

Space Meets Time: Local Spacetime Neural Network For Traffic Flow Forecasting

Song Yang, Jiamou Liu, and Kaiqi Zhao

School of Computer Science

The University of Auckland, Auckland, New Zealand

syan382@aucklanduni.ac.nz, {jiamou.liu, kaiqi.zhao}@auckland.ac.nz

Abstract—Traffic flow forecasting is a crucial task in urban computing. The challenge arises as traffic flows often exhibit intrinsic and latent spatio-temporal correlations that cannot be identified by extracting the spatial and temporal patterns of traffic data separately. We argue that such correlations are universal and play a pivotal role in traffic flow. We put forward spacetime interval learning as a paradigm to explicitly capture these correlations through a unified analysis of both spatial and temporal features. Unlike the state-of-the-art methods, which are restricted to a particular road network, we model the universal spatio-temporal correlations that are transferable from cities to cities. To this end, we propose a new spacetime interval learning framework that constructs a local-spacetime context of a traffic sensor comprising the data from its neighbors within close time points. Based on this idea, we introduce local spacetime neural network (STNN), which employs novel spacetime convolution and attention mechanism to learn the universal spatio-temporal correlations. The proposed STNN captures local traffic patterns, which does not depend on a specific network structure. As a result, a trained STNN model can be applied on any unseen traffic networks. We evaluate the proposed STNN on two public real-world traffic datasets and a simulated dataset on dynamic networks. The experiment results show that STNN not only improves prediction accuracy by 4% over state-of-the-art methods, but is also effective in handling the case when the traffic network undergoes dynamic changes as well as the superior generalization capability.

I. INTRODUCTION

With the increasing volume of traffic data and the growing impacts of data-driven technologies in modern transportation systems, the traffic flow forecasting problem is drawing increasing attention [1]. Reliable and timely predictions of the traffic dynamics can assist transportation management, help alleviate traffic congestion, and enhance traffic efficiency.

A traffic system is characterised by the changing flows at locations in a *road network*, i.e., nodes (sensors) interlinked by road segments, from which one may observe salient patterns: sudden bursts, drastic fluctuations, and periodic shifts. One can develop an understanding of these patterns from two perspectives. On the outside, traffic is the accumulation of various extrinsic factors, from road layout, to geographic features of the environment, to traffic laws, and erratic behaviours of drivers, all playing a role in shaping traffic conditions. On the inside, intrinsic principles are governing the flow of traffic, as vehicles – particles in the road system – maintain stable directional, speed and concentration features as they travel through space and time, giving rise to some *universal*




Fig. 1: An illustration of spacetime interval. The figure shows three snapshots of the network. The target node’s state at t_3 is strongly influenced by the state of a at t_1 and the state of b at t_2 , but only weakly influenced by the state of a at t_2 and the state of b at t_1 .

patterns of the traffic flow. These universal patterns are what is behind well-established traffic flow models, such as Wardrop’s equilibria, and Kerner’s three-phase traffic theory, that describe traffic flow from a mathematical point of view [2].

We argue that accurate predictions of traffic flow – based on macro-level traffic data – rest upon a model’s ability to grasp not only extrinsic features from the given data set, but also intrinsic principles that are universal to any traffic systems. At its core, traffic can be seen as a physical system [3] where localised changes cascade like waves along roads, to some other locations, some time later [4]. Therefore the key to capturing the intrinsic principles of traffic flow lies in extracting the latent correlations between locations’ present states and surrounding locations’ past.

Recent breakthroughs in neural-based techniques, in particular those based on graph neural networks (GNNs) [5], [6], [7], [8], represent a significant step towards representing non-linear traffic features by leveraging topological constraints while achieving better prediction results as compared to earlier statistical methods such as autoregressive models [9]. However, existing GNN-based traffic forecasting models have three common limitations: (1) These models heavily rely on the graph structure which varies across different road networks. As such, they are restricted to a particular road network, rather than revealing intrinsic properties of traffic system; (2) They apply computationally expensive feature aggregation operations (e.g., graph convolutions takes) on the entire net-

work, and thus are not scalable to large road networks that contain hundreds of thousands of sensors. For instance, graph convolution network (GCN) scales quadratically w.r.t. the number of nodes in the network. (3) These models all consist of *separate components* to extract features from the spatial and the temporal dimensions, respectively, before integrating them into the same feature map to derive a prediction. The extracted feature map represents “aggregated” correlations between locations of the network across time. This setup implicitly assumes that the correlations between locations stay uniform over time. However, two locations may have different correlations at different times: As indicated in Figure 1, against the current (t_3) target node, node a has a stronger correlation at an earlier time (i.e., t_1) than a later time (i.e., t_2), while b has a weaker correlation at t_1 than at t_2 .

To break the above limitations in existing GNN-based models, we propose a new spatio-temporal correlation learning paradigm, called *Spacetime Interval Learning*, which fuses the spatial dimensions and the temporal dimension into a single manifold, i.e., a *spacetime*, as coined in special relativity theory. The correlations that we aim to capture reflects a form of “distance” between two traffic events in spacetime, where the distance is referred to as *interval* in spacetime terminologies [10]. Specifically, the proposed spacetime interval learning paradigm extracts the traffic data of nearby sensors within a fixed time window, regarded as the *local-spacetime context* of each target location. The local-spacetime context is analogous to the “receptive field” of sensors to ignore the sensors that do not contribute much to the final prediction. Then, the spacetime interval learning paradigm learns to correlate a node x at a given time t with another node y at a different time $t' \geq t$ in the local-spacetime context.

The proposed learning paradigm addresses the limitations of the GNN-based models as follows. Firstly, it models the spatio-temporal correlations within the local-spacetime context of each target location. In this way, the model we build is independent of the graph structure, and thus is *universal* for traffic flows at large, as opposed to for any specific network. Secondly, our method focuses on the local-spacetime, such that it changes the traffic prediction from *network-level* to *node-level*. That is, our model can be predict the traffic for multiple locations of interest in parallel on different machines. Thirdly, unlike existing methods, we fuse the spatial and the temporal dimensions into a single manifold to capture the varying spatial correlations between locations across time.

Under the spacetime interval learning paradigm, we propose an implementation called Spacetime Neural Network (STNN), which combines novel *spacetime attention blocks* and *spacetime convolutional blocks*. The former highlights the interval between events, while the latter aggregates the learned features from spatial, temporal, and spatio-temporal aspects. The main advantage of the model are as follows. The code is available on <https://github.com/songyangco/STNN>.

1) *Accuracy*: The learned spacetime intervals explicitly capture the correlations between locations across different time instances, making the learned features highly informative

in traffic prediction. As validated through experiments on two real-world datasets, METR-LA and PeMS-Bay, when trained on data from a single network, our model outperforms existing methods, with, e.g., 4% improvement over current state-of-the-art complicated models for 15-min predictions. See Section VI-B.

