

Enhancing Traffic Flow Prediction using Outlier-Weighted AutoEncoders

Citation for published version (APA):

Choudhary, H., & Hassani, M. (2023). *Enhancing Traffic Flow Prediction using Outlier-Weighted AutoEncoders: Handling Real-Time Changes*. (arXiv).

Document status and date:

Published: 01/12/2023

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne


Take down policy

If you believe that this document breaches copyright please contact us at:


openaccess@tue.nl

providing details and we will investigate your claim.

Enhancing Traffic Flow Prediction using Outlier-Weighted AutoEncoders: Handling Real-Time Changes

Himanshu Choudhary 

Department of Mathematics and Computer Science
Eindhoven University of Technology
Eindhoven, Netherlands
himanshu.dce12@gmail.com

Marwan Hassani 

Department of Mathematics and Computer Science
Eindhoven University of Technology
Eindhoven, Netherlands
m.hassani@tue.nl

ABSTRACT

In today's urban landscape, traffic congestion poses a critical challenge, especially during outlier scenarios. These outliers can indicate abrupt traffic peaks, drops, or irregular trends, often arising from factors such as accidents, events, or roadwork. Moreover, given the dynamic nature of traffic, the need for real-time traffic modeling also becomes crucial to ensure accurate and up-to-date traffic predictions. To address these challenges, we introduce the Outlier Weighted Autoencoder Modeling (OWAM) framework. OWAM employs autoencoders for local outlier detection and generates correlation scores to assess neighboring traffic's influence. These scores serve as a weighted factor for neighboring sensors, before fusing them into the model. This information enhances the traffic model's performance and supports effective real-time updates, a crucial aspect for capturing dynamic traffic patterns. OWAM demonstrates a favorable trade-off between accuracy and efficiency, rendering it highly suitable for real-world applications. The research findings contribute significantly to the development of more efficient and adaptive traffic prediction models, advancing the field of transportation management for the future. The code and datasets of our framework is publicly available under <https://github.com/himanshudce/OWAM>.

CCS CONCEPTS

• **Information systems** → **Data streaming**.

KEYWORDS

Real-Time Traffic Prediction, Outliers, AutoEncoders, Correlation

1 INTRODUCTION

Traffic congestion is a well-known problem impacting public health, economic losses, and health problems [1]. It also contributes to environmental issues due to increased emissions and air pollution [29]. As increasing road capacity may not always be feasible, optimizing traffic controllers becomes crucial for enhancing traffic flow. These systems can help alleviate congestion, reduce travel times, and enhance overall traffic flow.

Traffic modeling has advanced considerably, resulting in effective traffic prediction models in general. Yet, challenges emerge when these models encounter outlier traffic scenarios. In the realm of traffic, outliers refer to situations where data substantially deviate from the usual traffic patterns. These outliers can indicate abrupt traffic peaks, drops, or irregular trends, often arising from factors such as accidents, events, or roadwork. Furthermore, these traffic prediction models often lack generalizability for real-time

predictions. Traffic is inherently dynamic, experiencing continuous fluctuations due to factors like accidents and special events. Emphasis should be placed on reducing training and evaluation time, ensuring the model's robustness for real-time updates and efficiency in handling real-world conditions.

To illustrate these challenges, let's consider a city, where we aim to predict traffic conditions at Intersection A (Figure 1), and, challenges arise due to abrupt changes in traffic patterns (Figure 1a & 1b). The first case (Figure 1a) presents a clear outlier situation, while another case (Figure 1b) shows a total shift in traffic dynamics. Such instances can lead to long congestion at Intersection A if prediction models fail to accurately anticipate and identify future traffic flows. Current research falls short of offering a holistic solution that can simultaneously manage both outlier scenarios and real-time adjustments. This study seeks to bridge this gap by introducing the Outlier Weighted Autoencoder Modeling (OWAM) framework, a real-time model that aims to improve traffic predictions in both dynamic and outlier traffic scenarios simultaneously.

Building upon the groundwork laid by T. Mertens [26], our research basically aims to identify "relevant" intersections by assessing the influence of outliers from nearby sensors on the target sensor. We begin the process by leveraging Autoencoders (AE) under real-time settings to detect local outliers at each sensor. To enhance AE's performance further, we integrated Earth Mover's distance [2], a well-known distance metric, as a loss function in AE. These localized outlier scores then serve as the basis for calculating correlations between the target sensor and its neighboring counterparts. Subsequently, we merge the neighboring data with the target data, assigning weights based on their correlation scores to ensure precise influence. Furthermore, we incorporate outlier-based model updates through OWAM to effectively address both outlier scenarios and the dynamic changes simultaneously. To check the effectiveness of our proposed architecture we tested it on three real-world datasets, i.e. Hague (Section 5.2.1), METR-LA (Section 5.2.2) and PEMS-BAY (Section 5.2.3).

The OWAM framework demonstrates superior performance compared to previous outlier-based approaches and the LSTM baseline. Though it falls short of surpassing the heavy graph-based techniques in terms of RMSE, it notably stands out in terms of runtime efficiency, rendering it highly suitable for real-world traffic management systems. This work extensively extends the presentation and the experimental study presented in [9]. Overall, the following are the main contributions of our work:

- (1) An innovative Outlier Weighted Autoencoder Modelling (OWAM) framework, that enhances traffic flow prediction in both outlier and real-time scenarios.

