

Mitigating Spatial Disparity in Urban Prediction Using Residual-Aware Spatiotemporal Graph Neural Networks: A Chicago Case Study

Dingyi Zhuang
Massachusetts Institute of Technology
Cambridge, MA, USA
dingyi@mit.edu

Hanyong Xu
Massachusetts Institute of Technology
Cambridge, MA, USA
hanyongx@mit.edu

Xiaotong Guo
Massachusetts Institute of Technology
Cambridge, MA, USA
xtguo@mit.edu

Yunhan Zheng
Singapore-MIT Alliance for Research
and Technology Centre
Singapore
yunhan@mit.edu

Shen hao Wang
University of Florida
Gainesville, FL, USA
shen hao wang@ufl.edu

Jinhua Zhao
Massachusetts Institute of Technology
Cambridge, MA, USA
jinhua@mit.edu

Abstract

Urban prediction tasks, such as forecasting traffic flow, temperature, and crime rates, are crucial for efficient urban planning and management. However, existing Spatiotemporal Graph Neural Networks (ST-GNNs) often rely solely on accuracy, overlooking spatial and demographic disparities in their predictions. This oversight can lead to imbalanced resource allocation and exacerbate existing inequities in urban areas. This study introduces a Residual-Aware Attention (RAA) Block and an equality-enhancing loss function to address these disparities. By adapting the adjacency matrix during training and incorporating spatial disparity metrics, our approach aims to reduce local segregation of residuals and errors. We applied our methodology to urban prediction tasks in Chicago, utilizing travel demand datasets as an example. Our model achieved a 48% significant improvement in fairness metrics with only a 9% increase in error metrics. Spatial analysis of residual distributions revealed that models with RAA Blocks produced more equitable prediction results, particularly by reducing errors clustered in central regions, supporting more balanced and equitable urban planning and policy-making.

CCS Concepts

• **Information systems** → **Data mining**; • **Computing methodologies** → *Machine learning approaches*.

Keywords

Urban Prediction, Deep Learning, Algorithmic Fairness, Equity

ACM Reference Format:

Dingyi Zhuang, Hanyong Xu, Xiaotong Guo, Yunhan Zheng, Shen hao Wang, and Jinhua Zhao. 2025. Mitigating Spatial Disparity in Urban Prediction

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebST '25, April 28-29, 2025, Sydney, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Using Residual-Aware Spatiotemporal Graph Neural Networks: A Chicago Case Study. In *Proceedings of the International World Wide Web Conference (WebST '25)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>


1 Introduction

Urban prediction tasks, such as forecasting traffic flow, crime rates, and temperature, play a critical role in urban planning and management [26]. These tasks benefit from machine learning and deep learning techniques, particularly Spatio-Temporal Graph Neural Networks (ST-GNNs), which effectively capture both spatial and temporal dependencies. By combining Graph Neural Networks (GNNs) with temporal processing architectures, ST-GNNs model complex relationships between urban regions and time points, delivering more accurate and reliable forecasts.

However, previous ST-GNN models focus exclusively on accuracy, neglecting the social implications of their predictions. Figure 1 shows that STGCN predictions in Chicago correlate with minority rates, resulting in over-predictions in low-minority areas and under-predictions in high-minority areas [34]. These patterns highlight spatial disparities, intrinsically tied to demographic segregation, as illustrated in Figure 1b. Such disparities often stem from biases in data collection (e.g., over-policing in minority neighborhoods) and inherent model biases, both of which can amplify inequities [3, 7].

Existing methods mitigate disparities by incorporating demographic information into model design [36]. However, practical limitations, such as privacy concerns, restrict the collection and use of demographic data, leaving disparities within ST-GNNs unaddressed [38]. This raises a critical question: **How can we improve equality in ST-GNN predictions without access to protected group memberships?**

This study addresses spatial disparities in ST-GNN models, aiming to reduce spatial segregation and its links to demographic disparities. We propose a Residual-Aware Attention (RAA) Block and an equality-enhancing loss function. The RAA Block adapts the adjacency matrix during training, dynamically adjusting spatial relationships based on residuals to reduce local error segregation. The loss function integrates mean squared error with metrics like Moran's I and Generalized Entropy Index (GEI), penalizing spatial



(a) STGCN prediction residual distribution in Chicago. (b) Minority rate distribution in Chicago.

Figure 1: Comparison between prediction residual and demographic distributions. (a): The red colors represent under-prediction while the blue colors represent over-prediction. (b): Distribution of the racial minority groups in Chicago, defined as the percentage of the non-white population in each community area.

clustering of residuals and redundancy. Our method mitigates disparities within the model, supporting equitable urban planning and policy-making.

Applied to urban prediction tasks in Chicago using datasets including travel demand, crime, and accident reports, our approach achieved a 48% improvement in fairness metrics with only a 9% increase in error metrics. Spatial analysis demonstrated that RAA Blocks reduce clustered prediction errors, particularly in central regions. Overall, our methodology effectively mitigates prediction disparities with minimal accuracy trade-offs, fostering fairer resource allocation and equitable urban development.

Our contributions are summarized as follows:

- (1) We introduce residuals as indicators of fairness in predictions, demonstrating their utility in highlighting disparities in ST-GNN model outputs.
- (2) We design the RAA Block and propose an equality-enhancing loss function that incorporates spatial and demographic disparity metrics. This function balances prediction accuracy with fairness, ensuring that urban prediction models do not entrench existing spatial disparities. We modify the message-passing mechanism in GNNs to ensure the propagation of information that reduces bias.
- (3) We conduct a comprehensive case study applying our method to urban prediction tasks in Chicago using the travel demand as an example. Our model identifies and mitigates prediction disparities without relying on demographic data, leading to more balanced resource allocation and fairer urban planning decisions.

2 Literature Review

2.1 Urban Prediction with GNNs

Recent advancements in Graph Neural Networks (GNNs) have made them highly effective in representing spatial data as graphs and

capturing both spatial and temporal dependencies. These capabilities have driven their adoption in urban computing tasks [13, 33]. One prominent application domain is transportation, where GNNs have been extensively used for tasks such as ride-hailing demand prediction [8, 29]. For example, Wang et al. [28] tackled the Origin-Destination Matrix Problem by integrating grid embedding and multi-task learning modules to capture spatial and temporal attributes. Similarly, Jin et al. [12] used graph convolutional neural networks with pixel-level representations to model joint latent distributions of ride-hailing demands.

Beyond ride-hailing, GNNs have been applied to traffic incident prediction [24], travel time estimation [11], human mobility analysis [30], traffic flow prediction [5], and trajectory forecasting [22]. For instance, Yu et al. [35] proposed a framework combining spatial graph convolutional networks with spatiotemporal convolutions to predict traffic accidents, while Fang et al. [5] enhanced GNNs’ depth to extract long-range spatiotemporal dependencies for traffic flow prediction.