- 2) *Transferrability* amounts to a main advantage of our universal model. For this, we train a single STNN model on one dataset and use it without fine-tuning on unseen traffic datasets. In all cases, the results achieve competitive, if not better, performance compared to state-of-the-art approaches, validating the applicability of our universal model. See Section VI-B.
- 3) *Ability to handle dynamic networks* is another natural consequence of converting raw input to a local-spacetime for every target node. This refers to, e.g., the network undergoes structural changes such as the blocking or addition of roads. We perform an experiment on a synthetic dynamic network. Our model is much better than baselines. See Section VI-C.
- 4) *Scalability*: Previous GNN-based models all assume a small fixed network whose size is limited by computational capability. By performing prediction on node-level, we break the limit on the network size. The complexity of our model is independent from the network size, making our model applicable to networks of arbitrary size. See Section VI-F.

II. RELATED WORKS

Studies on traffic forecasting through spatio-temporal data analysis can roughly be divided into four types.

Statistical methods. Early studies in the 2000s use various statistical methods, such as historical average (HA), autoregressive integrated moving average (ARIMA) [9], and vector autoregression (VAR) [11] for the traffic forecasting problem. However, these approaches assume the input time series to be stationary and univariate. This assumption may not hold in a complex traffic system. Thus, statistical methods often generate inaccurate predictions.

Machine learning methods. To account for more complex traffic data, various classical machine learning methods, such as kNN [12], SVM [13], SVR [14], are employed in the early 2010s. One shortcoming of these methods is that they fail to capture the highly non-linear spatial and temporal patterns present in the data.

Deep learning methods. After 2015, capitalising on the success of deep learning [15], several deep learning models are proposed to capture the underlying patterns of traffic data. Recurrent neural networks (RNNs) are adopted to analyse time series patterns [16], [17], while convolutional neural networks (CNNs) are used to capture spatial dependencies by treating a traffic network as a grid [18], [19]. However, these methods omit the road network, which is an essential spatial constraint on how traffic moves in the spatial dimension. This limitation inhibits traditional deep learning methods from accurate traffic forecasting.

TABLE I: Existing methods fall into the same paradigm

Studies	Spatial layers	Temporal layers
DCRNN (2017) [20]	Diffusion Conv	GRU
STGCN (2017) [5]	GCN	Gated TCN
DMVST-Net (2018) [6]	CNN	LSTM
STDN (2019) [23]	CNN	LSTM + Attn
GraphWaveNet (2019) [24]	GCN	Gated TCN
DSTGCN (2019) [7]	GCN	TCN
ASTGCN (2019) [22]	GCN + Attn	TCN + Attn
SLCNN (2020) [8]	GCN	Gated TCN
GMAN (2020) [25]	Spatial Attn	Temporal Attn
MTGCN (2020) [26]	GCN	Gated TCN
STNN (ours)	ST Conv + Attn	

Graph neural networks. More recently, graph-based deep learning models are increasingly used for spatio-temporal data analysis. Yu et al. [5] propose a spatio-temporal graph convolutional network based on spectral graph theory to extract features from the road network and the historical time series data. In particular, the spatial patterns are captured by graph convolutional layers while the temporal patterns are extracted by gated convolutional layers. Li et al. [20] model the traffic flow as a diffusion process and capture the spatial dependencies by random walks on a graph followed by an RNN to extract temporal patterns. Yao et al. [6] use graph embeddings to capture the road network information and integrate Long-Short Term Memory (LSTM) and CNN with the graph embeddings for traffic prediction. Attention mechanism has also been adopted to learn spatial and temporal patterns with improved performance [21], [22]. However, all of these methods fall into the same paradigm that separate layers are applied to extract spatial patterns and temporal patterns respectively. We summarize current methods under this paradigm in Table I. Particularly, CNN, Graph Convolutional Network (GCN), or spatial attention are often used to learn spatial patterns. Gated Recurrent Unit (GRU), LSTM, gated temporal convolution (gated TCN), or temporal attention are used to learn temporal patterns.

Unlike current methods, this work proposes a novel learning paradigm. Instead of using separate components or layers to learn spatial and temporal patterns, we design a novel space-time module to learn the local spatio-temporal correlations, and capture both spatial and temporal patterns at the same time. Our model does not fall into any existing categories. Specially, instead of exploiting GNN to learn spatial patterns, we fuse spatial features with temporal features together and learned by a unified layer.

III. PROBLEM FORMULATION

We summarise the list of important notations in Table II.

Traffic condition like vehicle speed is often recorded by the road-side sensor station v along with other auxiliary information including time, sensor location, etc. Aggregating N sensors together forms a sensor set $V \in \mathbb{R}^N$ which monitor the traffic flow of a certain area.

TABLE II: Table of Notations

Notation	Description
$V^{(t)}$	sensor set at time t
V	union sensor set of all time steps
$Q^{(t)}$	distance matrix at time t
\mathcal{Q}	stack $Q^{(t)}$ along time axis
$\mathbf{x}_i^{(t)}$	a vector of features recorded by i -th sensor at time t
$\mathbf{X}^{(t)}$	a matrix of features of all sensors at time t
\mathcal{X}	stack $\mathbf{X}^{(t)}$ along time axis
V_{tar}	a set of target nodes to predict
v_p	a target node
V_p	a set of neighboring nodes that close to v_p
F	number of features captured by each sensor station
$y_i^{(t)}$	future traffic measure of i -th sensor at time t
\mathbf{y}_i	future traffic measures of i -th sensor
\mathcal{Y}	future traffic measures of all nodes in V_{tar}
$\mathbf{A}^{(t)}$	normalized connectivity matrix at time t
$\mathbf{D}_p^{(t)}$	a local-spacetime snapshot for v_p at time t
\mathcal{D}_p	a local-spacetime for v_p

At each time step t , the i -th sensor station v_i records local traffic measures such as vehicle average speed in a fixed time duration and supplementary data. These measurements form a feature vector $\mathbf{x}_i^{(t)} \in \mathbb{R}^F$. In particular, we have two features in our input data: average speed, timestamp. Collecting the feature vectors of all sensor stations together, we obtain the overall observation of the traffic network at time step t as $\mathbf{X}^{(t)} = (\mathbf{x}_1^{(t)}, \mathbf{x}_2^{(t)}, \dots, \mathbf{x}_i^{(t)}) \in \mathbb{R}^{N \times F}$. A traffic flow dataset contains the measures of a sequence of T_h time steps, where $T_h > 0$, and is thus presented as a *sensor feature tensor* $\mathcal{X} := (\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots, \mathbf{X}^{(T_h)}) \in \mathbb{R}^{N \times F \times T_h}$.

