

FairTP: A Prolonged Fairness Framework for Traffic Prediction

Jiangnan Xia^{1*}, Yu Yang^{2*}, Jiaxing Shen³, Senzhang Wang^{1†}, Jiannong Cao⁴

¹School of Computer Science and Engineering, Central South University

²Centre for Learning, Teaching, and Technology, The Education University of Hong Kong

³School of Data Science, Lingnan University

⁴The Department of Computing, The Hong Kong Polytechnic University

jiangnanx129@gmail.com, yangyy@eduhk.hk, jiaxingshen@LN.edu.hk, szwang@csu.edu.cn, jiannong.cao@polyu.edu.hk

Abstract

Traffic prediction is pivotal in intelligent transportation systems. Existing works focus mainly on improving overall accuracy, overlooking a crucial problem of whether prediction results will lead to biased decisions by transportation authorities. In practice, the uneven deployment of traffic sensors in different urban areas produces imbalanced data, making the traffic prediction model fail in some urban areas and leading to unfair regional decision-making that eventually severely affects equity and quality of residents' life. Existing fairness machine learning models struggle to maintain fair traffic prediction over prolonged periods. Although these models might achieve fairness at certain time slots, this static fairness will break down as traffic conditions change. To fill this research gap, we investigate prolonged fair traffic prediction, introducing two novel fairness metrics, i.e., region-based static fairness and sensor-based dynamic fairness, tailored to fairness fluctuations over time and across areas. An innovative prolonged fairness traffic prediction framework, namely FairTP, is then proposed. FairTP achieves prolonged fairness by alternating between "sacrifice" and "benefit" the prediction accuracy of each traffic sensor or area, ensuring that the number of these two actions are balanced over time. Specifically, FairTP incorporates a state identification module to discriminate whether the traffic sensors or areas are in a "sacrifice" or "benefit" state, thereby enabling prolonged fairness-aware traffic predictions. Additionally, we devise a state-guided balanced sampling strategy to select training examples to further enhance prediction fairness by mitigating the performance disparities among areas with uneven sensor distribution over time. Extensive experiments in two real-world datasets show that FairTP significantly improves prediction fairness without causing significant accuracy degradation.

Code — <https://github.com/jiangnanx129/FairTP>

Extended version —

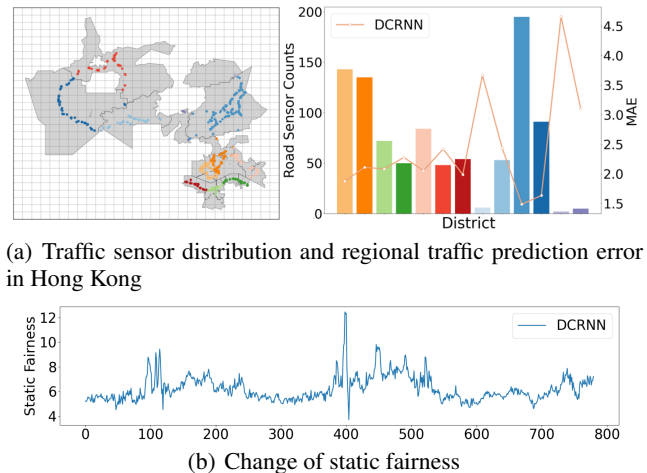
<https://github.com/jiangnanx129/FairTP>

Introduction

Traffic prediction is crucial to transportation planning, infrastructure management and optimizing resource allocation

*These authors contributed equally.

†Corresponding author



(a) Traffic sensor distribution and regional traffic prediction error in Hong Kong

(b) Change of static fairness

Figure 1: Illustration of fairness issues in traffic prediction.

and service provision (Miao et al. 2022, 2024, 2025; Xia et al. 2022; Wang et al. 2024). However, existing models merely focus on prediction accuracy while overlooking the key social impacts of traffic forecasting. In practice, achieving equitable (or fair) traffic prediction across city areas is essential, which enables transportation authorities to make unbiased decisions and significantly improve the life quality of urban residents.

The uneven distribution of traffic sensors installed in a city creates data volume disparities, making the traffic prediction model fail to produce precise predictions in some underprivileged regions with fewer sensors (Tedjopurnomo et al. 2020). Figure 1(a) visualizes the uneven sensor distribution in Hong Kong and the regional traffic prediction performance produced by DCRNN (Li et al. 2018), where the color gradients highlight different areas. The prediction error in regions with sparse sensors is notably higher versus the regions with dense sensor deployment. The forecast bias raises a fairness issue across city regions. Systematic underpredicting the traffic in certain areas can easily mislead transportation departments to provide insufficient services, thus forming an inequality treatment against privileged regions.

Existing fairness machine learning models fail to preserve

fair traffic prediction for a prolonged time (Dong et al. 2023; Wan et al. 2023). Although existing approaches can achieve fairness at certain time slots (Chai and Wang 2022; Li et al. 2024; Guo et al. 2023; Yang et al. 2023; Caton and Haas 2024), such static fairness may be broken as the traffic conditions change over time. Figure 1(b) presents the fluctuating static fairness, clearly reflecting its temporal dynamics. This metric is calculated based on a specific definition of group fairness (Yan and Howe 2020). Moreover, existing methods achieve static fairness but significantly reduce the accuracy in privileged areas. Enhancing fairness while achieving high prediction accuracy remains an open research problem.

