Estimating the impact of high-fidelity rainfall data on traffic conditions and traffic prediction

Prokhorchuk, Anatolii; Mitrovic, Nikola; Muhammad, Usman; Stevanovic, Aleksandar; Asif, Muhammad Tayyab; Dauwels, Justin; Jaillet, Patrick

DOI

[10.1177/03611981211026309](https://doi.org/10.1177/03611981211026309)

Publication date

2021

Document Version

Final published version

Published in

Transportation Research Record

Citation (APA)

Prokhorchuk, A., Mitrovic, N., Muhammad, U., Stevanovic, A., Asif, M. T., Dauwels, J., & Jaillet, P. (2021). Estimating the impact of high-fidelity rainfall data on traffic conditions and traffic prediction. In *Transportation Research Record* (11 ed., Vol. 2675, pp. 1285-1300). (Transportation Research Record; Vol. 2675, No. 11). SAGE Publications. <https://doi.org/10.1177/03611981211026309>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Estimating the Impact of High-Fidelity Rainfall Data on Traffic Conditions and Traffic Prediction

Transportation Research Record
2021, Vol. 2675(11) 1285–1300
© National Academy of Sciences:
Transportation Research Board 2021
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/03611981211026309
journals.sagepub.com/home/trr
 SAGE

Anatolii Prokhorchuk¹, Nikola Mitrovic², Usman Muhammad^{*3},
Aleksandar Stevanovic⁴, Muhammad Tayyab Asif^{*5},
Justin Dauwels⁶, and Patrick Jaillet⁷

Abstract

Accurate prediction of network-level traffic parameters during inclement weather conditions can greatly help in many transportation applications. Rainfall tends to have a quantifiable impact on driving behavior and traffic network performance. This impact is often studied for low-resolution rainfall data on small road networks, whereas this study investigates it in the context of a large traffic network and high-resolution rainfall radar images. First, the impact of rainfall intensity on traffic performance throughout the day and for different road categories is analyzed. Next, it is investigated whether including rainfall information can improve the predictive accuracy of the state-of-the-art traffic forecasting methods. Numerical results show that the impact of rainfall on traffic varies for different rainfall intensities as well as for different times of the day and days of the week. The results also show that incorporating rainfall data into prediction models improves their overall performance. The average reduction in mean absolute percentage error (MAPE) for models with rainfall data is 4.5%. Experiments with downsampled rainfall data were also performed, and it was concluded that incorporating higher resolution weather data does indeed lead to an increase in performance of traffic prediction models.

Weather conditions tend to have a measurable impact on traffic conditions on the roads. Many efforts have been made to incorporate weather-related variables in traffic modeling and to estimate the impact of adverse weather conditions on road traffic and driving behavior (1, 2). In particular, rainfall intensity has been frequently identified as an external factor that has the highest impact on driving conditions and network performance (2–7). Quantifying this impact could help transportation operators to better manage the safety and the efficiency of transportation systems during intense rainfall (8, 9).

In the estimation of the impact of rainfall on traffic conditions, the existing studies restrict their focus to a short segment or a single intersection only (3, 10–13). Moreover, the relevant literature often deals with a limited database where the information about the rainfall is frequently collected for a short period of time (e.g., a few weeks or days) and often aggregated in hourly or even daily intervals (3, 10–16). Such an approach fails to capture spatial and temporal dynamics of rainfall impact on traffic.

With the recent improvements in sensor technology, traffic-related information (e.g., rainfall intensity) is now

collected routinely from multiple sources with high temporal resolution. This study relies on rainfall and traffic datasets with high resolution and coverage to: (i) quantitatively investigate the impact of rainfall on the driving conditions in a large traffic network; and (ii) determine whether including information about the rainfall can improve the accuracy of traffic speed prediction algorithms for short- to long-term prediction horizons.

¹School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore

²Siemens, Austin, TX

³Panasonic, Singapore

⁴Department of Civil and Environmental Engineering, University of Pittsburgh, Pittsburgh, PA

⁵Amazon, Seattle, WA

⁶Department of Microelectronics, Delft University of Technology, Delft, The Netherlands

⁷Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA

*Work was done while at NTU

Corresponding Author:

Anatolii Prokhorchuk, anatolii001@e.ntu.edu.sg

For assessing the impact of rainfall on driving conditions, the average network traffic speed for rainfall and non-rainfall weather scenarios is evaluated. Since the rainfall and traffic intensities might significantly vary across time and space, the traffic speed for all 5 min time instances across the day is evaluated. This differentiates the current study from related papers that deal with the low- or medium-resolution weather information and traffic data for short road segments or a single corridor (16–22). It is important to verify whether similar findings might be achieved if high-resolution weather data is relied on, as well as heterogeneous and large traffic networks.

As a next step, it is explored how rainfall information may help to improve traffic speed prediction. To this end, speed information is combined with the rainfall data to train and test state-of-the-art forecasting methods involving ensemble methods and recurrent neural networks. It is important to point out that the focus of this study is to show how incorporating rainfall data can improve the prediction performance of traffic forecasting methods, whether they are ensemble methods or deep learning models. These predictions methods are compared for three different prediction horizons (15, 30, and 45 min) with and without rainfall data. The flowchart illustrating the preprocessing steps is presented in Figure 1.

The contributions of this study are as follows:

- To the best of the authors' knowledge, this is the first study concerned with the integration of high-resolution weather radar data into state-of-the-art algorithms for short- to medium-term traffic forecasting on the city scale network.
- To achieve this goal, 6 months of high-resolution rainfall intensity data for the Singapore road network was extracted using image processing.
- It is investigated how the rainfall affects the average traffic speed depending on various parameters: the hour of the day, the day of the week, the type of the road, and the intensity of the rain.
- A way of incorporating rainfall information in various prediction models (gradient boosting and recurrent neural networks) is proposed, and how it affects the model performance is studied. To provide a further comparison with the current literature, it is also investigated how the rainfall resolution affects the prediction performance by downsampling available high-resolution rainfall data.

The rest of the paper is structured as follows. First, the relevant literature is briefly reviewed. Next, the dataset analyzed in this paper is described. The following section presents the methodology for investigating the impact of rainfall on traffic conditions and for assessing

Figure 1. Flowchart of data extraction and fusion process.

the prediction algorithms. Then the results of the analysis and experiments are presented and discussed. In the final section, concluding remarks are offered, and topics for future research are suggested.

Related Work

It has been demonstrated that different rainfall intensities can have varying impact on the key traffic state parameters and their relationships (3–7, 14).