Based on the location of sensor stations, we can get the travel distance on road network between any two sensors. Consequently, we built a matrix \mathbf{Q} to reflect the distances of all sensor pairs. One can think \mathbf{Q} as the weighted adjacent matrix of a directed complete graph. The edge weight is affected by the actual distance and complete graph is because we can always calculate the distance between two sensors. Furthermore, we do not require \mathbf{Q} remain unchanged all the time. On the contrary, the travel distance between two sensors might be variable because of events like traffic accidents or road constructions which alter the underlying topology of road network, leading to a further travel distance. Thus, to better incorporate the network dynamics in our model, we use $\mathbf{Q}^{(t)}$ to just indicate the spatial context at the particular time step t . $\mathcal{Q} := (\mathbf{Q}^{(1)}, \mathbf{Q}^{(2)}, \dots, \mathbf{Q}^{(T_h)}) \in \mathbb{R}^{N \times N \times T_h}$ reflects the underlying road network structure dynamics over T_h time steps, where $T_h > 0$. For a static network, \mathcal{Q} reduces to \mathbf{Q} .

Traffic flow forecasting: Given a sensor feature tensor \mathcal{X} and a spatial context tensor \mathcal{Q} over past T_h time steps, as well as a set of nominated *target nodes* $V_{\text{tar}} = \{v_1, v_2, \dots, v_M\} \subseteq V$. The traffic flow forecasting problem aims to predict the traffic flow of the next T_r time steps for every target node in V_{tar} . The desired output is thus $\mathcal{Y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M) \in \mathbb{R}^{M \times T_r}$ where $\mathbf{y}_p = (y_p^{T_h+1}, y_p^{T_h+2}, \dots, y_p^{T_h+T_r}) \in \mathbb{R}^{T_r}$ denotes the

traffic conditions of the next T_r time steps at node $v_p \in V_{\text{tar}}$.

IV. SPACETIME INTERVAL LEARNING

Spacetime interval learning aims to discover the influence between traffic events in a single *spacetime manifold* instead of modeling the spatial and temporal dimensions separately. We define traffic events and spacetime as follows:

Definition 1 (Traffic Event): Given a traffic measurement s (e.g., speed) observed at sensor v_i and time t , a traffic event is a tuple consists of the measurement, time, and location, namely, (s, t, v_i) .

A traffic event, defined analogously to the notion of events in physics [27], embodies both spatial and temporal information of a traffic measurement by a sensor station.

Definition 2 (Spacetime): Given a subset of sensors $V' \subset V$, the sensor features \mathbf{X}' , the related spatial context \mathbf{Q}' , and the T time steps historical data, the spacetime \mathcal{D} is a manifold consists of all traffic events produced by $\{\mathbf{X}'^{(1)}, \mathbf{X}'^{(2)}, \dots, \mathbf{X}'^{(T)}\}$ and $\{\mathbf{Q}'^{(1)}, \mathbf{Q}'^{(2)}, \dots, \mathbf{Q}'^{(T)}\}$.

A spacetime manifold can be organized as a 3D structure where first dimension corresponding to the time, the second dimension corresponding to the space, and the third is the observed traffic measurement. The time dimension is easy to understand but the space data needs more work. In order to squeeze the spatial information into a single dimension, we transfer the sensor coordinates to the travel distance with the anchor/target sensor. As a result, we construct a local-spacetime around a target sensor and use that to predict the future traffic flow of that particular sensor.

Definition 3 (Local-spacetime): The local-spacetime \mathcal{D}_p for sensor v_p is a subset of \mathcal{D} , which only contains the traffic events occur at nearby location of v_p and recent time steps. Furthermore, the sensor location in all traffic events from \mathcal{D}_p is replaced by its travel distance with v_p .

The *spacetime interval* between two traffic events denotes the extent to which one event influences the other; a smaller interval means a closer association between two traffic events. In a local-spacetime, we only care about the intervals between traffic events of the target sensor with other sensors.

Definition 4 (Spacetime Interval): Spacetime interval is the quantified influence of a traffic event imposed on another traffic event regarding to the traffic measurement.

A crucial step in the proposed learning paradigm is to build the local-spacetime of a target node v_p , which is then used for spacetime interval learning. In this way, the trained model not only captures the spacetime interval explicitly, but also be able to generalize learned patterns to other local-spacetimes in the same road network or event different city. We now give details of how to construct a local-spacetime for an arbitrary node $v_p \in V_{\text{tar}}$ where V_{tar} is a set of nodes we want to predict. The pseudocode is outlined in Algorithm 1.

Given a time step t , we define a connectivity matrix $\mathbf{A}^{(t)} \in V^{(t)^2} \rightarrow [0, 1]$ by applying the Gaussian kernel on the distance matrix $\mathbf{Q}^{(t)}$ to convert travel distance to a weight reflect the

Algorithm 1: Local-spacetime construction

Input: $\mathcal{X} = (\mathbf{X}^{(1)}, \mathbf{X}^{(2)}, \dots, \mathbf{X}^{(T_h)}) \in \mathbb{R}^{N \times F \times T_h}$
 $\mathcal{Q} = (\mathbf{Q}^{(1)}, \mathbf{Q}^{(2)}, \dots, \mathbf{Q}^{(T_h)}) \in \mathbb{R}^{N \times N \times T_h}$
Target nodes: $V_{\text{tar}} = \{v_1, v_2, \dots, v_M\}$
Parameter: $\epsilon > 0, \theta > 0$
Output: Local-spacetime list $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_M\}$

- 1 Set $\mathbf{A}^{(t)} \leftarrow \exp\left(-\left(\mathbf{Q}^{(t)}\right)^2 / \theta^2\right) \forall t = 1, \dots, T_h;$
- 2 **for** $v_p \in V_{\text{tar}}$ **do**
- 3 **for** $t \leq T_h$ **do**
- 4 **if** $\mathbf{A}^{(t)}(i, p) > \epsilon$ **or** $\mathbf{A}^{(t)}(p, i) > \epsilon$ **then**
- 5 Set $V_p \leftarrow V_p \cup \{v_i\};$
- 6 **end**
- 7 **end**
- 8 **for** $t \leq T_h$ **do**
- 9 **for** $v_i \in V_p$ **do**
- 10 Set $\mathbf{D}_p^{(t)}(i) \leftarrow [\mathbf{x}_i^{(t)}; \mathbf{A}^{(t)}(i, p)] \in \mathbb{R}^{F+1};$
- 11 **end**
- 12 **end**
- 13 $\mathcal{D}_p \leftarrow (\mathbf{D}_p^{(1)}, \dots, \mathbf{D}_p^{(T_h)}) \in \mathbb{R}^{|V_p| \times (F+1) \times T_h};$
- 14 **end**

connectivity of two sensors. Longer distance corresponds to smaller weight, namely,

$$\mathbf{A}^{(t)}(i, j) := \exp\left(-\frac{\mathbf{Q}^{(t)}(i, j)^2}{\theta^2}\right) \in [0, 1] \quad (1)$$

where i, j are any two nodes in $V^{(t)}$ and θ is a hyper-parameter. Empirically, we set θ as the standard deviation among all $\mathbf{Q}^{(t)}(i, j)$. Intuitively, $\mathbf{A}^{(t)}(i, j)$ is a normalized value that expresses how easy to travel from v_i to v_j in the network at time step t . Note that Equation (1) guarantees $\mathbf{A}^{(t)}(i, i) = 1$. The superscript t used to depict the variational travel distance caused by the underlying networks structure change.