To fill this research gap, this paper investigates prolonged algorithmic fairness in traffic prediction, a challenging task due to several factors. First, there lacks of clear definitions of fairness to measure and quantify equity in dynamic traffic environments. Second, the scarcity of data in underprivileged areas poses a significant challenge for improving traffic prediction accuracy while maintaining fairness, as efforts to enhance performance in these areas can inadvertently harm the performance of privileged regions. Third, achieving both short-dated static and prolonged dynamic fairness remains difficult.

To tackle these challenges, this paper introduces two pioneering fairness definitions for dynamic traffic environments and proposes a prolonged fairness traffic prediction framework, namely FairTP. We argue that fairness in traffic prediction changes over time and across urban areas. The prediction accuracy of each traffic sensor or area should alternately “sacrifice” and “benefit” to pursue fairness. Performance sacrifice means that the model intentionally reduces the prediction accuracy of nodes that could have been made accurately, while the performance benefit is the opposite. Over time, prolonged fairness is achieved when these fluctuations are similar within a given period. Accordingly, we first define region-based static fairness (RSF) and sensor-based dynamic fairness (SDF) to measure performance disparities across regions at each time slot and state disparities among road sensors over a period, respectively. Next, FairTP is proposed that consists of the following two key modules. The state identification module discriminates “sacrifice” and “benefit” states, enabling prolonged fairness-aware predictions, SDF calculation, and guiding the sampling module. Then, the state-guided balanced sampling module is designed to adjust training examples by increasing the sampling frequency for sensors in a “sacrifice” state. It improves prediction accuracy in underprivileged areas and reduces performance disparities between regions. Lastly, both RSF and SDF are integrated into FairTP to achieve predictive fairness at shortdated static and prolonged dynamic levels. FairTP can plug into existing traffic prediction models and obtain equitable results.

To summarize, our main contributions are as follows.

- We systematically explore prolonged fairness in traffic prediction and propose two novel fairness definitions RSF and SDF for dynamic traffic scenarios.
- We propose the novel FairTP framework that can seamlessly integrate with existing traffic prediction models to

enhance fairness with minimal accuracy degradation.

- Extensive experiments on two real-world traffic datasets demonstrate that FairTP maintains high predictive accuracy while achieving both shortdated static and prolonged dynamic fairness.

Related Work

Traffic prediction. It has garnered significant research attention, attributable to the availability of urban data and its wide range of applications (Choi et al. 2022; Xia et al. 2024). Deep neural networks, notably Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have gained popularity in traffic prediction due to their superior learning capabilities (Yao et al. 2019). However, these models are designed for spatio-temporal grid data and are not suitable for graph-based data, which is prevalent in road networks. Recently, there has been a rising research interest in leveraging Graph Neural Networks (GNNs) for spatio-temporal data prediction (Zheng et al. 2023b; Yang et al. 2021; Ye et al. 2022). Existing models have combined GNN with RNN, Temporal Convolutional Networks (TCN), or attention mechanisms to capture the complex spatial and temporal dependencies in traffic data. These models, like DCRNN (Li et al. 2018), DGCRN (Li et al. 2023), and DSTAGNN (Lan et al. 2022), have made advancements in capturing the dynamics of road networks.

Different from existing models, this paper pioneers a systematic study on fair traffic prediction, introduces two novel fairness definitions suitable for dynamic traffic scenarios, and develops FairTP that can seamlessly integrate with existing traffic prediction models and make their predictions fair with very slight accuracy degradation.

Algorithmic Fairness. A considerable volume of research in machine learning highlights that models can exhibit discriminatory behavior towards certain groups across different domains (Boratto et al. 2023; Mahapatra, Dong, and Momma 2023; Hua et al. 2023). Previous fairness research primarily focused on identifying and mitigating biases towards specific sensitive groups, such as race, gender, incomes (Ghani et al. 2023; Mehrabi et al. 2021). Various fairness metrics, including group and individual fairness, have been proposed (Dong et al. 2023).

Fairness research in transportation is in its early stages, with recent efforts exploring fairness in mobility demand prediction (Du et al. 2024). For example, Yan et al. proposed FairST, introducing two fairness metrics to promote equity across demographic groups (Yan and Howe 2020). However, their approach relies on sensitive features like race or gender and applies regularization at time slot (static). Similarly, Zheng et al. developed SA-Net that uses socially-aware convolution operations to integrate socio-demographic and ridership data for fair demand prediction (Zheng et al. 2023a). This method also depends on external data and static fairness regularization. Dynamic fairness is a form of prolonged fairness. It has been explored in decision-making but lacks a clear definition. Song et al. defined a dynamic fairness based on general static individual fairness (Song, Ma, and King 2022). However, their method relies on an oracle similarity

matrix created with domain knowledge or human judgment (Song et al. 2022), making it unsuitable for traffic scenarios.

Significantly different from prior works that focus on static fairness at specific time slots, we propose RSF and SDF, two novel fairness measures for dynamic traffic. They are free from sensitive attributes and can enhance predictive fairness at shortdated static and prolonged dynamic levels.

Problem Statement

Before defining the problem of prolonged fair traffic prediction, we first introduce terminologies and notations.

$G = \{V, E\}$ is a road network where the node V represents on-road sensors, and the link E indicates the connections of sensors along the road. $(X^{t-T+1}, X^{t-T+2}, \dots, X^t)$ is historical traffic observations over T time steps, where $X^t = \{x_{v_i}^t | v_i \in V\}$ and $x_{v_i}^t$ is the observation of sensor v_i at time t . We divide a city into m regions $Re = (r_1, r_2, \dots, r_m)$ and the regional traffic condition is denoted as $X_{Re}^t = \{x_{r_p}^t | r_p \in Re\}$ where $x_{r_p}^t$ is the average traffic observations of the nodes in the region r_p at time t .

