

A Visual Analytics System for Improving Attention-based Traffic Forecasting Models

Seungmin Jin*, Hyunwook Lee*, Cheonbok Park, Hyeshin Chu, Yunwon Tae, Jaegul Choo, and Sungahn Ko**

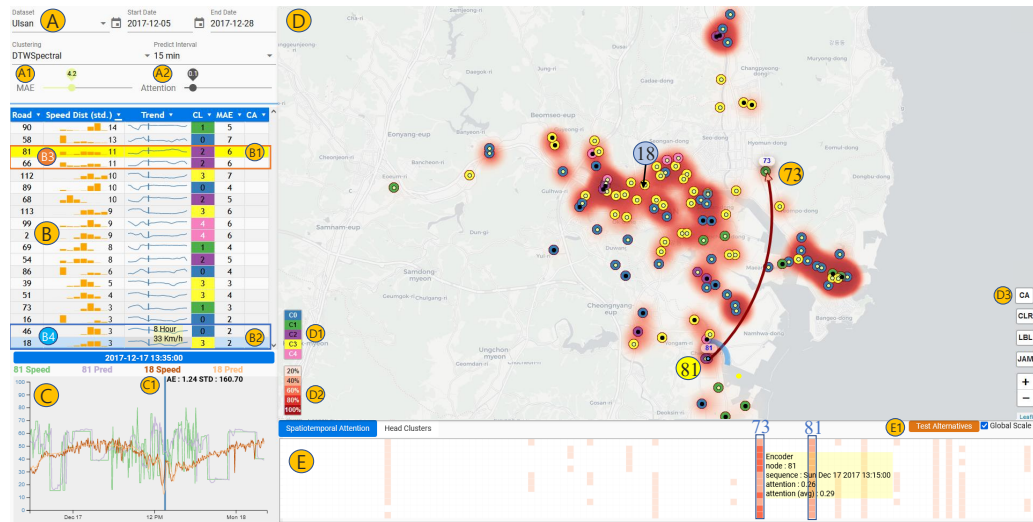


Fig. 1. Our system for analyzing how an attention-based deep learning model predicts traffic congestion. (a) filter view, (b) table view, (c) ground-truth&prediction result comparison view, (d) a map with attention curves and clusters, and (e) attention heatmap view.

Abstract—With deep learning (DL) outperforming conventional methods for different tasks, much effort has been devoted to utilizing DL in various domains. Researchers and developers in the traffic domain have also designed and improved DL models for forecasting tasks such as estimation of traffic speed and time of arrival. However, there exist many challenges in analyzing DL models due to the black-box property of DL models and complexity of traffic data (i.e., spatio-temporal dependencies). Collaborating with domain experts, we design a visual analytics system, AttnAnalyzer, that enables users to explore how DL models make predictions by allowing effective spatio-temporal dependency analysis. The system incorporates dynamic time warping (DTW) and Granger causality tests for computational spatio-temporal dependency analysis while providing map, table, line chart, and pixel views to assist user to perform dependency and model behavior analysis. For the evaluation, we present three case studies showing how AttnAnalyzer can effectively explore model behaviors and improve model performance in two different road networks. We also provide domain expert feedback.

Index Terms—Traffic Visualization, Deep Learning, Attention Model, Speed Prediction, Explainable Artificial Intelligence

1 INTRODUCTION

To mitigate traffic congestion, which is a major issue in modern cities, a lot of effort has been devoted to developing methods for predicting future congestion to support drivers' route decisions [2]. In particular, as many recent studies show that deep learning (DL) outperforms conventional methods in forecasting traffic congestion [62] and the estimated time of arrival (ETA) [75], there have been numerous practical approaches for developing and deploying new DL models. For example, reporters in broadcasting centers announce which roads will be congested in the next 15, 30, and 60 minutes to help drivers avoid the possibly congested roads or change their driving schedule (i.e.,

delayed departure time) [36]. Mobility industries, such as navigation service providers, also report that they have achieved greater than 40% accuracy in estimating travel time of drivers with DL [13].

But, many obstacles exist in improving DL models in the traffic domain. First, DL methods are black-box in nature [18], which means they do not show what they have learned during training or how they make predictions given input features. Second, traffic domain data is spatio-temporal, which is hard to analyze, as both space and time should be considered simultaneously. For example, when a road is congested because of an accident, there is a high chance that the roads linked to the road with the accident will soon become congested, but it is hard to estimate when the congestion happens and which neighboring roads would be affected [39]. Third, traffic data is heterogeneous with extreme cases [39, 62, 78] and easily affected by uncontrollable external factors (e.g., accidents). Thus the traffic prediction task is especially challenging in that the models need to learn not only spatio-temporal features from the data, but also how to respond to implicit external events on roads. The external factors could even vary by region [36], which further complicates to the models. Indeed, spatio-temporal analysis with a large space (e.g., major city's road network) and different time units (e.g., hours, days, months) is a nontrivial task [3].

In this work, we collaborate with domain experts to design a visual analytics (VA) system that supports users in effectively exploring how a DL model predicts future traffic conditions and in finding insights

*Equal contribution

**Corresponding author

- Seungmin Jin is with HSE (Higher School of Economics), Russia, and also with the Computer Science Department, UNIST, South Korea. E-mail: skjin@unist.ac.kr.
- Hyunwook Lee, Hyeshin Chu, and Sungahn Ko are with UNIST. E-mail: {gusdnr0916, hyeshinchu, sako}@unist.ac.kr
- Cheonbok Park is with NAVER. E-mail: cbok.park@navercorp.com
- Yunwon tae and Jaegul Choo are with KAIST. E-mail: {tyj204, jchoo}@kaist.ac.kr

Manuscript received 31 March 2022; revised 1 July 2022; accepted 8 August 2022.
Date of publication 26 September 2022; date of current version 2 December 2022.
This article has supplementary downloadable material available at <https://doi.org/10.1109/TVCG.2022.3209462>, provided by the authors.
Digital Object Identifier no. 10.1109/TVCG.2022.3209462

for performance improvements. We first perform task analysis with three domain experts and extract the system requirements needed to support the experts in answering questions that should be answered for their work. Then, we design a VA system, AttnAnalyzer, that provides three automated methods and visual interfaces for exploration. We incorporate the dynamic time warping (DTW) [57] and Granger causality tests to support users in exploring spatio-temporal dependencies of roads, such as proceeding and lagging speed patterns. We also utilize a spectral clustering algorithm for grouping roads with similar speed patterns for automated methods. There are several visualizations in AttnAnalyzer, including table, line chart, map and pixel views. The table view presents information of roads, such as speed trends, cluster index, and prediction accuracy. The map view allows spatial dependency analysis, while the line chart view reveals the temporal dependencies of roads. The pixel views have two sub-views to analyze attention information, extracted from an attention-based traffic forecasting model, called ST-GRAT. To validate hypotheses created from the system, we devise an attention enforcement method, which replaces the attention of problematic roads with those from the roads with low MAE error. For demonstration purposes, we use an attention-based traffic forecasting model because they show state-of-the-art performance and because experts have expressed that they extensively use attention-based DL models in navigation service.

To evaluate the system, we provide three case studies. In the first two case studies, we showcase how domain experts use the system to explore the spatio-temporal dependencies of two different large road networks. In the last case study, we show that the insights derived from the two case studies are meaningful for improving model performance, confirming with attention enforcement. Lastly, we provide domain expert feedback and discuss limitations and future work.

The main contributions of this work include the following: 1) a VA system design for exploring traffic forecasting model's behaviors from a spatio-temporal perspective, 2) incorporation of automated methods (DTW, Granger causality test, clustering) for visual temporal analysis, 3) development of an attention enforcement method, and 4) quantitative and qualitative evaluations of the system with three case studies, proven model's accuracy improvements with the attention enforcement method, and domain experts' feedback. Although we show how model designers can improve model performance using our tool, the main objective of this work is not to propose a new model, but to demonstrate how to explain deep learning models with attention. To our knowledge, this work the first attempt to exploring and attention-based traffic forecasting models' prediction process and improving performance in the traffic domain, demonstrating the power of visual analytics approaches [63]. Traffic data is heterogeneous with extreme cases [39, 62, 78] and affected by uncontrollable external factors, such as accidents. Thus the traffic prediction task is especially challenging in that the models in the domain need to learn not only spatio-temporal features from the data, but also how to respond to implicit external events on roads. The external factors could even vary by region [36], which further complicates to the models.

2 RELATED WORK

2.1 Spatio-Temporal Models for Traffic Prediction

As deep learning models are effective in many domains with large data sets, researchers have developed various models for the traffic domain. In particular, many models have been proposed for speed prediction that can be used in the field to relieve traffic congestion issues [36]. In order to accurately forecast road speeds, most of studies focus on modeling spatial and temporal dependencies that in general affect the dynamics of road speeds [39]. For example, if there is a vehicle accident on a road, it is highly possible that the other roads; linked to that road will become congested (i.e., spatial dependency). As time passes, the congestion beginning at the road with the accident may propagate to other roads (i.e., temporal dependency). As such, existing models aim to effectively model road network dependencies for speed prediction.

To effectively capture spatial dependency, graph convolution neural networks (GCNNs) [29] has been proposed that apply the convolution technique to graphs of roads. As using diffusion on graphs turns

out to be effective for modeling spatial dependency [40] in road networks, many researchers have employed diffusion convolution neural networks (DCNNs) for accurate prediction [40, 50, 72]. However, the models based on GCNNs and DCNNs are vulnerable, as they only consider spatial dependency as fixed values, regardless of inputs and traffic condition changes. To alleviate this issue, attention-based models have been employed to better model spatial dependency with considerations of road distances, showing state-of-the-art performance for speed prediction tasks [51, 51, 82].