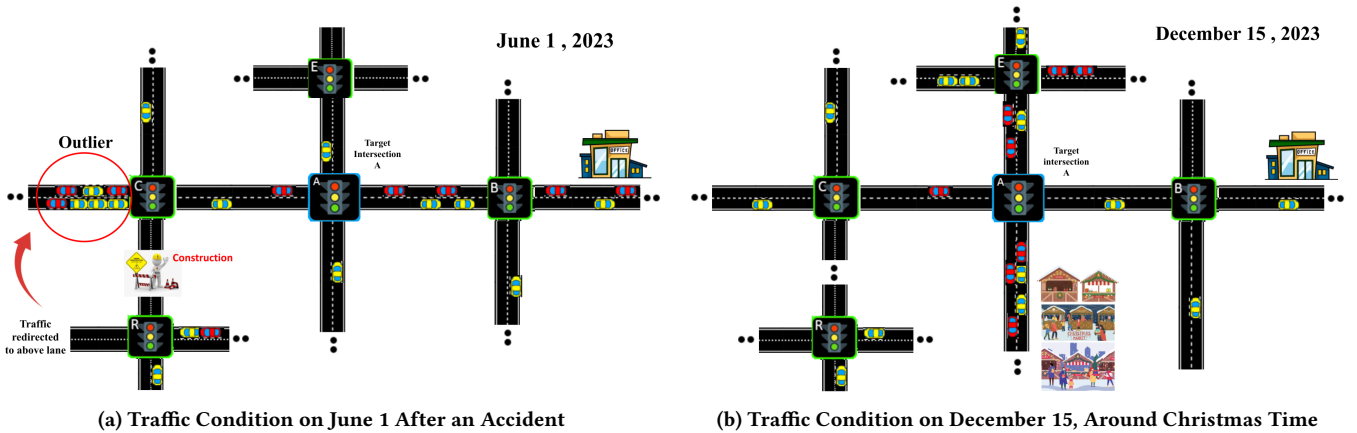


Figure 1: Traffic Conditions of the City on Different Dates

- (2) A novel outlier detection architecture that integrates Earth mover’s distance with Autoencoders to improve the accuracy of outlier identification.
- (3) An extensive experimental evaluation over three real-world datasets to demonstrate the effectiveness of our proposed architecture on traffic prediction tasks utilizing traffic flow and traffic speed.
- (4) Providing an open-source implementation of OWAM framework for reproducibility and further research.

The paper is structured as follows: In Section 2, we review relevant literature on outlier detection techniques and traffic modeling. Section 3 provides essential background and mathematical formulation. Section 4 presents the OWAM framework’s design. Section 5 describes datasets and metrics. Section 6 discuss our findings on traffic prediction. Section 7 summarizes the discussion and outlines future research directions.

2 RELATED WORK

This Section explores relevant literature, focusing on outlier detection techniques and traffic modeling, both of which are central to our research objective. Outlier detection involves identifying data points that significantly deviate from the norm, with various techniques applied across domains. These methods include One-Class Support Vector Machines (OC-SVM) [34], Local Outlier Factors (LOF) [7] and Isolation Forests [24]. However, most of these approaches are not well-suited for real-time applications like traffic flow forecasting due to high response times, scalability issues, and difficulties in handling high-dimensional real-time data. To handle such dynamic tasks, several contributions have been made in the realm of online anomaly detection for data streams. Notable approaches include Half-Space Trees (HST) [35], Robust Random Cut Forest (RRCF) algorithm [16] and the extension of the Local Outlier Factor (LOF) method for data streams [30].

Nevertheless, Autoencoders (AEs) have proven to be the most effective method in both offline and online scenarios. The KitNET algorithm [27] efficiently identifies network intrusions through an ensemble of autoencoders. MemStream [4] introduces a novel approach by combining denoising autoencoders with a memory

module to address concept drift and memory poisoning in real-time outlier detection for streaming data. Additionally, to tackle concentrated bursts and a higher proportion of anomalies in data streams, Probability Weighted Autoencoders (PW-AE) [8] are proposed, where the loss function is optimized based on the probability of an instance being an anomaly. Not only does it handle evolving data streams and concept drift, but it also outperforms previous models in terms of accuracy and efficiency. In this paper, we build upon this promising prior work [8] and further modify it for enhanced outlier detection, as detailed in subsequent sections.

In the domain of traffic modeling, various approaches have been explored to enhance prediction accuracy and efficiency. Broadly categorized, traffic flow prediction models fall into three main groups: parametric techniques, machine learning techniques, and deep learning techniques [25]. Parametric techniques use statistical methods with parameters to model data like Kalman filters [13] and ARIMA models [6] for time series analysis. However, these methods may not capture complex data patterns as effectively as machine learning or deep learning techniques. Machine learning techniques, such as Support Vector Machines (SVM) [28] and K-nearest neighbors (KNN) [19], consider spatial characteristics for congestion prediction. Some incorporate dynamic distance measures and state matrices for accuracy improvement. Deep learning employs neural networks with multiple layers to learn complex data patterns. They extract space-time dependencies from heterogeneous data sources, enhancing accuracy. For example, Convolution Neural Networks (CNNs) extract spatial features from traffic images, while LSTMs capture temporal dependencies in traffic flow data. Encoders reduce data dimensionality and Graph Convolution Networks (GCNs) model relationships between regions. Moreover, hybrid models, combining multiple techniques also boost performance [21, 23].

Outlier-based hybrid approaches offer valuable insights into traffic modeling by focusing on outlier detection in traffic flow data. Polson et al. [31] demonstrated the effectiveness of neural networks in handling outlier scenarios in traffic data. On the other hand, LOF-based models have been proposed to detect outliers within Probability Distributions of traffic flows (FPDs) [11, 12, 14]. Previous studies have also investigated different FPD-based distance

metrics for effective outlier detection, with Earth Mover’s distance [2] proving to be the most effective.

Prior outlier-based approaches have been constrained by traditional outlier detection methods, predominantly tested on similar datasets, and have been limited in their ability to intelligently integrate information. Additionally, researchers largely overlooked the real-time aspect of traffic, where traffic consistently evolves and adapts patterns in response to various ongoing events, a critical aspect demanding further investigation. Regular model updates are vital to ensure accurate and up-to-date traffic models that reflect real-time changes in traffic flow. These models should strike a balance between accuracy and efficiency to facilitate real-world scenarios.

3 PRELIMINARIES AND PROBLEM FORMULATION

Traffic controllers (*TC*) manage traffic signal timing at intersections, optimizing various traffic movements. They rely on data from multiple traffic monitoring and management (*TM*) systems, which collect real-time traffic information at each sensor l , including vehicle volume, speed, and occupancy. Our objective is to enhance the prediction accuracy of the traffic management systems TM_l , assisting *TCs* in decision-making and optimizing traffic flow. We trained a model M with both offline and online objectives for comparison and real-time applicability assessment.

Let $L^+ = \{(x(i), y(i)) \mid \forall i \in S\}$ be a sequence of S tuples, where each tuple consists of a feature-vector $x(i) \in R^d$ and a corresponding real value to be predicted $y(i) \in R$. Additionally, let $L_{train} = \{(x(i), y(i)) \mid \forall i \in \{1, \dots, P\}\}$ be a sub-sequence used for training the model M . Let $L_{test} = \{(x(i), y(i)) \mid \forall i \in \{1, \dots, m\}\}$ be the sub-sequence used for testing the model.