In certain cases, even light drizzle can make drivers cautious, with driving behavior matching that of a lower-capacity carriageway. This decrease in throughput leads to further degradation of other key traffic flow parameters such as speed and travel time (2). Xu et al. studied the effect of rainfall on the macroscopic fundamental diagram (MFD) (23). They found that rainfall decreases production, capacity, and speed, with the most pronounced effect during PM peak.

This quantifiable effect of rainfall is frequently utilized to develop mathematical models that can predict traffic conditions based on the current traffic and weather information (3, 4, 16, 24, 25). For instance, Tsirigotis et al. assess the effects of rainfall intensity on the predictability of traffic speed using vector autoregressive moving average models with explanatory variables (16). They analyze traffic and rainfall data spanned over 1 week for a single section of an urban motorway with a resolution of 10 min intervals. They conclude that information about the rainfall only marginally improves the prediction performance. Furthermore, they emphasize the need for considering a large database with intense rainfall incidents in traffic forecasting applications. In another study, Butler et al. apply neural networks to predict traffic volume by including weather conditions as exogenous variables (19). The test road network comprised of two traffic lanes approaching a junction in Dublin city, Ireland. The data was collected for 4 years with a sampling interval of 1 hour. They reveal that additional rainfall information, given in a coarse format, decreased the

forecasting accuracy for traffic volumes. In addition, they suggest that weather information should be observed at intervals finer than 1 hour in traffic applications. Hou et al. studied the impact of adverse weather on traffic to properly calibrate traffic-flow-related parameters during the inclement weather (26). They performed experiments based on rainfall data from automated surface observing system (ASOS) stations located at several U.S. airports. This data is sampled every 5 min; however, it is combined with data from loop detectors within 10 mi of the station. Based on simulations, they found that incorporating rainfall, visibility, and snow data can lead to producing more realistic traffic conditions during the adverse weather.

Peng et al. performed experiments on data from Georgia Department of Transportation (27). Both traffic and weather data are sampled with 1 h interval. They compared auto regressive integrated moving average (ARIMA) and neural network approaches, and observed that rainfall addition provided a marginal improvement in mean absolute percentage error (MAPE) (0.17% in absolute value for neural networks). Jia et al. applied deep learning models such as deep belief network (DBN) and long short-term memory (LSTM) to predict traffic speed and traffic flow (17, 18). They consider data from a single arterial road for the experiments sampled every 2 min, while the rainfall data is sampled hourly. They found that inclusion of the rainfall data did not improve prediction for 2 min horizon, while the slight improvement of 1.5% in MAPE was observed for 10 and 30 min horizons. Zhang and Kabuka provided results for a different deep learning model, that is, the gated recurrent unit (GRU) (21). They performed experiments on the Performance Measurement System (PeMS) dataset combined with hourly available weather data. They concluded that this can improve prediction accuracy up to 25% compared with the baseline (no weather data) case. Most recent relevant papers are summarized in Table 1.

Incorporating weather data has been explored in other applications intelligent transportation systems (ITS) as well. Patnaik et al. investigated adding rainfall data into the bus arrival time prediction model (28). They found that rainfall data did not increase the model accuracy and attributed it to the low resolution of the weather data. A similar problem was considered by Chen et al.; however, they did not investigate the improvement resulting from the addition of the precipitation data explicitly (29). A more recent study by Noor et al., similar to Patnaik et al., has not been able to significantly improve the bus arrival time prediction model by incorporating the weather data (28, 30). They suggested that the small sample size could have been the cause.

In summary, the scope of existing studies remains limited; either only small portions of the network were

considered, or the analyzed rainfall information is collected from a limited number of stations and/or aggregated in 1 h intervals.

Data Set

This section explains the traffic- and rainfall-related data sets analyzed in this study. The traffic network of Singapore is considered for analysis, since Singapore exhibits a tropical climate where the rainfall can be particularly heavy and persistent throughout the year. Table 2 summarizes the notation employed throughout this paper.

Traffic Data Set

Figure 2 shows the road network considered in the analysis. The Land Transportation Authority (LTA) of Singapore collects raw traffic data in the network from multiple sources such as loop detectors and probe vehicles. The Traffic Management of LTA analyzes these inputs and matches them to the corresponding road segments. Then, estimation models infer the speed along each segment during the 5 min sampling interval. These values are then discretized into speed bands. Each speed band corresponds to a range of 10 km/h: speed band one (SB1) corresponds to average speed of 0–10 km/h, SB2 corresponds to average speed of 10–20 km/h, and so on. While highly accurate speed data, associated with high data collection costs, is desired among practitioners, it is still important to understand the value of easily available discrete information. This is especially relevant for the prediction and similar applications where the forecast speed values are often used as inputs for complex routing and traffic control mechanisms.

The network being considered consists of $n = 2,896$ road segments: $v_1, \dots, v_n$. The observations are available for 6 months covering period from August 2016 to January 2017. For each 5 min interval j and for each road segment v_i , $s_{i,j}$ is defined as the discretized average speed during this interval for that segment. Every segment v_i has less than 10% missing data.

Rainfall Data Set

The National Environmental Agency (NEA) of Singapore administers automated weather instruments to collect information about the rainfall (<http://www.nea.gov.sg/>). The rainfall data from the meteorological radar is published on the NEA website in the form of an image every 5 to 15 min (see Figure 2). This image contains the information about the rainfall intensity for each location in the network. In addition to the precipitation scale, the

Table 1. Brief Summary of the Relevant Studies

Study	Traffic data	Weather data	Time period	Prediction method	Findings
Qiu et al. (20)	14 Roads, traffic flow reported every 2 min	Nine rain stations, data reported every 5 min	2 months	LRR, NN, PCA-SVR, model fusion	For most roads, relative improvement was at least 10%; however the proposed fusion model performance has not improved compared with the single model for more than half of the heavy rain scenarios.
Zhang and Kabuka (21)	PEMS dataset (http://pems.dot.ca.gov/), 1 h sampling interval	Weather data (precipitation, wind, temperature) from NOAA, 1 h sampling interval	2 years	Deep GRU recurrent neural network	Prediction accuracy improved up to 25%.
Koesdwiady et al. (22)	PEMS dataset (47 freeways) (http://pems.dot.ca.gov/), 15 min resolution	Rainfall data from 16 stations available hourly, resampled to 15 min	4 months	ARIMA, ANN and DBN	Reported improvement of 16% in MAE by incorporating rainfall data
Jia et al. (17, 18)	One arterial road, 2 min sampling interval	Hourly rainfall intensity	3 months	DBN, LSTM	For 10 and 30 min prediction horizon the accuracy is improved with the addition of the rainfall by 1.5% in terms of absolute MAPE values.
Lu (4)	183 road segments in Singapore	Radar images obtained every 5 to 10 min	3 months	Impulse response function and linear system	An improvement was observed for 15 min prediction horizon; however, the approach was only compared with trivial baselines such as the historical average.
This study	Singapore highway network (2,896 segments), speed bands reported every 5 min	Radar images covering the whole network, obtained every 5 to 10 min	6 months	Gradient boosting, LSTM, GRU	Achieved relative improvement of 5% in MAPE across different prediction horizon (or 0.5% in absolute values). Investigated the effect of the rainfall data resolution on the prediction accuracy.