Using $\mathbf{A}^{(t)}$, we extract a set of nodes, denoted as V_p , that could benefit the future traffic flow prediction for the target node v_p as

$$V_p := \left\{ v_i \mid \max\{\mathbf{A}^{(t)}(i, p), \mathbf{A}^{(t)}(p, i) \mid 1 \leq t \leq T_h\} > \epsilon \right\} \quad (2)$$

where ϵ is a pre-determined threshold parameter that indicates how close a node should be from (or to) v_p to be considered relevant to the traffic flow at v_p . Note that $\mathbf{A}^{(t)}(i, p)$ can be different from $\mathbf{A}^{(t)}(p, i)$ because $\mathbf{Q}^{(t)}$ is not symmetric. ϵ control the trade-off between computational cost and prediction accuracy. Moreover, to have the fixed shape input data for training, we require the size of V_p to be α , namely, $|V_p| = \alpha$. This is done by keeping the nearest α neighbors only if $|V_p| > \alpha$ or add dummy nodes if $|V_p| < \alpha$. Dummy nodes have no connections with other nodes, and their features are all zero. As a results, nodes in V_p are sorted based on the travel distance to v_p in ascending order.

We now construct a local-spacetime \mathcal{D}_p of v_p . For $t \in \{1, \dots, T_h\}$, each row of the matrix $\mathbf{D}_p^{(t)} \in \mathbb{R}^{|V_p| \times (F+1)}$ is defined as:

$$\mathbf{D}_p^{(t)}(v_i) \leftarrow [\mathbf{x}_i^{(t)}; \mathbf{A}^{(t)}(i, p)], v_i \in V_p \quad (3)$$

where $\mathbf{x}_i^{(t)}$ is the traffic measurement recorded at node v_i and time step t , and $[\cdot; \cdot]$ denotes concatenation. The matrix $\mathbf{D}_p^{(t)}$ can be regarded as a snapshot of local-spacetime \mathcal{D}_p . It encodes spatial relationship between v_p and those nodes in V_p , as well as traffic measurements of all these nodes at time step t . Finally, define the local-spacetime \mathcal{D}_p by

$$\mathcal{D}_p := (\mathbf{D}_p^{(1)}, \mathbf{D}_p^{(2)}, \dots, \mathbf{D}_p^{(T_h)}) \in \mathbb{R}^{|V_p| \times (F+1) \times T_h} \quad (4)$$

Since \mathcal{D}_p contains all the information we need to train a predictive model of the future traffic flow at v_p , the learning process is independent on the size of the entire network, thus resolving the scalability issue. Moreover, the incorporation of the network information at all time steps $t \in \{1, \dots, T_h\}$ makes the model capable of handling dynamic network topology.

In summary, the spacetime interval learning framework reduces the traffic flow forecasting problem to learning a universal function $g(\cdot)$ that maps \mathcal{D}_p to the future traffic flow of v_p and $p \in \{1, \dots, n\}$:

$$(\mathbf{D}_p^{(1)}, \mathbf{D}_p^{(2)}, \dots, \mathbf{D}_p^{(T_h)}) \xrightarrow{g(\cdot)} (y_p^{T_h+1}, \dots, y_p^{T_h+T_r}) \quad (5)$$

In next section, we propose a novel model as a realization of the mapping function $g(\cdot)$, namely, the spacetime neural network.

V. SPACETIME NEURAL NETWORK

The core task of spacetime interval learning is to adaptively learn the influence between traffic events and use them to predict the future traffic flow. To this end, we propose *SpaceTime Neural Network (STNN)*, which is an end-to-end deep learning model that realizes the function $g(\cdot)$ in Equation 5.

Design Principles. To quantify the influence of traffic events in the local-spacetime to the prediction of future traffic flow on the target sensor, a natural idea is to learn how each single traffic event influences the prediction. However, we observe that traffic events in local-spacetime are not independent and they may interfere each other. For example, congestion in an arterial road may increase the traffic in a surrounding region and then diffuse to further road segments. As such, we design our model with two main principles: (1) the model should be able to capture the pair-wise influence, and (2) it needs to be aware of the conditions of nearby regions, capturing the many-to-one influence.

Model Architecture. Following the above two principles, we propose a novel spacetime module, on top of which we build the proposed STNN model. A spacetime module comprises a *Spacetime Attention Block* (Section V-A) to capture the pair-wise influence, and a *Spacetime Convolution Block* (Section V-B) to capture the many-to-one influence.

The overall architecture of STNN is shown in Figure 2. The input is a local-spacetime of the target (central) node represented by the blue cylinder. STNN employs several spacetime modules to predict the future traffic of the target node.

More formally, convert the raw data to the local-spacetime \mathcal{D}_p and permute \mathcal{D}_p to the channel-first format, i.e., $\mathcal{D}_p \in \mathbb{R}^{(F+1) \times N \times T_h}$. Then the input \mathcal{D}_p is transformed using 1×1 convolution to map traffic events to a high-dimensional feature space. Next, we stack k spacetime modules where k controls the trade-off between model complexity and performance. A smaller k leads to a model with fewer parameters and unlikely to overfits the data. A larger k gives better performance but may jeopardize the generalization ability. Residual connections are also applied for each module. Last, we use a fully connected layer with linear activation to present the prediction results $(\hat{y}_p^{T_h+1}, \dots, \hat{y}_p^{T_h+T_r})$.

A. Spacetime Attention Block

Inspired by self-attention mechanism [28], we propose the *spacetime attention* block to automatically discover the pairwise influence between traffic events. Specifically, for the attention block at the l th layer, the input is a 3D tensor which is the output of previous convolutional layer or the initial input data, i.e., $\mathcal{D}_p^{[l-1]} \in \mathbb{R}^{C^{[l-1]} \times |V_p| \times T_h}$, where $C^{[l-1]}$ denotes the number of channels for feature maps output by the previous layer. Note that $\mathcal{D}_p^{[0]} = \mathcal{D}_p$. Let $U = |V_p| \times T_h$ and we calculate the attention map as

$$\mathbf{S} = (\mathbf{W}_q \mathcal{D}_p^{[l-1]})^T \mathbf{W}_k \mathcal{D}_p^{[l-1]} \in \mathbb{R}^{U \times U} \quad (6)$$

$$\mathbf{S}'_{i,j} = \frac{\exp(\mathbf{S}_{i,j})}{\sum_{j=1}^U \exp(\mathbf{S}_{i,j})} \quad (7)$$

where $\mathbf{W}_q \in \mathbb{R}^{C^{[l]} \times C^{[l-1]}}$ and $\mathbf{W}_k \in \mathbb{R}^{C^{[l]} \times C^{[l-1]}}$ are learnable weight matrices that project each traffic event into feature spaces corresponding to “query” and “key”, respectively. As such, the dot product of the “query” and “key”, i.e., $\mathbf{S}_{i,j}$ indicates the relevance between the i -th and j -th traffic events. The softmax function is used to guarantee the attention weights sum to one. Then, we apply another learnable weight matrix $\mathbf{W}_v \in \mathbb{R}^{C^{[l]} \times C^{[l-1]}}$ to project input data into an output feature space and use the attention matrix to weight the contribution from different traffic events:

$$\mathcal{D}_p^{[l]} = \mathbf{W}_v \mathcal{D}_p^{[l-1]} \mathbf{S}'^T \in \mathbb{R}^{C^{[l]} \times |V_p| \times T_h} \quad (8)$$

The attention block can dynamically adjust the impact of different traffic events to target events w.r.t. the features from the previous layer.