Here, L_{train} and L_{test} are disjoint subsets of L^+ . Given a function $f : R^d \rightarrow R$ that was trained on L_{train} with P samples, the objective can be formulated with $min(z_1)$ such that -

$$z_1 = \frac{1}{N} \sum_{i=1}^N (\mathcal{L}_r, L_{test}, f(L_{train})) \quad (1)$$

Where N denotes the total number of sensors and \mathcal{L}_r represents a regression loss.

In addition to the regular (offline) settings, we incorporated a real-time incremental updating strategy [32] with time windows T of different lengths. Let L_{train} contain the data up to time point t , and $L_{test} = \{(x(i), y(i)) \mid \forall i \in \{t+1, \dots, t_n\}\}$ represent infinite sub-sequences with interval length T . For a given validation protocol \mathcal{V} , the problem can be formulated by $min(z_2)$ such that -

$$z_2 = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{t_n} \mathcal{V} (\mathcal{L}_r, L_{test}, f(L_{train})) \quad (2)$$

For the validation protocol \mathcal{V} , we employ the test-then-train [5] approach. In this evaluation method, models are first validated based on a performance metric and subsequently trained on each instance within the time window T of L_{test} at its arrival time.

To achieve this objective, we introduce the OWAM (Outlier Weighted Autoencoder Model) framework. We outline its key components in details below for through understanding.

Flow Probability Distribution (FPD). Flow Probability Distributions (FPDs) [26] represent the probability of traffic occurrences within specific time intervals, derived from real traffic data. FPDs encapsulate traffic distribution patterns and can be created for vehicle flow, average speed or any relevant traffic data, to get a comprehensive understanding of traffic dynamics.

Let us denote the FPD for a specific sensor B during the time period H as $FPD(HB)$. To calculate $FPD(HB)$, we start with a collection of aggregated traffic flow values, denoted as X_{HB} :

$$X_{HB} = (x_{hB1}, x_{hB2}, \dots, x_{hBH}) \quad (3)$$

Here, x_{hB} represents the aggregated traffic flow or speed value for the intersection B during the h -th time interval, with h ranging from 1 to H .

To calculate the FPDs, we apply the formula as created by Zimek and Djenouri in [7]:

$$FPD(HB) = \frac{\sum_{x_h \in X_H} P(x_h)}{P(x_{hB})}, \quad \forall x_{hB} \in X_{HB}, \quad h = 1, \dots, H \quad (4)$$

Outlier Detection using Autoencoders. This research explores various outlier detection techniques, including well-established methods like Half-Space-Tree (HST) and Kit-net (see Section 2). However, our primary focus is on the Autoencoders, recognized for their superior online anomaly detection capabilities.




Figure 2: Autoencoders (Online): Outlier Detection

Autoencoders, commonly used for anomaly detection, aim to reconstruct \hat{D} input data D by encoding it into a lower-dimensional representation and then decoding it back to its original format (Equation 5). Anomalies, deviating significantly from the training data, result in higher reconstruction errors/loss (Equation 6). As shown in Figure 2, our research includes training AE on continuous traffic stream of FPDs in real-time and if faced with an anomaly, say, characterized by a left-skewed distribution with substantial magnitude (Figure 2), the model generates a higher reconstruction loss, indicating the presence of an outlier.

$$\hat{D} = Decoder(Encoder(D)) \quad (5)$$

$$Reconstruction Loss(\mathcal{L}) = \mathcal{L}_r / e(\hat{D}, D) \quad (6)$$

To train the autoencoder, we aim to minimize the reconstruction loss (\mathcal{L}_r), typically measured using Root Mean Square Error

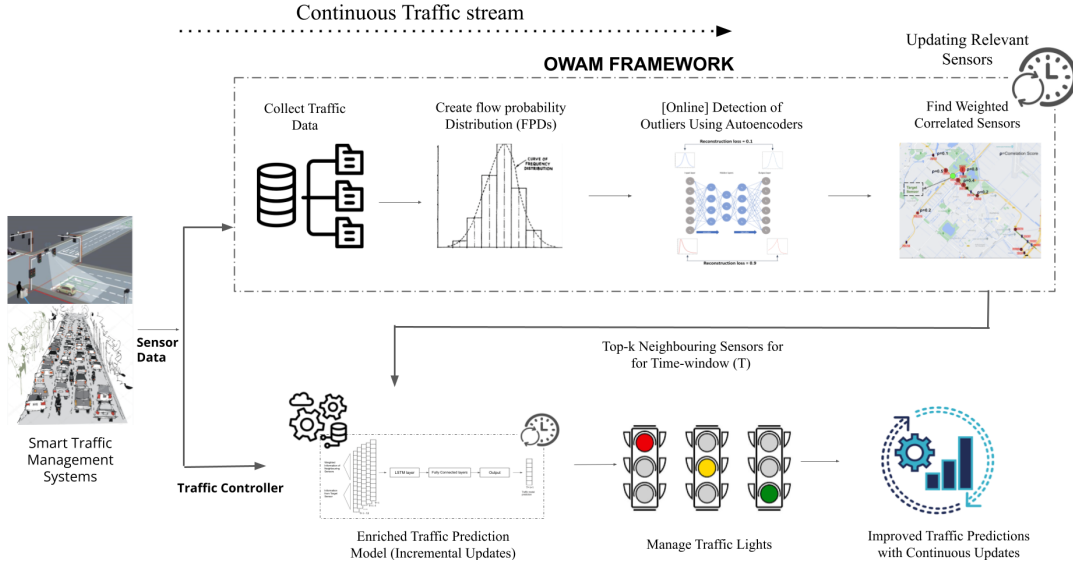


Figure 3: Overview of OWAM Framework with Real-Time Updates

(RMSE). Building upon the previous work by Bazan et al. [3], where they found that Earth Mover’s distance (EMDs) exhibited superior performance in comparing two distributions, we utilized this knowledge and integrated EMD as a loss function \mathcal{L}_e in Autoencoders as can be depicted in the Figure 2, where loss can be used either as RMSE (\mathcal{L}_r) or EMD (\mathcal{L}_e).

Earth’s Mover Distance (EMD) [2] is well-suited for capturing complex differences between entire distributions, in contrast to Euclidean distance (RMSE), which emphasizes point-wise errors. EMD quantifies the minimum effort required to transform one distribution into another, considering their overall structure, and is less influenced by individual extreme points.

Finally, this research incorporates the Long Short-Term Memory (LSTM) model [18], a widely used sequence prediction model, in its existing form. It is important to highlight that our proposed framework can be integrated with any traffic modeling framework, as we do not alter the original model but enhance its performance by passing only crucial information.