Note: ANN=Artificial Neural Network; GRU=Gated Recurrent Unit; LR=Linear Regression; LSTM=Long short-term memory; MAE=Mean Absolute Error; MAPE=Mean Absolute Percentage Error; NN=Neural Network; PCA-SVR=Principal Component Analysis - Support Vector Regression.

Table 2. Notation

Notation	Definition
n	Number of segments in the road network
v_i	i -th segment in the network
s_{ij}	Discretized speed value for i -th segment during j -th time period
r_{ij}	Computed rainfall intensity value for i -th segment during j -th time period
a, b	Parameters for rainfall intensity estimation
$s_j^{\text{non-rainfall}}$	Average network speed during j -th time among segments with zero rainfall intensity during that period
s_j^{rainfall}	Average network speed during j -th time period among segments with non-zero rainfall intensity during that period
$d(j), h(j)$	Day of the week and hour of the day for j -th time period
k	Prediction horizon

Figure 2. Singapore road network consisting of 2,139 highways, 10 arterial, and 747 slip roads. **Figure 3.** Locations of rainfall stations in Singapore (upper left). Format of rain map acquired from the National Environmental Agency (NEA) (bottom). Corresponding precipitation scale is also provided on NEA's website (upper right corner).

time stamp and the map of the highway network help to precisely integrate weather and traffic data (see Figures 2 and 3). Rainfall information was collected during the 6-month period that corresponds to the speed data set. Similarly to the traffic data, r_{ij} is defined as the estimated rainfall intensity for segment v_i during time period j .

Rainfall intensity at each segment v_i is estimated as follows: first the rainfall and traffic road maps are overlaid to estimate the location of the segment v_i on the rainfall map (see Figure 3). The nearest pixel of the rainfall map is assigned to each segment v_i . To derive the rainfall rate at a single pixel of the map visualizations, Doppler Radar reflectivity is employed:

$$r = a(10^{\frac{d}{10}})^b, \quad (1)$$

where:

r is rainfall rate in millimeters per hour;

d is the reading from the image visualization (see the upper right corner of Figure 3); and

a and b are empirically estimated parameters.

This equation is known as the Marshall-Palmer formula; it allows to convert dBZ (decibel relative to Z) values from the radar into rainfall intensity. The coefficients computed in Lu are employed (4). The rainfall intensity at a segment v_i is approximated by the mean of rainfall intensities of the assigned pixel and its neighbors from the 3×3 box (see the right side of Figure 4).

Since the rainfall data is irregularly reported (especially during the heavy showers), optical character recognition (OCR) is employed to extract the time components from each of the collected images (see the top right corner of Figure 4). OCR is the process of converting text images into text in ASCII format. As a pre-processing step, the text images are binarized and scaled (31). Then OCR functionality of Matlab is applied, and further analysis is only carried out on those images that fulfill the following requirement: $0 \leq (t_i - t_i^*) \leq \sigma$, where t_i is the time when the image appeared in the database (or time instant when the picture is collected from the

Figure 4. Integrated rain and roadway map (left) with the estimated rainfall intensity for 3×3 pixels region (right). Time instance of the map is extracted with the help of optical character recognition (OCR) (top right).

internet), t_i^* is the corresponding output from the OCR, and σ is an empirically determined parameter (15 min). This condition ensures that outdated rainfall data is not incorporated, since sometimes the image on the NEA website will not be updated for an extended period of time. Since the text being extracted has consistent formatting and is clearly visible in the image, it is found that OCR performs quite well. The extracted time has been compared with the image time stamp and no inconsistencies were found. Note, OCR was only applied to extract time stamps. By contrast, the actual rainfall data is obtained from the raw images via the formula (1).

The occasional irregularities in reporting rainfall maps lead to the unavoidable missing values in weather data sets. If there is no data for the previous 5–10 min, then the missing data is imputed via linear interpolation of the available rainfall information.

Data Analysis

This section describes the approach to determining the impact of rainfall on traffic speed across the network, and incorporating the rainfall information into traffic speed prediction algorithms.

Impact of Rainfall on the Network Speed

First, the average driving conditions (or the average speed-band) in the network for rainfall and non-rainfall weather conditions are analyzed. As mentioned previously, for a given road segment v_i and time period j , $s_{i,j}$ and $r_{i,j}$ are defined as average speed and average rainfall intensity, respectively. The interest is in overall network conditions during both scenarios. The average network speed is defined in both cases as:

$$S_j^{\text{non-rainfall}} = \frac{\sum_{i=1, r_{i,j}=0}^n s_{i,j}}{|\{r_{i,j} : r_{i,j}=0\}|}, S_j^{\text{rainfall}} = \frac{\sum_{i=1, r_{i,j}>0}^n s_{i,j}}{|\{r_{i,j} : r_{i,j}>0\}|}. \quad (2)$$

Since the rainfall data is available on road link level, it is known that, during a particular time frame, only a subset of links experience rainfall. For a fair comparison, average speed is computed only for links either with or without rainfall, and not for the whole network at once.

Impact of Rainfall on Traffic Prediction

This section explains how the information about the rainfall is incorporated, as an exogenous variable, in the prediction of speed of a certain segment. To predict the future driving conditions $s_{i,j+k}$ at time $j+k$ for segment v_i , a feature vector is constructed that consists of past driving conditions $(s_{i,j-P}, \dots, s_{i,j})$, past weather conditions $(r_{i,j-P}, \dots, r_{i,j})$, time of the day, and day of the week:

$$s_{i,j+k} = f_k(d(t_i), h(t_i), s_{i,j}, r_{i,j}, \dots, s_{i,j-P}, r_{i,j-P}), \quad (3)$$

where:

$d(j)$ and $h(j)$ are the day and hour, respectively, of the particular time period j ;

P represents the number of past speed band values of the segment that are used;

k is the prediction horizon; and

f_k is the predictor function.

