMSE} = \sqrt{\frac{1}{N}\sum(y-\hat{y})^2}, \quad \mathrm{MAPE} = \frac{100}{N}\sum\left|\frac{y-\hat{y}}{y}\right|.$$

## 4 Results and Analysis

Table 2 summarises forecasting performance in Glasgow and Edinburgh. Across both cities, the road graph provides the strongest single-graph baseline, achieving lower errors than either geographic adjacency or mobility interaction. However, the adaptive fusion model further improves predictive performance, achieving the lowest RMSE and MAPE while maintaining comparable MAE. This indicates that combining multiple spatial perspectives can provide additional predictive benefit when the relative contribution of each relational graph is learned adaptively rather than fixed a priori. In contrast, equal-weight fusion produces substantially higher errors, particularly in Glasgow, suggesting that improvement comes not from graph combination alone, but from how the relational graphs are comparatively weighted.

**Table 2** Prediction performance in Glasgow (GLA) and Edinburgh (EDI) (mean $\pm$ std over five runs). Lower values indicate better performance.

| City | Graph | MAE | RMSE | MAPE (%) |
|---|---|---|---|---|
| GLA | Road | 21.8 $\pm$ 0.9 | 31.5 $\pm$ 0.2 | 68.7 $\pm$ 6.2 |
| | Space | 23.2 $\pm$ 2.0 | 32.3 $\pm$ 1.1 | 78.4 $\pm$ 12.1 |
| | Mobility | 22.4 $\pm$ 1.5 | 32.0 $\pm$ 0.6 | 71.6 $\pm$ 9.8 |
| | Fusion (Equal weight) | 27.4 $\pm$ 9.6 | 36.9 $\pm$ 8.9 | 92.1 $\pm$ 49.8 |
| | **Fusion Adaptive** | **21.8 $\pm$ 1.2** | **31.3 $\pm$ 0.4$\downarrow$** | **67.8 $\pm$ 8.3$\downarrow$** |
| EDI | Road | 21.9 $\pm$ 1.7 | 31.8 $\pm$ 0.7 | 73.8 $\pm$ 12.7 |
| | Space | 23.6 $\pm$ 1.9 | 32.9 $\pm$ 0.8 | 81.2 $\pm$ 13.2 |
| | Mobility | 22.5 $\pm$ 2.6 | 32.1 $\pm$ 1.7 | 78.4 $\pm$ 15.1 |
| | Fusion (Equal weight) | 21.3 $\pm$ 0.8 | 33.2 $\pm$ 1.7 | 69.0 $\pm$ 4.2 |
| | **Fusion Adaptive** | **20.9 $\pm$ 1.5$\downarrow$** | **31.6 $\pm$ 0.8$\downarrow$** | **66.0 $\pm$ 8.4$\downarrow$** |

**Table 3** Learned adaptive spatial weights ($\alpha$ values) in Glasgow (GLA) and Edinburgh (EDI) (mean $\pm$ standard deviation over five runs).

| Component | GLA | EDI |
|---|---|---|
| Road | 0.35 $\pm$ 0.02 | 0.46 $\pm$ 0.05 |
| Space | 0.36 $\pm$ 0.02 | 0.23 $\pm$ 0.02 |
| Mobility | 0.29 $\pm$ 0.00 | 0.31 $\pm$ 0.03 |