4 METHODOLOGY

This section presents a comprehensive overview of our novel framework, Outlier Weightage Autoencoder Modeling (OWAM), along with an elaboration on the rationale behind its design choices.

Figure 3 outlines our comprehensive traffic modeling approach. Initially, traffic data is collected at 5-minute intervals. Subsequently, Flow Probability Distributions (FPDs) are generated from this data, and these distributions undergo online AE-based outlier detection to compute local anomaly scores for each sensor. A Pearson’s correlation coefficient is then employed to calculate the correlation score between the target sensor and its neighboring sensors to determine their influence. These correlation scores are used as weights to adjust the importance of the neighboring sensor information before integrating it with the data from the target sensor. This integrated information is subsequently fed into an LSTM model for efficient

traffic prediction. It is important to note that not only LSTM but any other modeling framework can benefit from the weighted information generated by OWAM. To address real-time changes, we continually recalculate outlier information using online AE for every T time window and incrementally update the LSTM model with this revised information. Each component is discussed below to provide a comprehensive understanding of the OWAM framework.

4.0.1 Creating Flow Probability Distribution (FPDs). The first step involves creating Flow Probability Distributions (FPDs) from the collected data (Equation 4), which offers several advantages. FPDs help smooth out short-term fluctuations and noise, simplifying analysis, and capturing higher-level traffic patterns and trends. They enhance the identification of peak hours, daily variations, and weekly trends, providing a better understanding of traffic dynamics. Additionally, FPDs are more robust to outliers compared to individual 5-minute intervals, improving analysis reliability. Choosing the right time window for FPDs is crucial, balancing pattern capture and noise reduction. Very short intervals can introduce excessive noise, while very long intervals may over smooth data, potentially missing patterns.

Figure 4 illustrates these interpretations. Figure 4a displays data aggregated every 5 minutes with noticeable noise and fluctuations. Although a 15-minute interval aggregation (Figure 4b) reduces some noise, it still remains. In contrast, the 1-hour window (Figure 4c) makes it easier to identify outliers and patterns. So we aggregated data into 1-hour FPDs using 5-minute intervals, for optimizing outlier detection.

4.0.2 Autoencoders for Outlier Detection. In the second step, the continuous stream of Flow Probability Distributions (FPDs) undergoes online anomaly detection to identify outliers at each sensor. Utilizing online autoencoders, which adapt to incoming data in real-time, allows continuous updates of weights and parameters to capture evolving patterns and anomalies without the need for

retraining on the entire dataset. Probability-weighted autoencoders (PW-AE) were primarily employed for their superior performance, as detailed in Section 2. Additionally, we introduce a novel PW-AE-EMD model, integrating Earth Mover’s distance (EMD) loss (\mathcal{L}_e) to measure dissimilarity between the input FPD (D) and the reconstructed FPD (\hat{D}) for enhanced outlier detection (Equation 5 & 6).

We preserve the outlier scores for each sensor without applying additional threshold filtering. The use of filtering can alter the correlation strength by discarding valuable information about how neighboring sensors were influenced by outliers and their subsequent return to normal behavior. By maintaining the raw outlier scores, we gain a comprehensive understanding of the dynamics of outliers and their effects on nearby sensors. This approach enhances our analysis of traffic patterns and behavior.

4.0.3 Weighted Correlation Analysis for Relevant Sensor Identification. To proceed, it is essential to have target sensors for correlation analysis and prediction. In a real-world scenario, we would need to predict each sensor individually, making every sensor a potential target. However, due to computational constraints, we opted for a more manageable approach by selecting representative samples from the data.

For the Hague dataset (refer to Section 5.2.1), we employed the same target sensors as used in a previous outlier-based study [26] due to their strategic positioning on the lane. However, in the case of the remaining two datasets i.e. METR-LA (Section 5.2.2), and PEMS-BAY (Section 5.2.3), the exact sensor locations were unavailable, making it impossible to apply this domain-specific knowledge. Consequently, in these cases, we randomly selected five target sensors from each dataset to act as representatives for our study.

For each target sensor, we compute correlation scores with other sensors using their outlier scores. Pearson’s correlation coefficient serves as the metric for this purpose. While, in an ideal scenario, we could consider introducing a time lag to account for traffic propagation, but the datasets used in our study contained closely situated sensors. Consequently, traffic could traverse between them within relatively short time intervals. In the Hague dataset, for instance, the average travel time between any two sensors was consistently less than 1 hour. Although we lacked precise travel time information for the PEMS-BAY and METR-LA datasets, available distance data suggested their relative proximity. In fact, the maximum distance between the farthest sensors in these datasets was approximately 20 km.

Utilizing correlation scores, we introduced the concept of outlier-weighted correlation integration. Instead of directly incorporating information from correlated neighboring sensors, we assign weights based on their correlation coefficients. These weights quantify the influence of each neighboring sensor on the target sensor. In Figure 5a, the green circle represents the target sensor, while the neighboring sensors are depicted as red circles with varying intensities corresponding to their correlation scores. When transmitting information from these sensors, along with the target sensor’s data, to the LSTM model, we apply a weighted approach. Specifically, we multiply the transmitted information by weights determined by

their respective correlation scores. This weighting mechanism allows the model to emphasize more influential sensors based on their outlier characteristics, resulting in enhanced traffic predictions.

4.0.4 Traffic Prediction. After obtaining the weighted correlation information from neighboring sensors, we concatenated their data with the target sensor’s data and inputted it into the LSTM model, as illustrated in Figure 5b, for predicting the traffic flow of the target sensor. As mentioned earlier, this weighted information from neighboring sensors can be utilized by any traffic prediction model, not exclusive to LSTM. To explore the impact of dimensionality, our experiments involved various sensor subsets, ranging from selecting only the top k sensors to utilizing all neighboring sensors. The threshold (θ) determined the selection of the top $k\%$ of sensors. For instance, choosing the top 5% ($\theta = 0.05$) in the METR-LA dataset with 207 sensors implied selecting the top 10 neighboring sensors for each target sensor.

In offline settings, we partitioned the data into an 80/20 ratio for model training and evaluation. The LSTM model processed the data sequentially, with input vectors (x_i) representing the previous 1-hour window (12 timestamps, 5-minute intervals), and output vectors (y_i) representing the subsequent 5-minute traffic situation (Speed/Number of vehicles) to be predicted. We conducted the comparative analysis in offline settings for fair comparison purposes, primarily because there were no online versions available for the other models being evaluated.