GNN applications extend beyond transportation to public safety and environmental monitoring [13]. In crime prediction, where spatiotemporal correlations play a vital role, Zhang and Cheng [37] developed a framework combining gated networks and localized diffusion networks to explain temporal and spatial crime propagations. Similarly, Wang et al. [25] introduced a model with adaptive region graph learning and homophily-aware constraints to predict crime patterns. Other applications include urban change detection [40] and disaster forecasting [6].

Despite their success, most GNN-based urban computation models prioritize improving prediction accuracy while neglecting fairness issues. This oversight risks propagating algorithmic biases and exacerbating social disparities.

2.2 Algorithmic Fairness

Fairness in algorithmic systems typically aligns with two key principles: *equality* and *equity* [20]. *Equality*, or horizontal equity, ensures equal treatment for all, while *equity*, or vertical equity, focuses on providing resources to those in need [18, 32]. These goals are evaluated through Disparate Treatment Analysis, which seeks fair treatment, or Disparate Impact Analysis, which ensures fair outcomes [20, 21]. This study adopts *equality* as the primary goal, aiming for equitable resource allocation across urban regions.

Prediction residuals offer a straightforward measure of algorithmic fairness by reflecting under- and over-prediction patterns [14]. This approach aligns with *fairness through unawareness*, which avoids using protected attributes in decision-making [4, 15]. In contrast, *fairness through awareness* integrates protected attributes directly into the model [9].

In transportation systems, studies have explored mitigating algorithmic disparities. Zheng et al. [39] proposed using absolute correlation regularization to address disparities in neural networks. Similarly, Zheng et al. [38], Zhang et al. [36], and Guo et al. [10] addressed group fairness in ride-hailing demand predictions but relied on demographic information, limiting their applicability due to privacy concerns. These methods fail to address the root causes of disparities, such as biases in the message-passing mechanisms of GNNs.

To overcome these limitations, we pursue *fairness through un-awareness* by focusing on the inherent biases in ST-GNN models without relying on demographic data. By targeting spatially segregated residual patterns, we aim to achieve equality in urban predictions and address disparities rooted in model design.

3 Problem Description and Preliminaries

3.1 ST-GNN for Forecasting

ST-GNNs have been extensively utilized in urban prediction tasks [31, 42]. We formulate the ST-GNN framework as follows: Let the graph be denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where \mathcal{V} represents the set of nodes (locations/regions), \mathcal{E} denotes the set of edges, and $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the adjacency matrix that encodes the relationships between nodes.

In this study, we define nodes as regions or census tracts and establish edges based on the geographic distances between regions. The adjacency matrix \mathbf{A} captures these geographical relationships, with larger values indicating closer proximities.

The urban spatiotemporal dataset inputs are denoted as $\mathcal{X} \in \mathbb{R}^{|\mathcal{V}| \times t}$, where t is the number of time steps. The objective is to predict the target values $\mathcal{Y}_{1:|\mathcal{V}|,t:t+k}$ for the future k time steps, given historical data up to time t , $\mathcal{X}_{1:|\mathcal{V}|,1:t}$. ST-GNN models aim to learn a mapping function f_θ , parameterized by θ , such that:

$$\hat{\mathcal{Y}}_{1:|\mathcal{V}|,t:t+k} = f_\theta(\mathcal{X}_{1:|\mathcal{V}|,1:t}; \mathcal{G}) = f_\theta(\mathcal{X}_{1:|\mathcal{V}|,1:t}; \mathcal{V}, \mathcal{E}, \mathbf{A}), \quad (1)$$

where $\hat{\mathcal{Y}}_{1:|\mathcal{V}|,t:t+k}$ is the predicted target value, typically corresponding to $\mathcal{X}_{1:|\mathcal{V}|,t:t+k}$ in forecasting tasks. Prediction performance is measured using residual-based metrics such as Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE). The prediction residuals are defined as $\mathbf{r}_{1:|\mathcal{V}|,t:t+k} = \mathcal{Y}_{1:|\mathcal{V}|,t:t+k} - \hat{\mathcal{Y}}_{1:|\mathcal{V}|,t:t+k}$, which we denote as \mathbf{r} for brevity.

The loss function used during training typically minimizes RMSE or MAE to optimize prediction accuracy:

$$\mathcal{L}_{\text{accuracy}} = \text{RMSE}(\mathcal{Y}, \hat{\mathcal{Y}}) \quad \text{or} \quad \text{MAE}(\mathcal{Y}, \hat{\mathcal{Y}}). \quad (2)$$

While these approaches focus on accuracy, few studies have explored whether they achieve equity in urban planning. Prediction biases may vary across spatial regions or demographic groups, potentially resulting in inequitable resource allocation and urban management outcomes.

3.2 Measurement of Disparity in ST-GNN Prediction Results

The most straightforward measurement of disparities in ST-GNN prediction results is the prediction residuals. Disparities are defined in terms of spatial and demographic factors.

Our models consider spatial disparities to address demographic disparities. Spatial disparity refers to prediction residuals exhibiting similar over-prediction or under-prediction patterns within local neighborhoods, resulting in significant differences across the global space. For instance, the central region of Chicago is consistently over-predicted while the southern part is under-predicted, as shown in Figure 1a. This leads to structural differences that can affect resource allocation, where over-predicted areas might receive more resources, and under-predicted areas receive less, raising equality issues.

Mathematically, multiplying a vector by the adjacency matrix \mathbf{A} aggregates information from neighboring nodes. We define the spatial disparity D_s based on the sign-aware residual variance of the neighboring nodes' predictions, as shown below:

$$D_s = \frac{1}{|\mathcal{V}|} \sum_{i=1}^{|\mathcal{V}|} \left[(\mathbf{A}\mathbf{r}_i^+ - \bar{\mathbf{s}}^+)^2 + (\mathbf{A}\mathbf{r}_i^- - \bar{\mathbf{s}}^-)^2 \right], \quad (3)$$

where \mathbf{r}^+ and \mathbf{r}^- represent the positive and negative residuals, \mathbf{s}^+ and \mathbf{s}^- are the spatially weighted residuals, and $\bar{\mathbf{s}}^+$ and $\bar{\mathbf{s}}^-$ are their respective averages. The adjacency matrix \mathbf{A} captures the influence of neighboring nodes, and D_s summarizes variance in residuals across neighbors.