Given a road network G and the historical traffic observations $(X^{t-T+1}, X^{t-T+2}, \dots, X^t)$, prolonged fair traffic prediction aims to accurately estimate the regional traffic conditions $(X_{Re}^{t+1}, X_{Re}^{t+2}, \dots, X_{Re}^{t+T})$ in the future T time steps while minimizing the performance disparity across regions at each time slot and the number of times each node's prediction performance is "sacrificed" or "benefited" to achieve former during T . Performance sacrifice means that the model intentionally reduces the prediction accuracy of nodes that could have been made accurately, while the performance benefit is the opposite.

Methodology

This section introduces RSF and SDF, two metrics designed to evaluate shortdated static and prolonged dynamic fairness in traffic prediction scenarios. Additionally, we present the FairTP framework.

Region-based Static Fairness

Uneven road sensor placement leads to imbalanced traffic prediction accuracy, creating fairness issues. Existing fairness models depend on sensitive features and focus on fairness at specific time slots, neglecting the dynamic nature of traffic. To address this, we propose RSF, a fairness metric that avoids sensitive features and evaluates shortdated static fairness in dynamic traffic scenarios. Unlike group equity approaches relying on sensitive attributes (e.g. race) (Dong et al. 2023), RSF measures disparities in predictive performance between regions without requiring additional sensitive data.

Areas with fewer sensors experience higher prediction errors due to limited data, resulting in unfair regional decisions that affect residents' equity. RSF reduces these disparities by minimizing performance gaps across regions.

RSF. $RSF(r_p, r_q)$ between region r_p and r_q at time slot t is defined as:

$$RSF(r_p, r_q) = |\mathcal{M}[\hat{y}_{r_p}^t] - \mathcal{M}[\hat{y}_{r_q}^t]|, \quad (1)$$

where $\hat{y}_{r_p}^t$ represents the predicted region traffic condition for region r_p at time t and $\mathcal{M}[\hat{y}_{r_p}^t]$ is the mean absolute percentage error for region r_p at time t . A smaller RSF value indicates higher fairness.

RSF Loss. To enhance fairness, we define the RSF loss at time t :

$$L_{RSF} = \frac{2}{m(m-1)} \sum_{r_p, r_q \in Re} |\mathcal{M}[\hat{y}_{r_p}^t] - \mathcal{M}[\hat{y}_{r_q}^t]|, \quad (2)$$

where m is the number of regions. The average of the difference in predictive performance between all regional pairings is measured by L_{RSF} . It is calculated at each time slot. To maximize shortdated static fairness in predictions, we aim to minimize L_{RSF} during training.

Sensor-based Dynamic Fairness

While RSF addresses immediate prediction disparities, it does not ensure prolonged fairness over time. Directly optimizing RSF may degrade accuracy in well-instrumented regions without significantly benefiting underprivileged areas. To address this limitation, we propose SDF, a sensor-based metric for dynamic traffic scenarios that evaluates fairness over a duration T_d . T_d is the time length of a batch of training samples used to adjust the sampling frequency for sensor-based dynamic fairness.

SDF ensures prolonged fairness by assessing discrepancies in sensor states over time. Fairness evolves dynamically, requiring sensors to alternate between "sacrifice" and "benefit" states based on their prediction accuracy. These states are determined by a state identification module. The prolonged fairness is achieved when the overall states of road sensors are similar within a defined period T_d .

SDF. For road sensors v_i and v_j , the $SDF(v_i, v_j)$ between them is defined as follows

$$SDF(v_i, v_j) = |\mathcal{D}_{T_d}[v_i] - \mathcal{D}_{T_d}[v_j]|, \quad (3)$$

where $d_{v_i}^{t_k}$ is the state of the road sensor v_i at time slot t_k . Its value is given by the state identification module. And $\mathcal{D}_{T_d}[v_i] = d_{v_i}^{t_1} + d_{v_i}^{t_2} + \dots + d_{v_i}^{t_{T_d}}$ represents the overall state of road v_i over a period of time T_d . SDF calculates the sum of state differences between all pairs of road sensors over a period. A smaller SDF value indicates higher fairness, meaning sensors are treated more consistently in T_d .

SDF Loss. We define the SDF loss during a period T_d as:

$$L_{SDF} = \frac{2}{n(n-1)} \sum_{v_i, v_j \in V} |\mathcal{D}_{T_d}[v_i] - \mathcal{D}_{T_d}[v_j]|, \quad (4)$$

where n is the number of sensors used in T_d . V is the set of road sensors that have been sampled in T_d . The average of the overall state difference between all sensor pairings in the training data is measured by L_{SDF} . It is calculated every T_d batches. To maximize prolonged dynamic fairness in predictions, we aim to minimize L_{SDF} during training.

Prolonged Fairness Traffic Prediction Framework

In this section, we introduce FairTP, a framework for achieving prolonged fair traffic prediction as shown in Figure 2.

FairTP consists of three parts: (a) a state-guided balanced sampling module for sensor selection, (b) a ST dependencies learning module for learning spatio-temporal correlations, and (c) a state identification module for sensor status assessment.