Figure 4: Aggregation Plot at 5, 15, and 60 Minutes - Sensor K504

4.0.5 Real-Time Traffic Modelling. To accommodate the dynamic nature of traffic, we have implemented an incremental updating strategy, as depicted in Figure 3. Correlated sensors are periodically refreshed at specific intervals referred to as time windows (T), and this updated information is utilized for incremental updates to the LSTM model. Initially, the base model is trained using the first 50% of the data, with the target sensor having access to this historical data up to time point t . Subsequently, we incrementally update the model with the next t_n time windows of interval T . For each window at $t + 1$, we obtain the outlier scores using online AE (Section 4.0.2), recalculate the correlation score (weights), and adjust the significance of neighboring sensors with incoming data. Prior to updating the LSTM model, we evaluate its performance and then train it with the data from the current window, following the test-then-train methodology [5].

We experimented with various time windows, ranging from 1 hour to 1 month, to assess both short-term and long-term impacts. We compare three scenarios: outlier-weighted updates (incremental updates using OWAM), standard real-time updates (incremental updates without OWAM, keeping neighboring sensors fixed), and an




Figure 5: Correlated Weighted Sensor and LSTM Overview

offline model with no updates post-training. Through this comparison, our goal is to demonstrate the superiority of real-time updates and underscore the advantages of incorporating outlier-weighted correlations in traffic prediction.

5 EXPERIMENTAL SETUP

This section presents the following key components: performance metrics, an overview of the datasets along with their statistics, and system details. These elements serve as the foundation for conducting our experiments.

5.1 Performance Metric

5.1.1 RMSE Metric. We employed the Root Mean Squared Error (RMSE) metric, a common measure for regression problem evaluation. RMSE quantifies the average difference between predicted values (\hat{y}_i) and true values (y_i) in a dataset. It is computed as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}.$$

5.1.2 Temporal Analysis. In addition to score metrics, we investigated model convergence and prediction speed during training and testing. Understanding how quickly models learn and make predictions is vital for real-time applications, where timely decisions are essential. Assessing convergence time also gauges efficiency and effectiveness in learning from data.

5.2 Dataset

This research utilizes three datasets to investigate traffic patterns: The Hague dataset, which provides information on the aggregated number of vehicles every 5 minutes, and two widely referenced datasets in the literature, METR-LA, and PEMS-BAY, which capture

the average speed of traffic in 5-minute intervals. All three datasets considered as traffic flow measurements in this study.

5.2.1 Hague Dataset. The Hague dataset, obtained from 23 intersections in The Hague, Netherlands, spans from January 1, 2018, to March 31, 2020, with 5-minute intervals (Table 1). The dataset had minimal missing values (less than 0.01%), which were imputed with linear interpolation. Sensor distances and average travel times were obtained from the Google Maps API and used as input for Graph Neural Networks [22].

5.2.2 METR-LA Dataset. The METR-LA dataset is a widely used benchmark dataset for traffic modeling and forecasting. The data is gathered from the highway of Los Angeles containing 207 selected sensors indicating speed in miles per hour and ranging from March 1st, 2012 to June 30th, 2012 as indicated in Table 1. It covers multiple sensor locations across different highways, providing a diverse and representative set of traffic patterns for analysis.

5.2.3 PEMS-BAY Dataset. Similar to the METR-LA dataset, the PEMS-BAY is also a widely used benchmark dataset for traffic modeling and forecasting. This public traffic speed dataset collected by California Transportation Agencies (CalTrans) contains 325 sensors in the Bay Area ranging from Jan 1st 2017 to May 31th 2017 denoting speed in miles/hour (Table 1).

Table 1: Dataset Information

Dataset	Hague	METR-LA	PEMS-BAY
#Samples	236,736	34,272	52,116
#Sensors	23	207	325
Sample Rate	5 minutes	5 minutes	5 minutes
Time Range	2 Years	4 months	6 months
Traffic Indicator	#Vehicles	Speed (mi/h)	Speed (mi/h)


In order to get higher-level insights into these datasets, we also computed the correlations among all pairs of sensors and examined their distributions. The results revealed significant inter-node (spatial) correlations in both METR-LA and Hague datasets, while the correlations in PEMS-BAY dataset are notably weaker. This observation is depicted in Figure 6d, 6e and 6f.

By analyzing the velocity distributions presented in Figure 6c, 6b, and 6a, notable differences can be observed among the datasets. PEMS-BAY dataset demonstrates a more uniform and closer-to-free-flow velocity distribution compared to METR-LA. This observation aligns with the lower inter-node correlations depicted in Figure 6f, suggesting simpler and non-dependent traffic conditions in PEMS-BAY as compared to METR-LA. Furthermore, the vehicle distribution in Hague dataset reveals that most intersections experience low volumes of vehicles at any given time, indicating light and smooth traffic conditions in the area.

5.3 System Details

The experiments were performed on a Mac M1 machine with 16 GB RAM running on macOS 13. To leverage GPU capabilities, the experiments utilized the MPS backend. Python 3.10 served as the programming language for conducting the experiments.

Figure 6: Distributions of Speed, Number of Vehicles, and Inter-node Correlations



6 RESULTS

To evaluate OWAM’s performance comprehensively, we employed several criteria. Firstly, we investigated the impact of OWAM on the dimensionality of the traffic prediction model. Secondly, we examined the integration of Earth’s Mover Distance loss in Autoencoders and its potential to enhance model predictions. Thirdly, we compared OWAM’s performance with that of an LSTM Baseline and established state-of-the-art techniques. Lastly, we assessed OWAM’s effectiveness in handling real-time scenarios, specifically its capability to adapt to dynamic traffic changes.

When considering dimensionality in traffic flow prediction, the challenge is to determine the optimal number of sensors to include in the model for predicting traffic flow at the target sensor. Figure 7b illustrates the RMSE scores obtained from the LSTM-based OWAM framework for three datasets. Each dataset exhibits an optimal threshold value for achieving peak performance. The Hague dataset, comprising only 8 to 15 sensors, shows consistent performance when including the top 50% or all neighboring sensors, making the dimensionality impact less pronounced. In contrast, METR-LA and PEMS-BAY datasets, which feature more sensors, experience performance degradation when considering all sensors. Instead, incorporating the top 10% and 50% of neighboring sensors yields improved predictions in these datasets. This highlights the dataset-specific nature of sensor selection.