Demographic disparity arises when over-prediction correlates with specific demographic groups, favoring the majority over the minority group. For instance, spatial disparities often align with demographic segregation, as seen in Chicago (Figure 1b). We define demographic disparity D_d as:

$$D_d = |\text{Corr}(\mathbf{r}, \text{Pop}_{\text{minor}})| + |\text{Corr}(\mathbf{r}, \text{Pop}_{\text{major}})|, \quad (4)$$

where $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{minor}})$ and $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{major}})$ are Pearson correlation coefficients between residuals \mathbf{r} and the percentages of minority ($\text{Pop}_{\text{minor}}$) and majority ($\text{Pop}_{\text{major}}$) populations, respectively.

3.3 Attention Mechanism Preliminary

The attention mechanism enhances neural networks by focusing on relevant parts of the input, enabling models to effectively capture dependencies in sequential and spatial data.

Given an input data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, where n is the number of elements (e.g., time steps, nodes) and d is the feature dimension, the query (\mathbf{Q}), key (\mathbf{K}), and value (\mathbf{V}) matrices are computed as $\mathbf{Q} = \mathbf{X}\mathbf{W}_Q$, $\mathbf{K} = \mathbf{X}\mathbf{W}_K$, and $\mathbf{V} = \mathbf{X}\mathbf{W}_V$, where \mathbf{W}_Q , \mathbf{W}_K , $\mathbf{W}_V \in \mathbb{R}^{d \times d_k}$ are learnable parameters, and d_k is the feature dimension of the query, key, and value vectors. The attention scores are calculated using the scaled dot-product of the query and key matrices as:

$$\mathbf{S} = \frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}}, \quad (5)$$

where $\mathbf{S} \in \mathbb{R}^{n \times n}$ quantifies the relevance between input elements, with higher scores indicating stronger relationships.

The attention weights are obtained by applying the softmax function to normalize the scores: $\mathbf{H} = \text{softmax}(\mathbf{S})$, where \mathbf{H} represents the importance of each element in the input.

This mechanism has become a cornerstone of modern neural networks, significantly improving performance in tasks requiring an understanding of complex spatial and sequential relationships.

4 Methodology

Our overall framework, depicted in Figure 2, enhances existing ST-GNN architectures for urban prediction by integrating two key components to address disparities: 1) The Residual-Aware Attention (RAA) Block, which dynamically adapts the adjacency matrix during training to mitigate spatial disparities. This mechanism leverages residuals to refine spatial relationships within the graph, reducing local segregation of residuals and errors; 2) An equality-enhancing loss function that penalizes similar residual patterns in

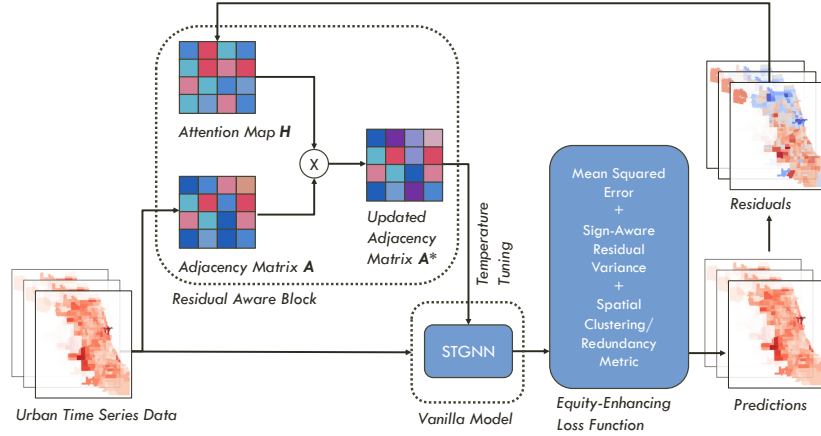


Figure 2: Prediction framework with residual aware attention block.

adjacent neighborhoods and incorporates a spatial clustering redundancy metric. By balancing prediction accuracy with spatial and demographic equality, this function addresses spatial disparities and, in turn, mitigates demographic disparities.

4.1 RAA Block

The adjacency matrix in GNN structures is crucial for shaping GNN results by ensuring that neighboring information is shared, leading to similar feature representations for adjacent nodes. However, this emphasis on local information can exacerbate spatial disparities, as discussed in Equation 3.

To adaptively mitigate these disparities, we incorporate a Residual-Aware Attention (RAA) Layer that uses residuals from each training step as inputs, leveraging the attention mechanism described in Section 3.3. During training, residuals r are used to compute the query (Q), key (K), and value (V) matrices through linear transformations. The \tanh activation function introduces non-linearity and maps values to the range $[-1, 1]$:

$$Q = \tanh(W_q r), \quad K = \tanh(W_k r), \quad V = \tanh(W_v r), \quad (6)$$

where W_q, W_k, W_v are learnable weight matrices. Based on Equation 5, the attention scores and attention weights are:

$$S = \frac{QK^T}{\sqrt{|K|}}, \quad H = \text{softmax}(S), \quad (7)$$

where S represents the relative importance between elements, and H normalizes these scores into probabilities.

The adapted adjacency matrix is then formed by applying an element-wise Hadamard product between the original adjacency matrix and the attention weights:

$$A_{\text{adapted}} = A \odot H. \quad (8)$$

This adapted adjacency matrix is used in the next training epoch, enabling the model to dynamically adjust spatial relationships within the graph and reduce the segregation of residuals and errors.

By integrating this adaptive mechanism, the RAA Block effectively mitigates spatial and demographic disparities in urban predictions.

4.2 Equality-enhancing Loss Function

To address disparities and promote equality in urban prediction, we propose an equality-enhancing loss function that integrates the mean squared error (MSE) loss with additional terms accounting for spatial disparities. The overall loss function is:

$$\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{prediction}} + \lambda_s D_s + \lambda_d D_d, \quad (9)$$

where $\mathcal{L}_{\text{prediction}}$ represents the MSE loss, and λ_s and λ_d are regularization parameters balancing the contribution of each term. These parameters are set to 0.05 by default but can be adjusted for different datasets.

The term D_s measures spatial disparity using the sign-aware residual variance as defined in Equation 3, ensuring that prediction residuals are evenly distributed across spatial regions. The D_d term incorporates fairness metrics, such as Moran’s I or the Generalized Entropy Index (GEI), to assess spatial clustering or redundancy in information. These metrics aim to reduce spatial unevenness in prediction residuals, promoting more equitable outcomes.

By including D_s and D_d in the loss function, this approach ensures that urban prediction models address disparities at both spatial and demographic levels, supporting equitable resource allocation and informed urban planning. Details of the D_d metric formulations are provided in Section 5.3.