B. Spacetime Convolution Block

After the attention layer, we design a *spacetime convolution block* that contains three kernels to capture the many-to-one influence from three different perspectives corresponding to space, time, and spacetime, respectively. The *spacetime kernel* is the main perspective for uncovering the spacetime correlations from a local local-spacetime to the target events.




Fig. 2: The architecture of STNN with an example local-spacetime constructed from the input data. STNN consists of k spacetime modules (ST-Modules) and a fully-connected output layer. Each ST-Module contains a spacetime attention block (ST-Attn block) and a spacetime convolution block (ST-Conv block). The ST-Attn block uses self-attention mechanism to spotlight the most contributive traffic events. In each ST-Conv block, three different convolution kernels are employed to aggregate the spatio-temporal correlations in different perspectives. Then, the extracted features are stacked, and condensed by the 1×1 convolution.




Fig. 3: Spatial kernel, temporal kernel and spacetime kernel on the sub-spacetime

The *temporal kernel* finds correlations between traffic events of the same location along time, and the *spatial kernel* captures the influence from nearby locations on the same time step. The motivation for adopting the temporal kernel is that each sensor may have unique periodic temporal traffic patterns (peak/off-peak, weekday/weekend, etc.), which may be undervalued from the spacetime perspective. Similarly, we use the spatial kernel to keep the underlying geo-spatial influence from neighbors which is invariant to the time. Figure 3 demonstrates these three kernels on the spacetime manifold. Each spacetime convolution block takes the output of former attention block as the input, e.g., $\mathcal{D}_p^{[l]} \in \mathbb{R}^{C^{[l]} \times |V_p| \times T_h}$. The output $\mathcal{D}_p^{[l+1]} \in \mathbb{R}^{C^{[l+1]} \times |V_p| \times T_h}$ is computed by

$$H = \text{LeakyReLU} \left[\Theta_{st}^{[l+1]} * \mathcal{D}_p^{[l]}; \Theta_t^{[l+1]} * \mathcal{D}_p^{[l]}; \Theta_s^{[l+1]} * \mathcal{D}_p^{[l]} \right] \quad (9)$$

$$\mathcal{D}_p^{[l+1]} = \text{LeakyReLU} \left(\Theta_o^{[l+1]} * H \right), \quad (10)$$

where $\Theta_{st}^{[l+1]}$, $\Theta_t^{[l+1]}$, and $\Theta_s^{[l+1]}$ are the $f \times f$ spacetime kernel, the $f \times 1$ temporal kernel and the $1 \times f$ spatial kernel, respectively; $\text{LeakyReLU}(\cdot)$ denotes the leaky rectified linear units function; and $*$ represents the convolution operation. For the above three convolution kernels, we take $f = 3$ in the experiments. Additionally, padding is used to make sure the

output has the same size as the input. Last, we concatenate the output of the three kernels and use an 1×1 convolution, $\Theta_o^{[l+1]}$, to condense features and restrict the number of channels.

At last, a fully connected layer is adopted to predict the future traffic flow from learned features.

VI. EXPERIMENTS

To evaluate the performance of the proposed model, we compare STNN with the state-of-the-art traffic prediction models on two public traffic datasets in Section VI-B. In addition, a simulated dynamic traffic network is used to demonstrate our model's capability of handling dynamic graphs in Section VI-C. We further investigate the learned spacetime interval via a case study (Section VI-D) and the impacts of different components of STNN (Section VI-E).

A. Experimental Setup

1) *Real-world Networks*: Two public real-world datasets METR-LA and PeMS-Bay are used for evaluation. METR-LA were collected from the loop detectors in the highway of Los Angeles County [29], and PeMS-Bay were collected from California Caltrans Performance Measurement System (PeMS) [30]. These datasets were released by [20] and have been widely used to evaluate traffic prediction models. METR-LA records four months of traffic data in early 2012 in Los Angeles County with 207 sensors. PeMS-Bay captures six months of traffic data in early 2017 in the Bay Area with 325 sensors. The speed reading of sensors in METR-LA and PeMS-Bay are aggregated to 5-minute windows, resulting in 288 data points per day. Missing values are filled with linear interpolation. Statistics of the datasets are summarized in Table III.

2) *Simulated Network*: Apart from the two public datasets with static networks, we simulate a traffic network to demonstrate our model's ability to handle dynamic networks. We

TABLE III: Statistics of datasets

Data	Nodes	Time steps	Traffic events	Dynamics
METR-LA	207	34,272	7,094,304	No
PeMS-Bay	325	52,116	16,937,700	No
Simulated	84	2,000	16,800	Yes

synthesize the dynamic traffic data using CityFlow [31] which is a multi-agent reinforcement learning tool for city traffic scenarios. It is worthy to note that the goal of this experiment is not to show how the simulated traffic network resembles reality but to evaluate how well our model predicts traffic flow as the network topology changes. The simulated dataset contains 84 nodes and 2000 time steps. Each node denotes a sensor station located in the middle of a road segment as shown in Figure 4. Each sensor station records the total number of vehicles passing it in the fixed time interval instead of the average speed. Dynamics are represented as road closures to simulate traffic accidents or road constructions. For example, road segment 8 in Figure 4 is closed from 400 to 600 and 1500 to 1900 time steps. Such closure will alter the travel distance between nearby nodes resulting in modified matrix $\mathbf{Q}^{(t)}$. Besides the evaluation on the entire dataset, we also highlight the results of node 8 and its nearby roads 7 and 9, whose surroundings changed the most.




Fig. 4: Simulated road network illustration

3) *Baselines*: We compare our model with the following baselines on the two real-world datasets. The default settings of these methods are used.

- HA: Historical average method. We use the moving average with window size 12 for the forecasting.
- ARIMA_{kal} [9]: Auto-Regressive integrated moving average model with Kalman filter for time series analysis.
- FC-LSTM [32]: Recurrent neural network with fully connected LSTM hidden units.
- DCRNN [20] combines diffusion convolution and gated recurrent units in an encoder-decoder manner.
- STGCN [5] employs graph convolution for spatial patterns and 1D convolution for temporal features.
- Graph WaveNet [24] uses node embeddings to learn a dependency matrix, and the dilated 1D convolution for prediction.
- GMAN [25]: Graph multi-attention network equipped with spatial, temporal and transform attention.