It’s worth noting that utilizing only the target sensor ($\theta = 0$) or incorporating all available sensor information ($\theta = 1$) results in lower prediction performance, as indicated by higher RMSE scores, across all datasets. This underscores the importance of selecting an optimal number of correlated sensors for both prediction accuracy and model efficiency, enabling the model to focus on influential factors while minimizing noise by including a limited subset of relevant sensors.

Regarding training time (Figure 7a), there is no consistent trend observed across datasets. However, an increase in training time is noticeable for the higher-dimensional PEMS-BAY dataset. It is

important to mention that we use model convergence time instead of total training time due to computational limitations, as this choice aligns with efficiency considerations.

Table 2: Comparison of EMD Loss and RMSE Loss on Integration with Autoencoders

Model	Dataset	Train Time	RMSE
OWAM (EMD Loss)	Hague	677.739	13.479
	METR-LA	37.903	7.525
	PEMS-BAY	313.028	4.915
OWAM (RMSE Loss)	Hague	245.465	13.539
	METR-LA	81.914	7.918
	PEMS-BAY	289.274	5.537

Table 3: Distribution of Outlier Score for RMSE and EMD Loss for a Random Sensor

	Mean	Standard Deviation
EMD Loss	0.70	0.13
RMSE Loss	0.81	0.11

Table 2 highlights the advantages of using Earth Mover’s Distance (EMD) as the loss function. The OWAM model with EMD loss consistently outperforms the model relying on the RMSE loss (Euclidean Distance) function. While the improvements may not be substantial in the Hague scenario, an overall enhancement is still evident. The RMSE score decreases by 5% and 11% for METR-LA and PEMS-BAY, respectively. Although no clear pattern emerges in training times, it’s important to note that RMSE is computationally more efficient, featuring $O(n)$ time complexity, whereas EMD involves solving a linear programming problem, often less efficient than $O(n)$, especially for large datasets.

To evaluate the advantages of incorporating EMD loss into the outlier model, refer to Table 3. When EMD loss is used, it leads to a




Figure 7: Model Training Time and RMSE Score for Different Thresholds for Hague, METR-LA, and PEMS-BAY Datasets

higher standard deviation in outlier scores compared to RMSE loss. This increased variance signifies the model’s capacity to differentiate highly relevant sensors from less relevant ones, allowing it to accurately capture complex data patterns. Conversely, a lower standard deviation indicates less differentiation among sensors, potentially resulting in the oversight of important correlations and data patterns.

Table 4: Results for the Hague dataset

Model Name	RMSE	Train Time (s)	Instance Pred Time (ms)	Eval Time (s)
OWAM	13.479	677.74	1.006	47.617
HST[35]	13.631	244.31	2.14	101.32
Kit-Net[27]	13.540	305.37	1.019	48.267
DGCRN[22]	13.110	5043.31	5.91	280.00
OBIS[26]	13.547	682.81	1.08	51.30
LSTM Baseline	13.547	682.81	0.63	29.83

Table 5: Results for the METR-LA dataset

Model Name	RMSE	Train Time (s)	Instance Pred Time (ms)	Eval Time (s)
OWAM	5.880	46.18	1.18	8.09
HST[35]	6.815	47.53	3.99	27.37
Kit-Net[27]	5.883	55.44	1.69	11.58
DGCRN[22]	4.161	11263.23	45.23	310.04
OBIS[26]	7.421	91.37	1.57	10.78
LSTM Baseline	8.831	106.06	1.62	11.11

Table 7 offers a comprehensive overview of the features associated with the compared approaches. The "LSTM Baseline" represents the standard modeling approach without any specific preprocessing steps. "DGCRN" (Diffusion Graph Convolution Recurrent Network) [22] stands as a representative state-of-the-art graph-based approach in traffic analysis. The choice of DGCRN is driven by its recognition in the literature and the availability of open-source code, enabling valuable insights into the contributions of outlier-based research compared to cutting-edge graph-based techniques in traffic management. This comparison aids in assessing the

Table 6: Results for the PEMS-BAY dataset

Model Name	RMSE	Train Time (s)	Instance Pred Time (ms)	Eval Time (s)
OWAM	4.036	303.94	14.23	97.56
HST[35]	4.816	372.86	16.15	110.73
Kit-Net[27]	4.312	356.96	15.23	104.393
DGCRN[22]	1.591	25900.67	113.79	780.04
OBIS[26]	4.375	377.49	14.48	99.27
LSTM Baseline	5.434	405.05	16.30	111.74

practical potential of our framework in intelligent traffic management systems. "OBIS" [26] represents the latest advancements in outlier-based approaches. Furthermore, we also compared OWAM to other online outlier-based methods (HST and Kit-Net), where sensors are selected similarly based on weighted correlation with only differentiation being the outlier technique. It should be noted that all state-of-the-art and other methods, including DGCRN, were applied using the same settings as OWAM to ensure a fair comparison. This approach guarantees the reproducibility of the results and provides a clear basis for evaluating OWAM’s performance relative to other models.

Tables 4, 5, and 6 provide summaries of the comparative analysis among different models. OWAM demonstrates superior performance compared to the previously developed outlier-based model and the Baseline LSTM in terms of RMSE scores. While the difference may not be highly pronounced for the Hague scenario, OWAM still enhances the average RMSE by 10% compared to OBIS and 12% compared to Baseline LSTM. Additionally, OWAM achieves similar or better training and evaluation times compared to OBIS and 15% less training time than Baseline LSTM. Furthermore, OWAM surpasses both HST and Kit-Net in terms of RMSE and prediction time, attributed to its advanced outlier detection techniques and the weighted inclusion of top neighboring sensors. These enhancements contribute to improved prediction accuracy and computational efficiency.

OWAM delivers competitive performance to DGCRN in the Hague dataset, which boasts a relatively large sample size but fewer sensors. However, in the METR-LA and PEMS-BAY datasets, featuring significantly more sensors but relatively smaller sample sizes,

Table 7: Feature Comparison of Traffic Prediction Models

Model	Spatial Dependency	Temporal Dependency	Dimensionality Reduction	Neighboring Sensor Integration
LSTM Baseline	None	LSTM	No	Binary
DGCRN [22]	Distance Matrix + Graph Convolutions	GRU	No	Binary
OBIS [26]	Outlier Correlation (LOF)	LSTM	Yes	Binary
OWAM	Outlier Correlation (AE)	LSTM	Yes	Weighted

DGCRN notably outperforms OWAM in terms of RMSE. This performance gap highlights DGCRN’s prowess in handling complex traffic networks. DGCRN’s superiority can be attributed to its effective utilization of spatial dependencies within the traffic network, leveraging sensor distances to construct adjacency matrices. In contrast, OWAM and the previous outlier-based model do not fully exploit sensor distances. However, DGCRN’s advantages come at the cost of longer training times, as it requires a substantial amount of time to converge compared to OWAM and other models. Furthermore, Graph Neural Networks (GNNs) demand significantly higher memory consumption and a more extensive runtime than other models, making their prediction phase considerably longer. This limits their suitability for real-time traffic applications where prompt decisions are essential.