Temporal dependency also plays a crucial role in encoding traffic congestion patterns. Examples include congestion propagation on roads over time. To effectively capture the dependency, prior studies have mainly used conventional modeling methods, such as recurrent neural networks (RNNs) [40, 77, 82]. However, such RNN-based temporal dependency modeling methods have a limitation in that they cannot effectively capture long-range temporal trends in a given sequence. To alleviate this weakness, recent models have employed other advanced techniques, such as convolution neural networks [72] or self-attention networks [51]. In this work, we use ST-GRAT to model both spatial and temporal dependencies due to its enhanced interpretability and better performance over other models for traffic prediction [51].

2.2 Visual Analytics for Deep Learning Models

Although deep learning models are effective, it is difficult to understand how the models work due to their black-box characteristic [48]. Existing visual analytics approaches and systems for the matter fall into global model analysis and instance-based analysis [23]. The global model analysis systems focus on how to visualize internal model structures through a graph structure, whose nodes and links are mapped to neural network neurons and weights between two connected neurons, respectively (e.g., CNNVis [42]). In general, this global model analysis category includes visual analytics systems for understanding convolution neural networks (CNNs) applied to computer vision tasks [10, 42, 54, 76, 79] and recurrent neural networks (RNNs) for natural language processing tasks [46, 61]. For example, Liu et al. [42] proposed a visual analytics system, called CNNVis to visualize a CNNs via a directed cyclic graph. They additionally use the edge bundling technique to visualize the learned filters and connections between layers. Zeng et al. [81] study a visual analytics approach to understanding the behaviors of CNNs, solving the modifiable areal unit problem. Shen et al. [59] suggest a visual analytics system that allows users to visually explore the model behavior from global and individual levels on a multi-dimensional time-series forecast. Compared to the work in the global model analysis category, the main objective of our work is different. For example, while their approach focuses on the analysis of input data with errors, aiming at stabilizing feature embedding and predictions, our work allows for the investigation of spatio-temporal dependencies that a model learns during training. Also, our automated approach is different in that it enables users to directly estimate the reasons for model behaviors using time-series causality analysis.

Fred et al. [22] present SUMMIT, which summarizes CNNs at scale by generating an attribute graph from a trained model and finding which filters influence on performance. Inspired by these studies, many VA systems have been proposed to solve similar problems with various models, including generative adversarial networks [25, 68], deep reinforcement learning [67], and sequence-to-sequence models with attention [60] and self-attention networks [12, 52].

The VA systems for instance-based analysis aim to show how an input instance influences a model and vice versa to help users identify models' robustness and internal mechanisms in given scenarios [7, 26, 41, 70]. For example, Strobel et al. [60] proposed Seq2Seq-VIS to help users manipulate models' internal parameters, and observe how models react to different inputs. SANVis [51] visualizes self-attention networks with a given sentence, displaying how the model captures the linguistic relationships between the words in the sentence.

2.3 Visual Analytics for Traffic Congestion

As traffic congestion is a crucial problem; affecting life quality, many VA systems have been proposed to find the causes and possible solu-

tions. Many computational and strategical approaches are employed in designing VA systems for traffic analysis. The systems often integrate new visualizations to better present traffic patterns and anomalies, such as T-Watcher with fingerprint visualization [56]. Wang et al. [69] showcases an automated method for detecting traffic congestion, enabling users to explore different traffic congestion propagation graphs of a large city. Zeng et al. [80] proposed an interchange circus diagram to present interchange patterns at a junction road and help users identify multi-spatial scales and temporal change patterns. Recently, Pi et al. [55] have built cumulative vehicle count curves (N-curve) and classified patterns of congestion cause by using entropy from information theory to analyze causes of congestion across intersections.

While much research has been conducted, few studies on visual tools exist for predicting traffic speeds and volumes, which is an important task for experts in the traffic domain, such as reporters at broadcasting centers. Lee et al. [36] proposed a visualization system to help the experts whose tasks include broadcasting traffic conditions across a city. They use and train Long Short-Term Memory (LSTM) [21] with three features—road network structure, neighbor roads' speed, and rush-hour—to improve prediction performance. However, as the system does not explain the prediction results, experts may not be confident when they broadcast possible congested roads to citizens based on the prediction results. In this work we design a visual tool that can help users understand the prediction results.

3 TASK DESCRIPTION AND REQUIREMENTS

Our system has been designed with input from three domain experts of a corporate, which has 19 million users and processes 100 million map and navigation services per day, as of 2021. We have met with experts over a period of 18 months to extract task requirements and discuss design considerations. The first expert (E1, 20 years of experience) is a director of the map service group, supervising all map-related developments and services. The second expert (E2, 12 years of experience) leads the traffic intelligence team of the group and has developed algorithms for navigation services with the third expert (E3, 6 years of experience). E2 and E3 are the machine learning and software engineers in the traffic domain, and both have software development skills and machine learning knowledge and have applied machine learning algorithms to the domain area (i.e., trainers [61]).

According to the experts, the group's main mission is to provide the best optimized navigation services to drivers and pedestrians. To fulfill this mission, they collect traffic speeds and paths from users (e.g., drivers who use their navigation services). They also collaborate with traffic departments of the national government, police agencies, and city halls to collect and utilize additional road events, such as accidents and construction in their services. The collected data is used for pathfinding and travel time estimation algorithms.

One of the concerns of the group is that although they have done well in the domain, it is possible that their competitors can take a technical advantage with the recent rise of machine learning models, which mean a loss in their market share in the country. To overcome this, they have utilized DL technologies [37, 50, 72, 77, 82], including LSTM [36], attention-based models for predicting future traffic speed [51] and arrival time [13, 16]. But, they often obtain unsatisfactory performance from the models, when they directly apply the models to their data. To find possible reasons for low performance, they conventionally investigate the model behaviors of roads with low accuracy, using the what-if method [17, 49, 70], in which they change the input traffic speed or features of a road in a model and compare forecasting performance. They also compare a road's traffic behavior to that of other roads. They know that a city has multiple regions, and each region shows different traffic and congestion patterns due to the various roles and capacities of the roads across the regions (e.g., residential and university areas, industrial complexes [37]). In terms of a navigation service point of view, finding out which time and locations where model produces high accuracy is important information for deployment.

But, after starting their investigation, they soon find that this case-based, individual prediction performance analysis with Python is onerous and ineffective. In particular, the approach takes a great deal of time

and effort to answer all of the questions that they create at work for achieving their goal and deployment criteria due to a large number of road links in cities [36] and complex dependencies among roads [19].

The questions that need to be answered for their tasks are as follows: **(Q1) Where and when does a model produce inaccurate predictions?; (Q2) What could be the causes of inaccurate predictions?; (Q3) Which past (combinations of) observations does a model use for speed prediction?; (Q4) Which road does a model refer to for speed prediction?; (Q5) How can the accuracy of a model be improved?; and (Q6) How does a model work across different road topologies?** Among the questions, Q1 and Q2 are the main questions, while Q3 and Q4 are auxiliary questions that need to be clarified to better answer Q1 and Q2. Q5 can be considered a question for increasing their competence in the market and Q6 is the last question that should be addressed before in-the-wild deployment.

From the discussions, we have derived the following requirements for a visual analytics system designed for exploring traffic forecasting model's behavior and performance improvements. First, as it is important to pop out problematic roads with a given model, **(R1) a VA system should highlight the roads with low accuracy and provide information on when a model has low accuracy (Q1)**. As models predict future speed based on spatio-temporal dependencies among roads, **(R2) a system should provide a method for effectively exploring encoded dependencies (Q2–Q4)** so that users can find evidence on the relationship between high errors and speed patterns [9, 39]. Example information for supporting spatio-temporal dependency exploration includes **(R2-1) historical traffic patterns of roads, speed distribution of data, standard deviation, daily speed trends (Q3, Q4), and (R2-2) data on model behaviors (Q2–Q5)**, such as which roads a model refers to for forecasting (i.e., which roads influence the prediction for a target road [19]), and which input sequences are importantly used for the prediction. In many cases, it is difficult for users to identify which roads affect the prediction performance of a certain road [39]. While prior work has revealed [19] that there is a strong dependency among neighboring roads and second and third linked roads by cross validation, methods for investigating whether a target road refers to appropriate neighboring roads or not have been rare. Therefore, **a system should also provide information on the similarity of temporal data and causality among roads (Q2–Q4)**. Lastly, to support users' formulation and validation of hypotheses, **(R3) a system should provide a method that shows how much improvement users can expect (Q5, Q6)**.

4 DATA DESCRIPTION AND ATTENTION MECHANISM

4.1 Data Description

In this work, we use the traffic data of two different road networks—the urban and highway road networks to explore the model's inference process for speed prediction. For the urban road network, we use dedicated short range communication (DSRC) data [28] generated from Ulsan, South Korea (range: 9/1/2017~12/28/2017), where more than 1.1 million people live with more than 540,000 registered vehicles as of 2017. A total of 116 DSRC sensors are used for data collection, which are installed every 5.7km and cover a total of 68 main roads. For the highway road network, we use the METR-LA data [24], which were collected from 207 loop detectors (range: 3/1/2012~6/27/2012) on the highways of Los Angeles. Note that the highway network data we use are the standard benchmark data for traffic forecasting tasks [40, 77, 82]. After discussing with domain experts and reviewing training results, we replace the missing data and explicit errors with historical data. We also use 5-minute aggregated data to mitigate possible effects of outliers, as performed in many previous studies (e.g., [40]).

