

Received May 4, 2020, accepted May 22, 2020, date of publication June 3, 2020, date of current version June 16, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2999662

Traffic Data Imputation Using Deep Convolutional Neural Networks

OUAFA BENKRAOUDA^{1,4}, BILAL THONNAM THODI^{2,4}, Hwasoo YEO³, (Member, IEEE),
MÓNICA MENÉNDEZ^{1,2,4}, (Member, IEEE), AND SAIF EDDIN JABARI^{1,2,4}

¹Department of Urban and Regional Planning, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA

²Department of Civil and Urban Engineering, NYU Tandon School of Engineering, Brooklyn, NY 11201, USA

³Department of Civil and Urban Engineering, Korea Advanced Institute of Science and Technology, Daejeon 34141, South Korea

⁴Division of Engineering, New York University Abu Dhabi, Abu Dhabi, United Arab Emirates

Corresponding author: Saif Eddin Jabari (sej7@nyu.edu)

This work was supported by the New York University Abu Dhabi (NYUAD) Center for Interacting Urban Networks (CITIES), funded by Tamkeen, through the New York University Abu Dhabi (NYUAD) Research Institute Award under Grant CG001, and in part by the Swiss Re Institute through the Quantum CitiesTM Initiative.

ABSTRACT We propose a statistical learning-based traffic speed estimation method that uses sparse vehicle trajectory information. Using a convolutional encoder-decoder based architecture, we show that a well trained neural network can learn spatio-temporal traffic speed dynamics from time-space diagrams. We demonstrate this for a homogeneous road section using simulated vehicle trajectories and then validate it using real-world data from the Next Generation Simulation (NGSIM) program. Our results show that with probe vehicle penetration levels as low as 5%, the proposed estimation method can provide a sound reconstruction of macroscopic traffic speeds and reproduce realistic shockwave patterns, implying applicability in a variety of traffic conditions. We further discuss the model's reconstruction mechanisms and confirm its ability to differentiate various traffic behaviors such as congested and free-flow traffic states, transition dynamics, and shockwave propagation. We also provide a comparison against a widely used adaptive smoothing technique used for the same purpose and demonstrate the superiority of the proposed approach, even with probe vehicle lower penetration levels.

INDEX TERMS Convolutional neural networks, data expansion, data imputation, estimation, filtering, traffic dynamics, traffic state estimation.

I. INTRODUCTION

Recent studies have shown that connected and autonomous vehicles (CAVs) can ease traffic flow instabilities at low CAV penetration levels [1]–[4]. To do so, these systems require accurate knowledge of traffic conditions. In general, most advanced traffic management and control tools require accurate inputs. However, measurements remain sparse in today's road traffic networks and high penetration levels of CAVs are not expected in the coming decade or two. In the meantime, traffic state reconstruction methods are needed to fill this gap. Most traffic state estimation techniques in the literature have focused on reproducing traffic conditions at aggregate levels (averaged over segment lengths greater than 100 meters and time intervals greater than one minute [5], [6]). Advanced traffic management tools (e.g., adaptive traffic signal control) require knowledge of traffic conditions at a much finer

scale [7]–[13]; see [14] for a comprehensive review of past traffic state estimation methods.

We present here a data driven methodology to estimate traffic speeds across a road section at fine spatio-temporal resolutions and using limited probe vehicle information. Characterizing traffic speed across small intervals of space and time is especially challenging because of the inherent stochastic nature of traffic flow and noisy sparse traffic observations [15]–[17]. Researchers often resort to different statistical learning methods to model traffic speed variations from historical datasets [18]–[20]. However, the majority of these methods only succeed in capturing recurring traffic patterns such as daily variations or within-day variations, and do not necessarily learn the actual flow dynamics. A few studies have focused on modeling the sharp discontinuities found in spatio-temporal traffic dynamics (e.g., caused by transition from free flow to congested state). Nonetheless, the low interpretability of these non-linear models poses reliability and robustness issues when applied [5], [19], [21].

The associate editor coordinating the review of this manuscript and approving it for publication was Chao Chen.

Various kernel-based estimation schemes have also been proposed to interpolate spatio-temporal macroscopic traffic speeds from heterogeneous data sources, mainly, loop detectors and probe vehicles [22]–[26]. However, these interpolation methods need field calibration of static model parameters (offline or online) such as shockwave speeds and free flow speeds, which can dynamically change depending on the local-traffic conditions, leading to biased results.

In this paper, we propose a method for traffic state reconstruction of traffic speeds dynamics using vehicle trajectory information obtained from sparse probe data (with as low as 5% penetration rate). We employ a deep convolutional neural network (CNN) [27], [28] to learn the traffic speed dynamics from a time-space diagram plotted using partial observations of vehicle trajectories. Using an encoder-decoder architecture [29], we develop a CNN that learns the local spatial and temporal correlations (or features) from the time-space traffic speed profiles, and reconstructs the full (macroscopic) traffic speed dynamics for any given sparse vehicle trajectories. We show that a well trained deep CNN can learn to reproduce short-term non-recurring traffic features such as congestion waves and free flow speed dynamics, which state-of-the-art estimation methods often fail to capture. The proposed learning method can adapt to local traffic speed conditions, and reproduce observed shock-wave patterns and stop-and-go traffic. This is highly relevant for freeway monitoring, especially incident response applications [30], [31]. We further provide critical insights into the learned neural network model to understand the essential features detected while reconstructing the full speed map. Such insights can help unlock some of the black-box features of neural networks when applied to traffic state estimation problems.

In the rest of the paper, we formally state the traffic speed estimation problem and motivate the use of an encoder-decoder based convolutional neural network architecture to learn the spatio-temporal traffic speeds in Sec. II. We perform numerical experiments using the proposed model on a homogeneous road section and discuss its efficacy in reconstructing the traffic speed maps on a simulated scenario, as well as on a dataset obtained from the Next Generation Simulation (NGSIM) program in Sec. III. This is followed by a brief discussion of the model's interpretability and the role played by the hidden layers in the estimation of speed dynamics. We close Sec. III with a comparison against an adaptive smoothing technique that is widely employed in this context. We conclude the paper with a discussion of the main insights and suggest potential future research directions in Sec. IV.

II. METHODOLOGY

A. TRAFFIC SPEED ESTIMATION PROBLEM

Consider a set of vehicles \mathcal{N} passing through a given road section of length $X = |\mathcal{X}|$ over time period $T = |\mathcal{T}|$. Let (x_i, t_i, v_i) denote the coordinates of vehicle $i \in \mathcal{N}$, where x_i and v_i represent its relative longitudinal position (with respect

to the starting position of the road section) and speed at time t_i . We denote by $\mathcal{G} := \{(x_i, t_i, v_i) \mid i \in \mathcal{N}\}$ the set of all \mathcal{N} vehicle trajectories. Let $\mathcal{G}_p \subseteq \mathcal{G}$ represent the trajectories of sampled vehicles in \mathcal{N} (with sampling percentage p), within the space-time domain $X \times T$. Note that the cardinalities of \mathcal{G} and \mathcal{G}_p can vary depending on the number of vehicles in the system and the resolution of vehicle coordinates (the sampling cadence).

The estimation problem is to determine an arbitrary function from the domain of partial trajectories \mathcal{G}_p to the full trajectories \mathcal{G} (or some function of \mathcal{G} , such as macroscopic traffic speeds or traffic density). However, the variable (and not necessarily known) sized nature of \mathcal{G} (and \mathcal{G}_p) presents a challenge, as a unique mapping function may not exist. Hence, we represent \mathcal{G} (and \mathcal{G}_p) on a two-dimensional space-time plane as plotted vehicle trajectories, and color code it with the respective speed data; see Fig. 1 for an example.