- MTGCN [26]: Multivariate time series forecasting with graph neural networks.

4) *Experiment Settings*: We implemented our model via the PyTorch framework and on the following hardware platform: (CPU: Intel(R) Xeon(R) Gold 6128 CPU @ 3.40GHz, GPU: NVIDIA Quadro RTX 8000). For the STNN model, we set $k = 2$, i.e., two ST-Modules are stacked, which is enough to produce satisfying results. The number of output channels for two ST-Modules are 32 and 64, respectively. The convolutional kernel size is set to be $f = 3$. The parameter used in the LeakyReLU is 0.2. In the experiments, datasets are split as 70% training, 20% validation and 10% testing in chronological order. The length of input and output time sequence are both 12. In the training phase, the batch size is 80, and the number of epochs is 50, which takes about 5 hours to train the model on METR-LA dataset. Adam optimizer is adopted with learning rate 1×0.001 to update model parameters towards the minimized L1 loss between the predicted value and the grounded truth. Dropout is enabled with rate 0.3. We apply the commonly used metrics: mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percentage error (MAPE) to evaluate the mean prediction accuracy of future T_r time steps. We report the average performance of five runs. Random seed of PyTorch is set to be 1 for reproducibility purpose.

We perform a grid search for several important hyper-parameters to determine the appropriate values regarding complexity and performance. The first hyper-parameter is the number of nodes in each local-spacetime. The second is the proportion of data used for training, ranges from 0.05 to 1. The third is the number of spacetime modules employed, ranges from 1 to 4. Results show that the model performance improves drastically with the increasing number of neighbors in a local-spacetime and reach the plateau after $\epsilon = 15$. Similar pattern shows for the ratio of training set and the best-tradeoff point is 0.2. In other words, in 70% training data, we random select $0.2 \times 70\%$ for the final training. Last, model is not very sensitive to the number of ST modules after stack two modules.

B. Evaluation on Real-World Data

One salient feature of this work is that a trained STNN model can be used directly on unseen traffic networks. None of the existing state-of-the-art methods can do this. Current models train and predict on the same dataset, and the trained model cannot be applied to other networks with different sizes. Thus, we evaluate the performance of STNN in three settings: (1) train and test on the same network; (2) train and test on different networks; (3) train on multiple networks and test on each network.

1) *Train and Test on the Same Network*: This experiment follows the conventional settings where we train and test the STNN model on the same network. We train two STNN models separately on the two real-world networks, namely STNN₁ for METR-LA dataset and STNN₂ for PeMS-bay dataset, respectively. The prediction accuracy on test sets are

TABLE IV: Overall performance of short-term (15 mins), mid-term (30 mins) and long-term (60 mins) traffic forecasting.

Training Data	Test Data	Models	15 mins			30 mins			60 mins		
			MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
METR-LA	METR-LA	HA	4.16	7.80	13.02%	4.16	7.80	13.02%	4.16	7.80	13.02%
		ARMIA	3.60	7.18	8.95%	4.37	8.72	10.42%	5.56	10.01	13.50%
		FC-LSTM	3.02	5.27	8.92%	3.29	6.78	9.10%	3.71	7.32	9.54%
		DCRNN	2.52	4.96	6.81%	2.71	5.57	7.85%	3.03	6.15	8.84%
		STGCN	2.76	5.01	6.90%	2.94	6.17	8.26%	3.85	7.50	9.73%
		Graph WaveNet	2.36	4.54	6.04%	2.59	5.25	7.05%	2.98	6.07	8.21%
		GMAN	2.39	4.63	5.95%	2.58	5.31	6.93%	2.94	6.02	8.16%
		MTGCN	2.41	4.58	6.01%	2.59	5.23	6.98%	2.96	6.04	8.18%
		STNN (ours)	2.27	4.46	5.80%	2.56	5.29	6.84%	3.01	6.23	8.50%
PeMS-Bay	PeMS-Bay	HA	2.88	5.59	6.85%	2.88	5.59	6.85%	2.88	5.59	6.85%
		ARMIA	1.55	3.01	3.12%	1.98	4.12	4.76%	2.71	5.34	6.62%
		FC-LSTM	1.89	3.88	4.40%	2.01	4.19	4.75%	2.16	4.38	5.24%
		DCRNN	1.28	2.63	2.57%	1.51	3.49	3.40%	1.78	4.19	4.28%
		STGCN	1.27	2.65	2.60%	1.58	3.88	3.70%	2.20	4.63	4.87%
		Graph WaveNet	1.22	2.45	2.50%	1.47	3.31	3.21%	1.68	3.56	3.88%
		GMAN	1.26	2.50	2.58%	1.47	3.32	3.18%	1.65	3.53	3.61%
		MTGCN	1.24	2.46	2.53%	1.50	3.34	3.30%	1.69	3.65	3.90%
		STNN (ours)	1.20	2.41	2.50%	1.49	3.23	3.26%	1.86	4.17	4.30%
PeMS-Bay	METR-LA	STNN (ours)	2.60	5.02	6.75%	3.01	5.94	8.37%	3.68	7.33	11.13%
METR-LA	PeMS-Bay		1.36	2.73	2.92%	1.68	3.58	3.73%	2.15	4.72	5.01%
Combined	METR-LA		2.30	4.50	5.83%	2.60	5.29	6.90%	3.07	6.37	8.80%
	PeMS-Bay		1.21	2.43	2.51%	1.49	3.22	3.26%	1.87	4.17	4.40%

shown in the first two rows of Table IV. STNN surpasses the baselines in both datasets. For short period predictions like the METR-LA 15 mins case, STNN₁ achieves 2.27 MAE and 4.46 RMSE. Compared with the baselines, where the best MAE and RMSE are 2.36 and 4.54, respectively, STNN₁ produces up to 4% performance improvement. For mid-term predictions (30 mins), STNN can reach the SOTA performance. For long period predictions like 60 mins case, our model is slightly worse than the baselines like GMAN which designed specifically for long-term traffic prediction.

2) *Train and Test on Different Networks*: Since STNN is designed to discover universal traffic patterns. An interesting question is to apply a trained model on an unseen dataset. This is a huge challenge for methods that rely on the graph adjacency/Laplacian matrix because these matrices can be dramatically different in various networks. The proposed STNN concentrates on the local spatial and temporal context such that the learned model can be easily used on unseen datasets. In previous experiments, we have trained the STNN₁ on METR-LA and STNN₂ on PeMS-bay. Without any fine-tuning, we take STNN₁ to make predictions on PeMS-bay data and the performance is comparative to state-of-the-art methods trained directly on PeMS-bay. Similar results for STNN₂ on METR-LA. The details are shown in the middle two rows of Table IV. The MAE of STNN₁ on PeMS-bay 15 mins prediction period is 1.36, that is comparable to existing methods which train and test on the same dataset. The MAE of STNN₂ on METR-LA 15 mins prediction period is 2.60 while the MAEs of baselines range from 2.39 to 2.76.