5 Experimental Setup

5.1 Urban Data Collections in Chicago

The demographic dataset for this research is sourced from the American Community Survey (ACS) for the years 2017-2018. The ACS provides detailed sociodemographic information, including income brackets, age demographics, racial compositions, commuting modes, and average commuting durations. This data is specific to

each of the 811 census tracts within Chicago [27, 41]. Key attributes include total population, age groups, racial composition, education levels, economic status, and travel information such as commuting times¹.

We use **Chicago Data Portal (CDP)**² to evaluate the effects of our model. CDP contains the trip records of Transportation Network Providers (ride-sharing companies) in the Chicago area. The city of Chicago is divided into 77 zones and the trip requests with pick-up and drop-off zones are recorded every 15 minutes. We use 4-month observations from September 1st, 2019 to December 30th, 2019.

We have chosen Chicago as a case study since its heavy spatial segregation of different racial groups, for which unfair algorithmic prediction results may lead to even worse segregation. Fig 1b shows the racial minority rate in Chicago in 2019. It is clear that minority races cluster around the south and middle-west of the city, while regions in the north has far less minorities than the other regions.

5.2 Model Comparison

We selected four prevalent ST-GNN models as base models to demonstrate the effectiveness of our proposed method:

- **DCRNN** [17] models traffic dynamics using diffusion convolution. It captures spatial dependencies through bidirectional random walks on the graph and temporal dependencies with an encoder-decoder architecture and scheduled sampling.
- **DSTAGNN** [16] learns dynamic association attributes from data to represent the graph. It employs a multi-head attention mechanism for spatial variances and handles temporal dependencies with features of multi-receptive fields.
- **STGCN** [34] includes two spatio-temporal convolutional blocks that utilize graph convolutional layers to capture spatial dependencies and temporal gated convolution layers for temporal dynamics.
- **AGCRN** [1] comprises two modules: Node Adaptive Parameter Learning, which learns node-specific parameters from node embeddings, and Data Adaptive Graph Generation, which generates a graph from the training data. This architecture effectively captures fine-grained variability in space and time.

For each model, the performance of the vanilla version will be compared with the performance of the models with the RAA block and enhanced loss functions added. The vanilla version model implementations were based on the repository by Liu et al. [19]³.

5.3 Fairness Metrics

Fairness metrics are crucial for evaluating equality in urban prediction models, addressing both spatial and demographic disparities. These metrics ensure that models do not exacerbate existing inequalities, thereby promoting fair and equitable urban predictions. In this study, we selected the following fairness metrics to evaluate three key aspects of residual distributions: entropy, correlation with demographic features, and spatial clustering.

¹<https://www.census.gov/programs-surveys/acs/data.html>

²<https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Trips/m6dm-c72p>

³<https://github.com/liuxu77/LargeST/>

5.3.1 Generalized Entropy Index (GEI). It is developed by Speicher et al. [23] and originates from economic inequality indices and is used to compare unfair algorithmic treatments across a population. In this study, GEI reflects spatial disparity rather than demographic disparity.

We calculate the GEI of prediction residuals, where a smaller value indicates a more even distribution of residuals across all units of analysis, resulting in fairer predictions. Denote a non-negative transformation of the residuals $b_i = r_i + m$, where $m \in \mathbb{R}^+$ such that $b_i \geq 0 \forall i$, and r_i represents the residual at node i . The formulation is as follows:

$$GEI = \frac{1}{|\mathcal{V}|^{\alpha(\alpha-1)}} \sum_{i=1}^{|\mathcal{V}|} \left[\left(\frac{b_i}{\bar{b}} \right)^{\alpha} - 1 \right], \quad (10)$$

The parameter α controls the emphasis on larger residuals, and we set $\alpha = 2$ in this research.

5.3.2 Moran's I. It measures spatial autocorrelation, indicating how similar values are clustered or dispersed across space. It ranges from -1 (dispersion) to 1 (clustering), with 0 indicating random distribution. Lower values are preferred to reduce spatial clustering.

The weight W is computed as: $W = \sum_{i=1}^N \sum_{j=1}^N a_{ij}$, where a_{ij} are elements of the adjacency matrix A , and N is the number of nodes. Moran's I is then calculated as:

$$I = \frac{N}{W} \frac{\sum_{i=1}^N \sum_{j=1}^N a_{ij} (r_i - \bar{r})(r_j - \bar{r})}{\sum_{i=1}^N (r_i - \bar{r})^2}, \quad (11)$$

where r_i and r_j are residuals at nodes i and j , and \bar{r} is the residual mean. In order to ensure non-negative values for the loss function, we adjust as: $I^* = I + 1$. This adjustment supports penalizing spatial clustering in our loss function.

5.3.3 Scaled Disparity Index (SDI). While we do not incorporate demographic disparities in model training, we still need to measure if mitigating spatial disparities also reduces demographic disparities. Hence, we propose using the Pearson correlation coefficients between prediction residuals and the minority populations ($\text{Corr}(\mathbf{r}, \text{Pop}_{\text{minor}})$) and residuals and the majority populations ($\text{Corr}(\mathbf{r}, \text{Pop}_{\text{major}})$) to assess fairness.

To quantify disparity, we introduce the SDI:

$$SDI = \frac{|\Delta \text{Corr}|}{|\text{Corr}_{\text{minor}}| + |\text{Corr}_{\text{major}}|} \cdot \sqrt{|\text{Corr}_{\text{minor}} \cdot \text{Corr}_{\text{major}}|}, \quad (12)$$

The first term of the SDI normalizes the disparity by taking the difference between $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{minor}})$ and $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{major}})$ and scaling it by the sum of their absolute values. The geometric mean of $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{minor}})$ and $\text{Corr}(\mathbf{r}, \text{Pop}_{\text{major}})$ then adjusts this normalized disparity by accounting for the magnitude of the correlations. This scaling is crucial, as it distinguishes between scenarios where correlations have different magnitudes but similar patterns of disparity.

Ultimately, the SDI yields non-negative values, with smaller values indicating greater fairness in our context. The SDI can also be extended to compare other advantaged and disadvantaged groups, such as evaluating equality between poor and rich communities.