Note, the rainfall has a near-immediate impact on the traffic, which is in alignment with the findings from other relevant studies (4). The function f_k can be approximated by means of any state-of-the-art prediction algorithm. This study applies an ensemble method of extreme gradient boosting (XGBoost) and several variants of recurrent neural networks (RNNs) for regression. A brief review of these methods is given below.

Ensemble Methods. Ensemble methods are data-driven algorithms that have been widely applied for a variety of applications (32, 33). These techniques heavily rely on the underlying assumption that collection of weak classifiers can be a better predictor than the single classifier. This study applied XGBoost, as it is the most popular choice for both classification and prediction problems because of its highly scalable nature (34). In XGBoost, additive training is implemented whereby a new tree is added to the model that has been learned so far. The general objective function can thus be written as:

$$O^{(t)} = \sum_{i=1}^n \ell(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)) + \Omega(f_t), \quad (4)$$

where:

ℓ is the loss function;

f_t is the function that contains the structure of tree t ; and

Ω is the regularization for the complexity of the tree structure defined in terms of number of leaves T and weights w_j of leaves as:

$$\Omega(f) = \gamma T + \frac{1}{2} \lambda \sum_{j=1}^T w_j^2.$$

The objective function (4) is optimized through first- and second-order gradient of the loss function ℓ which allows XGBoost to support a wide range of custom loss functions. Further details of the method can be found in Chen and Guestrin (34).

Recurrent Neural Networks (RNNs)

Previously, traditional machine learning models such as support vector machines (SVM) have been applied to traffic prediction (35, 36). However, currently deep learning approaches have become widespread (37–39). LSTM networks are a variation of RNNs that are able to capture long-term dependencies (40). GRU is another network architecture that is commonly applied for time-series forecasting (41). Both architectures are characterized by gating mechanisms that aim to overcome the vanishing gradient problems of common RNN. Another RNN modification is the bidirectional recurrent network, first proposed in Schuster and Paliwal (42). In this case, the recurrent layer is duplicated and the input sequence is fed both forwards and backward. In Cui et al. it was observed that stacked bidirectional LSTM outperforms plain LSTM and GRU networks for traffic prediction (39). This study investigates whether this approach is still beneficial with the addition of rainfall information.

In the case of gradient boosting, a separate model is trained for each of the segments v_i . For recurrent networks, a single model is trained for the whole network.

The input to the network in non-rainfall case is a tensor with the dimensions $(n_{\text{samples}}, n_{\text{segments}}, n_{\text{timestamps}})$. There are several ways of including the rainfall data in the tensor. The most straightforward one is to treat rainfall time series for each segment as another input feature, so that the tensor has the following dimensions: $(n_{\text{samples}}, 2n_{\text{segments}}, n_{\text{timestamps}})$.

A similar structure is employed for all tested RNN models (GRU, LSTM, BLSTM). One recurrent layer is followed by two fully connected layers with n_{segments} neurons in all cases. Dropout regularization is applied after each layer. A hyperparameter search is performed to find the optimal layer size and dropout rate. The RNN results are reported only for the best hyperparameters values: LSTM and BLSTM layers consist of 1,024 neurons, while the GRU layer consists of 512; the dropout rate is set to 0.1. Rectified linear unit is used as an activation function for all layers, and Adam is employed as an optimizer with mean squared loss (43). The networks are implemented via Keras library (44).

Evaluation

For training the prediction algorithm, all those data points when the corresponding rainfall intensity data is available ($r_{ij} \geq 0$) are considered; whereas, for evaluation, only those data points with rainfall intensity above zero ($r_{ij} > 0$) are considered. The prediction performance is evaluated by three-fold rolling cross-validation. The dataset contains approximately 180 days of data: for the first fold the training data consists of observations from day 1 to day 45 and test data from day 46 to day 90; for the second fold, the training data consists of observation from day 1 to day 90 and test data from day 90 to day 135; and for the last fold, the training data consists of observations from day 1 to day 135 and test data from day 135 to day 180. This is one of several approaches for time-series cross-validation described in the literature (45). The usual cross-validation procedure is not always suitable for time-series data, especially if it is non-stationary.

Several common evaluation metrics are employed, namely mean absolute error (MAE), mean squared error (MSE) and MAPE. MAPE for segment v_i and k^{th} prediction horizon is defined as:

$$\text{MAPE}(v_i, k) = \frac{1}{d} \sum_{j=1}^d \frac{|\hat{s}_{i,j}^k - s_{i,j}^k|}{s_{i,j}^k}, \quad (5)$$

where:

$\hat{s}_{i,j}^k$ is the predicted speed band at time t_j ; and

d is the number of test samples.

MAPE for the whole network G containing p segments for the k^{th} prediction horizon is defined as:

Figure 5. Histogram of rainfall intensities across the available dataset:

$$\text{MAPE}(k) = \frac{1}{P} \sum_{i=1}^P \text{MAPE}(v_i, k). \quad (6)$$

Impact of Rainfall on Network Speed

This section discusses and analyzes the effects of rainfall on traffic speed. Figure 5 shows the distribution of rainfall intensities across the available dataset. It can be observed that the dataset contains the whole range of intensities, with a slight overrepresentation of high-intensity time periods (as can be expected from the tropical climate of Singapore). Figure 6 shows the average speed during rainy and corresponding non-rainy time instants. Reduction in average speed (across the network), influenced by the rainfall, can be seen even in speed-band data (see Figure 6). The reduction is higher for the periods with congested traffic, which is in alignment with the results from other studies (14). In the case of weekdays, as expected, these periods are evening and morning peak hours (see Figures 6 and 7a). For the analysis, AM peak is defined to be from 06:00 to 09:59, and PM peak to be from 17:00 to 20:59. Concerning weekends, the highest impact of rainfall is during the evening and night hours (see Figures 6b and 7a) when an increasing number of recreational- and shopping-related trips are realized in a vibrant city such as Singapore. Figure 6 further shows that the lines of average network speed during weekdays are smoother than for weekends (see blue and red lines in Figure 6) as a consequence of a larger dataset. Figure 7a shows the reduction in speed band because of rainfall for different times of the day and road categories. From Figure 7a, it can be seen that reduction in speed band is greater for highways than for arterial roads (see also Figure 7b). One possible reason for this difference is that vehicles tend to travel faster on

highways, and therefore they tend to slow down more when roads are slippery. Figure 7b shows the reduction in speed band because of rainfall for different rainfall intensities and road categories in the network. As expected, heavy rainfall has a greater effect on traffic than light rainfall. Heavy and even light rain might have significant impact even during the weekends. This impact has been frequently ignored by other studies, as they often are limited to peak hours.