3) *Train on Multiple Networks*: Finally, we trained a single model on the combined dataset by mixing the training ex-

amples from PeMS-Bay and METR-LA. The learned STNN model captures the general traffic patterns from both datasets. The results are shown in the last two rows of Table IV. The performance of the uniform STNN is very close to STNNs that train and test on the same network, and outperforms baselines on both datasets of short-term prediction.

C. Evaluation on Simulated Dynamic Network

Table V shows the traffic volume prediction error for the next three time steps of nodes 7, 8, 9, and the average error on the whole network. We only compare with HA, ARMIA and FC-LSTM because the other methods cannot be applied on networks with dynamic topology. In general, STNN outperforms the baselines by large margin. The overall MAE and RMSE of STNN are 5.31 and 19.83 for the entire dataset, which is 42% better than baselines. Moreover, for a specific location like node 7, the MAE and RMSE of the baseline are 9.75 and 19.39, respectively. STNN achieves 5.27 MAE and 10.31 RMSE, leading to the 46% and 47% improvement against the baseline. For the road segmentation monitored by sensor 8, the traffic volume is always zero during the closure period. It is not very sensible to predict the traffic flow for such time steps. Thus, we exclude them during evaluation. Similar to other nodes, STNN still outperforms baselines with a 35% better MAE. It is worthy to note that the prediction error on node 7 is larger than node 9, mainly because node 7 is one of the traffic sources where new vehicles keep entering the traffic network continuously and randomly makes it hard to predict.

D. Case Study

To develop a better insight of our model, we present a real-world case study. The aim is to reveal the spacetime


TABLE V: Performance of traffic volume prediction on the simulated dynamic network. We highlight sensor 7,8,9 as they are most affected by the changing network topology.

Data	Models	MAE	RMSE	MAPE
Node 7	HA	26.69	64.57	12.53 %
	ARMIA	12.08	29.16	9.84 %
	FC-LSTM	9.75	19.39	7.74 %
	STNN (ours)	5.27	10.31	3.97 %
Node 8	HA	20.75	43.82	28.79 %
	ARMIA	17.37	35.22	23.54 %
	FC-LSTM	15.19	30.46	18.69 %
	STNN (ours)	9.78	26.31	9.63 %
Node 9	HA	11.75	29.63	10.55 %
	ARMIA	6.35	11.35	6.10 %
	FC-LSTM	5.23	9.61	5.19 %
	STNN (ours)	2.65	6.14	2.88 %
Overall	HA	15.53	42.19	10.01 %
	ARMIA	11.02	31.75	7.81 %
	FC-LSTM	9.17	25.65	5.89 %
	STNN (ours)	5.31	19.83	3.79 %


intervals that are extracted by our model. Figure 5 illustrates a small part of the METR-LA dataset that contains a number of nodes around the Los Angeles Zoo. These nodes are used for the prediction of a target node, as indicated by node 111. We set the prediction period to 15 minutes (3 time steps). The color scale in Figure 5a, which is transformed from the learned attention map, indicates the different influence (i.e., spacetime interval) made by those surrounding traffic events have towards the target node 111. The vertical axis represents the spatial dimension where nodes are listed in ascending order by their distance with the target node 111. The horizontal axis represents the temporal dimension where 0–11 time steps data are used for predicting the next 3 time steps of node 111. Figures 5b displays the spatial context of all nodes involved. From these two figures, we can observe that the most significant nodes for prediction are 111, 42, 54, 37, 142, 145. Among them, 54, 37, 142, 145 are the nearest upstream neighbors; and 42 is the immediate downstream neighbor. It is remarkable that downstream locations could be useful for traffic forecast, while the correlation fades out rapidly as distance increases as for node 24. Besides the spatial context, the temporal dimension also plays an important role. The time steps of node 111 that are key to the prediction are the immediate last four while this is not true for other nodes. Despite close proximity of nodes 125 and 58 to 111, they demonstrate a smaller influence as 58 is on the opposite direction of the same road while 125 is on another road. This case study demonstrates that our model can extract meaningful correlations between the traffic events along both the spatial and temporal dimensions.

E. Ablation Study

We perform an ablation study to validate the impact of each proposed component to the model performance. In each experiment, a block/layer is removed and the rest remain unchanged. Particularly, six types of experiments are conducted: full STNN, STNN without ST-Attn block, STNN without ST-Conv block, STNN without spatial convolution layer, STNN



(a) Impact of traffic events for the prediction of future traffic flow at node 111. Nodes are sorted by travel distance with 111 on road network.



(b) Spatial context of node 111. Red markers represent roads from west to east, blue markers represent roads from east to west, and green markers stand for north to south.

Fig. 5: Case study

without temporal layer, and STNN without spacetime convolution layer. We report the performance of the above models for short-term (15 mins) forecasting in Table VI.

TABLE VI: Ablation study

	METR-LA			PeMS-Bay		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE
Full STNN	2.27	4.46	5.80	1.20	2.41	2.53
No ST-Attn block	2.35	4.56	5.85	1.23	2.43	2.55
No ST-Conv block	2.41	4.69	5.97	1.30	2.58	2.71
No spatial conv	2.31	4.51	5.74	1.22	2.42	2.50
No temporal conv	2.33	4.58	5.90	1.22	2.47	2.55
No spacetime conv	2.34	4.59	5.98	1.23	2.49	2.58

The effect of proposed spacetime attention block is evident since it highlights the contributive traffic events. The spacetime convolution block improves the prediction accuracy significantly as it aggregated the traffic events in three perspectives. Moreover, the spacetime convolutional layer appears to be more important than the other two convolutional layers.

F. Complexity Analysis

We analyze the complexity of the proposed model to justify its scalability by showing the computation of STNN is independent of the road network size. In addition, the total trainable parameters in STNN is 308,786, which is smaller than STGCN (454,358), MTGCN (405,452), and other state-of-the-art methods.

The time complexity of the spacetime convolutional block is $O(4\alpha T d_i d_o)$ where α is the number of sensors in local-spacetime, T is the length of input time steps, d_i and d_o are channels of input feature maps and output feature

maps. Additionally, four different convolutional layers (space-time, spatial, temporal, 1*1) were utilized in the spacetime convolutional block. The spacetime attention block incurs $O(c_i c_o (\alpha^2 T^2 + \alpha T))$ time complexity where c_i is the channels of input data and c_o is the channels of output data. α and T still represents the size of the spacetime manifold. In short, the computational complexity of the proposed STNN depends on the local-spacetime size, input sequence length, and channels. However, it is independent of the entire network.