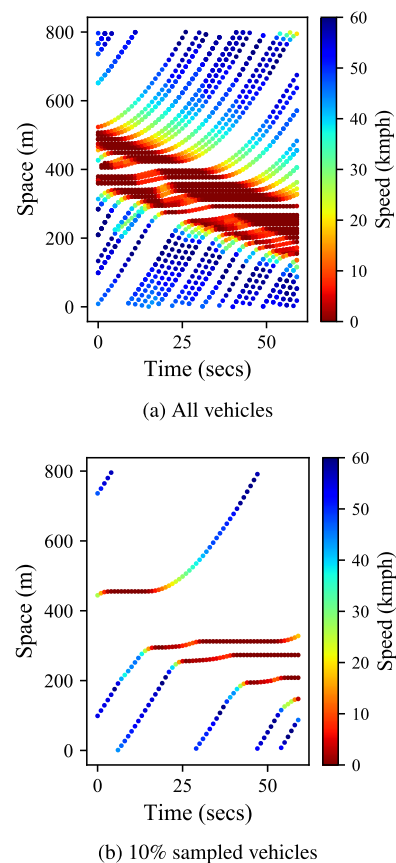


FIGURE 1. Vehicle trajectories plotted on a space-time diagram and color coded with speeds ($X = 800$ m, $T = 60$ sec).

To interpret this graphical representation, we discretize the entire space-time domain into finite regions called cells, each of which has dimensions x and t , the cell's length and time duration, respectively. The values in each cell represent the traffic state in that cell, defined by a three-scale color value based on the RGB format. The three-scale color representation of traffic states allows us to clearly differentiate between the empty cells (where there are no vehicles),

free flowing cells (where vehicles travel at the free-flow speed), and heavily congested cells (where vehicles cannot move). For example, empty cells are represented by a white color with RGB value (255, 255, 255) whereas heavily congested cells are represented by a red color with RGB value (231, 11, 5). Colors ranging between red and blue represent traffic states in between free-flow and heavily congested. Using this representation, we encode the vehicle trajectory information as a three-dimensional tensor, denoted by $\mathbf{z} \in \{1, \dots, 255\}^{X \times T \times C}$, where C is the number of color channel arrays. Note here that when treating the time-space diagrams as images using RGB, the tensor values are restricted to the set $\{1, \dots, 255\}$. Let \mathbf{z}^p denote the respective tensor with partial trajectory information, and \mathbf{z}^f the tensor with full vehicle trajectory information. The traffic state estimation problem can be stated as one that seeks to determine a mapping function g , such that $g(\mathbf{z}^p) = \mathbf{z}^f$, where \mathbf{z}^p and \mathbf{z}^f are now fixed sized tensors. We refer to the mapping function g as the traffic speed reconstruction model.


B. REPRESENTATION OF SPEED RECONSTRUCTION MODEL

We approximate the speed reconstruction function g using a neural network model. We chose an encoder-decoder convolutional neural network architecture [27], [29] to represent the reconstruction model as shown in Fig. 2a. The encoder model g^e extracts the primary features from the sparse input vehicle trajectory tensor \mathbf{z}^p and maps them into a latent space representation h , which can be thought of as an abstract projection of the sparse-vehicle trajectory information contained in \mathbf{z}^p onto some higher dimensional space (often uninterpretable). The decoder model, g^d , uses the information contained in h to reconstruct the output tensor \mathbf{z}^f , which in this case is the macroscopic traffic speed map.

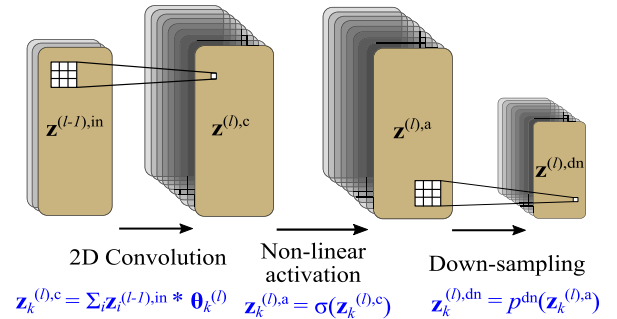
The encoder model (g^e) is a sequence of CNN layers Λ stacked horizontally, where each layer $l \in \Lambda$ consists of three operations - convolution, non-linear activation and down-sampling (up-sampling for the decoder model); see Fig. 2b. The convolution operation is performed using a set of layer-specific kernels $\Theta^{(l)} \equiv [\Theta_k^{(l)}]_{k \in \mathcal{K}^{(l)}}$, where $\mathcal{K}^{(l)}$ denotes the set of kernels in layer l . Each kernel is a tensor $\Theta_k^{(l)} \in \mathbb{R}^{M \times N \times C}$, where $M \times N$ is the kernel size (a hyper-parameter) and C is the number of color channels used in the feature maps. The convolution operation is given by

$$\mathbf{z}_k^{(l),c} = \sum_{i \in \mathcal{M}^{(l-1)}} \mathbf{z}_i^{(l-1),in} * \Theta_k^{(l)}, \quad k = 1, \dots, K, \quad l \in \Lambda, \quad (1)$$


where $\mathbf{z}_i^{(l-1),in}$ is the i th input feature map and $\mathcal{M}^{(l-1)}$ is the set of feature maps in layer $l - 1$, while $\mathbf{z}_k^{(l),c}$ is the convolved feature map produced using the k kernel in layer l . Note that the dimensions of $\mathbf{z}_k^{(l),c}$ are the same as the dimensions of $\mathbf{z}_i^{(l-1),in}$. This is achieved by padding the input matrix with zeros so that the information around the boundary of the time-space plane is effectively utilized during the convolution. However, the number of feature maps in the convolved



(a) Convolutional Encoder-decoder architecture



(b) A single CNN layer in the encoder model



(c) A single CNN layer in the decoder model

FIGURE 2. Representation of the speed reconstruction model using a convolutional encoder-decoder neural network architecture.

layer l , $|\mathcal{M}^{(l)}|$, and those in the input layer $l - 1$, $|\mathcal{M}^{(l-1)}|$, can differ. The convolution is followed by a clipping operation using a non-linear activation function σ :

$$\mathbf{z}_k^{(l),a} = \sigma(\mathbf{z}_k^{(l),c}), \quad k = 1, \dots, K, \quad l \in \Lambda, \quad (2)$$

where $\mathbf{z}_k^{(l),a}$ is the k th activated feature map in layer l . The kernels $\Theta_k^{(l)}$ represent the primary feature detectors in layer l , where each kernel $\Theta_k^{(l)}$ detects distinct patterns (or features) from the input feature maps $\mathbf{z}^{(l),in} \equiv [z_i^{(l)}]_{i \in \mathcal{M}^{(l)}}$. For instance, the initial layer can detect discrete traffic states such as free flow and congested traffic. Successive layers detect more complex patterns such as shockwave propagation and transient dynamics. We provide a more detailed interpretation of the kernels in Sec. III-B.

In the encoder model (Fig. 2b), the activated feature maps undergo a down-sampling operation to produce an abstract (lower dimensional) representation of the feature activations

as shown below:

$$\begin{aligned} \mathbf{z}^{(l),\text{dn}} &= p^{\text{dn}}(\mathbf{z}^{(l),\text{a}}) \\ &= \max_{m,n} \left(\mathbf{z}^{(l),\text{a}} \cdot \mathbf{u}(m,n) \right), \quad l \in \Lambda, \end{aligned} \quad (3)$$

where $\mathbf{u}(m,n)$ is a two-dimensional spatial-window applied to the input tensor $\mathbf{z}^{(l),\text{a}}$. Note that these depend on the kernel dimensions $M \times N$. The down-sampling (3) returns the largest value of $\mathbf{z}^{(l),\text{a}}$ (a.k.a. max pooling function [32]). Note that (3) reduces the dimension of $\mathbf{z}^{(l),\text{a}}$ by a factor of $(m \times n)$, where (m,n) represents the max-pooling filter size. The max-pooling operation summarizes the primary features detected over various spatial regions in the time-space diagram using the information contained in its smaller sub-regions.