4.2 Attention Mechanism

We describe how we utilize ST-GRAT to demonstrate our VA approach. We choose ST-GRAT because 1) it has demonstrated state-of-the-art performance and because 2) we can produce attention matrices on spatio-temporal dependency (e.g., Eq. 1, Fig. 3D) [51]. We also consider that the domain experts have expressed that attention models are extensively used at work for not only speed prediction but also other tasks, such as travel time [16] and taxi demand prediction [74].

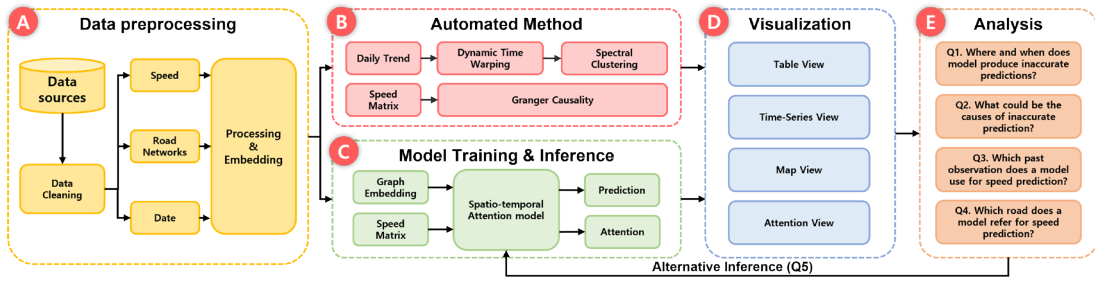


Fig. 2. An overview of the system workflow with functional modules and questions. (A) Data pre-processing, (B) automated methods for supporting spatio-temporal analysis, (C) model training and inference, (D) visualization modules, and (E) answering questions using the system.

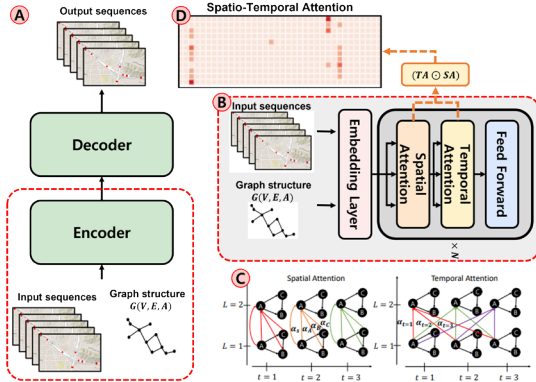


Fig. 3. Overall architecture of ST-GRAT (A). Each layer in the encoder is composed of a stack of identical layers (B). $G(V, E$ and $A)$ in (A) indicates a set of nodes, links, and a weighted adjacency matrix, respectively. We create a spatio-temporal attention matrix (D) by calculating point-wise multiplication between spatial and temporal attention weights (C).

ST-GRAT (Fig. 3) is a variation of the transformer [65] that uses the encoder-decoder architecture with self-attention (i.e., temporal attention). Additionally, ST-GRAT utilizes graph attention as spatial attention before temporal attention with a sentinel vector. The sentinel vector acts as weights for skip connection within the same road. ST-GRAT utilizes 12-length sequential historical speed with encoded features for each road and predicts 12 sequential speed predictions.

There are three types of layers in the encoder and decoder: embedding, spatial attention, and temporal attention layers (Fig. 3B). To allow the model to extract the spatio-temporal dependencies, we provide a road network, speed, and observed time as the input features for the embedding layer (Fig. 3B). The road network graph is directed graph $G(V, E, A)$, where a road is represented as a node (i.e., V) and the connection between roads is shown as a link (i.e., E). Note that road network G is directed graph, which allow model to distinguish the road directions. Note that we provide the order of a given sequence using the position embedding method [65].

The model captures spatial dependencies among neighbor roads in the spatial attention layer (Fig. 3B) by using a graph attention network [66, 82], a well-known graph modeling method. In short, the spatial attention layer integrates information among neighboring roads by directed graph attention. This directed spatial attention improves dependency modeling and helps developers interpret ST-GRAT.

The temporal attention layer (Fig. 3B) models the temporal dependency and trends of given sequences. For modeling, temporal attention performs multi-head attention to compute temporal correlations. The attention type is decided by attention axes; spatial attention aggregates features among the spatial axis (i.e., neighboring roads), while temporal attention aggregates input features within the temporal axis (i.e., different time steps of the same roads (Fig. 3C).

To sum up, there are two important attention layers—spatial and temporal attention—in ST-GRAT. Vaswani et al. [65] describe an attention function as mapping a query and a set of key-value pairs to an output, where the query, key, values, and output are all vectors. From this perspective, from each key-value pair, we can derive an attention matrix from each attention layer, called **SA** and **TA**, and create a **spatio-**

temporal matrix (ST matrix). Specifically, for each attention head, given X before spatial attention and H after temporal attention, their relationship can be written as follows:

$$H = (TA \odot SA)X, \quad (1)$$

where $TA \odot SA$, named *spatio-temporal attention*, is the attention that users mainly explore for understanding model behaviors.

5 VISUAL ANALYTICS ENVIRONMENT

We describe our visual analytics system, AttnAnalyzer, which we have designed to effectively explore the deep learning model's internal process for traffic forecasting using attentions. Fig. 2 shows the pipeline of the analysis with AttnAnalyzer. We first prepare the data (A) for automated methods (B) and model training (C). After model training, we extract ST attention from the model (C), which is used for visualizations (D). Users explore model behaviors using the automated methods, (E) answering the questions (Q1–Q4) and test and confirm performance improvements (Q5).

5.1 Automated Methods

In this work, we incorporate two automated methods—**dynamic time warping (DTW)** with spectral clustering [6, 57, 73] and **Granger causality tests** [14], to allow users to determine which references are appropriate for predictions and whether a model employs proper references for the predictions (R2-3) [38].

We use daily trends of roads as input for DTW computation, which should be explored since they are encoded as temporal dependencies among roads, which a model learns during training [74]. However, it is difficult to analyze road trends for two reasons. First, there are too many trends (i.e., roads) in a traffic dataset. Second, there are different time gaps between the trends of roads due to different levels of dependencies, which are presented with either lagging or preceding trends [39]. For example, when congestion occurs on a road, the roads linked to the congested road also become congested, but there are no concrete patterns when the neighboring roads are congested. As such, we incorporate dynamic time warping (DTW) [6] in this work, which calculates the similarity in two time series over time gaps after finding the best matching alignment that minimizes the distance, as Le Guen and Thome do for their time-series forecasting model design [35]. For efficiency, we use FastDTW [57], whose complexity is $O(N)$.

Then we perform spectral clustering [47] with the computed DTW scores to further allow cluster-based analysis, which has shown its effectiveness in analyzing traffic data [5, 73]. For example, users can easily confirm whether referred roads of a target road have similar daily trends as the target road using clusters. If the target and referred roads are in the same cluster, it is highly possible that the model learns the dependencies of the target and referred roads during training [39]. Note that we use the ELBOW method [1] with visual inspection to choose the number of clusters.

As Wu et al. [71] show in their study, catching preceding patterns in time-series enables effective analysis of deep learning models. Therefore, to help users better explore the preceding patterns in the traffic data, and to complement DTW that distorts the time during its computation, we incorporate the Granger causality test [14] in this work, a well-known temporal dependency investigation method. We first describe the definition of Granger causality based on two principles: (1)

the cause happens prior to its effect and (2) has unique information about the future values of its effect. Given these two principles about causality, Granger causality proposes testing the following hypothesis for the identification of a causal effect.

$$\text{GrangerCausality} = \mathbb{P}[Y(t+1) \in A | \mathcal{I}(t)] \neq \mathbb{P}[Y(t+1) \in A | \mathcal{I}_{-X}(t)]$$

Here, P is probability, A is an arbitrary non-empty set, and $\mathcal{I}(t)$ and $\mathcal{I}_{-X}(t)$ denote the information available as of time t in the entire universe, and in the modified universe where X is excluded, respectively. If the above hypothesis is accepted, we say X Granger-causes Y .

5.2 Visual Interfaces

AttnAnalyzer consists of five views: (A) filter, (B) table, (C) line, (D) map, and (E) attention, as shown in Fig. 1. The attention view (E) has two sub views—spatio-temporal and head cluster (Fig. 5) views. In the filter view (A), users can select a dataset, date range, and target prediction time (e.g., 15, 30, 45, and 60 minutes; default: 15 minutes).

Users can filter out the roads using Mean Absolute Error (A1: ● scaled: min–max values) and spatio-temporal attention (A2: ● scaled: 0–1) values. When roads' MAEs are higher than the threshold value, the center circle of the roads are colored in black (R1). If a road's MAE value is below the value at Q1 (i.e., top 25% in accuracy), the road is included in the low error group. If the MAE of a road is above the value at Q3 (i.e., bottom 25% in accuracy), it belongs to the high error group. When users hover over the filters, a tooltip pops up to show the MAE values at the Q1 and Q3 boundaries (Fig. 1, A1). The attention filter is applied to AttnArrows (e.g., Fig. 4) on the map, when a road is selected for investigation. There are two legends to represent clusters (Fig. 1 D1) and ratio values (D2). We use the color sets from ColorBrewer [20].