Models	Original model					Residual-Aware Attention Variants					
	MAE	SMAPE	GEI	SDI	Moran's I	Variants	MAE	SMAPE	GEI	SDI	Moran's I
DCRNN	8.092	0.458	1.28	0.2	0.182	RAA block + D_s	7.492 \downarrow 7%	0.478 \uparrow 4%	1.106 \downarrow 13%	0.16 \downarrow 20%	0.068 \downarrow 62%
						+ loss with Moran's I	8.911 \uparrow 10%	0.555 \uparrow 21%	1.148 \downarrow 10%	0.183 \downarrow 8%	-0.135 \downarrow 174%
						+ loss with GEI	7.608 \downarrow 5%	0.483 \uparrow 5%	1.2 \downarrow 6%	0.072 \downarrow 64%	0.046 \downarrow 74%
DSTAGNN	8.564	0.425	1.383	0.308	0.542	RAA block + D_s	8.185 \downarrow 4%	0.456 \uparrow 7%	1.2 \downarrow 13%	0.269 \downarrow 12%	0.436 \downarrow 19%
						+ loss with Moran's I	8.509 \downarrow 0%	0.448 \uparrow 5%	2.862 \uparrow 106%	0.281 \downarrow 8%	0.558 \uparrow 2%
						+ loss with GEI	9.555 \uparrow 11%	0.447 \uparrow 5%	0.434 \downarrow 68%	0.238 \downarrow 22%	-0.021 \downarrow 103%
STGCN	6.948	0.428	1.506	0.337	0.394	RAA block + D_s	7.329 \uparrow 5%	0.538 \uparrow 25%	1.075 \downarrow 28%	0.235 \downarrow 30%	0.01 \downarrow 97%
						+ loss with Moran's I	7.885 \uparrow 13%	0.506 \uparrow 18%	1.692 \uparrow 12%	0.106 \downarrow 68%	0.201 \downarrow 48%
						+ loss with GEI	9.349 \uparrow 34%	0.476 \uparrow 11%	0.35 \downarrow 76%	0.056 \downarrow 83%	-0.108 \downarrow 127%
AGCRN	6.988	0.425	1.265	0.316	0.064	RAA block + D_s	7.039 \downarrow 0%	0.465 \uparrow 9%	0.975 \downarrow 22%	0.007 \downarrow 97%	0.019 \downarrow 70%
						+ loss with Moran's I	7.393 \uparrow 5%	0.49 \uparrow 15%	0.978 \downarrow 22%	0.07 \downarrow 77%	-0.034 \downarrow 153%
						+ loss with GEI	8.419 \uparrow 20%	0.49 \uparrow 15%	0.373 \downarrow 70%	0.084 \downarrow 73%	0.04 \downarrow 37%

Table 1: Comparison of Models with different Residual-Aware Attention versions

5.4 Error Metrics

We also evaluate prediction accuracy using common error metrics: the Symmetric Mean Absolute Percentage Error (SMAPE) and the Mean Absolute Error (MAE). The SMAPE is given by:

$$\text{SMAPE} = \frac{1}{N} \sum_{i=1}^N \frac{2|\hat{y}_i - y_i| + \epsilon}{|\hat{y}_i| + |y_i| + \epsilon}, \quad (13)$$

where \hat{y}_i and y_i are the predicted and true values, N is the number of observations, and ϵ is a small constant to prevent division by zero.

Similarly, the MAE is defined as:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i|, \quad (14)$$

providing the average magnitude of prediction errors.

For all metrics (GEI, Moran's I, SDI, SMAPE, MAE), lower values indicate better performance.

6 Results and Analysis

6.1 Overall Performance

The results are presented in two parts: the first evaluates the accuracy and fairness performance of the original vanilla models compared to the RAA-enhanced model variants, while the second explores the impact of the RAA block, sign-aware residual variance, and Moran's I or GEI metrics through an ablation study.

We evaluate three RAA-enhanced variants for all four base models in Section 5.2, differentiated by the D_d term in the loss function:

- (1) RAA block with D_s in the loss function;
- (2) RAA block with D_s and Moran's I as D_d ;
- (3) RAA block with D_s and GEI as D_d .

Table 1 summarizes the accuracy and fairness metrics for each model and its variants. The base model metrics are shown on the left, while RAA-enhanced variants are on the right. Improvements are highlighted in green, and declines in red, with percentages indicating the magnitude of change compared to the base model. All experiments were performed on a machine running Ubuntu 22.04, equipped with an Intel(R) Core(TM) i9-10980XE CPU @ 3.00GHz, 128GB RAM, and an NVIDIA GeForce RTX 4080 GPU.

In general, introducing the RAA blocks and the equality-enhancing loss functions improved fairness metrics compared to the base models. On average, GEI, SDI, and Moran's I decreased by 18%, 47%, and 80%, respectively. Across 12 RAA experiments, 10 showed reduced values for all fairness metrics, and all achieved lower SDI. These results highlight the effectiveness of the RAA modules. Averaging the percentage changes for each RAA variant shows reductions of 40%, 37%, and 67%, respectively. The variant combining the RAA block with GEI as D_d exhibited the best fairness improvement. Among the four GNN models, DCRNN and AGCRN with RAA blocks demonstrated the most consistent improvements, achieving reductions across all fairness metrics in all three variants. AGCRN, in particular, showed the highest overall fairness improvements, demonstrating that reducing spatial disparities can also mitigate demographic disparities.

We observed a trade-off between accuracy and fairness. Among the 12 experiments, all achieved improvements in at least one fairness metric while experiencing slight decreases in error metrics. This aligns with the project's objective of reducing prediction error variances at a minimal cost to accuracy. Notably, the percentage increases in MAE (7%) and SMAPE (12%) were significantly lower than the fairness improvements. These results suggest that fairness improvements can be achieved with minimal accuracy trade-offs. Furthermore, in several cases, adding RAA blocks and regularization terms improved both accuracy and fairness. This phenomenon is attributed to *Model Multiplicity* or *Under-Specification*, where the complex nature of the loss function enables the model to balance multiple objectives during training [2].

6.2 Residual Spatial Distribution

To understand the effect of the RAA modules on the residual spatial distributions, we have plotted the average residuals at each community area on the map, as shown in Figure 3. The red colors represent positive residuals, or under-prediction, while the blue colors representing the opposite. Darker colors demonstrate larger residuals.