The above comprehensive analysis indicates OWAM’s superior performance compared to previous outlier-based and baseline techniques, coupled with significantly faster processing times than GNN-based techniques like DGCRN, making it the preferred choice for real-world applications. This speed is crucial for scalability and timeliness in large-scale, real-time traffic modeling, where rapid data processing and decision-making are essential. OWAM’s ability to efficiently manage larger datasets and more complex scenarios with less computational burden than DGCRN positions it as a scalable, time-efficient solution for traffic management and similar real-time applications.

Regarding real-time (online) settings, Figure 8 demonstrates that dynamic updates based on outlier-weighted correlations (incremental updates using OWAM) significantly enhance the model’s performance compared to static updates (incremental updates without OWAM, keeping neighboring sensors fixed) and no update models. This improvement is especially pronounced for very short (1-hour) or longer time windows (weekly or monthly), indicating their effectiveness in adapting to extended outlier situations and improving predictions. Contrasting results for the hague dataset at a 1-hour window indicate very small and frequent fluctuations in the traffic traffic.

The selection of an optimal update window is crucial in real-time modeling, influenced by specific requirements and computational constraints. While it might be assumed that smaller time windows always lead to better performance and improved model understanding due to more frequent updates, this is not universally true. In our experiments (Figure 8), we observed that updating the model every 3, 6, or 12 hours results in a decrease in performance compared to not updating it at all. This can be attributed to changes in traffic before even adapting to it. Therefore, it is essential to analyze the

data and study the patterns before selecting a specific time window for model updates.

Overall, OWAM offers a promising solution for traffic flow prediction, demonstrating competitive performance when compared to baseline and state-of-the-art models. Its advanced outlier detection techniques and the weighted inclusion of top neighboring sensors contribute to enhanced accuracy and computational efficiency. However, the selection of the ideal model should be tailored to specific application requirements. Graph Neural Networks (GNNs) like DGCRN excel in accuracy-driven tasks, while OWAM provides a more efficient option for real-time applications.

7 CONCLUSION AND FUTURE WORK

In conclusion, this paper introduces the Outlier Weighted Autoencoder Modeling (OWAM) framework, representing a pioneering approach to outlier-driven traffic modeling that addresses crucial gaps in the field. OWAM effectively integrates advanced outlier techniques into traffic modeling, intelligently incorporates outlier insights, and prioritizes real-time modeling. By combining autoencoders with Earth Mover’s Distance loss, it enhances outlier detection and seamlessly incorporates anomaly information into traffic prediction models. Additionally, OWAM provides dynamic model updates for efficient real-time traffic modeling. Experimental findings showcase OWAM’s remarkable ability to significantly improve prediction accuracy. It consistently outperforms other models, achieving an average RMSE score improvement of 10% compared to OBIS and 12% compared to the baseline LSTM model. While Graph Neural Networks (GNNs) demonstrate robust RMSE performance, OWAM excels in computational efficiency, positioning it as an attractive choice for real-time applications. The versatility and adaptability of OWAM make it a promising tool for real-world traffic analysis and management systems. Its contributions extend to the advancement of traffic modeling techniques and the enhancement of our understanding of urban traffic dynamics.

Future research should explore the use of more intricate datasets, such as NE-BJ [22], to better capture the complexities of urban traffic patterns. Conducting comparative analyses with a variety of traffic modeling approaches, including recent advanced Graph Neural Networks (GNNs) techniques [10], [15], can further enhance our understanding of traffic prediction. Efforts to enhance the interpretability of GNNs are essential for making informed data-driven decisions in traffic management. The incorporation of elements from GNN architecture can be considered to achieve an equilibrium between the superior accuracy of GNNs and the rapid processing capabilities of outlier-based models. Additionally, merging weight




Figure 8: Comparison Between Real-Time Dynamic and Static Updates over Different Datasets

allocation and model training into one unified network through the use of an attention mechanism could yield a robust and interpretable model. Moreover, OWAM’s applicability extends to various domains that require anomaly detection and real-time responsiveness after immediate concept drift detection [17, 20, 33], such as environmental monitoring, fraud detection [36], and healthcare.