5.2.1 Table View

The sortable view (Fig. 1 B) allows users to explore roads' characteristics in a compact form with information on road ID (Road), speed distribution (Speed Dist. (std.)), speed trends (Trend), cluster index (CLS), mean absolute error (MAE), and causality analysis (CA). The view is initially sorted by clusters, but users can sort differently using a small triangle (∇) at column titles. The column title is underscored to represent which column is used for current sorting in the table. For example, the table is currently sorted by the standard deviation (i.e., std) of roads' speed (Fig. 1, Speed Dist. (std)). We show the speed distribution of a road (Speed Dist.) in a histogram with the x-axis representing different speed ranges with bins (e.g., 10 miles/bin) and the y-axis meaning the frequency. The bins' height is normalized by the maximum count. At the right side of the distribution, we also place each road's standard deviation value. If a road has a high deviation score, it means the road speed often significantly changes, while the low deviation implies that significant speed changes rarely happen on the road (e.g., highways).

To allow speed trend analysis, we provide normalized speed trends using spark lines [64] (Fig. 1B, Trend column), with the x-axis representing the time range (0–23) and the y-axis denoting average speeds of roads with the 5-minute interval. As users hover over a trend line, a vertical line and tooltip appear, showing time stamps and speeds. After the trend column, we provide cluster indices that each road belongs to and the MAE values of roads to help users investigate the relationship between speed trends, road clusters, and model accuracy (Q2, Q3). When users click on a trend line, a new window pops up, and shows speeds and MAE values.

We present Granger causality test results (F values) in the last column (CA) to support users inspecting spatio-temporal dependency in terms of causality. When users click on a road on the table, the F-values between the clicked road and other roads are presented. If the F-value of a pair is higher than of other pairs, we can say that the pair with a higher F-value has a more significant Granger causality relationship than other pairs. For example, we can see from Fig. 6 B1 that the Road 113 and Road 112 pair has a more significant Granger causality relationship than Road 113 with other roads. Note that if the p-value of a pair is less than 0.05, we do not show the results. We provide




Fig. 4. AttnArrows for presenting which roads a target road refers to for inference. The donut chart surrounding the target road is filled in blue (clockwise) based on the self-reference intensity.

the test results to help users make and validate hypotheses on model behaviors. For example, users can review F-values to determine if the model sufficiently refers the information of the road with a significant Granger causality relationship for performance analysis (Q2, Q5).

Users can perform interactions in the table view for coordinating multiple views for detailed investigation. When users hover over a road in the table, corresponding road on the map and the attention matrix are highlighted. When users click a road id, (1) table view highlights (light-blue) the row of the road (e.g., Fig. 1 B2), (2) map view places the road at the center of the map with highlighting in light blue, and (3) line view shows the collected and predicted speed lines (Fig. 1 C).

5.2.2 Time-Series Views

To help users investigate roads' temporal dependencies (R2), comparing road speeds, we provide a line chart view (Fig. 1C), where the x-axis and y-axis represent time (5-minute intervals) and road's speed, respectively. Users can add raw and predicted speeds by selecting roads from other views, such as table, map, and attention views, and remove the lines by toggling road labels at the top. There is a blue vertical bar in the view (Fig. 1 C1), indicating the current data point in visualizations (e.g., "2017-12-17 13:35:00," Fig. 1C, top blue panel). It also shows the absolute error (AE) and STD values in the past one hour of the time specified by the bar (e.g., "AE: 1.24 STD:160.70" Fig. 1 C1). Users can move the bar to anywhere to update map and attention views with the data specified by the bar.

5.2.3 Map View

As models learn dependencies among roads and forecast based on the learned dependencies, it is important to explore what dependencies a model learns with road speed patterns and which roads a model refers to for performance investigation. To help users perform such exploration (R2-1, R2-2), we first present each road as a white dot. Each circle surrounding the dots shows the cluster that the road belongs to. We provide three visualizations in the map view: heatmap, attention arrows (AttnArrows), and cluster visualization. First, to help users overview the relationship between roads' congestion levels and model performance (Q1), we visualize roads' congestion levels using heatmaps (Fig. 1D). Here, the redder the heatmaps are, the slower the roads are (heatmap legend: Fig. 1 D2). We also have considered providing heatmaps with a time filter so that users can explore regions with high model error but decided against this because this approach requires many interactions for filtering and memorizing heatmaps that have changed due to filtering. Using Bezier curves, AttnArrows (Fig. 4) link a target road and the roads that the target road attends for prediction. Here, the color represents the amount of attention given to the roads. The darker the head color, the more attention the reference road is given. For example, the head near Fig. 1, 73 shows that the attention level is about 40% of the maximum attention value (legend: Fig. 1 D2). To reduce visual clutter, users can hide unimportant AttnArrows by using attention filters (Fig. 1 A2), or click AttnArrows to turn off [30]. By default, the AttnArrows head visualizes the encoder's attention data.

When users click a road to set a target road for investigating its reference roads, the dot is highlighted steelblue (e.g., Fig. 1D, road 81). To represent the selected target road's self-reference intensity (i.e., how much a road attends to its own speed pattern), we present a donut chart around the road, filling the chart according to the self-reference intensity (scale: 0–1, clockwise). For example, the donut chart of road 81 (Fig. 1D) has self-reference intensity of 0.29. There are buttons for interactions on the map (Fig. 1 D3)—CA (Causality Analysis), CLR (Clear map), LBL (Label display), JAM (Traffic Jam Heatmap), zoom-in, and zoom-out.

When users click the CA button, AttnAnalyzer turns on the Causality Analysis mode. Then if users click a road, AttnAnalyzer removes the roads on the map, whose p-value in the test is higher than 0.05 to allow users to analyze with statistically meaningful roads. Users can draw a selection area by creating a polygon on the map where a set of roads can be included for investigation and un-select the selected roads using CLR. The speed heatmap and road labels are turned on and off with JAM and LBL, respectively.

To help users find which road belongs to which cluster, we initially implemented and compared BubbleSet [11] and Kelfusion [45], which are effective for set visualization. From our observations, although effective with a small number of clusters, they produced several overlaps among clusters when roads in different clusters are located closely. We have also observed that they may not be helpful, as a large number of overlaps means that there rarely is a dominant clustering pattern among the neighboring roads. Appendix Figure 11 shows our implementation results with 5 and 8 clusters. After the investigation, we decide to encode the clusters by placing an outer, colored circle on each road [33]. Figure 4 shows two example roads, where both the target (left) and reference (right) roads belong to cluster C0 (legend: Fig. 1 D1).

5.2.4 Attention View

There are two visualization tabs in the attention views: spatio-temporal (ST) and head-cluster views for **R2-2**. We use pixel-based visualization [27, 31] to present the model's spatio-temporal attention in a matrix form, because of its scalability on the number of items [27, 31]. In the ST attention matrix view (e.g., Fig. 1E), the x-axis indicates different roads, and the y-axis indicates past time steps (5-minute intervals, 12 steps, top: 5 minutes ago, bottom: 60 minutes ago). In the view, users can analyze which roads and time steps the model focuses on with attention intensity, represented by the color (legend: Fig. 1 D2). For example, when users click road 81 on the map, the spatio-temporal view (Fig. 1E) shows that road 81 attends itself (i.e., self-reference) and 73 more than other roads for its prediction.

There are 8 matrices in the head-cluster view (e.g., Fig. 5 top, first row) to show how much each attention head refers to reference roads for making predictions for target roads in a cluster point of view. In the view, the first four matrices visualize the attention patterns of the four heads with the roads for which the model records high error (top 25%), while the other four matrices display those of the four heads with the roads for which the model shows low error (lower 25%).

In each matrix, the column represents the clusters of reference roads, while the row means clusters of target roads in ascending order. The color of each cluster cell represents the intensity of the attention that each head assigns, so if a cell color is the darkest red, it means an attention head heavily refers to the roads in the cell (i.e., cluster) for making predictions (legend: Fig. 1 D2). For example, many target clusters attend cluster 0 and 3 (Fig. 5 top, low error) and record low error. This indicates that the roads in cluster 0 and 3 have strong similarity with the roads in the target clusters in terms of daily speed trend.

There are global and local scales to normalize the attention values (i.e., intensity) in the matrices differently. All attention values across matrices are divided by the largest attention value to show relative attention intensity in the global normalization (Fig. 5 top, first row), but in the local normalization, each cell is divided by the sum of each row to make row sums of the matrices equal to 1. We show the attention patterns with the two scales to help users analyze the attention patterns by individual heads and across heads (**R2-2**) and find the reason for the failed inference, using the level of attention intensity information.

When users hover over a cell on a matrix, a tooltip pops up, showing detailed information, such as head index, relative attention ratio, and the average of their attentions. The interaction also highlights the road in the map and table views associated with the hovered item.

Inspired by Strobelt et al.'s approach that nullifies attention [60], we present an **attention enforcement method** in the head cluster view that replaces the attention values of the roads with high error with those from the roads with low error. Fig. 5 shows the process. For example, as users select four clusters (blue outline) in the view, AttnAnalyzer searches top-k target roads with high error in each cluster (Fig. 5 [2]).