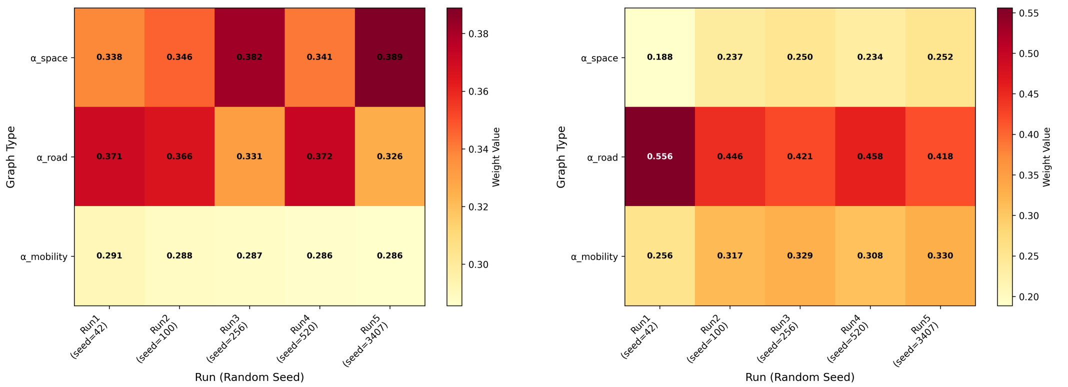


**(a)** GLA. **(b)** EDI.

**Figure 2** Learned spatial weights ($\alpha$) under the adaptive graph construction in Glasgow and Edinburgh.

The learned spatial weights reveal clear differences between the two cities (Table 3 and Figure 2). In Glasgow, the weighting pattern is relatively balanced: road connectivity (0.35) and geographic adjacency (0.36) contribute at similar levels, while mobility interaction (0.29) plays a smaller role. In Edinburgh, by contrast, road connectivity is more strongly prioritised (0.46), followed by mobility interaction (0.31) and geographic adjacency (0.23). This indicates greater reliance on infrastructural connectivity in Edinburgh, compared with the more balanced relational structure observed in Glasgow.

Although road connectivity provides the strongest individual graph in both cities, the adaptive model assigns non-negligible weights to all three relational graphs. This suggests that no single relational representation is sufficient on its own, and that multiple spatial perspectives contribute to forecasting performance in different ways.

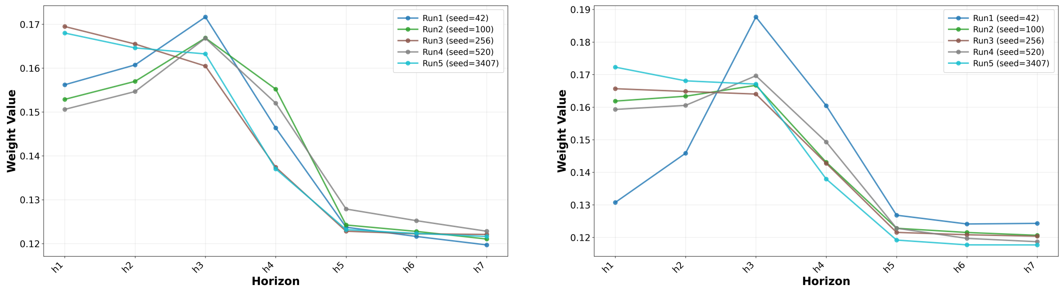


**(a)** GLA. **(b)** EDI.

**Figure 3** Learned temporal weights ($\beta$) across prediction horizons for (a) Glasgow and (b) Edinburgh.

Temporal weights (Figure 3) follow a similar pattern in both cities, placing greater emphasis on short- and mid-term horizons and less emphasis on longer horizons. Compared with the spatial weights, which vary across cities, the temporal weighting patterns are more consistent. Taken together, these results suggest that adaptive weighting captures context-dependent spatial structure, while temporal weighting remains comparatively stable across the two urban settings.

## 5 Discussion and Conclusion

This study suggests that multi-graph spatiotemporal modelling can be understood not only as a technical problem of graph fusion, but also as a problem of spatial perspective selection. By representing multiple relational encodings within a unified framework and learning their relative contributions through adaptive spatial weighting, the proposed approach makes explicit how different relational representations are prioritised within the forecasting process.

The empirical results indicate context-dependent weighting patterns across the two cities. More broadly, the findings suggest that adaptive spatial weighting may provide a useful mechanism for exploring how alternative spatial perspectives contribute to spatiotemporal modelling in different contexts. Rather than identifying a universally optimal spatial representation, the results suggest that the relevance of different relational structures may vary across urban settings.