State-guided balanced sampling module. Spatial imbalances in traffic data often lead to unfair predictions and biased decisions by transportation authorities. Effective sampling of real-world traffic data is essential for building fair machine learning models. It helps counteract biases from uneven sensor deployment and reduces spatio-temporal redundancies. Moreover, it can improve models' generalization, enhance training efficiency, and balance performance between well-instrumented and under-instrumented areas. Therefore, we propose a state-guided sampling strategy to provide balanced training data and reduce performance disparities caused by uneven sensor distribution. The strategy enhances predictions in underprivileged areas by sampling N_{sam} sensors from the road network. N_{sam} indicates the number of road sensors to be sampled for training. The sampling scheme is periodically adjusted based on sensor states, giving more training opportunities to those sensors identified as "sacrifice" (elaborated in the state identification module). This helps to improve prediction performance in these underprivileged areas.

Specifically, based on the sampled number N_{sam} , we start with stratified sampling to initialize the sampling. Then, after each time interval T_d , we calculate the sampling probabilities for all the sensors according to the results of the training feedback, and perform the next round of sampling. A greedy algorithm is employed to adjust the sampled data every T_d batches, prioritizing sensors with lower sampling probabilities until N_{sam} is reached. Notably, each sampling round is influenced by the results of the previous round. Next, we describe this process in detail.

To begin the sample data collection, we use stratified sampling. This method calculates the number of road sensors in different city regions and selects a proportionate number of sensors from each region based on N_{sam} . It effectively reduces sampling errors and improves the representativeness and reliability of the sample. The initial sampled sensors are denoted as $Sam_0(v)$.

Second, we train the model using $Sam_0(v)$ to determine the states d of the sampled sensors via the state identification module. Over a period T_d , these states are aggregated into an overall state D . And the unsampled sensors are assigned $D=0$. Then, new sampling probabilities for all sensors in V are calculated based on D . It guides the next sampling round. The process for sensor $v_i \in V$ is shown as follows

$$\begin{aligned} d_{v_i}^{t_k} &= \begin{cases} d_{v_i}^{t_k} - 0.5, & v_i \in \text{Sam}_l(v), \\ 0, & v_i \in V \text{ and } v_i \notin \text{Sam}_l(v), \end{cases} \\ \mathcal{D}_{T_d}[v_i] &= d_{v_i}^{t_1} + \dots + d_{v_i}^{t_k} + \dots + d_{v_i}^{T_d}, \\ P_{v_i}^l &= \text{sigmoid}(\mathcal{D}_{T_d}[v_i]), v_i \in V, \end{aligned} \quad (5)$$

where $d_{v_i}^{t_k}$ is the state of sensor v_i at time t_k , determined by the state identification module. $\text{Sam}_l(v)$ represents the sensors sampled in the l -th round, and $\mathcal{D}_{T_d}[v_i]$ accumulates the

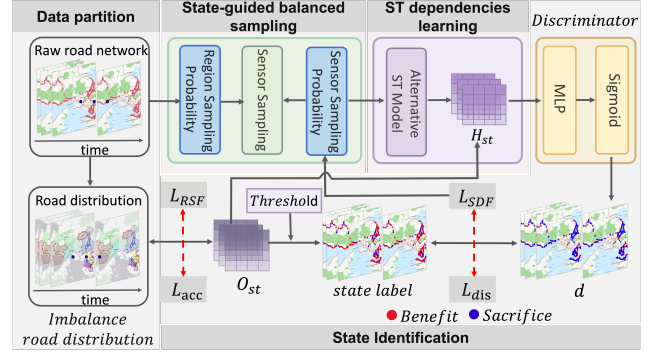


Figure 2: The framework of the proposed FairTP.

overall state of sensor v_i over the period T_d . The sampling probability P_{v_i} for sensor v_i is calculated using a sigmoid function and is used to guide the sampling strategy for the $(l+1)$ -th round.

Third, to ensure balanced sampling among different regions in the $(l+1)$ -th round, we calculate the sampling probabilities for each region. This encourages an equal number of sensors to be selected from each region during the process. The formula is shown as follows

$$\begin{aligned} (P_{r_1}^l, P_{r_2}^l, \dots, P_{r_m}^l) \\ = \text{softmax}(C_{r_1} - C_a, C_{r_2} - C_a, \dots, C_{r_m} - C_a), \end{aligned} \quad (6)$$

where $\{C_{r_p} | p = 1, 2, \dots, m\}$ denotes the number of sensors sampled in region r_p . $C_a = N_{sam}/m$ is the target number of sensors to be sampled in each region under balanced sampling. P_{r_p} represents the sampling probability for region r_p . A smaller P_{r_p} indicates that the number of sampling nodes in r_p is farther away from the balance.

Fourth, we combine the region sampling probabilities and the sensor sampling probabilities to update the sampling probabilities for all sensors in the road network. The process is shown as follows

$$P_{r_p}^l * P_{v_i}^l \xrightarrow{\text{Partition}} P_{v_i}^{l+1}, r_p \in Re, v_i \in V, \quad (7)$$

where Partition refers to the zoning of the city, and $P_{v_i}^{l+1}$, $v_i \in V$ represents the updated sampling probability for sensor v_i in the $(l+1)$ -th round.

Finally, based on $P_{v_i}^{l+1}$, $v_i \in V$, we use a greedy algorithm to select sensors with the lowest sampling probabilities as training samples. Steps 3 and 4 are repeated, updating sampling probabilities and selecting sensors iteratively until the desired number N_{sam} is reached. The new sample data $Sam_{l+1}(v)$ is then used for the next training period T_d . This approach ensures a representative and balanced sample across different road network regions.

ST dependencies learning module. After data sampling, we utilize a ST model to capture spatial relationships and temporal trends within the sampled data for accurate traffic forecasting. The process can be represented as follows

$$O_{st}, H_{st} = ST(\text{Sam}(v)), \quad (8)$$

where $Sam(v)$ denotes the sampled data. O_{st} represents the predicted traffic results, and H_{st} is the hidden representation that contains the spatio-temporal dependencies learned by the traffic model. In this paper, the ST model is replaceable, as FairTP can be extended to any traffic prediction model.