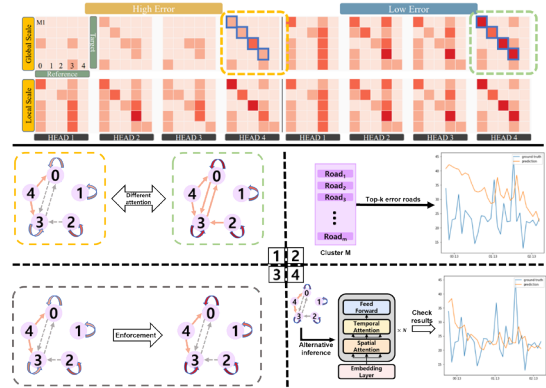


Fig. 5. (Top) The head-cluster view with four attention heads, (Bottom) The enforcement process—1) comparing attention between low (left) and high (right) error cases, 2) selecting k highest error roads in each cluster, 3) replacing the attention of the selected roads with that from low error roads, and 4) testing alternatives.

Then, the system automatically finds the most appropriate reference roads for each selected road in the chosen clusters, using the DTW distance matrix and Granger causality tests with equal importance (0.5). Lastly, the system extracts the attention values of the reference roads (i.e., alternative inference, Fig. 5 [3]), applies alternative inference with the replaced attention to the target roads, and expresses new prediction results (Fig. 5 [4]). Fig. 5 [1] and [3] show the differences in attention distribution among clusters (i.e., color darkness of the arrows) in the initial and final cases. When users click on the “Test Alternatives” in the view (Fig. 1 E1), a new view pops up, showing two line charts to describe the model's original and resulting accuracy (e.g., Fig. 10). Here, the x-axis corresponds to the degree of error, while the y-axis presents the frequency of the roads with specified error. In sum, if the original performance graph is moved to the left, there is an improvement in model performance. Detailed analysis with the result of using the method is presented in Sec. 6.3.

6 CASE STUDY

We present two case studies with two different road networks—the urban and highway road networks—exploring the model's speed prediction process. The case studies are archetypal use-cases that drove the design along with the feedback from our domain experts. The third case study shows how model performance can be improved based on the findings derived from the two case studies. As described in Sec. 4.1, we use DSRC data generated from Ulsan, South Korea as urban road networks and METR-LA data as highway road networks. Note that we define high error (low performance) and low error (high performance) groups of roads. The high error group includes the roads whose MAE is higher than the third MAE quartile (Q3), while the low error group include roads whose MAE is lower than the first MAE quartile (Q1).

We set DTW's window size as 4 to capture up to 20-minute time lags and use 5 and 6 clusters for Ulsan and LA, respectively, based on the elbow methods [53]. We performed pre-processing so that both data sets have a 5-minute interval of speeds and timestamps and replace missing values with averages of past data. We used 70% of data for training, 10% of the data for validation, and 20% of the data for testing.

6.1 Analysis with Urban Roads

A machine learning model developer, Jane, has recently developed a new deep learning model for traffic forecasting using the attention mechanism and found that the new model shows higher accuracy in experiments, compared to existing models (Appendix, Table 2); however, she believed the accuracy can be improved more, if she understands how the model works. So, she decided to use AttnAnalyzer to explore how the model predicts road speeds.

After the traffic data of Ulsan was loaded, AttnAnalyzer showed the roads on the map and the detailed information on the roads in the table (Fig. 1B). Jane then turned on the speed heatmap to obtain an overview of typical traffic congestion across the city by clusters and to analyze error patterns by cluster. For example, she found that

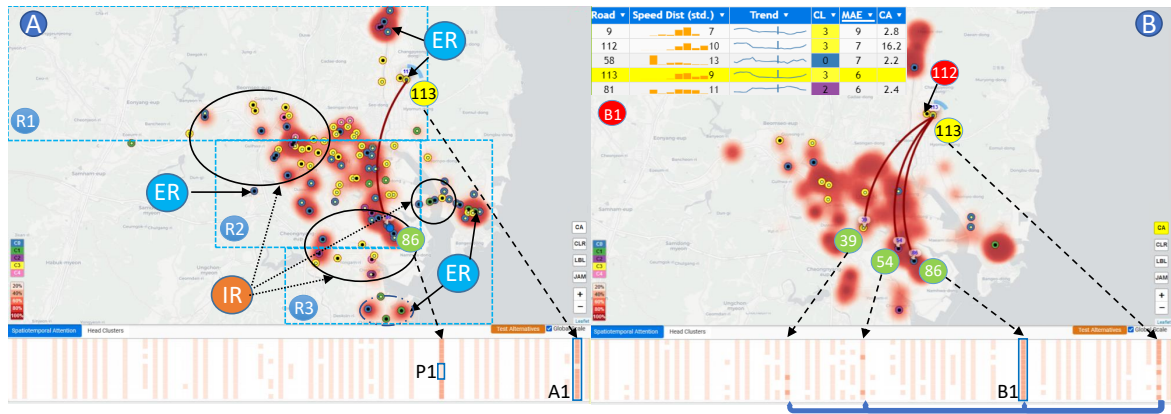


Fig. 6. Two example visualizations during analysis. (A) There are three regions—R1: residential, R2: downtown, and R3: industry complex. R2 has heavy congestion during rush-hours due to in-and-out traffic from R1 and R3. High error roads show inner black dots as a filtering result (Sec. 5.2). (B) Roads with low causality values and attention for Road 113 are filtered out in the causality mode.




Fig. 7. Inaccuracy is related to a high standard deviation of the speed trend in a given window. The model records a low error at T1 (low deviation) and a high error at T2 (high deviation). The traffic pattern of Road 112 precedes that of Road 113 (P2).

Cluster 3 usually experienced traffic congestion during the morning and evening rush hours as it includes many roads in the downtown region (Fig. 6A, R2), while Cluster 4, which is in a residential region (Fig. 6A, R1) is congested in early morning [36]. Next, as she set the MAE filter's threshold at 4.2, which is about one-fourth of 11 (the largest MAE, Fig. 1 A1), AttnAnalyzer highlighted the roads with low accuracy on the map. Interestingly, there seems to be no strong correlation between traffic congestion and model performance. Instead, it was shown that high error roads tend to be located at edges of regions (ER) and intersections of regions (IR), such as commercial and industry complexes, and residential regions [36] as shown in Fig. 6A.

As the roads with similar characteristics may have similar error levels, she found it is important to explore the roads of high errors (Q1). To identify the characteristics in common, she sorted the table, making hypotheses on a relationship among the deviation, distribution, and accuracy. She quickly noticed many cases seem to imply a relation between speed distribution patterns and the road error, MAE in the table view (Fig. 1B). For example, she found that road 81 and road 66 (Fig. 1 B3) had similar speed distribution, deviation, and error. They both had high deviation (e.g., 11, the speeds were similarly distributed in multiple ranges, and high error was recorded (e.g., MAE: 6). In contrast, road 16 and road 46 (Fig. 1 B4) had low standard deviation value (e.g., 3) with a pointy and unbalanced speed distribution and a low error (e.g., MAE: 2).

Then, she began further investigation on the reason for such a relation, assuming that (Finding 1) **the model makes inaccurate predictions (i.e., high error) for roads with high speed fluctuations (i.e., large speed changes)**. During her investigation, she found many cases that could confirm this assumption. For example, Fig. 1C shows two example roads—road 81: high fluctuation and, road 18: low fluctuation. At road 81, the model did not accurately forecast when the speed suddenly dropped or soared (MAE: 6), but it accurately predicted speeds for road 18 (MAE: 2). She thought this finding could also explain why there are many roads with higher error in the intersections of regions (IR) that experience severe speed changes (Q1).

As she understood which roads tend to have high error and when the high error occurs, she investigated how the model makes prediction with the roads and possible causes of inaccuracy (Q2–Q4). To do so, she first set the attention filter (Fig. 1 A2) as 0.1 to pass the roads that

the model did not give much attention. Then, she clicked road 113, one of the highlighted roads with the highest error on the map (Fig. 6B, yellow-fill circle) to add it to the line view and turned on AttnArrows. Once the speed data line was added, she moved the blue bar in the line view to the point that had a low absolute error (Fig. 7 T1), updating all other visualization views. Based on the updated map and attention views, (Q4) she found that the model mainly refers to road 113 and road 86 (Fig. 6 A). She also noticed several vertical light-pink lines in the attention view, including that of road 113 and 86. The reference to road 113 is a self-reference, implying that it is the model used in the road's past speed data (12 steps, 1 hour) for forecasting (Q3). This self-reference at the low speed deviation is advantageous when the current road speed follows a past speed trend, or periodic pattern, because the speed pattern does not change much from the previous steps [74].

It is interesting that the model referred to road 86 among many other candidates. Initially, she confirmed that road 86's speed trend was a similar to that of road 113 (Fig. 7 A). In addition, she found that the model effectively attended road 86's information. For example, there was a time when road 86 had a speed drop (Fig. 7 P1), which is different from its average speed trend. When this sudden drop occurred, the model did not refer to road 86 (Fig. 6 P1) as it learned the temporal dependency that referring to road 86 in the drop timing is not helpful for accurate predictions for road 113.

Once she determined how the model makes accurate predictions, she became interested in why the model behaves differently for the road with a high speed fluctuation, making an inaccurate prediction (Q2–Q4). As she moved the blue bar in the line view to the point where road 113 had a sudden speed change and high error (Fig. 7, T2), she found that the model began attending other roads (e.g., road 39, 54, 86) and reduced existing self-reference intensity (Fig. 6 B, attention view). The head-cluster view also showed similar behaviors of the model in terms of cluster point of view. For example, when the model has sufficient self-reference, there are dark red diagonal patterns, similar to Head 2, 3, and 4 in the low error case (Fig. 5 top right). However, when the model lacks self-reference, these patterns vanish from the matrix, as shown in the high error case (Fig. 5, top left).