Understanding the relative importance of spatial perspectives is valuable because different relational graphs imply different explanations of how spatial processes operate. By making these differences explicit, the proposed framework provides interpretative value beyond predictive accuracy and links empirical modelling with broader GIScience questions concerning spatial representation and relational thinking. Its primary contribution therefore lies not in large predictive gains, but in making representational prioritisation observable within a spatiotemporal learning framework and reframing multi-graph learning as a process of selecting among alternative spatial perspectives rather than simply combining multiple graphs.

Although demonstrated using COVID-19 forecasting, the framework may be applicable to a broader range of spatiotemporal phenomena in which multiple relational structures coexist, including traffic forecasting, urban mobility analysis, and environmental monitoring. Nevertheless, the current study is limited to two urban contexts and does not include a full ablation analysis or comparisons with stronger multi-graph baselines. Future work will extend the framework to additional cities, relational layers, and application domains, while further examining how learned weighting patterns relate to different spatial processes. Overall, this work highlights that spatial representation is not fixed, but can be learned and adapted in a context-sensitive manner within spatiotemporal models.

## References


1 Hugo Barbosa, Marc Barthelemy, Gourab Ghoshal, Chris R. James, Maxime Lenormand, Thomas Louail, Ronaldo Menezes, José J. Ramasco, Filippo Simini, Matteo Tomasini, and Alessandro Vespignani. Human mobility: Models and applications. *Physics Reports*, 734:1–74, 2021. `doi:10.1016/j.physrep.2018.01.001`.

2 Michael Batty. *The New Science of Cities*. MIT Press, Cambridge, MA, 2013. `doi:10.7551/mitpress/9399.001.0001`.

3 Stefano Boccaletti, Ginestra Bianconi, Regino Criado, et al. The structure and dynamics of multilayer networks. *Physics Reports*, 544(1):1–122, 2014. `doi:10.1016/j.physrep.2014.07.001`.

4 Manlio De Domenico, Albert Solé-Ribalta, Emanuele Cozzo, Mikko Kivelä, Yamir Moreno, Mason A. Porter, Sergio Gómez, and Alex Arenas. Mathematical formulation of multilayer networks. *Physical Review X*, 3(4):041022, 2013. `doi:10.1103/PhysRevX.3.041022`.

5 Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiaowen He. Learning discrete structures for graph neural networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 1972–1982, 2019.

6 Shengnan Guo, Youfang Lin, Nan Feng, Chuan Song, and Huaiyu Wan. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):922–929, 2019. `doi:10.1609/aaai.v33i01.3301922`.

7 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. `doi:10.48550/arXiv.1412.6980`.

8 Mikko Kivelä, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. Multilayer networks. *Journal of the Royal Society Interface*, 11(98):20140343, 2014. `doi:10.1093/comnet/cnu016`.

9 Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2018. `doi:10.48550/arXiv.1707.01926`.

10 Daniel R. Montello. Scale and multiple psychologies of space. In *Spatial Information Theory: A Theoretical Basis for GIS*, volume 716 of *Lecture Notes in Computer Science*, pages 312–321. Springer, 1993. `doi:10.1007/3-540-57207-4_21`.

11 OpenStreetMap contributors. Geofabrik download server – scotland extract. `https://download.geofabrik.de/europe/united-kingdom/scotland.html`, 2023. Accessed 1 March 2023.

12 Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of Modern Physics*, 87(3):925–979, 2015. `doi:10.1103/RevModPhys.87.925`.

13 Public Health Scotland. Covid-19 in scotland dataset. `https://www.opendata.nhs.scot/dataset/covid-19-in-scotland`, 2023. Accessed 1 March 2023.

14 Urban Big Data Centre. Ubdc mobility dataset. `https://data.ubdc.ac.uk/datasets/6d87635c-96bb-4b7a-9a2e-676bb250032d`, 2023. Accessed 1 March 2023.

15 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2018. `doi:10.48550/arXiv.1710.10903`.

16 Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2021. `doi:10.1109/TNNLS.2020.2978386`.

**17** Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 753–763, 2020. `doi:10.1145/3394486.3403118`.

**18** Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3634–3640, 2018. `doi:10.24963/ijcai.2018/505`.