The down-sampling operation can handle the sparsity in the vehicle trajectory information contained in the input tensor. We illustrate this using an arbitrary (sparse) input vehicle trajectory dataset, and its corresponding activated feature map obtained before and after the down-sampling operation in Fig. 3. The traffic state information contained in cells (2, 3) and (3, 2) is absent from the activated feature maps (middle part of Fig. 3). However, the down-sampling operation provides an improved estimate of the states of these cells (as free flow and congested states, respectively) by utilizing information contained in neighboring cells (left part of Fig. 3). This partially explains why the spatial distribution of probe vehicles along a road section can affect state estimation accuracy [33], [34]. The output of the down-sampling




FIGURE 3. Illustrating the need for down-sampling to handle the sparsity in vehicle trajectory information.

operation, $\mathbf{z}^{(l),\text{dn}}$ serves as the input into the next layer in the encoder model g^e , or to the hidden state representation \mathbf{h} . The decoder model g^d also consists of a similar set of layers as g^e , except that the down-sampling operation is replaced by an up-sampling operation; see Fig. 2c. The most common up-sampling operation is the nearest neighbor function [32], defined as

$$\mathbf{z}^{(l),\text{up}} = p^{\text{up}}(\mathbf{z}^{(l),\text{a}}) = \mathbf{z}^{(l),\text{a}} \otimes \mathbf{1}_{m' \times n'}, \quad l \in \Lambda, \quad (4)$$

where \otimes is the Kronecker product and $\mathbf{1}_{m' \times n'}$ is an $m' \times n'$ matrix of ones. The dimensions $m' \times n'$ depend on the size of the full feature maps $X \times T$. The up-sampling operation in (4), besides scaling the activated feature maps by $(m' \times n')$, enforces a plausible and logical reconstruction of

traffic states by drawing inferences from super-regions. For example, in the left part of Fig. 3, one can reconstruct the traffic states in its sub-regions by logically looking at the states in the neighboring regions.

The motivation to use a CNN layer as opposed to a naive feed forward neural network in the encoder-decoder architecture (Fig. 2) is the parameter sharing and sparse connectivity properties of CNNs [27]. For traffic data, parameter sharing implies that a specific set of learned features $\Theta^{(l)}$ (such as for recognizing the free flow speed, congested traffic, and transition dynamics) can be used anywhere in the time-space plane (in other words, it is translation invariant). Sparse connectivity ensures that the speed estimates in certain regions in space and time heavily depend on the states in the immediate (local) surroundings, which is especially true for traffic dynamics [23]. Moreover, the estimation from sparse data problem can be thought of as a spatial imputation/interpolation problem, to which CNNs are well suited.

C. PARAMETER OPTIMIZATION

The output from the decoder model, g^d is the reconstructed speed map denoted as $\hat{\mathbf{z}}_{\Theta}^d$, and is a function of the kernel set $\Theta \equiv \{\Theta^{(1)}, \dots, \Theta^{(|\Lambda|)}\}$. Thus, the traffic speed estimation problem is cast as a kernel set learning problem:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}(\mathbf{z}^f, \hat{\mathbf{z}}_{\Theta}^d), \quad (5)$$

where \mathcal{L} denotes a loss function over the estimated and actual speed values. Here, (5) is a differentiable function in $\Theta^{(l)} \in \Theta$, and any gradient based optimization technique can be used to find the optimal kernel set [27]. Note that (5) only optimizes the kernel set Θ , and does not include other hyper-parameters such as the number of CNN layers, $|\Lambda|$, in the encoder and decoder model (g^e and g^d), the number of kernels in each layer ($\{|\mathcal{K}^{(l)}|\}_{l \in \Lambda}$), the kernel dimensions (M and N), the down-sampling filter dimensions (m and n), and the up-sampling filter dimensions (m' and n'). These hyper-parameters have to be fine tuned separately.

III. EXPERIMENTS AND RESULTS

Here we demonstrate the proposed speed reconstruction model for a homogeneous road segment (a freeway section) using both simulated data and real-world data. By looking at different initial and boundary conditions, we produce various traffic conditions corresponding to free-flow, congested, and stop-and-go traffic. The trajectories of all vehicles passing through the segment are recorded at a one second cadence.

A. LEARNING SPEED RECONSTRUCTION MODEL FROM DATA

Using the simulated vehicle trajectories (for a 3-lane road that is 800 m long), we generate input-output samples for training the speed reconstruction model, with the following parameters: $X = 800$ m, $T = 60$ sec, $x = 10$ m, and $\tau = 1$ sec. The input samples are generated using a $p = 5\%$ penetration rate sampled from an unknown probability distribution.

TABLE 1. Hyper-parameters for the speed reconstruction model.

Layers		Convolution operation	Non-linearity	Sampling operation
Input layer		Identity	-	-
Encoder model	CNN-1	Convolution ($5 \times 5 \times 8$)	Relu	Max-pooling (2×3)
	CNN-2	Convolution ($5 \times 5 \times 32$)	ReLU	Max-pooling (2×2)
	CNN-3	Convolution ($5 \times 5 \times 64$)	ReLU	Max-pooling (2×2)
Decoder model	CNN-4	Convolution ($3 \times 3 \times 64$)	ReLU	Nearest neighbour (2×2)
	CNN-5	Convolution ($5 \times 5 \times 32$)	ReLU	Nearest neighbour (2×2)
	CNN-6	Convolution ($5 \times 5 \times 8$)	ReLU	Nearest neighbour (2×3)
Output layer		Convolution ($5 \times 5 \times 3$)	σ	-

This ensures that the sampling replicates actual field observations and no bias in the spatial-distribution of probe vehicles is introduced. The output samples are generated using all vehicle trajectories. In each sample, the vehicle trajectories are color-coded using a fixed speed gradient of 0-60 kmph, where 0 kmph is represented by red and 60 kmph is represented by blue; see Fig. 1 for example. These figures are then converted to $80 \times 60 \times 3$ tensor, where 3 corresponds to the RGB channel arrays.

We now test our model using two different datasets. The first set is generated from the simulated vehicle trajectories for the same road section used to train the model but under different traffic demand conditions. The second set is produced from the NGSIM vehicle trajectory dataset from Highway I-80 in Emeryville, California, U.S.A. [35]. The NGSIM program provides real-world vehicle trajectory data collected at three different sites. The dataset includes positions and speeds of all vehicles that traversed each of the three sites over the study period (45 minutes). The time cadence of the data is 0.1 seconds and it is extracted from video data. This allows us not only to test the transferability of our model, but to validate it against real world field data, and check its applicability to traffic behaviors not seen in the training sets. For both datasets we sample 5% of vehicle trajectories as input data.

The encoder and decoder models used here consists of three CNN layers, and the kernel size, number of kernels, and sampling stride for each layer are shown in Table 1. We use a random search technique to determine the best number of layers, the number of kernels in each layer, and the kernel sizes. One can also resort to more sophisticated approaches such as Bayesian-optimization to determine the optimal hyper-parameters [36]. For all the layers we use a rectified linear unit (ReLU) for activation (defined as $\text{ReLU}(z) = \max(z, 0)$ for any given input z) except for the last layer which has a Sigmoid activation function (defined as $\sigma(z) = (1 + \exp(-z))^{-1}$ for any given input z). As we treat the reconstructed speed maps as images, the (normalized) model outputs are bounded between 0 and 1, and this can be obtained using a sigmoid activation function.

We use a larger kernel size (5×5) for the initial and the final CNN layers as opposed to the smaller (3×3) kernels commonly found in traditional image processing architectures (see [28] for instance). This is because our input tensor

is very sparse (only 5-10% of the pixels have values other than 1), and as the convolution operation primarily detects local correlations in the data, smaller kernels may fail to detect essential features in the input tensor, and can hence produce unrealistic results. We overcome this by resorting to larger kernels, which when convolved can better learn the spatio-temporal speed dynamics, especially the backward propagating shockwave patterns, even from sparsely observed speeds.