Comparing the spatial distribution of residuals between the original model and models with RAA modules, we observe that the latter produces more equitable prediction results, as indicated by the less prominent residual clusters on the map. Models with RAA

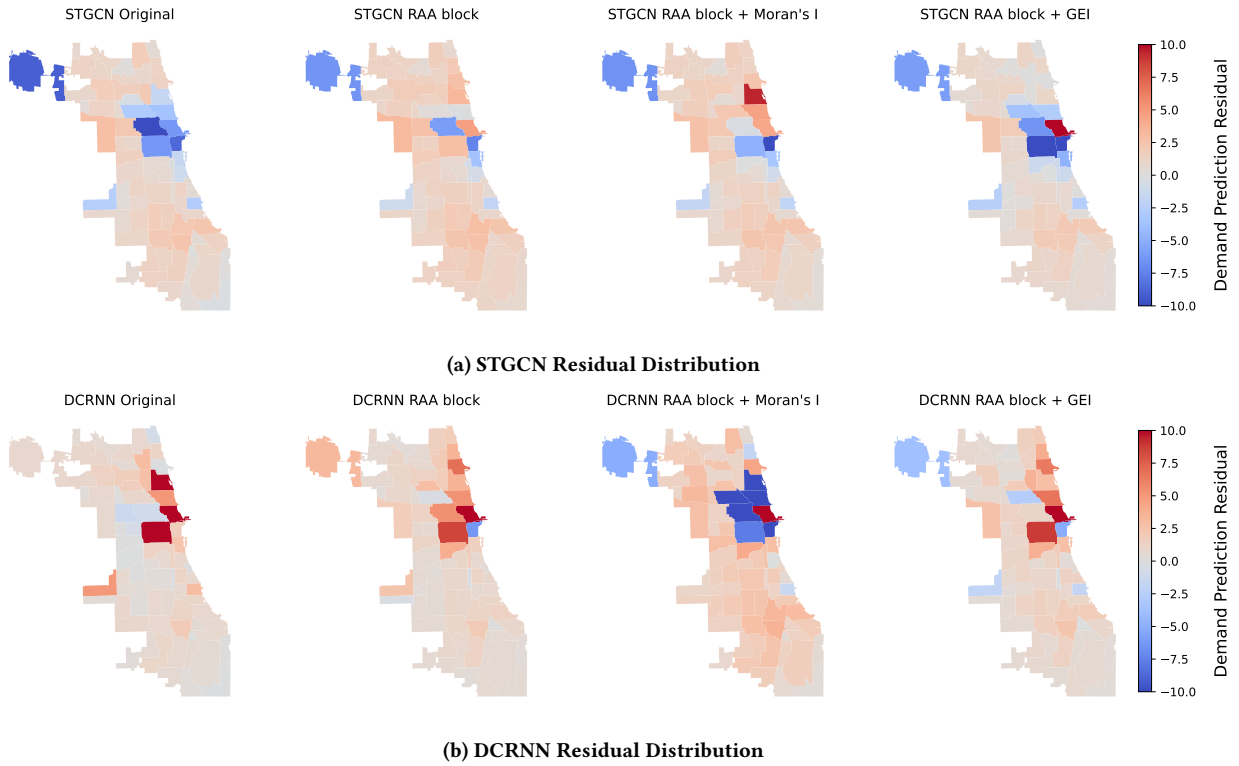


Figure 3: Residual spatial distribution in Chicago.

blocks are effective in reducing errors concentrated in central city regions. For instance, in the case of the STGCN model shown in Figure 3a, significant over-prediction occurs in central Chicago. However, models with the RAA block and Moran’s I in the loss function display a much smoother residual distribution. Although the STGCN model with GEI as a regularization term D_d still exhibits some clustering of large residuals in the central region, the residual signs are more heterogeneous, indicating reduced spatial autocorrelation.

It is also worth noting that, although RAA blocks demonstrate some effectiveness in reducing disparities in the residual spatial distribution, their effectiveness is model-dependent. There is a trade-off between residual clustering, variance, and the imbalance between over- and under-predictions. It is challenging for a single model to improve all three aspects simultaneously, which explains the varied performances of different models on the map and in terms of different fairness metrics.

6.3 Ablation Study

An ablation study is a method used in machine learning to evaluate the impact of individual components on a model’s performance. By systematically altering specific elements and observing changes in performance metrics, researchers can identify which components are essential. We conducted 6 ablation cases in total, corresponding to:

- **RAA block**: adding RAA block to the architecture alone;

- **RAA block + D_s** : adding the RAA block to the architecture and adding the D_s regularization term in the loss function;
- **RAA block + Moran’s I**: adding the RAA block to the architecture and adding the Moran’s I metric as the D_d term in the loss function;
- **RAA block + GEI**: adding the RAA block to the architecture and adding the GEI metric as the D_d term in the loss function;
- **RAA block + D_s + Moran’s I**: adding the RAA block to the architecture, adding the D_s term and the Moran’s I metric as the D_d term to form the Equation 9;
- **RAA block+ D_s + GEI**: adding the RAA block to the architecture, adding the D_s term and the GEI metric as the D_d term to form the Equation 9.

No.	Model Variants	MAE	SMAPE	GEI	SDI	Moran’s I
1	Original	6.948	0.428	1.506	0.337	0.394
2	RAA block	7.830	0.475	0.402	0.154	-0.000
3	RAA block + D_s	7.257	0.474	1.854	0.140	0.212
4	RAA block + Moran’s I	10.465	0.470	0.842	0.151	0.251
5	RAA block + GEI	8.028	0.525	0.525	0.044	-0.136
6	RAA block + D_s + Moran’s I	7.924	0.508	1.461	0.015	0.012
7	RAA block + D_s + GEI	9.349	0.476	0.350	0.056	-0.108

Table 2: Ablation study of our designed modules using the STGCN model.

To illustrate the effect of model architecture design choices, we conducted the aforementioned ablation study in the context of the

STGCN model, as shown in Table 2. The ablation study validates each strategy we have introduced.

First, the RAA block significantly improves fairness metrics compared to the original model, although it sacrifices accuracy. Comparing Model 1 and Model 2 from Table 2, we observe reductions in GEI, SDI, and Moran’s I, indicating reduced residual variances, correlations to demographics, and spatial autocorrelation. This supports the hypothesis that using residuals to calculate the Attention map helps direct the model toward more equitable results. Moreover, both spatial and demographic disparities could be mitigated using this method.

Second, adding the D_s term improves accuracy but trades off with fairness results. This is demonstrated by comparing three pairs of models: Model 2 and Model 3, Model 4 and Model 6, and Model 5 and Model 7. Models with D_s regularization show at least one lower accuracy metric compared to those without. The D_s term impacts error metrics because it separately calculates the variance of positive and negative residuals, effectively weighting errors more heavily in the loss function, causing the model to focus more on accuracy.

Lastly, adding D_s and D_d terms improves the corresponding fairness metrics. Comparing between Model 2 and Model 5, as well as Model 3 and Model 7, we see that introducing GEI in the loss function as a regularization term significantly enhances the GEI performance of model predictions. On the other hand, comparing between Model 2 and Model 4 as well as Model 3 and Model 6, we see that the Moran’s I metric as the D_d term is less effective in improving the output Moran’s I metric, but can help with other fairness metrics. Thus, over-emphasizing any one metric can lead to worse results in others. It is crucial to balance multiple objectives in the loss function to avoid such scenarios. In our case, this means ensuring similar regularization weights for both accuracy and fairness terms.