She speculated that this attention behavior is not helpful for an accurate prediction of road 113 for two reasons. First, as the referred roads (e.g., road 39, 54 and 86) are far from the target road (Fig. 6B), there is a low probability that the referred roads have similar traffic patterns as that of road 113 and any congestion is propagated among the roads in the near-future (e.g., 15 minutes) with given speed levels [19]. Second, a distant road can be helpful for prediction if there is any relationship between the traffic patterns of the roads. For example, if the traffic pattern of a referred road precedes that of a target road, referring to distant roads could be advantageous [14]. But, she could not find crucial evidence of preceding traffic patterns from the line chart view. She also noticed from the cluster visualization and table view they are not even in the same cluster, as shown Fig. 1B.

Rather, during this investigation, she found that road 112 is close

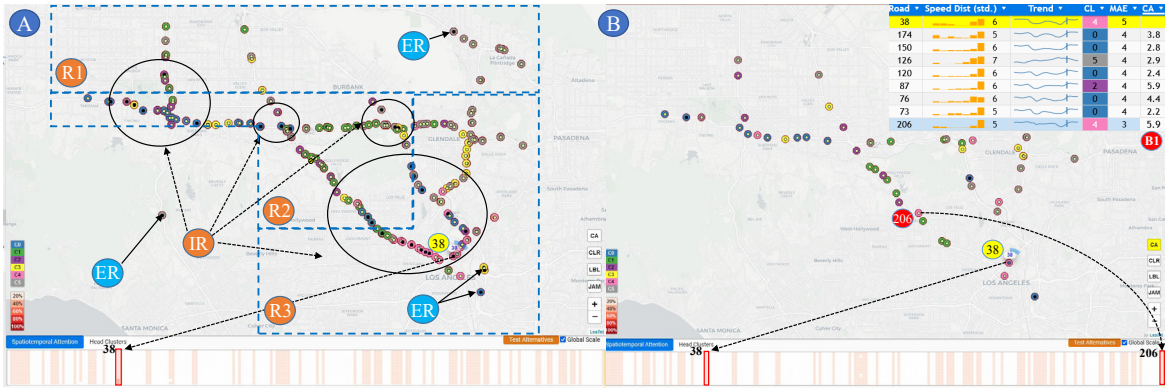


Fig. 8. An example visualization view with the highway speed data in L.A. (R1: residential, R2: downtown, R3 industrial regions). The two attention views show how the model attends in the low (A) and high error (B) cases.




Fig. 9. Summary of tables and Causality Analysis in LA. Overall there is a positive correlation between std. and MAE. Road 206's traffic pattern precedes that of road 38 (line chart, P1).

to road 113 and is also in the same cluster as road 113. The speed pattern also seemed to precede that of Road 113, as shown in Fig. 7 P2. To inspect this further, she ran a Granger causality test [14] with Road 112 and 113. The test result indicated that road 112's speed trend indeed precedes that of road 113 ($F[6,268]=16.2, p=0.001$), as shown Fig. 6 B1 (Table). These observations led her to think that **(Finding 2) the model effectively captures similar speed trends, but may not effectively recognize preceding speed patterns (Q2)**. She also noted that **(Finding 3) in the case of high speed fluctuations, the model may attend many other roads and lose the importance of self-reference (Q2, Q4)**. This behavior is also not appropriate, as existing autoregressive models have already shown that past temporal data acquired by self-reference is more crucial for short-term forecasting (e.g., 15 minutes) [8, 44].

6.2 Analysis with Highway Roads

To determine whether the model shows the same behaviors with different road networks (Q6), she decided to analyze the model with the speed data on highways in L.A., USA. As her purpose includes confirming the model behaviors on the roads with high and low errors (Q1, Q2), the analysis procedure is similar to that of Ulsan in the previous study. After she replaced existing data with L.A. data in the menu (Fig. 1A), she set the MAE filter as 3.6—the value at the third quartile and sorted the table view by the speed distribution. Fig. 9 (left) shows an example sorted table by the distribution.

She observed that most roads have similar speed trends across the clusters. This could be the characteristic of highway roads since the speed is rather faster than the urban road networks. Speed distributions are skewed right and pointy in the table (Fig. 9A), indicating typical highways' speed trends. She also obtained an overview of the positive correlation between standard deviation and MAE from the table. For example, roads 140, 61, 38 had a high error having wide range of speed (Fig. 9C), while roads 40 and 85 with a pointy speed distribution had a low error (Fig. 9B), which confirms **(Finding 1)**. She also noticed from the map view that high error roads exist at the intersections of regions (IR) and edges of regions (ER) [43] (Fig. 8A, roads with black dots inside). This pattern was also found in the previous study.

As she clicked on road 38 on the map view, one of the black dots (i.e., high error), she moved the blue bar to the point where the model recorded a high error (e.g., AE: 45.63, Fig. 9 T1). Then, she instantly noticed several light-pink vertical bars in the attention view (Fig. 8B, attention view) and found that none of the bars had an attention intensity

Table 1. Improved accuracy using the findings from our visual analytics approach. 10% of roads with the highest error are used.

Dataset	Method	15 Mins	30 Mins	60 Mins	Average
Ulsan	ST-GRAT	8.70	9.00	9.52	8.98
	Enforced w/ DTW	8.23	8.43	8.81	8.39
	Enforced w/ DTW + Granger	8.24	8.37	8.76	8.37
METR-LA	ST-GRAT	6.08	7.43	10.61	7.68
	Enforced w/ DTW	6.06	7.38	10.24	7.55
	Enforced w/ DTW + Granger	6.04	7.28	10.08	7.48

larger than 0.1. She thought that this was the same behavior shown in the previous study and may not be a proper behavior because the roads that encode preceding traffic patterns of other roads **(Finding 2)** or refer its own past speed pattern **(Finding 3)** tend to have high accuracy. The head-cluster view also showed similar behaviors from a cluster point of view (Appendix Fig. 13). The head-cluster view displayed many light pink clusters in the diagonal pattern (i.e., insufficient self-reference), when the model recorded high error. However, with sufficient self-reference, it had high accuracy with dark red diagonal patterns.

To further inspect the model's attention behavior, she performed a causality analysis (CA) and found that road 206 has a speed pattern that precedes that of road 38 (i.e., dropping and increasing traffic speed earlier than road 38's traffic speed), as shown in Fig. 9 P1. The Granger causality test confirmed that road 206's speed pattern indeed exceeds that of road 38 ($F[5, 271]=5.9, p<0.001$) as shown in Fig. 8 B1 (Table). But the model rarely attended road 206, nor performed self-referencing with attention intensity less than 0.1 **(Finding 2)**. The pie chart of road 38 (yellow) and attention matrix in Fig. 8B presented this observations. Note that the model had performed self-referencing for road 38's prediction, before the traffic speed fluctuated **(Finding 3)**, as shown in Fig. 8 A, pixel view.

6.3 Attention Enforcement Tests

From the two case studies with AttnAnalyzer, she had three main findings. **Finding 1**: when a road's speed highly fluctuates, the model's error increases; **Finding 2**: the model often fails to find appropriate references with preceding speed trends; and **Finding 3**: the model loses important past self-reference information by dispersing attention to other unimportant roads.

As a next step, she decided to improve the model performance with the findings, using the attention enforcement method, provided in the attention view (Sec. 5.2.4). First, she selected four clusters in the diagonal pattern at Head 4, as shown in Fig. 5, green dotted box. As she clicked the "Test Alternatives" button in the view, a new view popped up that presents two line charts for performance comparison (Fig. 10). From the chart, she observed a shift in the original model's chart, which indicates improved performance. For example, from the chart of the METR-LA data, she recognized that the number of roads with about 30 absolute error (AE) decreased, while those with about 5 AE increased. Similarly, the number of roads with about 20 AE decreased the same as those with about 5 AE increased with the Ulsan data. She thought this result came from the fact that the attention enforcement method had the model focus on self-reference and roads with preceding speed patterns in the same clusters, confirming and using **Finding 2 and 3** together.

7 EXPERT FEEDBACK

To evaluate our VA approach, we conducted a semi-structured interview session (2-hour long) with two domain experts, E3 and E4. E3 and E4 had 6 and 2 years of experience, respectively, as machine learning engineers in the traffic domain, developing deep learning models in the transportation domain. E3 had also participated in meetings to perform task analysis (Sec. 3). Before the session, we performed an online tutorial session on AttnAnalyzer and sent them online access to the system so that they could freely use and test the system. In the interview session, we asked about their challenges at work, showcased the system with the case studies, and had a Q&A session on the system. We then asked about the usefulness and impact of the system on their work.

First, when we asked them about the challenges in their work, they explained that improving models in the traffic domain is difficult, because the data have complicated spatio-temporal dependencies, and they do not have effective tools for analyzing the dependencies in the layers. Thus, the process is performed in a brute-force manner using scripting languages—*“...whenever we evaluate the importance of some features or layers, we have to iteratively retrain our model and check the raw outputs,”* E4 reported. E4 also commented that the attention view and AttnArrows are helpful for analyzing complex spatio-temporal attention, and they can easily discover important features and layers and debug models’ misdeavor in advance, which is critical for shortening their work time due to the reduced number of retraining in the end.