Next, to investigate whether the rainfall impact is statistically significant, a regression analysis is performed. Two scenarios are considered: first, the impact of any rainfall (i.e., intensity is greater than zero); second, the impact of heavy rainfall (the rainfall is considered heavy if the intensity is higher than the median value). Day of the week and time of the day are used as covariates to perform the regression on the speed value as a target. Figure 8 shows the histogram of estimated rainfall effects for both scenarios. It can be seen that, for almost all the links, the rainfall has a negative effect on speed. In addition to average effect values, the 95% confidence intervals are also computed. For the rainfall scenario, out of 2,896 links, only 57 links (or less than 2%) do not have the whole interval below zero, making the results statistically significant.

Results and Discussion

This section presents and discusses the results of using rainfall information with a high temporal resolution for short- to long-term speed prediction.

First, the prediction performance of XGBoost regression is investigated for three different prediction horizons: 15, 30, and 45 min. For testing, only those time points where the rainfall intensity is greater than zero are considered. Table 3 shows the results for varying values of P in (3) for 15, 30, and 45 min prediction horizon.

From Table 3 it can be observed that addition of rainfall information into the model decreases both the mean and standard deviation of MAPE for all three prediction horizons. The percentage decrease in mean for all P values and all prediction horizons is around 4%, while the decrease in standard deviation is about 7%. The performance of the model varies slightly for different P values and an optimum P value was obtained for each horizon. Utilizing more than five past values does not significantly increase MAPE. The authors believe that this is because the more recent speed data is a better predictor of short-term speed fluctuations.

There is an increase in both the mean and standard deviation on increasing the prediction horizon from 15 to 45 min, as expected, since the long-term prediction is more challenging. However, the effectiveness of incorporating rainfall information into the prediction model is

Figure 6. Reduction in average traffic speed because of rainfall for various times of the day: (a) during different times of the day, and (b) during light and heavy rain.

Figure 7. Impact of rainfall on the average traffic speed for different road types: (a) during different times of the day, and (b) during light and heavy rain.

Figure 8. Histogram of average effects of rainfall (left) and heavy rainfall (right) from the regression analysis.

Table 3. Mean and Standard Deviation of Mean Absolute Percentage Error (MAPE) for Varying Numbers of Previous Timestamps Used for Training (P) for Three Prediction Horizons

(a) 15 min Horizon				
P	Mean		Standard deviation	
	Without rain	With rain	Without rain	With rain
2	10.22	9.21	1.70	1.30
3	10.09	9.14	1.70	1.30
4	10.02	9.11	1.69	1.30
5	10.00	9.10	1.70	1.30
6	9.99	9.09	1.70	1.31
7	9.98	9.09	1.71	1.31
8	9.98	9.10	1.72	1.32

(b) 30 min Horizon				
P	Mean		Standard deviation	
	Without rain	With rain	Without rain	With rain
2	12.56	11.43	2.34	1.91
3	12.45	11.37	2.35	1.91
4	12.37	11.33	2.34	1.90
5	12.35	11.31	2.34	1.90
6	12.34	11.31	2.35	1.91
7	12.34	11.31	2.35	1.91
8	12.33	11.32	2.38	1.93

(c) 45 min Horizon				
P	Mean		Standard deviation	
	Without rain	With rain	Without rain	With rain
2	14.07	12.90	2.69	2.28
3	13.97	12.83	2.70	2.28
4	13.90	12.79	2.71	2.28
5	13.87	12.76	2.70	2.27
6	13.85	12.76	2.70	2.28
7	13.82	12.76	2.69	2.27
8	13.82	12.75	2.71	2.28

consistent for all cases. A Mann-Whitney U test is also performed to determine the significance of these results (46). This is a non-parametric test to determine whether the two samples come from the same distribution: for all prediction horizons, the p-value was smaller than $1e-3$. This is also illustrated in Figure 9 which shows the distribution of the difference in MAPE metric. It was obtained by subtracting MAPE for the prediction model without rainfall feature from the MAPE of the model with rainfall feature for each of the 2,896 segments. All the distributions reveal that the prediction model shows improved performance for around 80% of the segments in the network, and the decrease in MAPE for these segments is much higher than the increase in MAPE for the remaining segments that are adversely affected by the inclusion of rainfall intensity feature. These results show that traffic speed prediction during inclement weather can be

improved by the inclusion of rainfall information into the model. There are several possibilities why some links do not show improvement in MAPE: a larger share of missing values (either in the speed data or the rainfall data), speed profile which differs significantly from the rest of the network, or something particular to these links, that make them less sensitive to rainfall effects (e.g., tunnels).

Next, it is investigated how different models behave with the inclusion of rainfall data. Tables 4 to 6 show the results for several models. It can be observed that, in most cases, the addition of rainfall improves the performance of the model: this holds for linear regression, gradient boosting, and RNNs. Linear regression is added for comparison as a baseline model. For all prediction horizons, it performs worse than XGBoost and LSTM, as expected. It is observed that linear regression

Figure 9. Distribution of the difference in mean absolute percentage error (MAPE) between the baseline and the rainfall models for all segments in the network (positive numbers represent improvement because of the addition of rainfall data): (a) 15 min prediction horizon, (b) 30 min prediction horizon, and (c) 45 min prediction horizon.

predictions are usually close to the mean, which can explain the squared error showing the largest gap between the regression and other models (since it is more sensitive to larger errors). Next, it is noticed that for 15 and 30 min prediction horizons, gradient boosting

outperforms various RNN architectures. Only for the 45 min horizon is the difference in metrics between gradient boosting and best RNNs marginal, with LSTM showing the lowest MAPE among all models. While RNNs can seem like a better-suited model for time-series prediction because of their design, results show that gradient boosting can achieve better performance for some of the prediction horizons and metrics. Only a few traffic forecasting studies use gradient boosting as a baseline, with most focusing only on historical average and ARIMA models for this purpose. However, results in Yao et al. show that gradient boosting outperformed convolutional LSTM and deep spatio-temporal residual networks on two datasets; at the same time, both deep multi-view spatial-temporal network and spatial-temporal dynamic networks achieved smaller errors (47–50). While the results are not dissimilar to these studies, another possible reason for RNN performing worse in this case can be that discretized speed data is used.