Notice that the number of kernels in the middle layers of the CNN (the bottleneck region) is larger than the number of kernels in the boundary layers of the CNN (see Table 1). This is attributed to the feature learning capability of each layer in the network. For instance, kernels in the initial layers can be considered as primary feature detectors, useful for example, to detect colors (corresponding to congested and free-flowing traffic) and slow moving traffic, which are primitive and can be characterized using fewer kernels. However, subsequent layers use these primary features to capture more complex patterns such as free flowing regimes and transition dynamics, which may require a larger number of kernels to capture all relevant patterns. We use an iterative optimization algorithm known as ‘‘Adam’’ [37] to train the model. The Adam optimizer employs first-order gradient descent with an adaptive learning rate, and is found to have good convergence rates in practice. The test results and discussion are presented below.

B. RESULTS AND DISCUSSION

We first provide a visual comparison between the actual vehicle trajectories and the reconstructed speed map for the test datasets, and quantify the reconstruction error using a Structural Similarity Index. We then analyze and interpret the model’s reconstruction behavior by visualizing what the kernels learned from the data and its representation behavior in the latent space.

1) RECONSTRUCTION PERFORMANCE

The test results are presented in Figures 4– 6. From Fig. 4, we can see examples where the proposed approach reproduced the traffic speeds correctly. We observe that the model has learned to reconstruct the actual spatio-temporal speed maps and can clearly differentiate between the direction and

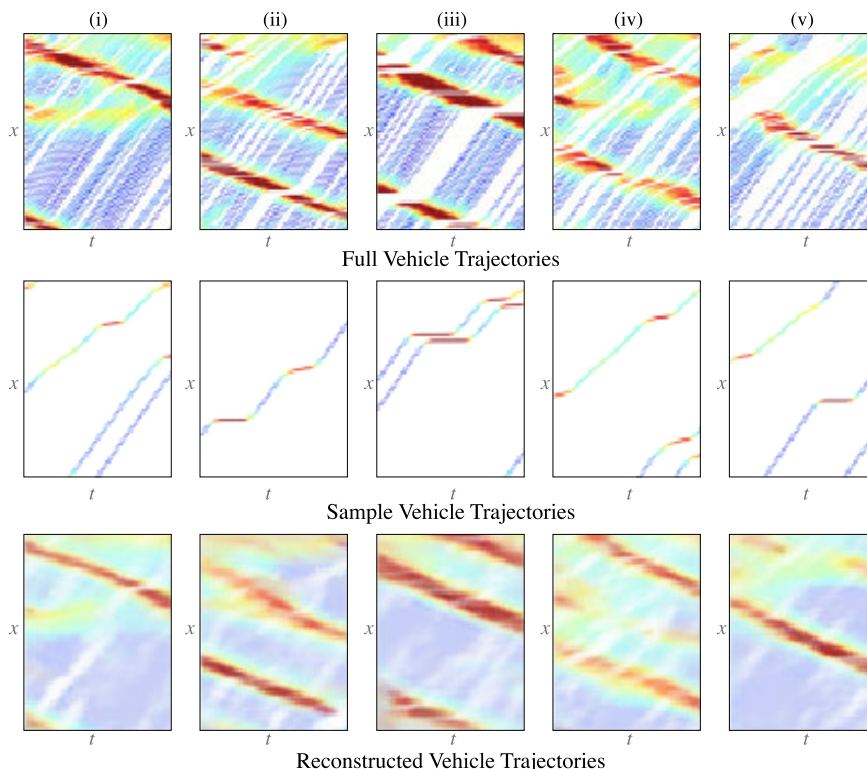


FIGURE 4. Selected examples where the model accurately reconstructed the macroscopic speed states (simulation data).

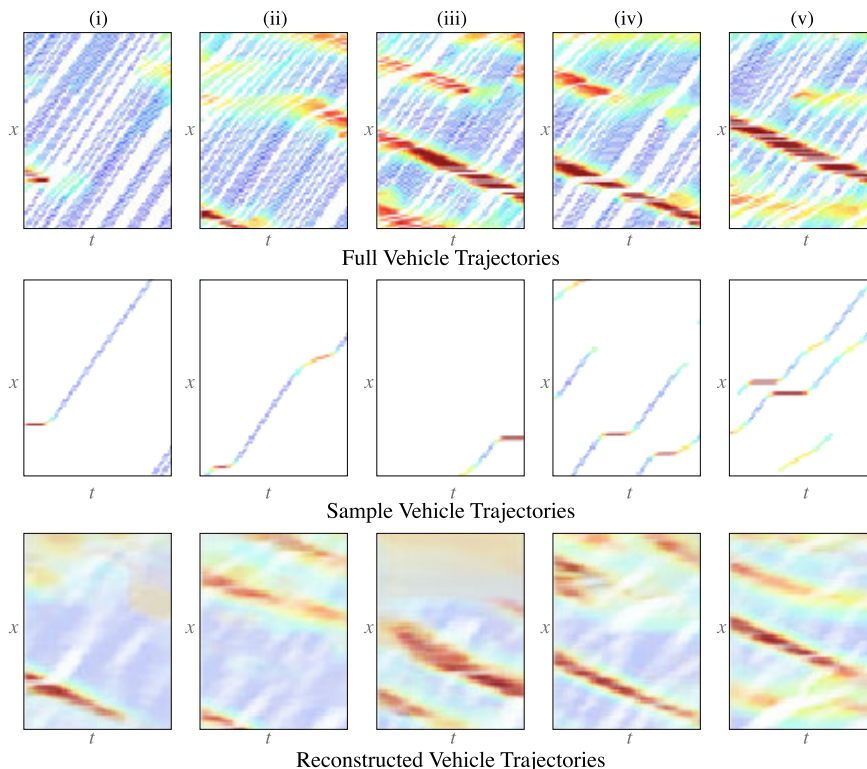


FIGURE 5. Selected examples where the model partially failed to provide a sound reconstruction of traffic states (simulation data).

regime of free flow (blue), congested (red), and transient dynamics (orange, yellow, and cyan). Interestingly, it succeeds in tracing backward propagating shockwave patterns

by tying together the limited information obtained from the sparse trajectories. In other words, the model is able to reconstruct shockwave patterns with varying sizes and

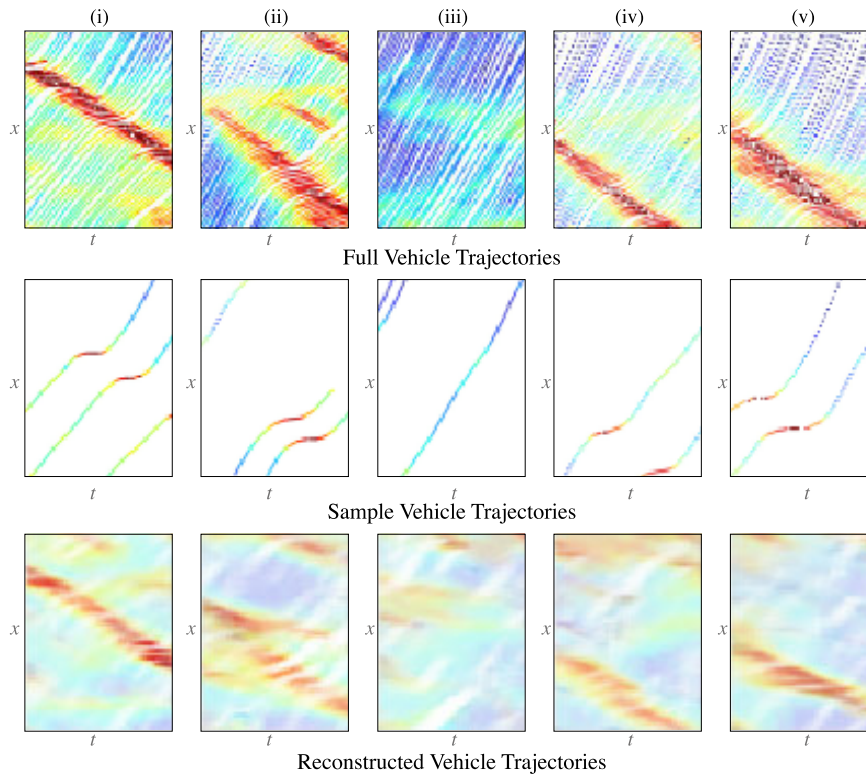


FIGURE 6. Selected examples from the reconstruction of the NSGIM dataset.

shapes depending on the local traffic conditions, rather than using a mere interpolation assuming a constant wave speed [4], [38]–[44]. Fig. 4 (i, iv, v) shows a scenario where the model captured the dynamic stop-and-go traffic regime (red regions), which cannot be characterized by a single shockwave speed.