State identification module. We design a state identification module that consists of a state marker and a discriminator for real-time sensor state evaluation. The discriminator is trained to infer sensor states with no ground truth during testing.

To begin with, we manually assign state labels to the sensors based on their performances. A state label of 1 indicates a “benefit” state, and a label of 0 indicates a “sacrifice” state. We use the output O_{st} of the ST model to calculate the MAPE for each sensor. This MAPE is then compared to a predefined threshold to determine the state of the sensor. And the threshold is obtained from the previous training round. Specifically, we argue that each sensor’s prediction accuracy varies in every training round. If a sensor’s MAPE is lower than the threshold, we label it as 1 (state “benefit”) for that round, indicating an improvement in performance. If the MAPE exceeds the threshold, we mark it as 0 (state “sacrifice”), indicating a drop in performance. Notably, the threshold is determined based on the selected ST model. We train the original ST model and record the MAPE for each round. This MAPE value is then used as the threshold when embedding the ST model into FairTP. A detailed example is provided in the extended version.

Next, we introduce a discriminator to classify sensor states at each time slot. The discriminator takes the hidden representation H_{st} as input and outputs a state prediction $d \in (0, 1)$. We denote the discriminator as $Dis_{\theta_{dis}} : H_{st} \rightarrow d$, where θ_{dis} represents the model parameters. The discrimination loss is computed as follows

$$L_{dis} = -\left(d \log(Y) + (1 - d) \log(1 - Y)\right), \quad (9)$$

where d represents the sensor states predicted by the discriminator $Dis_{\theta_{dis}}$. And Y denotes the corresponding state label. By minimizing the discrimination loss L_{dis} during training, the discriminator can accurately identify the sensor states during prediction. This, in turn, guides FairTP to perform appropriate sampling, leading to fair prediction results.

Overall Objective Function. FairTP is trained end-to-end by minimizing a composite loss function that combines accuracy and fairness objectives

$$L = L_{acc} + \lambda_1 L_{RSF} + \lambda_2 L_{SDF}, \quad (10)$$

where L_{acc} represents the mean absolute error (MAE). L_{RIF} and L_{SDF} correspond to formula (2) and (4), respectively. The latter is periodically included every T_d batches. Both L and L_{dis} are co-train to ensure that sampling is continuously adjusted in a fair direction, ultimately yielding fair traffic predictions.

Experiments

Experiment Setup

Dataset. We use two real-world datasets for regional traffic prediction: the HK and the SD datasets. The HK dataset contains six months of taxi trajectory data with 938 road sensors. The SD dataset includes data from 716 road sensors, sourced from the PeMS platform in 2019. Details are provided in the extended version.

Baseline. We select several types of representative traffic prediction models as underlying ST models, including **DCRNN**(Li et al. 2018), **AGCRN**(Bai et al. 2020), **GWNET**(Wu et al. 2019), **ASTGCN**(Guo et al. 2019), **DSTAGNN**(Lan et al. 2022), **DGCRN**(Li et al. 2023) and **D2STGNN**(Shao et al. 2022). These models help demonstrate the effectiveness and scalability of the proposed FairTP. Additionally, we compare with fairness mitigation baselines **FairST**(Yan and Howe 2020) and **SA-Net**(Zheng et al. 2023a) are compared as fairness mitigation baselines. Details of these models are provided in the extended version.

Implementation Details. We set the sampled number N_{sam} to 200 for both the SD and HK datasets. The dynamic time length T_d is fixed at 3, representing 3 batches. We set the hyperparameters λ_1 and λ_2 to 0.01 and 0.1, respectively. The proposed FairTP can be extended to various traffic prediction models. All models are implemented on the GeForce RTX 3090. Full implementation details are provided in the extended version.

Comparison With Baselines

First, we extend the proposed FairTP framework to multiple baselines, denoted as *FairTP-baseline*. The results are shown in Table.1, with the best performance highlighted in bold. Notably, the calculation of SDF relies on the sensor state predicted by FairTP’s state identification module. Since it is absent in the baselines, they cannot produce SDF outputs.

The newly proposed FairTP (*FairTP-baseline*) demonstrates superior performance in most scenarios. It outperforms the corresponding baselines in terms of both fairness and accuracy. On the HK dataset, the baselines show a reduction in MAE by 0.93% to 13.76%, and an improvement in RSF by 20.94% to 128.69%. While the MAE slightly gains (2.46% to 5.24%) over DGCRN, ASTGCN, and D2STGNN, the fairness improvements are substantial. On the SD dataset, *FairTP-baseline* significantly enhances the MAE performance of each baseline, with improvements ranging from 12.38% to 55.12%. RSF performance is enhanced by 15.89% to 90.58%. Notably, all *FairTP-baseline* models incorporate SDF, ensuring prolonged fairness. These results demonstrate FairTP’s adaptability to various traffic prediction models, achieving strong overall performance while significantly improving short-dated static and prolonged dynamic fairness. This improvement can be attributed to the effectiveness of data sampling, which boosts the predictive performance of underprivileged regions without significantly impacting privileged regions. Further details on the trade-off between accuracy and fairness are provided in the extended version.