Table 7 shows the computation times for XGBoost and LSTM, as well as linear regression, which is added as a baseline. The experiments were performed on an Intel(R) Xeon(R) CPU E5-2640 v4 @ 2.40 GHz server with 4 GeForce GTX 1080 GPUs. The slight difference in computation times between XGBoost and LSTMs is because one LSTM model is trained for all segments, while gradient boosting is trained for each segment individually. Depending on the technical requirements and the scale of the dataset, one or other model can be chosen. It was observed that gradient boosting outperforms RNNs on the shorter horizon forecasting; however, it also requires training a separate model for each segment, which comes at an additional computational cost. Linear regression shows the fastest training and inference times while producing the highest forecasting errors.

It can be seen from Figure 6 that the traffic speed reduction because of rainfall varies across time; therefore, the performance of the predictive models was investigated for different times of the day. For weekdays, four time categories were considered, to distinguish the performance during peak hours. Weekends were split into two categories, as no distinct pattern can be observed in Figure 6 for the speed during weekends. The prediction was performed for a horizon of 15 min by XGBoost. The results for weekdays in Figure 10 show that MAPE is higher during the peak hours. The traffic is most congested and rainfall has the greatest impact on traffic speed during peak hours, as is observed in Figure 7. It can also be seen that the inclusion of rainfall data provides the largest improvement in accuracy for AM and PM peak hours (see Figure 10).

Next, it is investigated how the resolution of the rainfall data affects the model performance. Several temporal resolutions often found in the literature (15, 30, and

Table 4. Comparison between Different Models: 15 min Prediction Horizon

Model	MAPE	MAE	MSE
Linear regression (no rain)	10.6	0.74	0.99
XGBoost (no rain)	9.35	0.49	0.61
GRU (no rain)	13.29	0.74	1.05
LSTM (no rain)	11.38	0.63	0.88
BLSTM (no rain)	13.34	0.74	1.19
Linear regression (rainfall)	9.71	0.70	0.89
XGBoost (rainfall)	9.03	0.49	0.59
GRU (rainfall)	12.39	0.67	0.92
LSTM (rainfall)	11.49	0.64	0.88
BLSTM (rainfall)	12.84	0.72	1.07

Note: BLSTM=Bidirectional Long Short-Term Memory; GRU=Gated Recurrent Unit; LSTM=Long short-term memory; MAE=Mean Absolute Error; MAPE=Mean Absolute Percentage Error; MSE=Mean Squared Error. Bold indicate the Lowest error values.

Table 5. Comparison between Different Models: 30 min Prediction Horizon

Model	MAPE	MAE	MSE
Linear regression (no rain)	13.63	0.70	1.11
XGBoost (no rain)	11.75	0.60	0.88
GRU (no rain)	13.03	0.71	1.04
LSTM (no rain)	12.23	0.67	0.98
BLSTM (no rain)	12.45	0.69	1.00
Linear regression (rainfall)	13.16	0.68	1.05
XGBoost (rainfall)	11.32	0.59	0.85
GRU (rainfall)	14.78	0.81	1.29
LSTM (rainfall)	13.89	0.88	1.21
BLSTM (rainfall)	13.21	0.73	1.13

Note: BLSTM=Bidirectional Long Short-Term Memory; GRU=Gated Recurrent Unit; LSTM=Long short-term memory; MAE=Mean Absolute Error; MAPE=Mean Absolute Percentage Error; MSE=Mean Squared Error. Bold indicate the Lowest error values.

Table 6. Comparison between Different Models: 45 Min Prediction Horizon

Model	MAPE	MAE	MSE
Linear regression (no rain)	14.98	0.74	1.20
XGBoost (no rain)	13.32	0.67	1.07
GRU (no rain)	13.36	0.70	1.06
LSTM (no rain)	12.96	0.69	1.04
BLSTM (no rain)	13.61	0.71	1.08
Linear regression (rainfall)	13.94	0.73	1.15
XGBoost (rainfall)	12.93	0.67	1.02
GRU (rainfall)	13.90	0.75	1.17
LSTM (rainfall)	12.90	0.70	1.07
BLSTM (rainfall)	14.19	0.73	1.17

Note: BLSTM=Bidirectional Long Short-Term Memory; GRU=Gated Recurrent Unit; LSTM=Long short-term memory; MAE=Mean Absolute Error; MAPE=Mean Absolute Percentage Error; MSE=Mean Squared Error. Bold indicate the Lowest error values.

Table 7. Average Training and Prediction Times Per Segment for Linear Regression, XGBoost and Long Short-Term Memory (LSTM)

Model	Training time per segment, s	Prediction time per segment, s
Linear regression	0.002	0.0001
XGBoost	0.14	0.001
LSTM	0.04	0.001

Figure 10. Time category-wise comparison of mean absolute percentage error (MAPE) and mean squared error (MSE) with and without rainfall data.

60 min) and several spatial resolutions are compared. A new temporal resolution is computed by averaging the intensities for the corresponding time period for each segment. Also, the original (segment-level) spatial resolution is compared with city-level resolution. To this end, rainfall intensities from all segments during the time period are averaged. In all cases, the test data remains the same as in the original experiments, for a fair evaluation. Results are presented in Figure 11. It can be observed that higher resolution does indeed produce better results. However, almost all resolution levels produce improvements compared with the baseline model. Only city-level data sampled at the 60 min interval decrease the model's accuracy, similar to the effect observed in Butler et al. (19). Additionally, the results show that the spatial resolution of rainfall data plays a significant role in short-term (15 min) prediction. In this case, the local information is required, as the rainfall is not homogenous across the country. At the same time, both segment-level and city-level data produce similar results for medium-term (45 min) prediction, since rain typically affects the entire country over a long time period.

Conclusions

This study has investigated the impact of rainfall on traffic conditions, and also how rainfall information can be leveraged to improve traffic predictions. Some of the previous studies were not successful in showing that the inclusion of lower-resolution precipitation data can improve the performance of short-term traffic estimation procedures. This study has considered high-resolution and high-volume rainfall and traffic datasets. This study has presented a method to merge rainfall information with traffic data for traffic prediction, and its prediction accuracy has been evaluated for horizons of 15 to

Figure 11. Impact of rainfall data spatiotemporal resolution on traffic speed prediction accuracy for three prediction horizon. Island-level resolution refers to city-level: (a) 15 min prediction horizon, (b) 30 min prediction horizon, and (c) 45 min prediction horizon.

45 min. The results have shown that, for predicting traffic conditions during inclement weather, the inclusion of rainfall information improves the prediction accuracy. It is observed that most models can benefit from the rainfall information; however, a better choice of the model architecture can provide a larger improvement than the inclusion of the weather data into some models.