Moreover, the model cleverly extrapolates the shockwave occurrence using few speed measurements obtained from the transient dynamics (or the onset of congestion). This can be clearly seen in the lower pane of Fig. 4 (i, iii), where the probe trajectories only capture slowing traffic (cyan color), yet the model still managed to detect the occurrence of a shockwave and reconstruct it accordingly. Also, in cases when there was a sufficient number of representative trajectories, the neural network learned to reconstruct smaller shockwaves and their dissipation behavior as seen from Fig. 4 (i, iv, v).

There are, however, a few cases where the method did not perform well. Where probe trajectories are very sparse, the method produces poor estimates of the propagation length and dissipation behavior of the shockwaves. As a result, the reconstructed upstream and downstream jams are either exaggerated or missed; see Fig. 5 (i, ii, v). Another shortcoming is the model's extrapolation of nonexistent shockwaves. The model learns that with the occurrence of any lower speeds (cyan or yellow), a congestion instance will take place, which is not always true (see Fig. 5 (i, iv)). This is caused by insufficient information about the local traffic conditions. The backward propagation of stop-and-go traffic depends on

the local densities, and if the observed trajectories failed to capture this, then the model may overestimate a shockwave occurrence. Also, if the sample trajectories are very sparse (as in Fig. 5 (iii)), the model only succeeds in predicting the represented parts and fails elsewhere by completely ignoring the section for which no information exists. Additionally, one can notice the random reconstruction of white areas which depict the absence of vehicle trajectories. If the areas are too large, the model almost always fails in reconstructing these areas as they don't follow any specific pattern. Notice, however, that the undesirable cases described above are very few, and not representative of the model's overall performance. They are just included here for completeness.

In order to assess the model's estimation precision with respect to real traffic behavior, we now validate it against NGSIM data in Fig. 6. The reason behind cross validating our model with unseen and real data is to ensure that the encoder-decoder neural network has in fact learned to mimic traffic behavior rather than memorizing images similar to the training data. Results show that the method successfully captures most of the shockwaves and their varying characteristics. However, since the neural network was only trained on simulation data, it performs less accurately when estimating shockwaves than in Fig. 4. In real traffic, congestion and free flow patterns depend on unobservable aspects that a simulator cannot always reproduce, and therefore we can see how the model didn't accurately reconstruct the shockwave propagation, for example in Fig. 6 (v). The same applies to

the free flow regimes. The method was not trained on cases with pure free flow behavior, which explains the encoder-decoder's poor estimation power in uncongested areas, see Fig. 6 (iii, iv).

To properly quantify the error difference between the reconstructed and the actual traffic speed maps, we use the Structural Similarity Index Method [45], SSIM. Unlike conventional distance functions such as mean squared error, SSIM captures changes occurring in the image's structural information. It is defined as:

$$\text{SSIM}(\mathbf{a}, \mathbf{b}) = \frac{(2\mu_{\mathbf{a}}\mu_{\mathbf{b}} + c_1)(2\sigma_{\mathbf{ab}} + c_2)}{(\mu_{\mathbf{a}}^2 + \mu_{\mathbf{b}}^2 + c_1)(\sigma_{\mathbf{a}}^2 + \sigma_{\mathbf{b}}^2 + c_2)}, \quad (6)$$

where \mathbf{a} and \mathbf{b} represent a sliding Gaussian window with a size of 11×11 , $\mu_{\mathbf{a}}$ and $\mu_{\mathbf{b}}$ are the means of \mathbf{a} and \mathbf{b} respectively, $\sigma_{\mathbf{a}}^2$ and $\sigma_{\mathbf{b}}^2$ are the variances of \mathbf{a} and \mathbf{b} , respectively, $\sigma_{\mathbf{ab}}$ is the covariance of \mathbf{a} and \mathbf{b} , c_1 and c_2 are stabilizing variables where $c_1 = (k_1L)^2$ and $c_2 = (k_2L)^2$, L is the dynamic range of pixels, k_1 and k_2 are default constants set to 0.01 and 0.03, respectively (as defined by [45]). The resulting SSIM index value ranges between $[-1, 1]$ where 1 and -1 indicate perfect similarity and 0 indicates no structural similarity. We compute the SSIM for the free flow and congested regimes separately and the results are shown in Fig. 7 for the validation datasets.

We observe that the average SSIM for simulation data in the congested and free flow regimes are 0.6 and 0.5 respectively (Fig. 7a), implying the model's strong reconstruction power. On the other hand, the average SSIM values for the NGSIM data are lower with 0.4 and 0.45 for congestion and free flow regimes respectively (Fig. 7b). However, having $\text{SSIM}_{\text{cong}} > \text{SSIM}_{\text{free}}$ in both cases, indicates that the model is better at estimating slow-moving traffic than free flowing traffic. This, however, could probably be changed if the model was better trained with pure free flow behavior.

2) INTERPRETING THE SPEED RECONSTRUCTION MODEL

We provide here a visual interpretation of the features or patterns that each kernel in the encoder model has learned to detect. Each kernel is activated if it sees a specific pattern in the input space, and this is determined as follows [46]:

$$\mathbf{z}_k = \arg \max_{\mathbf{z}_k^{(l),a}} \mathcal{L}'(\mathbf{z}_k^{(l),a}(\Theta_k^{(l)})), \quad (7)$$

where $\mathbf{z}_k^{(l),a}(\Theta_k^{(l)})$ is the activated feature map in layer (l) using kernel k , and $\mathcal{L}'(\cdot)$ can be any monotonically increasing function (such as L^1 -norm or L^2 -norm) defined over the activated feature map for the kernel k in layer (l) . Thus, (7) computes an input space \mathbf{z}_k , which produces the maximum activations for the kernel $\Theta_k^{(l)}$. This is shown below in Fig. 8 for a selected subset of kernels in the first 3 layers of our trained speed reconstruction model.

Interestingly, we see that the patterns or features learned by the model are more relevant to the macroscopic traffic speed states. For instance, the kernels in the first layer (Fig. 8a) detect discrete traffic states such as congestion