	HK					SD				
	MAE	RMSE	MAPE	RSF	SDF	MAE	RMSE	MAPE	RSF	SDF
DCRNN	2.353	3.694	0.048	1.779	-	28.437	44.359	0.131	2.427	-
FairTP-DCRNN	2.411	3.734	0.050	1.471	0.085	21.709	32.893	0.113	1.668	1.299
AGCRN	1.957	2.974	0.040	1.613	-	19.215	29.067	0.078	1.337	-
FairTP-AGCRN	1.939	3.243	0.039	0.907	0.057	14.798	20.782	0.082	1.002	1.602
GWNET	2.189	3.323	0.046	1.688	-	21.158	31.117	0.096	1.206	-
FairTP-GWNET	2.132	3.448	0.042	0.835	1.999	18.828	27.836	0.094	0.945	6.533
ASTGCN	2.005	3.065	0.042	1.632	-	23.689	35.812	0.103	1.320	-
FairTP-ASTGCN	2.110	3.264	0.043	0.945	0.067	17.965	25.269	0.103	1.048	0.614
DSTAGNN	2.373	3.662	0.049	1.697	-	22.005	31.298	0.113	1.113	-
FairTP-DSTAGNN	2.086	3.351	0.042	0.973	2.612	18.209	25.708	0.104	0.584	0.907
DGCRN	2.257	3.271	0.047	1.790	-	33.346	49.799	0.150	1.758	-
FairTP-DGCRN	2.399	3.606	0.049	1.241	3.795	21.497	31.182	0.109	1.517	4.944
D2STGNN	1.992	2.895	0.042	1.873	-	18.012	25.941	0.072	1.242	-
FairTP-D2STGNN	1.928	3.086	0.039	0.819	0.739	13.725	19.083	0.071	0.736	2.292

Table 1: Performance comparison of FairTP and traffic prediction model.




Figure 3: Regional performance on two datasets.

	HK			SD		
	MAE	RSF	SDF	MAE	RSF	SDF
FairST-AGCRN	1.86	1.87	-	19.33	1.19	-
SA-Net-AGCRN	1.88	1.65	-	18.91	1.21	-
FairTP-AGCRN	1.94	0.91	0.06	14.80	1.00	1.60
FairST-D2STGNN	2.06	1.86	-	19.80	1.26	-
SA-Net-D2STGNN	2.02	1.69	-	19.26	1.13	-
FairTP-D2STGNN	1.93	0.82	0.74	13.72	0.74	2.29

Table 2: Performance comparison of FairTP and fairness mitigation methods.

Next, we compare FairTP with existing fairness mitigation methods. Due to the strict 7-page limit, we present results based on two representative baselines, AGCRN and D2STGNN, both of which perform well.

As shown in Table.2, FairTP consistently achieves the highest fairness across all cases. On the HK dataset, FairTP improves RSF by 81.3% to 126.8%, with MAE reductions ranging from 4.6% to 6.7%. On the SD dataset, FairTP achieves 19.0% to 71.9% RSF improvements and enhances MAE performance by 27.8% to 44.3%. FairTP achieves 19.0% to 70.3% RSF improvements and enhances MAE performance by 27.8% to 44.3%. In summary, compared to FairST, which directly constrains predicted values, and SA-Net, which constrains MAPE values, our proposed FairTP focuses on minimizing regional MAPE differences to bal-

ance performance across areas for shortdated static fairness. Additionally, SDF is introduced to achieve prolonged dynamic fairness for road sensors.

Regional Performance Analysis

In this section, we provide the regional visualizations to verify the specific performance of *FairTP-baseline*.

The regional prediction performance are shown in Figure 3. In the HK dataset, *FairTP-baseline* significantly improves performance in underprivileged regions such as r_8 , r_{12} and r_{13} , which have fewer road sensors. At the same time, the performances in the privileged regions, such as r_1 , r_2 and r_{10} , with more sensors, remains largely unaffected. Similarly, in the SD dataset, *FairTP-baseline* shows noticeable improvements in underprivileged regions like r_6 , r_7 , r_{11} , and r_{12} . The performances in privileged regions like r_1 and r_2 do not observably deterioration. We also conducted similar experiments with six other baseline models, and the results are consistent. *FairTP-baseline* improves the predicted performance of the underprivileged regions on both data, while the performance of the privileged regions does not decline significantly. These results demonstrate that *FairTP* indeed improves the prediction performance of underprivileged regions, leading to the better overall city-wide performance and enhancing fairness by reducing performance disparities.

		HK			SD		
		MAE	RSF	SDF	MAE	RSF	SDF
FairTP-AGCRN	noS	2.02	1.01	0.18	14.81	1.12	2.35
	noD	1.99	0.88	188.20	15.50	1.05	344.12
	noAS	1.78	1.07	0.65	12.40	1.23	2.18
	ALL	1.94	0.91	0.06	14.80	1.00	1.60
FairTP-G2STGNN	noS	1.94	0.95	0.83	13.80	0.80	2.32
	noD	2.14	0.83	181.89	14.00	0.74	83.30
	noAS	1.73	1.04	1.38	11.42	0.98	2.47
	ALL	1.93	0.82	0.74	13.72	0.74	2.29

Table 3: Performance of ablation study.




Figure 4: Effect of dynamic time length T_d .




Figure 5: Effect of sampled number N_{sam} .

Ablation Study

We further conduct an ablation study to evaluate the contribution of each component in FairTP to the performance gain. We deactivate different components and form the following variants. **noS** removes the L_{RSF} , and it does not consider static fairness constraints. **noD** removes the L_{SDF} , and it does not consider constraints on fairness at the prolonged dynamic level. **noSA** removes the state-guided sampling module but uses the fixed stratified sampling.