Also, experiments were performed with rainfall data aggregated at different resolutions to study the impact of resolution. It was observed that higher-resolution rainfall data leads to higher model accuracy; however, most of the resolutions can lead to a marginal improvement over the baseline. In future work, the authors will consider network topology to improve the speed prediction models. The authors will also explore the impact of rainfall on the occurrence and duration of accidents in the context of a large traffic network and high-resolution rainfall data. Also, the authors will explore how these accidents affect traffic conditions and deteriorate the prediction performance.

Author Contributions

The authors confirm contribution to the paper as follows: Study conception and design: A. Prokhorchuk, N. Mitrovic, J. Dauwels, P. Jaillet, M. T. Asif; data collection: A. Prokhorchuk, N. Mitrovic, U. Muhammad, M. T. Asif; analysis and interpretation of results: A. Prokhorchuk, N. Mitrovic, J. Dauwels, P. Jaillet, A. Stevanovic; draft manuscript preparation: A. Prokhorchuk, N. Mitrovic, U. Muhammad. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was partially supported by the Singapore National Research Foundation through the Singapore-MIT Alliance for Research and Technology (SMART) Centre for Future Urban Mobility (FM).

ORCID iDs

Anatolii Prokhorchuk <https://orcid.org/0000-0001-6347-1107>
Aleksandar Stevanovic <https://orcid.org/0000-0003-1091-3340>

References

1. Bogren, J., and T. Gustavsson. RSI—Road Status Information a New Method for Detection of Road Conditions. *Proc., 17th International Road Weather Conference*, La Massana, Andorra, 2014.
2. Lam, W. H. K., M. L. Tam, X. Cao, and X. Li. Modeling the Effects of Rainfall Intensity on Traffic Speed, Flow, and Density Relationships for Urban Roads. *Journal of Transportation Engineering*, Vol. 139, No. 7, 2013, pp. 758–770.
3. Bartlett, A., W. Lao, Y. Zhao, and A. W. Sadek. *Impact of Inclement Weather on Hourly Traffic Volumes in Buffalo, New York*. Technical Report No. 13-3240. Transportation Research Board 92nd Annual Meeting, Washington D.C., 2013.
4. Lu, H. Short-Term Traffic Prediction Using Rainfall. *International Journal of Signal Processing Systems*, Vol. 2, No. 1, 2014, pp. 70–73.
5. Agarwal, M., T. H. Maze, and R. Souleyrette. The Weather and Its Impact on Urban Freeway Traffic Operations. *Proc., 2005 Mid Continent Transportation Research Symposium*, Washington D.C., 2006.
6. Billot, R., N.-E. El Faouzi, and F. De Vuyst. Multilevel Assessment of the Impact of Rain on Drivers' Behavior: Standardized Methodology and Empirical Analysis. *Transportation Research Record: Journal of the Transportation Research Board*, 2009. 2107: 134–142.
7. Rakha, H., M. Farzaneh, M. Arafteh, and E. Sterzin. Inclement Weather Impacts on Freeway Traffic Stream Behavior. *Transportation Research Record: Journal of the Transportation Research Board*, 2008. 2071: 8–18.
8. Perrin, H., P. Martin, and B. Hansen. Modifying Signal Timing During Inclement Weather. *Transportation Research Record: Journal of the Transportation Research Board*, 2001. 1748: 66–71.
9. Katz, B., C. O'Donnell, K. Donoughe, J. E. Atkinson, M. D. Finley, K. N. Balke, B. T. Kuhn, and D. Warren. *Guidelines for the Use of Variable Speed Limit Systems in Wet Weather*. Technical Report, Federal Highway Administration. Office of Safety, Washington, 2012.
10. Qiao, W., A. Haghani, and M. Hamedi. Short-Term Travel Time Prediction Considering the Effects of Weather. *Transportation Research Record: Journal of the Transportation Research Board*, 2012. 2308: 61–72.
11. Tselentis, D. I., E. I. Vlahogianni, and M. G. Karlaftis. Improving Short-Term Traffic Forecasts: To Combine Models or Not to Combine? *IET Intelligent Transport Systems*, Vol. 9, No. 2, 2014, pp. 193–201.
12. Dailey, D. J., and T. Trepanier. *The Use of Weather Data to Predict Non-Recurring Traffic Congestion*. Technical Report. Washington Department of Transportation, 2006.
13. Lin, L., M. Ni, Q. He, J. Gao, and A. W. Sadek. Modeling the Impacts of Inclement Weather on Freeway Traffic Speed: Exploratory Study with Social Media Data. *Transportation Research Record: Journal of the Transportation Research Board*, 2015. 2482: 82–89.
14. Liu, X., B. Wu, Y. Zhao, and L. Zheng. Analysis of Speed Characteristics of Urban Expressways under Rainy Conditions. *Proc., 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Yokohama, Japan, 2017, pp. 1–6. IEEE, New York.
15. Dunne, S., and B. Ghosh. Weather Adaptive Traffic Prediction Using Neurowavelet Models. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 14, No. 1, 2013, pp. 370–379.
16. Tsirigotis, L., E. I. Vlahogianni, and M. G. Karlaftis. Does Information on Weather Affect the Performance of Short-Term Traffic Forecasting Models? *International Journal of Intelligent Transportation Systems Research*, Vol. 10, No. 1, 2012, pp. 1–10.