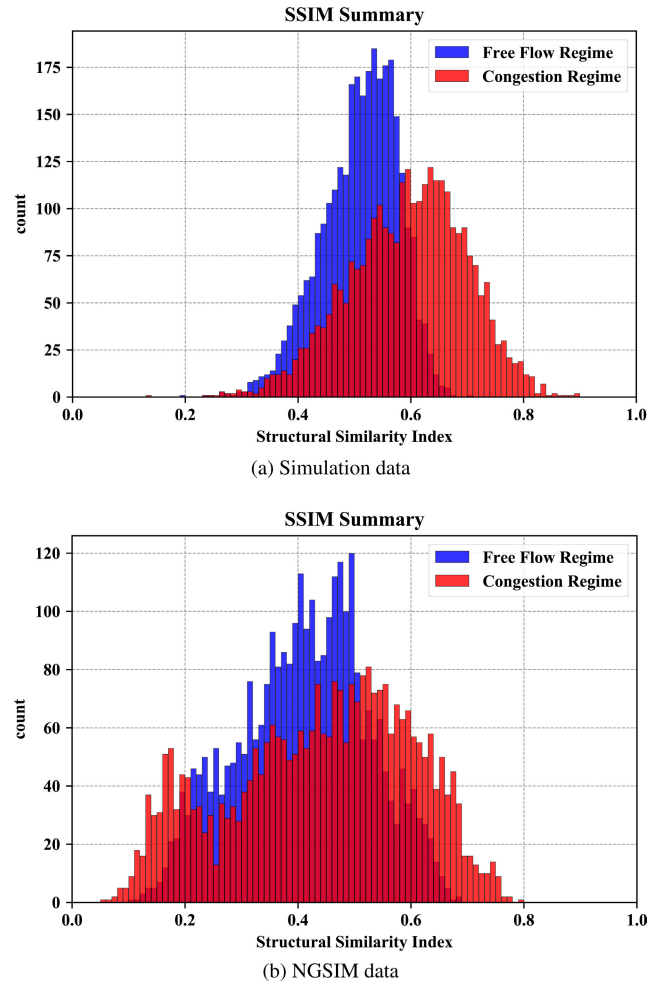


FIGURE 7. Histogram showing similarity indices between the actual and reconstructed spatiotemporal maps for the congested and free flow regimes.

(pale red color), free flow (purple color), slow moving traffic (red stripes) and vehicles in the transition regime (yellow). In the second layer (Fig. 8b), the kernels detect information propagation for congested and free flow traffic as seen from the sloped green, red, and purple regions. The third layer (Fig. 8c) combines these primitive features and produces more realistic spatio-temporal patterns such as a free flow regime, a congested regime, and transient dynamics (from free flow to congestion and vice-versa). This clearly illustrates the model's ability to recognize different traffic phenomena, and the mechanisms by which this information is used to reconstruct the complete spatio-temporal speed maps.

To further understand the reconstruction abilities of the model, we investigate the latent space representation (\mathbf{h}) for different probe levels. Here, $\mathbf{h} \in \mathbb{R}^{(10 \times 5 \times 64)}$ contains 64 feature maps, each of which has a dimension of (10×5) (see Table 1). We take the (normalized) sum of all these feature map activations in order to visualize them over a two-dimensional plane; this is shown in Fig. 9 for probe levels of 1%, 5%, 10%, and 100%.

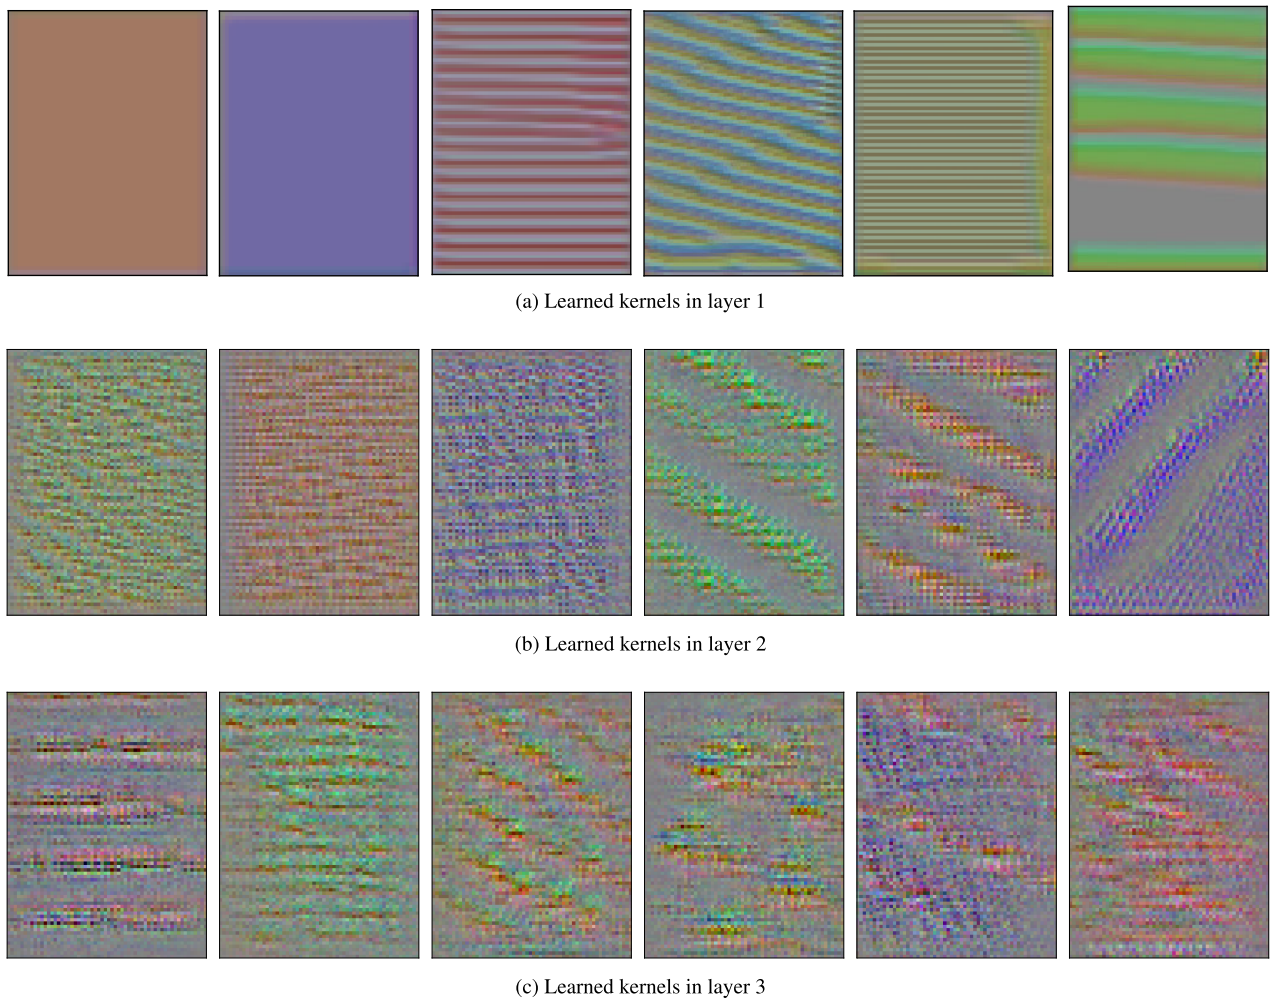


FIGURE 8. Primary features learned by each layer of the trained speed reconstruction model (only a subset of kernels in each layer are shown here).




FIGURE 9. Normalized latent space activations for different probe levels. To reduce the variance, the input trajectories are sampled 100 times and average activations are shown.

Recall that \mathbf{h} encodes the primary information from the sparse input vehicle trajectories, and the decoder model uses this encoded information to reconstruct the macroscopic speed states; i.e., the output from the decoder model

only depends on \mathbf{h} . Assuming that the visualized activations in Fig. 9 are a surrogate for the information contained in \mathbf{h} , we see similar patterns at 100% and 5% probe levels (in Fig. 9c and Fig. 9e), implying that the reconstructed speed