E3 expressed similar views to E4 regarding the system’s usefulness for analyzing attention, as well as exploring and debugging model behaviors. In particular, he mentioned that the system is effective in cases, where DL models do not have specific performance improvement, such as speed predictions for rush hours [36, 37]—*“...there are cases where deep learning models’ performance do not outperform existing approaches. ...using the attention enforcement, we can have improved performance, better responding to user requests with higher accuracy.”*

He then commented that the system allows great inspiration and has large potential in optimizing computing resources in a novel way. He explained that as deep learning models become heavier with many layers and input features, consuming many computing resources, it becomes a burden to utilize heavy models all the time. In this case, he suggested using the system at work with two types of models—one with high accuracy and resource-demanding and another with reasonable accuracy without resource-demanding. Then, the attention mechanism is used to decide which model to use for prediction with consideration of resource usage (i.e., switching models for prediction by conditions). Then they can explore the behaviors of the attention as described in the case studies and even improve performance by using the attention enforcement method, all of which could contribute to saving computing resources—*“...if we can interchange between lightweight and heavy models, [...] we can save tremendous resources for our navigation and prediction services.”*

E3 gave positive feedback on the automated methods. According to E3, DTW and Granger causality methods have already been used and have allowed notable results in their teams dealing with time series data analysis and prediction. E3 said, *“...DTW and the enforcement of attention with DTW is a very interesting and novel approach for the traffic forecasting.”* Lastly E3 spotted that the system can be easily extended for other cities, which is advantageous, as they deal with many large cities for services—*“Confirming models with different cities is also important [...] and this system can be frequently used for confirming model performance with many cities’ data.”*

Lastly, when asked how to strengthen the system, they answered that temporal analysis on the map could be further strengthened, such as direct time-series pattern visualization on the map with respect to roads, which could help developers better explore the dependencies. They also expressed their interest in using the system with DL models for travel time prediction, as it could reveal new types of insights into how deep learning models deal internally with sub-paths and local roads.

8 LIMITATIONS, DISCUSSION, AND FUTURE WORK

We encode clusters using colors because our algorithm experiment results show that there are less than eight meaningful traffic patterns




Fig. 10. A result of the attention enforcement presents error distribution in our urban and highway traffic data due to the inherent periodicity of the traffic on roads (e.g., existence of rush hour). Note that Kwon et al. [33] have demonstrated that a greater number of clusters can be identified with colors (e.g., 20 clusters). Although we assign colors to best distinguish the clusters with reduced visual complexity, assigning colors to clusters based on similarity levels could allow for effective cluster analysis. For example, clusters with high similarities could have a similar color [32]. By modifying Eq. 1, the attention matrix can be calculated from other attention-based spatio-temporal forecasting models. As Table 1 with top 10% high error roads shows, the attention enforcement method results in improved accuracy. While a negative relationship is known to exist between traffic volume and speed [4, 34, 39], it has not been investigated in this work; the model behavior with regard to this relationship could be the topic of a future study. The volume data for the LA dataset can be downloaded from the official website [15]. We place road information in the table view, but some information could be shown with visualization on the map for effective visual analysis, though this would also increase visual complexity.

We perform two experiments and acquire the best performance, when both DTW and Granger causality test are applied. Note that the improvement is meaningful, as most recent traffic forecasting models are recognized with a similar level of performance improvements [37]. The colored parts of the two inset visualizations in Fig. 10 mean improved performance gaps and some roads show significant performance improvement (e.g., road 32, 9). We show the usefulness of DTW and Granger causality tests for spatio-temporal attention analysis, but incorporating other methods (e.g. Bayesian inferences) also could produce additional insights. We use a fixed range of error using quartiles for the sake of simplicity; hence, the high (Q3) and low (Q1) error bounds are relative and different datasets would have different low and high error bounds with different error distributions. A filtering function for setting the error bound threshold could help users effectively investigate model behaviors with different error ranges. Although the experts in the interview have not reported any concerns regarding the usability of the system, a user study could be performed to investigate usability and complexity issues. For in-depth analysis, a longitudinal study [58] can be performed with the experts to show how the system can be used at work and to measure the impact of the system.

As a future work, we plan to research how to make an automated switch based on attention between the light weight and heavy forecasting models for DL resource optimization. It is of interest regarding what VA systems can support the exploration of such models to answer when and how the models are switched during the forecasting process.

9 CONCLUSION

We design a VA approach to help users explore the process of traffic forecasting and improve model performance. We perform task analysis with domain experts, which inform our system design. The system provides users with multiple views, including filter, line, map, and attention views, for effective model exploration in a spatio-temporal perspective. For evaluation, we perform two case studies, showing how users form and validate hypotheses and generate insights into model behaviors. We also show that the insights derived by using our VA approach are critical in improving the accuracy.

ACKNOWLEDGEMENTS

This work was supported by the Korean National Research Foundation (NRF) grant (No. 2021R1A2C1004542) and by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants (No. 2020-0-01336–Artificial Intelligence Graduate School Program. UNIST), funded by the Korea government (MSIT). This work was also partly supported by NAVER Corporation.

REFERENCES

- [1] M. Afzalan and F. Jazizadeh. An automated spectral clustering for multi-scale data. *Neurocomputing*, 347:94–108, 2019.
- [2] M. Akhtar and S. Moridpour. A review of traffic congestion prediction using artificial intelligence. *Journal of Advanced Transportation*, 2021:1–18, 2021.
- [3] N. Andrienko, G. Andrienko, and P. Gatalasky. Exploratory spatio-temporal visualization: an analytical review. *Journal of Visual Languages & Computing*, 14(6):503–541, 2003.
- [4] N. Andrienko, G. Andrienko, and S. Rinzivillo. Leveraging spatial abstraction in traffic analysis and forecasting with visual analytics. *Information Systems*, 57:172–194, 2016.
- [5] M. Aven. Daily traffic flow pattern recognition by spectral clustering. *CMC Senior Theses*, p. 1597, 2017.
- [6] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, p. 359–370, 1994.
- [7] A. A. Cabrera, W. Epperson, F. Hohman, M. Kahng, J. Morgenstern, and D. H. Chau. Fairvis: Visual analytics for discovering intersectional bias in machine learning. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 46–56, 2019.
- [8] C. Chen, J. Hu, Q. Meng, and Y. Zhang. Short-time traffic flow prediction with arima-garch model. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pp. 607–612, 2011.
- [9] C. Chen, Y. Wang, L. Li, J. Hu, and Z. Zhang. The retrieval of intra-day trend and its influence on traffic prediction. *Transportation research part C: emerging technologies*, 22:103–118, 2012.
- [10] S. Chung, C. Park, S. Suh, K. Kang, J. Choo, and B. C. Kwon. Revacnn: Steering convolutional neural network via real-time visual analytics. *Future of Interactive Learning Machines Workshop (FILM at NerIPS)*, 2016.
- [11] C. Collins, G. Penn, and S. Carpendale. Bubble sets: Revealing set relations with isocontours over existing visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1009–1016, 2009.
- [12] J. F. DeRose, J. Wang, and M. Berger. Attention flows: Analyzing and comparing attention mechanisms in language models. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1160–1170, 2021.
- [13] A. Derrow-Pinion, J. She, D. Wong, O. Lange, T. Hester, L. Perez, M. Nunkesser, S. Lee, X. Guo, B. Wiltshire, et al. Eta prediction with graph neural networks in google maps. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 3767–3776, 2021.
- [14] C. Diks and V. Panchenko. A new statistic and practical guidelines for nonparametric granger causality testing. *Journal of Economic Dynamics and Control*, 30(9-10):1647–1669, 2006.
- [15] D. C. Division. Traffic count data: Lac open data, <https://data.lacounty.gov/transportation/traffic-count-data/uview-g569>, Oct 2021.
- [16] X. Fang, J. Huang, F. Wang, L. Zeng, H. Liang, and H. Wang. Constgat: Contextual spatial-temporal graph attention network for travel time estimation at baidu maps. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2697–2705, 2020.
- [17] Google. Model understanding with the what-if tool dashboard. Available at https://www.tensorflow.org/tensorboard/what_if_tool/.
- [18] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):1–42, 2018.
- [19] S. Guo, D. Zhou, J. Fan, Q. Tong, T. Zhu, W. Lv, D. Li, and S. Havlin. Identifying the most influential roads based on traffic correlation networks. *EPJ Data Science*, 8(1):1–17, 2019.
- [20] M. Harrower and C. A. Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.
- [21] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.
- [22] F. Hohman, H. Park, C. Robinson, and D. H. Polo Chau. Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1096–1106, 2020.
- [23] F. M. Hohman, M. Kahng, R. Pienta, and D. H. Chau. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2674–2693, 2019.
- [24] H. V. Jagadish, J. Gehrke, A. Labrinidis, Y. Papakonstantinou, J. M. Patel, R. Ramakrishnan, and C. Shahabi. Big data and its technical challenges. *Commun. ACM*, 57:86–94, 2014.
- [25] M. Kahng, N. Thorat, D. H. Chau, F. Viégas, and M. Wattenberg. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation. *IEEE Transactions on Visualization and Computer Graphics*, 25:310–320, 2019.
- [26] M. Kahng, P. Y. Andrews, A. Kalro, and D. Hornig Polo Chau. Activis: Visual exploration of industry-scale deep neural network models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):88–97, 2017.
- [27] D. A. Keim. Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):59–78, 2000.
- [28] J. B. Kenney. Dedicated short-range communications (dsrc) standards in the united states. *Proceedings of the IEEE*, 99(7):1162–1182, 2011.
- [29] T. N. Kipf and M. Welling. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the International Conference on Learning Representations*, 2017.
- [30] S. Ko, S. Afzal, S. Walton, Y. Yang, J. Chae, A. Malik, Y. Jang, M. Chen, and D. Ebert. Analyzing high-dimensional multivariate network links with integrated anomaly detection, highlighting and exploration. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 83–92, 2014.
- [31] S. Ko, R. Maciejewski, Y. Jang, and D. S. Ebert. Marketanalyzer: An interactive visual analytics system for analyzing competitive advantage using point of sale data. *Computer Graphics Forum*, 31(3pt3):1245–1254, 2012.
- [32] A. Kumar, N. Timmermans, M. Burch, and K. Mueller. Clustered eye movement similarity matrices. In *Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications*, pp. 1–9, 2019.
- [33] B. C. Kwon, B. Eysenbach, J. Verma, K. Ng, C. deFilippi, W. F. Stewart, and A. Perer. Clustervision: Visual supervision of unsupervised clustering. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):142–151, 2018.
- [34] I. Lana, J. Del Ser, M. Velez, and E. I. Vlahogianni. Road traffic forecasting: Recent advances and new challenges. *IEEE Intelligent Transportation Systems Magazine*, 10(2):93–109, 2018.
- [35] V. Le Guen and N. Thome. Shape and time distortion loss for training deep time series forecasting models. *Advances in neural information processing systems*, 32, 2019.
- [36] C. Lee, Y. Kim, S. Jin, D. Kim, R. Maciejewski, D. Ebert, and S. Ko. A visual analytics system for exploring, monitoring, and forecasting road traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3133–3146, 2020.
- [37] H. Lee, C. Park, S. Jin, H. Chu, J. Choo, and S. Ko. An empirical experiment on deep learning models for predicting traffic data. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, pp. 1817–1822, 2021.
- [38] L. Li, X. Su, Y. Zhang, Y. Lin, and Z. Li. Trend modeling for traffic time series analysis: An integrated study. *IEEE Transactions on Intelligent Transportation Systems*, 16(6):3430–3439, 2015.
- [39] Y. Li and C. Shahabi. A brief overview of machine learning methods for short-term traffic forecasting and future directions. *SIGSPATIAL Special*, 10(1):3–9, 2018.
- [40] Y. Li, R. Yu, C. Shahabi, and Y. Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *Proceedings of the International Conference on Learning Representations*, pp. 1–16, 2018.
- [41] M. Liu, S. Liu, H. Su, K. Cao, and J. Zhu. Analyzing the noise robustness of deep neural networks. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 60–71, 2018.
- [42] M. Liu, J. Shi, Y. Li, C. Li, J. Zhu, and S. Liu. Towards better analysis of deep convolutional neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 23:91–100, 2016.
- [43] Los Angeles City Planning. Citywide maps. Available at <https://planning.lacity.org/>.
- [44] N. Madzlan, K. Ibrahim, et al. Arima models for bus travel time prediction. *Journal of the Institution of Engineers Malaysia*, 2010.
- [45] W. Meulemans, N. H. Riche, B. Speckmann, B. Alper, and T. Dwyer. Kelpfusion: A hybrid set visualization technique. *IEEE Transactions on Visualization and Computer Graphics*, 19(11):1846–1858, 2013.
- [46] Y. Ming, S. Cao, R. Zhang, Z. Li, Y. Chen, Y. Song, and H. Qu. Understanding hidden memories of recurrent neural networks. In *IEEE Conference on Visual Analytics Science and Technology*, pp. 13–24, 2017.