17. Jia, Y., J. Wu, and M. Xu. Traffic Flow Prediction with Rainfall Impact Using a Deep Learning Method. *Journal of Advanced Transportation*, Vol. 2017, 2017.
18. Jia, Y., J. Wu, M. Ben-Akiva, R. Seshadri, and Y. Du. Rainfall-Integrated Traffic Speed Prediction Using Deep Learning Method. *IET Intelligent Transport Systems*, Vol. 11, No. 9, 2017, pp. 531–536.
19. Butler, S., J. Ringwood, and D. Fay. Use of Weather Inputs in Traffic Volume Forecasting.
20. Qiu, H., R. Li, and H. Liu. Integrated Model for Traffic Flow Forecasting under Rainy Conditions. *Journal of Advanced Transportation*, Vol. 50, No. 8, 2016, pp. 1754–1769.
21. Zhang, D., and M. R. Kabuka. Combining Weather Condition Data to Predict Traffic Flow: A GRU-Based Deep Learning Approach. *IET Intelligent Transport Systems*, Vol. 12, No. 7, 2018, pp. 578–585.
22. Koesdwiady, A., R. Soua, and F. Karray. Improving Traffic Flow Prediction with Weather Information in Connected Cars: A Deep Learning Approach. *IEEE Transactions on Vehicular Technology*, Vol. 65, No. 12, 2016, pp. 9508–9517.
23. Xu, F., Z. He, Z. Sha, L. Zhuang, and W. Sun. Assessing the Impact of Rainfall on Traffic Operation of Urban Road Network. *Procedia-Social and Behavioral Sciences*, Vol. 96, 2013, pp. 82–89.
24. Zhang, J., G. Song, D. Gong, Y. Gao, L. Yu, and J. Guo. Analysis of Rainfall Effects on Road Travel Speed in Beijing, China. *IET Intelligent Transport Systems*, Vol. 12, No. 2, 2017, pp. 93–102.
25. Huang, S.-H., and B. Ran. *An Application of Neural Network on Traffic Speed Prediction under Adverse Weather Condition*. PhD thesis. University of Wisconsin, Madison, 2003.
26. Hou, T., H. S. Mahmassani, R. M. Alfelori, J. Kim, and M. Saberli. Calibration of Traffic Flow Models under Adverse Weather and Application in Mesoscopic Network Simulation. *Transportation Research Record: Journal of the Transportation Research Board*, 2013. 2391: 92–104.
27. Peng, H., S. U. Bobade, M. E. Cotterell, and J. A. Miller. Forecasting Traffic Flow: Short Term, Long Term, and When It Rains. In *Proc., International Conference on Big Data*, (F. Chin, C. Chen, L. Khan, K. Lee, L.J. Zhang LJ. eds.), Big Data – BigData 2018. BIGDATA 2018. Lecture Notes in Computer Science, Springer, Cham, 2018, pp. 57–71. Springer.
28. Patnaik, J., S. Chien, and A. Bladikas. Estimation of Bus Arrival Times Using APC Data. *Journal of Public Transportation*, Vol. 7, No. 1, 2004, p. 1.
29. Chen, M., X. Liu, J. Xia, and S. I. Chien. A Dynamic Bus-Arrival Time Prediction Model Based on APC Data. *Computer-Aided Civil and Infrastructure Engineering*, Vol. 19, No. 5, 2004, pp. 364–376.
30. Noor, R. Md., N. S. Yik, R. Kolandaisamy, I. Ahmedy, M. A. Hossain, K.-L. A. Yau, W. Md. Shah, and T. Nandy. Predict Arrival Time by Using Machine Learning Algorithm to Promote Utilization of Urban Smart Bus, 2020.
31. Breuel, T. M. *The OCROPUS Open Source OCR System*. Document Recognition and Retrieval XV, Vol. 6815, 2008, 68150F. International Society for Optics and Photonics.
32. Moretti, F., S. Pizzuti, S. Panzneri, and M. Annunziato. Urban Traffic Flow Forecasting Through Statistical and Neural Network Bagging Ensemble Hybrid Modeling. *Neurocomputing*, Vol. 167, 2015, pp. 3–7.
33. Wang, G.-W., C.-X. Zhang, and G. Guo. Investigating the Effect of Randomly Selected Feature Subsets on Bagging and Boosting. *Communications in Statistics-Simulation and Computation*, Vol. 44, No. 3, 2015, pp. 636–646.
34. Chen, T., and C. Guestrin. Xgboost: A Scalable Tree Boosting System. *Proc., 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Association for Computing Machinery, New York, NY, 2016, pp. 785–794. ACM.
35. Yang, Z. S., Y. Wang, and Q. Guan. Short-Term Traffic Flow Prediction Method Based on SVM. *Journal of Jilin University (Engineering and Technology Edition)*, Vol. 36, No. 6, 2006, pp. 881–884.
36. Asif, M. T., J. Dauwels, C. Y. Goh, A. Oran, E. Fathi, M. Xu, M. M. Dhanya, N. Mitrovic, and P. Jaillet. Spatiotemporal Patterns in Large-Scale Traffic Speed Prediction. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 15, No. 2, 2014, pp. 794–804.
37. Fu, R., Z. Zhang, and L. Li. Using LSTM and GRU Neural Network Methods for Traffic Flow Prediction. *Proc., 2016 31st Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, Wuhan, China, 2016, pp. 324–328. IEEE, New York.
38. Zhao, Z., W. Chen, X. Wu, P. C. Y. Chen, and J. Liu. LSTM Network: A Deep Learning Approach for Short-Term Traffic Forecast. *IET Intelligent Transport Systems*, Vol. 11, No. 2, 2017, pp. 68–75.
39. Cui, Z., R. Ke, Z. Pu, and Y. Wang. Deep Bidirectional and Unidirectional LSTM Recurrent Neural Network for Network-Wide Traffic Speed Prediction. *arXiv Preprint arXiv:1801.02143*, 2018.
40. Hochreiter, S., and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, Vol. 9, No. 8, 1997, pp. 1735–1780.
41. Chung, J., C. Gulcehre, K. Cho, and Y. Bengio. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv Preprint arXiv:1412.3555*, 2014.
42. Schuster, M., and K. K. Paliwal. Bidirectional Recurrent Neural Networks. *IEEE Transactions on Signal Processing*, Vol. 45, No. 11, 1997, pp. 2673–2681.
43. Kingma, D. P., and J. Ba. Adam: A Method for Stochastic Optimization. *arXiv Preprint arXiv:1412.6980*, 2014.
44. Chollet, F., et al. Keras, 2015. <https://keras.io>.
45. Bergmeir, C., and J. M. Benítez. On the Use of Cross-Validation for Time Series Predictor Evaluation. *Information Sciences*, Vol. 191, 2012, pp. 192–213.
46. Mann, H. B., and D. R. Whitney. On a Test of Whether One of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, Institute of Mathematical Statistics, 1947, pp. 50–60.
47. Yao, H., X. Tang, H. Wei, G. Zheng, and Z. Li. Revisiting Spatial-Temporal Similarity: A Deep Learning Framework for Traffic Prediction. *Proc., AAAI Conference on Artificial Intelligence*, Honolulu, Hawaii, Vol. 33, 2019, pp. 5668–5675.

48. Shi, X., Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W. Woo. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In *Advances in Neural Information Processing Systems*, MIT Press, Montreal, Canada 2015, pp. 802–810.
49. Zhang, J., Y. Zheng, and D. Qi. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. *Proc., 31st AAAI Conference on Artificial Intelligence*, San Francisco, CA, 2017.
50. Yao, H., F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, and Z. Li. Deep Multi-View Spatial-Temporal Network for Taxi Demand Prediction. *Proc., 32nd AAAI Conference on Artificial Intelligence*, New Orleans, LO 2018.