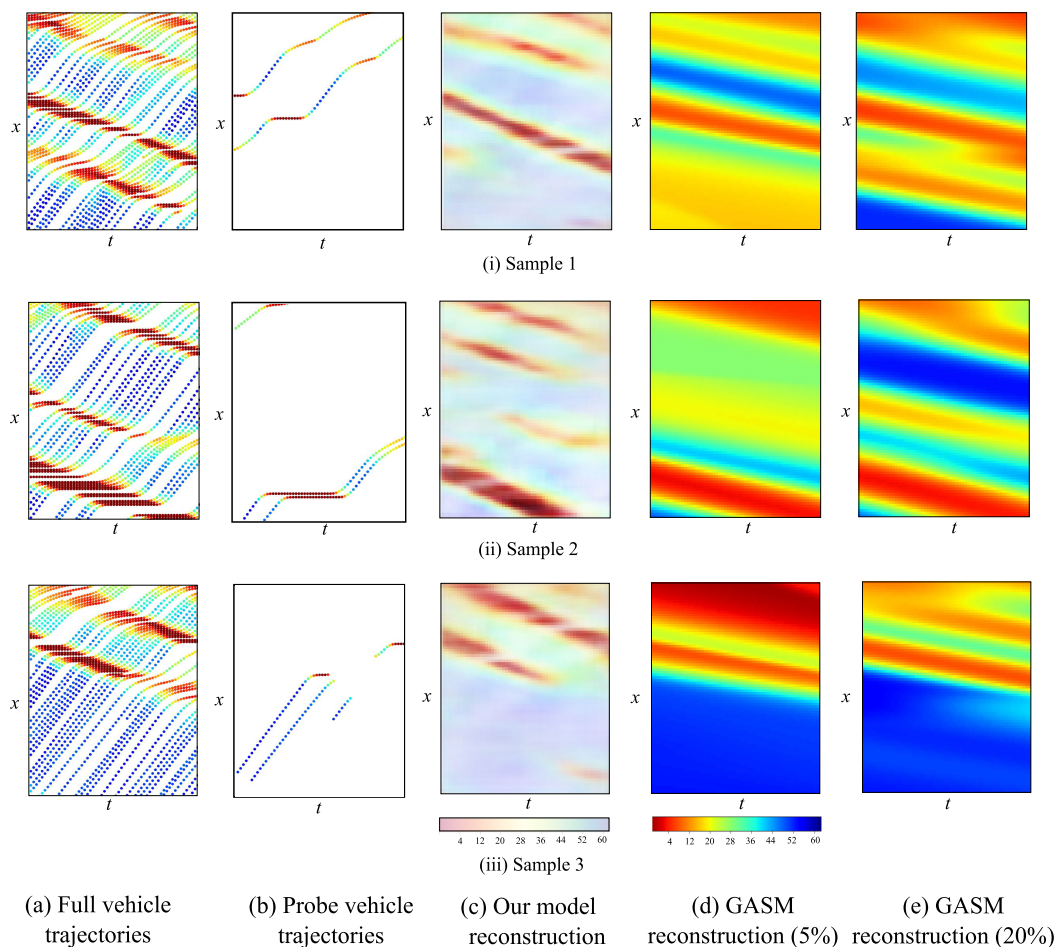


FIGURE 10. Comparison results between the proposed model and GASM (parameters used: $c_{free} = 60$ kmph, $c_{cong} = -15$ kmph, $v_{thr} = 25$ kmph, $\Delta V = 5$ kmph, $\lambda = 50$ m, $\tau = 15$ sec).

states will also be the same. This explains why the estimation model performs well even at 5% probe penetration rates.

C. COMPARISON RESULTS WITH ANOTHER TRAFFIC SPEED RECONSTRUCTION METHOD

We now provide a comparison of the proposed method with another popular traffic reconstruction methodology called the Generalized Adaptive Smoothing Method (GASM) [23]. It is a traffic speed interpolation technique over the two-dimensional time-space plane, and takes into account information propagation velocities in free flow and in congested traffic. GASM is different from the proposed method in that it uses a single asymmetric filter to interpolate unknown traffic speed on a time-space plane as opposed to a combination of several filters that are employed by our CNN. In addition, the asymmetric filter is fine-tuned manually using certain parameters, namely the free flow speed, the backward shockwave speed in congested traffic, and a threshold traffic congestion speed, whereas the CNN filters are learned from traffic speed data directly. In other words, GASM needs a priori information on three traffic parameters. More details of the GASM are provided in Appendix A.

In Fig. 10 we show a comparison of our method using 5% penetration rate and GASM reconstruction using 5% and 20% penetration rates. Using GASM, the results from a 20% penetration rate (Fig. 10e) are evidently better than those from a 5% rate (Fig. 10d). However, when compared to the real trajectories (Fig. 10a) and the results from the proposed approach (Fig. 10c), they still fail to capture the appropriate shockwave speeds. Below, we provide more insights from this comparison. (i) In general, GASM provides poor traffic speed estimation results at lower probe levels in comparison to the proposed method; see the GASM reconstruction results with 5% and 20% probe vehicle data. (ii) GASM predicts a constant shockwave speed for stop-and-go traffic, whereas the actual shockwave is dynamically varying and can be reproduced using our method (as demonstrated in Sec. III-B1). Moreover, the shockwave speed inferred by GASM differs from the actual speed (see Fig. 10d). Shockwave speeds in GASM are likely sensitive to the input parameter for the information propagation speed in congested traffic. A proper sensitivity analysis [47] could shed further light into these relations. CNN filters, on the other hand, learn these characteristics from the data, leading our model

to produce backward propagating stop-and-go traffic consistent with actual observations. (iii) When only a limited amount of information is available on the stop-and-go traffic, GASM produces implausible results. For instance, in sample 1, the second vehicle trajectory only witnesses a part of the shockwave (see Fig. 10b-i), and GASM extrapolated this observed lower speed to the rest of $x - t$ plane (see Fig. 10d-i). A similar result is also observed in the sample 3 reconstruction. This is in contrast with our method where the CNN filters learned the transition behaviors of congested to free flowing traffic (and vice-versa), even when the input probe trajectories do not completely capture these transitions; see our model reconstruction for sample 3. (iv) As previously mentioned, GASM reconstruction is likely very sensitive to the model parameters (free-flow speed, congestion wave speed) at this higher resolution ($10 \text{ m} \times 1 \text{ sec}$) traffic estimation, indicating that it needs very careful tuning; an additional step not required by the model proposed in this paper.

IV. CONCLUSION

In this paper, we address the problem of estimating dynamic traffic states (namely speeds) from limited probe vehicle data. We proposed a convolutional encoder-decoder neural network model to learn traffic speed dynamics from space-time diagrams. We illustrated this for a long road section and the results showed a sound and accurate reconstruction of macroscopic speed maps, even at 5% probe penetration rate. Notice that in the present setting of expressway traffic, network size has little effect on reconstruction quality. Data quality (e.g., probe vehicle penetration rate) has a more profound effect on the quality of the reconstruction. The model captured various spatio-temporal traffic behaviors such as backward propagating shockwaves, free-flow regimes and transient dynamics, which existing estimation methods fail to reproduce at low probe penetration rates. This was further confirmed with an analysis of the model's reconstruction mechanism where we visually observed that the kernels in the initial layers of the CNN learned to identify discrete traffic states: free flow, congested, and transient states. Successive layers recognized more compound patterns such as stop-and-go traffic, free flow propagation, and transient dynamics.

We also found that the model learned to reproduce dynamically varying shockwave patterns depending on the local traffic conditions, thus circumventing the constant wave speed assumption widely used in many of the existing estimation tools. Furthermore, when validated against the NGSIM dataset, the neural network successfully identified all shockwaves, despite being trained using simulated data. These insights attest to the viability of data-driven methods for real world applications. We also performed a comparison with another estimation method based on an asymmetric filter. The comparison demonstrated the superiority of the proposed method in speed reconstruction at lower probe levels. In the future, our efforts will continue along the lines of rigorous analysis of the model's reconstruction mechanism, as well as further applications to arterials with signalized intersections.

The latter will serve as a building block for the broader study of reconstruction of traffic states in a network setting, which is the ultimate aim of this line of research. The mathematical treatment of traffic across network intersections differs in a subtle way from that of expressway road segments. It is notable that traffic state reconstruction for urban networks has yet to be tackled effectively using classical techniques, the difficulty stemming from the intersection dynamics. We believe that the present approach using convolutional neural networks can be extended to networks with a clever choice of network architecture. We leave this to future research.