7 Discussion and conclusions

Previous urban prediction work prevalently built on ST-GNNs focuses solely on accuracy, neglecting social impacts. By focusing on residuals as indicators of fairness, we effectively highlight disparities in traditional ST-GNN outputs. This study addresses spatial and demographic disparities in urban prediction tasks by developing an RAA Block and an equality-enhancing loss function. Our approach, integrated into existing ST-GNNs, dynamically adjusts spatial relationships during training, mitigating spatial disparities.

Applied to urban prediction tasks in Chicago, our methodology demonstrates significant improvements both in fairness metrics and error metrics. This shows that reducing spatial disparities can also help mitigate demographic disparities. Moreover, our approach reduces the local segregation of residuals and errors. Spatial analysis of residual distributions shows that models with RAA Blocks effectively reduced clustered prediction errors in central regions. Through the comprehensive case study in Chicago, we demonstrate the effectiveness of our approach in mitigating prediction disparities for future equitable urban city management.

However, the RAA Block’s effectiveness is model-dependent, and there is a trade-off between residual clustering, variance, and the

balance of over- and under-predictions. Future work should optimize these factors to enhance performance. Although our approach does not rely on demographic data, incorporating such information when available could provide a more comprehensive understanding of fairness in urban predictions.

Acknowledgment

This research is based upon work supported by the U.S. Department of Energy’s Office of Energy Efficiency and Renewable Energy (EERE) under the Vehicle Technology Program Award Number DE-EE0009211.

References

- [1] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2020. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems* 33 (2020), 17804–17815.
- [2] Emily Black, Manish Raghavan, and Solon Barocas. 2022. Model Multiplicity: Opportunities, Concerns, and Solutions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’22)*. Association for Computing Machinery, New York, NY, USA, 850–863. <https://doi.org/10.1145/3531146.3533149>
- [3] Yushun Dong, Ninghao Liu, Brian Jalaian, and Jundong Li. 2022. EDITS: Modeling and Mitigating Data Bias for Graph Neural Networks. In *Proceedings of the ACM Web Conference 2022*. 1259–1269. <https://doi.org/10.1145/3485447.3512173> arXiv:2108.05233 [cs].
- [4] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS ’12)*. Association for Computing Machinery, New York, NY, USA, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [5] Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. 2021. Spatial-Temporal Graph ODE Networks for Traffic Flow Forecasting. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 364–373.
- [6] Hamed Farahmand, Yuanchang Xu, and Ali Mostafavi. 2023. A spatial-temporal graph deep learning model for urban flood nowcasting leveraging heterogeneous community features. *Scientific Reports* 13, 1 (April 2023), 6768. <https://doi.org/10.1038/s41598-023-32548-x>
- [7] Gillian Franklin, Rachel Stephens, Muhammad Piracha, Shmuel Tiosano, Frank Lehouillier, Ross Koppel, and Peter L Elkin. 2024. The Sociodemographic Biases in Machine Learning Algorithms: A Biomedical Informatics Perspective. *Life* 14, 6 (2024), 652.
- [8] Xu Geng, Yaguang Li, Leye Wang, Lingyu Zhang, Qiang Yang, Jieping Ye, and Yan Liu. 2019. Spatiotemporal Multi-Graph Convolution Network for Ride-Hailing Demand Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence* 33 (2019), 3656–3663. <https://doi.org/10.1609/aaai.v33i01.33013656>
- [9] Nina Grgic-Hlac̃a, Muhammad Bilal Zafar, Krishna P Gummedi, and Adrian Weller. 2016. The Case for Process Fairness in Learning: Feature Selection for Fair Decision Making. (2016).
- [10] Xiaotong Guo, Hanyong Xu, Dingyi Zhuang, Yunhan Zheng, and Jinhua Zhao. 2023. Fairness-Enhancing Vehicle Rebalancing in the Ride-hailing System. arXiv:2401.00093 (Dec. 2023). <https://doi.org/10.48550/arXiv.2401.00093> arXiv:2401.00093 [cs].
- [11] Jizhou Huang, Zhengjie Huang, Xiaomin Fang, Shikun Feng, Xuyi Chen, Jiexiang Liu, Haitao Yuan, and Haifeng Wang. 2022. DuETA: Traffic Congestion Propagation Pattern Modeling via Efficient Graph Learning for ETA Prediction at Baidu Maps. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 3172–3181. <https://doi.org/10.1145/3511808.3557091> arXiv:2208.06979 [cs].
- [12] Guangyin Jin, Yan Cui, Liang Zeng, Hanbo Tang, Yanghe Feng, and Jincui Huang. 2020. Urban ride-hailing demand prediction with multiple spatio-temporal information fusion network. *Transportation Research Part C: Emerging Technologies* 117 (Aug. 2020), 102665. <https://doi.org/10.1016/j.trc.2020.102665>
- [13] Guangyin Jin, Yuxuan Liang, Yuchen Fang, Zezhi Shao, Jincui Huang, Junbo Zhang, and Yu Zheng. 2023. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. *IEEE Transactions on Knowledge and Data Engineering* (2023), 1–20. <https://doi.org/10.1109/TKDE.2023.3333824>
- [14] Nathan Kallus and Angela Zhou. 2018. Residual unfairness in fair machine learning from prejudiced data. In *International Conference on Machine Learning*. PMLR, 2439–2448.
- [15] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual Fairness. In *Advances in Neural Information Processing Systems*, Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/hash/a486cd07e4ac3d270571622f4f316ec5-Abstract.html