Last, the local-spacetime construction is part of the proposed spacetime interval learning framework which is not included in the forward pass and backpropagation of STNN. The complexity of local-spacetime construction equals $O(NT)$ where N is the number of sensors in the network.

VII. CONCLUSION

In this paper, we propose a novel spacetime interval learning framework and spacetime neural network (STNN) for accurate traffic prediction. Our approach captures the intrinsic principles of traffic flow by learning the spacetime intervals between traffic events. The model works on the local-spacetime extracted for target nodes and thus can handle dynamic network topology. Experiments on two real-world datasets and a simulated dynamic network show that the proposed framework significantly outperforms state-of-the-art methods. This confirms the effectiveness of capturing spatio-temporal correlations directly. In the future, we will further explore the possibility of capturing the graph dynamics from the input data when the underlying graph structure is unknown. Another promising direction is to embed the local-spacetime construction into an end-to-end deep learning framework, making the parameters used in local-spacetime construction learnable. Last, we would like to investigate the spacetime interval learning paradigm in other spatio-temporal data modeling scenarios.

REFERENCES

- [1] E. I. Vlahogianni, M. G. Karlaftis, and J. C. Golias, "Short-term traffic forecasting: Where we are and where we're going," *Transportation Research Part C: Emerging Technologies*, vol. 43, pp. 3–19, 2014.
- [2] M. Treiber and A. Kesting, "Traffic flow dynamics," *Traffic Flow Dynamics: Data, Models and Simulation*, Springer-Verlag Berlin Heidelberg, 2013.
- [3] I. Yperman, S. Logghe, and B. Immers, "The link transmission model: An efficient implementation of the kinematic wave theory in traffic networks," in *Proceedings of the 10th EWGT Meeting*. Poznan Poland, 2005, pp. 122–127.
- [4] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, "Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam," *New journal of physics*, vol. 10, no. 3, p. 033001, 2008.
- [5] B. Yu, H. Yin, and Z. Zhu, "Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting," in *IJCAI*, 2018.
- [6] H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, and Z. Li, "Deep multi-view spatial-temporal network for taxi demand prediction," in *AAAI*, 2018.
- [7] Z. Diao, X. Wang, D. Zhang, Y. Liu, K. Xie, and S. He, "Dynamic spatial-temporal graph convolutional neural networks for traffic forecasting," in *AAAI*, vol. 33, 2019, pp. 890–897.

- [8] Q. Zhang, J. Chang, G. Meng, S. Xiang, and C. Pan, "Spatio-temporal graph structure learning for traffic forecasting," in *AAAI*, vol. 34, no. 01, 2020, pp. 1177–1185.
- [9] B. M. Williams and L. A. Hoel, "Modeling and forecasting vehicular traffic flow as a seasonal arima process: Theoretical basis and empirical results," *Journal of transportation engineering*, vol. 129, no. 6, pp. 664–672, 2003.
- [10] E. P. Rowe, *Geometrical physics in Minkowski spacetime*. Springer Science & Business Media, 2013.
- [11] E. Zivot and J. Wang, "Vector autoregressive models for multivariate time series," *Modeling financial time series with S-PLUS®*, pp. 385–429, 2006.
- [12] J. Van Lint and C. Van Hinsbergen, "Short-term traffic and travel time prediction models," *Artificial Intelligence Applications to Critical Transportation Issues*, vol. 22, no. 1, pp. 22–41, 2012.
- [13] N. Zhang, Y. Zhang, and H. Lu, "Seasonal autoregressive integrated moving average and support vector machine models: prediction of short-term traffic flow on freeways," *Transportation Research Record*, vol. 2215, no. 1, pp. 85–92, 2011.
- [14] R. Chen, C.-Y. Liang, W.-C. Hong, and D.-X. Gu, "Forecasting holiday daily tourist flow based on seasonal support vector regression with adaptive genetic algorithm," *Applied Soft Computing*, vol. 26, pp. 435–443, 2015.
- [15] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [16] N. Laptev, J. Yosinski, L. E. Li, and S. Smyl, "Time-series extreme event forecasting with neural networks at uber," in *International conference on machine learning*, vol. 34, 2017, pp. 1–5.
- [17] Z. Cui, R. Ke, Z. Pu, and Y. Wang, "Deep bidirectional and unidirectional lstm recurrent neural network for network-wide traffic speed prediction," *arXiv preprint arXiv:1801.02143*, 2018.
- [18] J. Zhang, Y. Zheng, D. Qi, R. Li, and X. Yi, "Dnn-based prediction model for spatio-temporal data," in *SIGSPATIAL*, 2016, pp. 1–4.
- [19] J. Zhang, Y. Zheng, and D. Qi, "Deep spatio-temporal residual networks for citywide crowd flows prediction," in *AAAI*, 2017.
- [20] Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," in *ICLR*, 2018.
- [21] J. Zhang, X. Shi, J. Xie, H. Ma, I. King, and D.-Y. Yeung, "Gaan: Gated attention networks for learning on large and spatiotemporal graphs," in *UAI*, 2018.
- [22] S. Guo, Y. Lin, N. Feng, C. Song, and H. Wan, "Attention based spatial-temporal graph convolutional networks for traffic flow forecasting," in *AAAI*, vol. 33, 2019, pp. 922–929.
- [23] H. Yao, X. Tang, H. Wei, G. Zheng, and Z. Li, "Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction," in *AAAI*, vol. 33, 2019, pp. 5668–5675.
- [24] Z. Wu, S. Pan, G. Long, J. Jiang, and C. Zhang, "Graph wavenet for deep spatial-temporal graph modeling," *IJCAI*, 2019.
- [25] C. Zheng, X. Fan, C. Wang, and J. Qi, "Gman: A graph multi-attention network for traffic prediction," in *AAAI*, vol. 34, no. 01, 2020, pp. 1234–1241.
- [26] 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 *KDD*, 2020.
- [27] S. M. Carroll, *Spacetime and geometry*. Cambridge University Press, 2019.
- [28] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-attention generative adversarial networks," in *International Conference on Machine Learning*, 2019, pp. 7354–7363.
- [29] H. V. Jagadish, J. Gehrke, A. Labrinidis, Y. Papanikolaou, J. M. Patel, R. Ramakrishnan, and C. Shahabi, "Big data and its technical challenges," *Communications of the ACM*, vol. 57, no. 7, pp. 86–94, 2014.
- [30] C. Chen, K. Petty, A. Skabardonis, P. Varaiya, and Z. Jia, "Freeway performance measurement system: mining loop detector data," *Transportation Research Record*, vol. 1748, no. 1, pp. 96–102, 2001.
- [31] H. Zhang, S. Feng, C. Liu, Y. Ding, Y. Zhu, Z. Zhou, W. Zhang, Y. Yu, H. Jin, and Z. Li, "Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario," in *The World Wide Web Conference*, 2019, pp. 3620–3624.
- [32] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *NeurIPS*, 2014.