The results of the ablation study are shown in Table.3. One can see that the L_{RSF} , L_{SDF} and the sampling module are all useful for FairTP as removing any one of them increases the prediction error or decreases the RSF or DSF. Among these, the L_{SDF} appears to have the most significant impact. When it is removed, the SDF decreases remarkably, demonstrating the importance of the prolonged fairness. Removing the L_{RSF} results in a decline in both RSF and SDF, highlighting the importance of static fairness constraints. When using fixed stratified sampling, the performance of the model improves. It may contribute to the

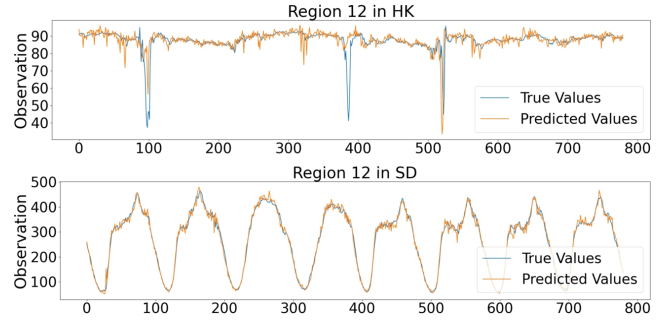


Figure 6: Case study.

adaptive matrix in the baseline that is able to learn more ST information based on the fixed road sensors. The improvement in RSF is slightly noticeable, likely because the predictive performance of most privileged regions is already quite close, but the SDF performance significantly decreases.

Parameter Analysis

We investigate the effects of dynamic time length T_d and sample size N_{sam} on the performance of FairTP-AGCRN. The results are shown in Figure 4 and Figure 5. For T_d , we vary it from 2 to 5 and find that the best performance, in terms of RSF and SDF, occurred at $T_d = 3$. Performance declined when T_d increased beyond 3, likely due to the greater fluctuations in sensor states. Similarly, we tune N_{sam} from 100 to 300. A small number of sensors result in insufficient data representation, while excessive sensors introduce noise. The optimal number of sensors for balanced performance is $N_{sam} = 200$.

Case Study

Comparative visualizations in Figure 6 illustrate the effectiveness of FairTP by comparing its predicted traffic flow to the ground truth. The case study is conducted on underprivileged regions (r_{12}) with minimal sensor coverage in both HK and SD, highlights the ability of FairTP-AGCRN to accurately capture the real traffic patterns, including sharp fluctuations. The close alignment between the orange predicted curves and the blue ground truth curves demonstrates the model's high predictive accuracy and its success in improving fairness at both shortdated static and prolonged dynamic levels.

Conclusion

We investigate prolonged fair traffic prediction for the first time. Two novel fairness definitions, i.e., region-based static fairness and sensor-based dynamic fairness, are proposed to characterize the fairness fluctuations over time and across areas. Moreover, we innovatively devise FairTP, a prolonged fairness traffic prediction framework. This framework can be easily integrated into existing traffic prediction models to enhance their short-term and long-term prediction fairness across urban areas without significantly sacrificing prediction performance.

Acknowledgments

This research was funded by the National Science Foundation of China (No.62172443), Hunan Provincial Natural Science Foundation of China (No.2022JJ30053), the Hong Kong Research Grants Council (RGC) under the Theme-based Research Scheme with grant No. T43-513/23-N and T41-603/20-R, Lingnan University (LU) (DB23A4) and Lam Woo Research Fund at LU (871236). This work is partially conducted at the Research Institute for Artificial Intelligence of Things (RIAIoT) at PolyU.