REFERENCES

- [1] Tanzina Afrin and Nita Yodo. 2020. A Survey of Road Traffic Congestion Measures towards a Sustainable and Resilient Transportation System. (2020).
- [2] Erik Andersen, Marco Chiarandini, Marwan Hassani, Stefan Jänicke, Panagiotis Tampakis, and Arthur Zimek. 2022. Evaluation of Probability Distribution Distance Metrics in Traffic Flow Outlier Detection. In *IEEE MDM*. <https://doi.org/10.1109/MDM55031.2022.00030>
- [3] Eric Bazan, Petr Dokladal, and Eva Dokladalova. 2019. Quantitative Analysis of Similarity Measures of Distributions. In *BMVC*. <https://api.semanticscholar.org/CorpusID:68048254>
- [4] Siddharth Bhatia, Arijit Jain, Shivin Srivastava, Kenji Kawaguchi, and Bryan Hooi. 2022. MemStream: Memory-Based Anomaly Detection in Multi-Aspect Streams with Concept Drift. (2022). <https://dl.acm.org/doi/fullHtml/10.1145/3485447.3512221>
- [5] Albert Bifet, Geoff Holmes, Bernhard Pfahringer, and Andreas Pfisterer. 2018. *Machine Learning for Data Streams: With Practical Examples in MOA*. MIT Press.
- [6] G.E.P. Box and G.M. Jenkins. 1970. *Time Series Analysis: Forecasting and Control*. Holden-Day.
- [7] Markus M. Breunig, Hans-Peter Kriegel, Raymond T. Ng, and Jörg Sander. 2000. LOF: Identifying Density-Based Local Outliers. (2000). <https://doi.org/10.1145/335191.335388>
- [8] Lucas Cazzonelli and Cedric Kulbach. 2023. Detecting Anomalies with Autoencoders on Data Streams. In *ECML PKDD*.
- [9] Himanshu Choudhary and Marwan Hassani. 2024. Autoencoder-based Continual Outlier Correlation Detection for Real-Time Traffic Flow Prediction. In *the 39th ACM/SIGAPP Symposium on Applied Computing SAC '24*. to appear. <https://doi.org/10.1145/3605098.3636162>
- [10] Shaojie Dai, Jinshuai Wang, Chao Huang, Yanwei Yu, and Junyu Dong. 2023. Dynamic Multi-View Graph Neural Networks for Citywide Traffic Inference. *ACM Trans. Knowl. Discov. Data* 17, 4 (2023). <https://doi.org/10.1145/3564754>
- [11] Youcef Djenouri and Arthur Zimek. 2018. Outlier Detection in Urban Traffic Data. In *WIMS*. ACM. <https://doi.org/10.1145/3227609.3227692>
- [12] Youcef Djenouri, Arthur Zimek, and Marco Chiarandini. 2018. Outlier Detection in Urban Traffic Flow Distributions. In *ICDM*. <https://doi.org/10.1109/ICDM.2018.00114>
- [13] Azadeh Emami, Majid Sarvi, and Saeed Asadi Bagloee. 2020. Short-term traffic flow prediction based on faded memory Kalman filter fusing data from connected vehicles and bluetooth sensors. (2020).
- [14] Wesley Fitters, Alfredo Cuzzocrea, and Marwan Hassani. 2021. Enhancing LSTM Prediction of Vehicle Traffic Flow Data via Outlier Correlations. In *COMPSAC*. 210–217. <https://doi.org/10.1109/COMPSAC51774.2021.00039>
- [15] Junhua Gu, Zhihao Jia, Taotao Cai, Xiangyu Song, and Adnan Mahmood. 2023. Dynamic Correlation Adjacency-Matrix-Based Graph Neural Networks for Traffic Flow Prediction. *Sensors* 23, 6 (2023).
- [16] Sudipto Guha, Nina Mishra, Gourav Roy, and Okke Schrijvers. 2016. Robust Random Cut Forest Based Anomaly Detection on Streams. In *ICML*. JMLR.
- [17] Marwan Hassani. 2019. Concept Drift Detection Of Event Streams Using An Adaptive Window. In *ECMS*. 230–239. <https://www.scs-europe.net/dlib/2019/2019-0230.htm>
- [18] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. (1997). <https://doi.org/10.1162/neco.1997.9.8.1735>
- [19] Victoria Hodge and Jim Austin. 2004. A Survey of Outlier Detection Methodologies. (2004). <https://doi.org/10.1023/B:AIRE.0000045502.10941.a9>
- [20] Jesús Huete, Abdulhakim Ali Qahtan, and Marwan Hassani. 2023. PrefixCDD: Effective Online Concept Drift Detection over Event Streams using Prefix Trees. In *COMPSAC*. 328–333. <https://ieeexplore.ieee.org/document/10196975>
- [21] Georgios N Kouziokas. 2021. Deep bidirectional and unidirectional LSTM neural networks in traffic flow forecasting from environmental factors. In *CSUM*.
- [22] Fuxian Li, Jie Feng, Huan Yan, Guangyin Jin, Depeng Jin, and Yong Li. 2022. Dynamic Graph Convolutional Recurrent Network for Traffic Prediction: Benchmark and Solution. (2022). <https://dl.acm.org/doi/10.1145/3532611>
- [23] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *ICLR*.
- [24] Fei Tony Liu, Kai Ting, and Zhi-Hua Zhou. 2009. Isolation Forest. In *ICDM*. <https://doi.org/10.1109/ICDM.2008.17>
- [25] Boris Medina-Salgado, Eddy Sánchez-DelaCruz, Pilar Pozos-Parra, and Javier E. Sierra. 2022. Urban traffic flow prediction techniques: A review. (2022). <https://doi.org/10.1016/j.suscom.2022.100739>
- [26] Tom Mertens and Marwan Hassani. 2022. Can we Learn from Outliers? Unsupervised Optimization of Intelligent Vehicle Traffic Management Systems. In *ECML PKDD*.
- [27] Yisroel Mirsky, Tomer Doitshman, Yuval Elovici, and Asaf Shabtai. 2018. Kitsune: An Ensemble of Autoencoders for Online Network Intrusion Detection. (2018).
- [28] Minh-Nghia Nguyen and Ngo Anh Vien. 2019. Scalable and Interpretable One-Class SVMs with Deep Learning and Random Fourier Features. In *ECML PKDD*.
- [29] C. Pasquale, I. Papamichail, C. Roncoli, S. Sacone, S. Siri, and M. Papageorgiou. 2015. Two-class freeway traffic regulation to reduce congestion and emissions via nonlinear optimal control. (2015).
- [30] Dragoljub Pokrajac, Aleksandar Lazarevic, and Longin Jan Latecki. 2007. Incremental Local Outlier Detection for Data Streams. In *IEEE-CIDM*. <https://doi.org/10.1109/CIDM.2007.368917>
- [31] Nicholas G. Polson and Vadim O. Sokolov. 2017. Deep learning for short-term traffic flow prediction. (2017). <https://doi.org/10.1016/j.trc.2017.02.024>
- [32] Williams Rizzi, Chiara Di Francescomarino, Chiara Ghidini, and Fabrizio Maria Maggi. 2022. How do I update my model? On the resilience of Predictive Process Monitoring models to change. (2022). <https://doi.org/10.1007/s10115-022-01666-9>
- [33] Erik Scharwächter, Emmanuel Müller, Jonathan Donges, Marwan Hassani, and Thomas Seidl. 2016. Detecting Change Processes in Dynamic Networks by Frequent Graph Evolution Rule Mining. In *ICDM*. 1191–1196. <https://ieeexplore.ieee.org/document/7837971>
- [34] Bernhard Schölkopf, Robert C Williamson, Alex Smola, John Shawe-Taylor, and John Platt. 1999. Support Vector Method for Novelty Detection. In *NIPS*. MIT Press.
- [35] Swee Chuan Tan, Kai Ming Ting, and Fei Tony Liu. 2011. Fast Anomaly Detection for Streaming Data. In *IJCAI*.
- [36] Haseeb Tariq and Marwan Hassani. 2023. Topology-Agnostic Detection of Temporal Money Laundering Flows in Billion-Scale Transactions. In *Machine Intelligence and Data Science Applications*. to appear.