- [47] S. Miyahara and S. Miyamoto. A family of algorithms using spectral clustering and dbscan. In *IEEE International Conference on Granular Computing*, pp. 196–200, 2014.
- [48] T. Mühlbacher, H. Piringer, S. Gratzl, M. Sedlmair, and M. Streit. Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1643–1652, 2014.
- [49] PAIR. What-if tool. Available at <https://pair-code.github.io/what-if-tool/>.
- [50] Z. Pan, Y. Liang, J. Zhang, X. Yi, Y. Yu, and Y. Zheng. Hyperst-net: Hypernetworks for spatio-temporal forecasting. *ArXiv*, abs/1809.10889, 2018.
- [51] C. Park, C. Lee, H. Bahng, Y. Tae, S. Jin, K. Kim, S. Ko, and J. Choo. St-grat: A novel spatio-temporal graph attention networks for accurately forecasting dynamically changing road speed. In *Proceedings of the ACM International Conference on Information & Knowledge Management*, p. 1215–1224, 2020.
- [52] C. Park, I. Na, Y. Jo, S. Shin, J. Yoo, B. C. Kwon, J. Zhao, H. Noh, Y. Lee, and J. Choo. Sanvis: Visual analytics for understanding self-attention networks. *IEEE Conference on Visual Analytics Science and Technology Short*, 2019.
- [53] M. Pensky and T. Zhang. Spectral clustering in the dynamic stochastic block model. *Electronic Journal of Statistics*, 13(1):678–709, 2019.
- [54] N. Pezzotti, T. Holtt, J. Van Gemert, B. Lelieveldt, E. Eiseemann, and A. Vilanova. Deepeyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):98, 2018.
- [55] M. Pi, H. Yeon, H. Son, and Y. Jang. Visual cause analytics for traffic congestion. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2186–2201, 2021.
- [56] J. Pu, S. Liu, Y. Ding, H. Qu, and L. Ni. T-watcher: A new visual analytic system for effective traffic surveillance. In *IEEE International Conference on Mobile Data Management*, vol. 1, pp. 127–136, 2013.
- [57] S. Salvador and P. Chan. Toward accurate dynamic time warping in linear time and space. *Intelligent Data Analysis*, 11(5):561–580, 2007.
- [58] P. Saraiya, C. North, V. Lam, and K. A. Duca. An insight-based longitudinal study of visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1511–1522, 2006.
- [59] Q. Shen, Y. Wu, Y. Jiang, W. Zeng, K. Alexis, A. Vianova, and H. Qu. Visual interpretation of recurrent neural network on multi-dimensional time-series forecast. In *2020 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 61–70. IEEE, 2020.
- [60] H. Strobelt, S. Gehrmann, M. Behrisch, A. Perer, H. Pfister, and A. M. Rush. SEQ2seq-VIS: A Visual Debugging Tool for Sequence-to-Sequence Models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):353–363, 2019.
- [61] H. Strobelt, S. Gehrmann, H. Pfister, and A. M. Rush. Lstmvis: A tool for visual analysis of hidden state dynamics in recurrent neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):667–676, 2017.
- [62] D. A. Tedjopurnomo, Z. Bao, B. Zheng, F. Choudhury, and A. Qin. A survey on modern deep neural network for traffic prediction: Trends, methods and challenges. *IEEE Transactions on Knowledge and Data Engineering*, 34(4):1544–1561, 2022.
- [63] J. Thomas and K. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Ctr, 2005.
- [64] E. R. Tufte. *Beautiful evidence*. Graphis Pr, 2006.
- [65] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, vol. 30, pp. 5998–6008, 2017.
- [66] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.
- [67] J. Wang, L. Gou, H. Shen, and H. Yang. Dqnviz: A visual analytics approach to understand deep q-networks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):288–298, 2019.
- [68] J. Wang, L. Gou, H. Yang, and H.-W. Shen. Ganviz: A visual analytics approach to understand the adversarial game. *IEEE Transactions on Visualization and Computer Graphics*, 24:1905–1917, 2018.
- [69] Z. Wang, M. Lu, X. Yuan, J. Zhang, and H. Van De Wetering. Visual traffic jam analysis based on trajectory data. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2159–2168, 2013.
- [70] J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viégas, and J. Wilson. The what-if tool: Interactive probing of machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):56–65, 2020.
- [71] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, and C. Zhang. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 753–763, 2020.
- [72] Z. Wu, S. Pan, G. Long, J. Jiang, and C. Zhang. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 1907–1913, 2019.
- [73] P. Xiao, N. Liu, Y. Li, Y. Lu, X.-j. Tang, H.-w. Wang, and M.-x. Li. A traffic classification method with spectral clustering in sdn. In *2016 17th International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT)*, pp. 391–394, 2016.
- [74] H. Yao, X. Tang, H. Wei, G. Zheng, and Z. Li. Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 5668–5675, 2019.
- [75] X. Yin, G. Wu, J. Wei, Y. Shen, H. Qi, and B. Yin. Deep learning on traffic prediction: Methods, analysis and future directions. *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–17, 2021.
- [76] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. In *Proceedings of the International Conference on Machine Learning*, 2015.
- [77] B. Yu, H. Yin, and Z. Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 3634–3640, 2018.
- [78] Z. Yuan, X. Zhou, and T. Yang. Hetero-convlstm: A deep learning approach to traffic accident prediction on heterogeneous spatio-temporal data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 984–992, 2018.
- [79] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *Proceedings of the European Conference on Computer Vision*, pp. 818–833, 2014.
- [80] W. Zeng, C.-W. Fu, S. M. Arisona, and H. Qu. Visualizing interchange patterns in massive movement data. *Computer Graphics Forum*, 32(3-3):271–280, 2013.
- [81] W. Zeng, C. Lin, J. Lin, J. Jiang, J. Xia, C. Turkyay, and W. Chen. Revisiting the modifiable areal unit problem in deep traffic prediction with visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):839–848, 2020.
- [82] J. Zhang, X. Shi, J. Xie, H. Ma, I. King, and D. Yeung. Gaan: Gated attention networks for learning on large and spatiotemporal graphs. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pp. 339–349, 2018.