References

- Bai, L.; Yao, L.; Li, C.; Wang, X.; and Wang, C. 2020. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems*, 33: 17804–17815.
- Boratto, L.; Fabbri, F.; Fenu, G.; Marras, M.; and Medda, G. 2023. Counterfactual graph augmentation for consumer unfairness mitigation in recommender systems. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3753–3757.
- Caton, S.; and Haas, C. 2024. Fairness in machine learning: A survey. *ACM Computing Surveys*, 56(7): 1–38.
- Chai, J.; and Wang, X. 2022. Fairness with adaptive weights. In *International Conference on Machine Learning*, 2853–2866. PMLR.
- Choi, J.; Choi, H.; Hwang, J.; and Park, N. 2022. Graph neural controlled differential equations for traffic forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 6367–6374.
- Dong, Y.; Ma, J.; Wang, S.; Chen, C.; and Li, J. 2023. Fairness in graph mining: A survey. *IEEE Transactions on Knowledge and Data Engineering*.
- Du, X.; Li, Z.; Long, C.; Xing, Y.; Philip, S. Y.; and Chen, H. 2024. FELight: Fairness-Aware Traffic Signal Control Via Sample-Efficient Reinforcement Learning. *IEEE Transactions on Knowledge and Data Engineering*.
- Ghani, R.; Rodolfa, K. T.; Saleiro, P.; and Jesus, S. 2023. Addressing bias and fairness in machine learning: A practical guide and hands-on tutorial. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5779–5780.
- Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 922–929.
- Guo, Z.; Li, J.; Xiao, T.; Ma, Y.; and Wang, S. 2023. Towards fair graph neural networks via graph counterfactual. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 669–678.
- Hua, W.; Li, L.; Xu, S.; Chen, L.; and Zhang, Y. 2023. Tutorial on large language models for recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, 1281–1283.
- Lan, S.; Ma, Y.; Huang, W.; Wang, W.; Yang, H.; and Li, P. 2022. Dstagnn: Dynamic spatial-temporal aware graph neural network for traffic flow forecasting. In *International conference on machine learning*, 11906–11917. PMLR.
- Li, F.; Feng, J.; Yan, H.; Jin, G.; Yang, F.; Sun, F.; Jin, D.; and Li, Y. 2023. Dynamic graph convolutional recurrent network for traffic prediction: Benchmark and solution. *ACM Transactions on Knowledge Discovery from Data*, 17(1): 1–21.
- Li, Y.; Yang, Y.; Cao, J.; Liu, S.; Tang, H.; and Xu, G. 2024. Toward Structure Fairness in Dynamic Graph Embedding: A Trend-aware Dual Debiasing Approach. In *SIGKDD*.
- Li, Y.; Yu, R.; Shahabi, C.; and Liu, Y. 2018. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *International Conference on Learning Representations*.
- Mahapatra, D.; Dong, C.; and Momma, M. 2023. Querywise fair learning to rank through multi-objective optimization. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1653–1664.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2021. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)*, 54(6): 1–35.
- Miao, H.; Liu, Z.; Zhao, Y.; Zheng, K.; Zhang, Y.; and Jensen, C. S. 2025. LightTR: A Lightweight Framework for Federated Trajectory Recovery. In *ICDE*.
- Miao, H.; Shen, J.; Cao, J.; Xia, J.; and Wang, S. 2022. MBA-STNet: Bayes-enhanced discriminative multi-task learning for flow prediction. *IEEE Transactions on Knowledge and Data Engineering*.
- Miao, H.; Zhao, Y.; Guo, C.; Yang, B.; Kai, Z.; Huang, F.; Xie, J.; and Jensen, C. S. 2024. A unified replay-based continuous learning framework for spatio-temporal prediction on streaming data. In *ICDE*.
- Shao, Z.; Zhang, Z.; Wei, W.; Wang, F.; Xu, Y.; Cao, X.; and Jensen, C. S. 2022. Decoupled dynamic spatial-temporal graph neural network for traffic forecasting. *arXiv preprint arXiv:2206.09112*.
- Song, W.; Dong, Y.; Liu, N.; and Li, J. 2022. Guide: Group equality informed individual fairness in graph neural networks. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1625–1634.
- Song, Z.; Ma, Y.; and King, I. 2022. Individual Fairness in Dynamic Financial Networks. In *NeurIPS 2022 Workshop: New Frontiers in Graph Learning*.
- Tedjopurnomo, D. A.; Bao, Z.; Zheng, B.; Choudhury, F. M.; and Qin, A. K. 2020. A survey on modern deep neural network for traffic prediction: Trends, methods and challenges. *IEEE Transactions on Knowledge and Data Engineering*, 34(4): 1544–1561.
- Wan, M.; Zha, D.; Liu, N.; and Zou, N. 2023. In-processing modeling techniques for machine learning fairness: A survey. *ACM Transactions on Knowledge Discovery from Data*, 17(3): 1–27.
- Wang, H.; Xia, J.; Yang, Y.; Wang, S.; and Cao, J. 2024. STS2ANet: Spatio-Temporal Synchronized Sliding Attention Network for Accurate Cross-Day Origin-Destination

Prediction. In *International Conference on Database Systems for Advanced Applications*, 186–202. Springer.

Wu, Z.; Pan, S.; Long, G.; Jiang, J.; and Zhang, C. 2019. Graph wavenet for deep spatial-temporal graph modeling. *arXiv preprint arXiv:1906.00121*.

Xia, J.; Wang, S.; Wang, X.; Xia, M.; Xie, K.; and Cao, J. 2022. Multi-view Bayesian spatio-temporal graph neural networks for reliable traffic flow prediction. *International Journal of Machine Learning and Cybernetics*, 1–14.

Xia, J.; Wang, S.; Wang, X.; Xia, M.; Xie, K.; and Cao, J. 2024. Multi-view Bayesian spatio-temporal graph neural networks for reliable traffic flow prediction. *International Journal of Machine Learning and Cybernetics*, 15(1): 65–78.

Yan, A.; and Howe, B. 2020. Fairness-aware demand prediction for new mobility. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1079–1087.

Yang, T.; Xu, Z.; Wang, Z.; and Ai, Q. 2023. FARA: Future-aware Ranking Algorithm for Fairness Optimization. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2906–2916.

Yang, Y.; Cao, J.; Stojmenovic, M.; Wang, S.; Cheng, Y.; Lum, C.; and Li, Z. 2021. Time-capturing dynamic graph embedding for temporal linkage evolution. *IEEE Transactions on Knowledge and Data Engineering*, 35(1): 958–971.

Yao, H.; Tang, X.; Wei, H.; Zheng, G.; and Li, Z. 2019. Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. In *Proceedings of the AAAI*, volume 33, 5668–5675.

Ye, J.; Liu, Z.; Du, B.; Sun, L.; Li, W.; Fu, Y.; and Xiong, H. 2022. Learning the evolutionary and multi-scale graph structure for multivariate time series forecasting. In *Proceedings of the 28th ACM SIGKDD Conference*, 2296–2306.

Zheng, Y.; Wang, Q.; Zhuang, D.; Wang, S.; and Zhao, J. 2023a. Fairness-enhancing deep learning for ride-hailing demand prediction. *arXiv preprint arXiv:2303.05698*.

Zheng, Y.; Zhong, L.; Wang, S.; Yang, Y.; Gu, W.; Zhang, J.; and Wang, J. 2023b. Diffuflow: Robust fine-grained urban flow inference with denoising diffusion model. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3505–3513.