- [16] Shiyong Lan, Yitong Ma, Weikang Huang, Wenwu Wang, Hongyu Yang, and Pyang Li. 2022. Dstagnn: Dynamic spatial-temporal aware graph neural network for traffic flow forecasting. In *International conference on machine learning*. PMLR, 11906–11917.
- [17] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings* (2018), 1–16.
- [18] Todd Alexander Litman. 2023. Evaluating Transportation Equity. (March 2023).
- [19] Xu Liu, Yutong Xia, Yuxuan Liang, Junfeng Hu, Yiwei Wang, Lei Bai, Chao Huang, Zhenguang Liu, Bryan Hooi, and Roger Zimmermann. 2023. LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. In *Advances in Neural Information Processing Systems*.
- [20] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2022. A Survey on Bias and Fairness in Machine Learning. <https://doi.org/10.48550/arXiv.1908.09635> arXiv:1908.09635 [cs].
- [21] Dana Pessach and Erez Shmueli. 2022. A Review on Fairness in Machine Learning. *Comput. Surveys* 55, 3 (Feb. 2022), 51:1–51:44. <https://doi.org/10.1145/3494672>
- [22] Bogdan Ilie Sighencea, Ion Rareș Stanciu, and Cătălin Daniel Căleanu. 2023. D-STGCN: Dynamic Pedestrian Trajectory Prediction Using Spatio-Temporal Graph Convolutional Networks. *Electronics* 12, 33 (Jan. 2023), 611. <https://doi.org/10.3390/electronics12030611>
- [23] Till Speicher, Hoda Heidari, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. A Unified Approach to Quantifying Algorithmic Unfairness: Measuring Individual & Group Unfairness via Inequality Indices. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, London United Kingdom, 2239–2248. <https://doi.org/10.1145/3219819.3220046>
- [24] Beibei Wang, Youfang Lin, Shengnan Guo, and Huaiyu Wan. 2021. GSNet: Learning Spatial-Temporal Correlations from Geographical and Semantic Aspects for Traffic Accident Risk Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 55 (May 2021), 4402–4409. <https://doi.org/10.1609/aaai.v35i5.16566>
- [25] Chenyu Wang, Zongyu Lin, Xiaochen Yang, Jiao Sun, Mingxuan Yue, and Cyrus Shahabi. 2022. HAGEN: Homophily-Aware Graph Convolutional Recurrent Network for Crime Forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 4 (June 2022), 4193–4200. <https://doi.org/10.1609/aaai.v36i4.20338>
- [26] Dongkun Wang, Jieyang Peng, Xiaoming Tao, and Yiping Duan. 2024. Boosting urban prediction tasks with domain-sharing knowledge via meta-learning. *Information Fusion* (2024), 102324.
- [27] Qingyi Wang, Shenhao Wang, Yunhan Zheng, Hongzhou Lin, Xiaohu Zhang, Jinhua Zhao, and Joan Walker. 2024. Deep hybrid model with satellite imagery: How to combine demand modeling and computer vision for travel behavior analysis? *Transportation Research Part B: Methodological* 179 (2024), 102869.
- [28] Yuandong Wang, Hongzhi Yin, Hongxu Chen, Tianyu Wo, Jie Xu, and Kai Zheng. 2019. Origin-Destination Matrix Prediction via Graph Convolution: a New Perspective of Passenger Demand Modeling. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19)*. Association for Computing Machinery, New York, NY, USA, 1227–1235. <https://doi.org/10.1145/3292500.3330877>
- [29] Yuandong Wang, Hongzhi Yin, Tong Chen, Chunyang Liu, Ben Wang, Tianyu Wo, and Jie Xu. 2021. Gallat: A Spatiotemporal Graph Attention Network for Passenger Demand Prediction. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. 2129–2134. <https://doi.org/10.1109/ICDE51399.2021.00212>
- [30] Zhaonan Wang, Renhe Jiang, Hao Xue, Flora D. Salim, Xuan Song, and Ryosuke Shibasaki. 2021. Event-Aware Multimodal Mobility Nowcasting. arXiv:2112.08443 (Dec. 2021). <https://doi.org/10.48550/arXiv.2112.08443> arXiv:2112.08443 [cs].
- [31] Yuankai Wu, Dingyi Zhuang, Aurelie Labbe, and Lijun Sun. 2021. Inductive Graph Neural Networks for Spatiotemporal Kriging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4478–4485.
- [32] An Yan and Bill Howe. 2020. Fairness-Aware Demand Prediction for New Mobility. *Proceedings of the AAAI Conference on Artificial Intelligence* 34, 01 (April 2020), 1079–1087. <https://doi.org/10.1609/aaai.v34i01.5458> Number: 01.
- [33] Jiexia Ye, Juanjuan Zhao, Kejiang Ye, and Chengzhong Xu. 2022. How to Build a Graph-Based Deep Learning Architecture in Traffic Domain: A Survey. *IEEE Transactions on Intelligent Transportation Systems* 23, 5 (May 2022), 3904–3924. <https://doi.org/10.1109/TITS.2020.3043250> arXiv:2005.11691 [cs, eess].
- [34] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-Temporal Graph Convolutional Networks : A Deep Learning Framework For Traffic Forecasting. *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)* (2018), 3634–3640. <https://aaafoundation.org/american-driving-survey-2014-2015/>
- [35] Le Yu, Bowen Du, Xiao Hu, Leilei Sun, Liangzhe Han, and Weifeng Lv. 2021. Deep spatio-temporal graph convolutional network for traffic accident prediction. *Neurocomputing* 423 (Jan. 2021), 135–147. <https://doi.org/10.1016/j.neucom.2020.09.043>
- [36] Xiaojian Zhang, Qian Ke, and Xilei Zhao. 2024. Travel Demand Forecasting: A Fair AI Approach. *IEEE Transactions on Intelligent Transportation Systems* 25, 10 (Oct. 2024), 14611–14627. <https://doi.org/10.1109/TITS.2024.3395061>
- [37] Yang Zhang and Tao Cheng. 2020. Graph deep learning model for network-based predictive hotspot mapping of sparse spatio-temporal events. *Computers, Environment and Urban Systems* 79 (Jan. 2020), 101403. <https://doi.org/10.1016/j.compenvurbsys.2019.101403>
- [38] Yunhan Zheng, Qingyi Wang, Dingyi Zhuang, Shenhao Wang, and Jinhua Zhao. 2023. Fairness-Enhancing Deep Learning for Ride-Hailing Demand Prediction. *IEEE Open Journal of Intelligent Transportation Systems* 4 (2023), 551–569. <https://doi.org/10.1109/OJITS.2023.3297517>
- [39] Yunhan Zheng, Shenhao Wang, and Jinhua Zhao. 2021. Equality of opportunity in travel behavior prediction with deep neural networks and discrete choice models. *Transportation Research Part C: Emerging Technologies* 132 (Nov. 2021), 103410. <https://doi.org/10.1016/j.trc.2021.103410>
- [40] Yanpeng Zhou, Jinjie Wang, Jianli Ding, Bohua Liu, Nan Weng, and Hongzhi Xiao. 2023. SIGNet: A Siamese Graph Convolutional Network for Multi-Class Urban Change Detection. *Remote Sensing* 15, 99 (Jan. 2023), 2464. <https://doi.org/10.3390/rs15092464>
- [41] Dingyi Zhuang, Qingyi Wang, Yunhan Zheng, Xiaotong Guo, Shenhao Wang, Haris N Koutsopoulos, and Jinhua Zhao. 2024. Advancing Transportation Mode Share Analysis with Built Environment: Deep Hybrid Models with Urban Road Network. *arXiv preprint arXiv:2405.14079* (2024).
- [42] Dingyi Zhuang, Shenhao Wang, Haris Koutsopoulos, and Jinhua Zhao. 2022. Uncertainty quantification of sparse travel demand prediction with spatial-temporal graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4639–4647.