APPENDIX A

GENERALIZED ADAPTIVE SMOOTHING METHOD (GASM)

Define the macroscopic speed field as a function over position (x) and time (t) that provides the speeds at every (x, t) pair. We denote the macroscopic speed field as $V(x, t)$; the GASM treats $V(x, t)$ as a convex combination of two other speed fields, the free-flow speed field $V_{\text{free}}(x, t)$ and a congested speed field $V_{\text{cong}}(x, t)$. Given a set of n discrete traffic measurements from probe vehicles $\{x_i, t_i, v_i\}_{i=1, \dots, n}$, these speed fields are calculated as

$$V_{\text{free}}(x, t) = \frac{1}{N(x, t)} \sum_i \phi \left(x - x_i, t - t_i - \frac{x - x_i}{c_{\text{free}}} \right) v_i \quad (\text{A.1})$$

and

$$V_{\text{cong}}(x, t) = \frac{1}{N(x, t)} \sum_i \phi \left(x - x_i, t - t_i - \frac{x - x_i}{c_{\text{cong}}} \right) v_i, \quad (\text{A.2})$$

where c_{free} and c_{cong} are the information propagation speeds in free flow and congested traffic, respectively. The smoothing kernels $\phi_i(x, t)$ are localized functions that decrease with increasing $|x|$ and $|t|$. The GASM uses a bivariate gaussian kernel, defined as

$$\phi(x, t) = \exp \left[- \left(\frac{\|x\|_2^2}{2\lambda^2} + \frac{\|t\|_2^2}{2\tau^2} \right) \right], \quad (\text{A.3})$$

where λ and τ are the widths of the kernel in the x and t directions, respectively. Finally, $N(x, t)$ is a normalization field that is given by

$$N(x, t) = \sum_i \phi_i(x - x_i, t - t_i). \quad (\text{A.4})$$

The macroscopic speed field is obtained via a convex combination as follows:

$$V(x, t) = w(x, t)V_{\text{cong}}(x, t) + (1 - w(x, t))V_{\text{free}}(x, t), \quad (\text{A.5})$$

where the $w(x, t) \in [0, 1]$ is a weight field that depends on the free-flow and congested fields in the following way:

$$w(x, t) = \frac{1}{2} \left[1 + \tanh \left(\frac{V_{\text{thr}} - \min\{V_{\text{free}}(x, t), V_{\text{cong}}(x, t)\}}{\Delta V} \right) \right], \quad (\text{A.6})$$

where V_{thr} is a threshold speed and ΔV is interpreted as a transition width. These parameters are chosen a priori

based on field traffic observations. If the smaller of the two predictions $V_{\text{free}}(x, t)$ and $V_{\text{cong}}(x, t)$ is large relative to the threshold parameter V_{thr} , $w(x, t) \rightarrow 0$ and, consequently, $V(x, t) \rightarrow V_{\text{free}}(x, t)$. On the other hand, if the smaller of the two is small relative to V_{thr} , then $w(x, t) \rightarrow 1$ and $V(x, t) \rightarrow V_{\text{cong}}(x, t)$.

REFERENCES

- [1] C. Wu, A. M. Bayen, and A. Mehta, "Stabilizing traffic with autonomous vehicles," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2018, pp. 1–7.
- [2] R. E. Stern, S. Cui, M. L. D. Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D. B. Work, "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transp. Res. C, Emerg. Technol.*, vol. 89, pp. 205–221, Apr. 2018.
- [3] S. Cui, B. Seibold, R. Stern, and D. B. Work, "Stabilizing traffic flow via a single autonomous vehicle: Possibilities and limitations," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 1336–1341.
- [4] G. Tilg, K. Yang, and M. Menendez, "Evaluating the effects of automated vehicle technology on the capacity of freeway weaving sections," *Transp. Res. C, Emerg. Technol.*, vol. 96, pp. 3–21, Nov. 2018.
- [5] N. G. Polson and V. O. Sokolov, "Deep learning for short-term traffic flow prediction," *Transp. Res. C, Emerg. Technol.*, vol. 79, pp. 1–17, Jun. 2017.
- [6] A. Koesdwiady, R. Soua, and F. Karray, "Improving traffic flow prediction with weather information in connected cars: A deep learning approach," *IEEE Trans. Veh. Technol.*, vol. 65, no. 12, pp. 9508–9517, Dec. 2016.
- [7] S. E. Jabari and L. Wynter, "Sensor placement with time-to-detection guarantees," *EURO J. Transp. Logistics*, vol. 5, no. 4, pp. 415–433, Dec. 2016.
- [8] S. Jabari, F. Zheng, H. Liu, and M. Filipovska, "Stochastic Lagrangian modeling of traffic dynamics," in *Proc. 97th Annu. Meeting Transp. Res. Board*, Washington, DC, USA, 2018, pp. 1–14, Paper 04-18170.
- [9] F. Zheng, S. E. Jabari, H. X. Liu, and D. Lin, "Traffic state estimation using stochastic Lagrangian dynamics," *Transp. Res. B, Methodol.*, vol. 115, pp. 143–165, Sep. 2018.
- [10] S. E. Jabari, N. M. Freris, and D. M. Dilip, "Sparse travel time estimation from streaming data," *Transp. Sci.*, vol. 54, no. 1, pp. 1–20, 2020.
- [11] K. Yang, S. I. Guler, and M. Menendez, "Isolated intersection control for various levels of vehicle technology: Conventional, connected, and automated vehicles," *Transp. Res. C, Emerg. Technol.*, vol. 72, pp. 109–129, Nov. 2016.
- [12] S. I. Guler, M. Menendez, and L. Meier, "Using connected vehicle technology to improve the efficiency of intersections," *Transp. Res. C, Emerg. Technol.*, vol. 46, pp. 121–131, Sep. 2014.
- [13] L. Li and S. E. Jabari, "Position weighted backpressure intersection control for urban networks," *Transp. Res. B, Methodol.*, vol. 128, pp. 435–461, Oct. 2019.
- [14] T. Seo, A. M. Bayen, T. Kusakabe, and Y. Asakura, "Traffic state estimation on highway: A comprehensive survey," *Annu. Rev. Control*, vol. 43, pp. 128–151, Jan. 2017.
- [15] S. E. Jabari, J. Zheng, and H. X. Liu, "A probabilistic stationary speed-density relation based on Newell's simplified car-following model," *Transp. Res. B, Methodol.*, vol. 68, pp. 205–223, Oct. 2014.
- [16] S. E. Jabari and H. X. Liu, "A stochastic model of traffic flow: Theoretical foundations," *Transp. Res. B, Methodol.*, vol. 46, no. 1, pp. 156–174, Jan. 2012.
- [17] K. Yang and M. Menendez, "Queue estimation in a connected vehicle environment: A convex approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 7, pp. 2480–2496, Jul. 2019.
- [18] X. Ma, Z. Tao, Y. Wang, H. Yu, and Y. Wang, "Long short-term memory neural network for traffic speed prediction using remote microwave sensor data," *Transp. Res. C, Emerg. Technol.*, vol. 54, pp. 187–197, May 2015.
- [19] Y. Jia, J. Wu, and Y. Du, "Traffic speed prediction using deep learning method," in *Proc. IEEE 19th Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2016, pp. 1217–1222.
- [20] A. Hofleitner, R. Herring, P. Abbeel, and A. Bayen, "Learning the dynamics of arterial traffic from probe data using a dynamic Bayesian network," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 4, pp. 1679–1693, Dec. 2012.
- [21] J. Wang, Q. Gu, J. Wu, G. Liu, and Z. Xiong, "Traffic speed prediction and congestion source exploration: A deep learning method," in *Proc. IEEE 16th Int. Conf. Data Mining (ICDM)*, Dec. 2016, pp. 499–508.
- [22] M. Treiber and D. Helbing, "Reconstructing the spatio-temporal traffic dynamics from stationary detector data," *Cooperat. Transp. Dyn.*, vol. 1, no. 3, pp. 1–3, 2002.
- [23] M. Treiber, A. Kesting, and R. E. Wilson, "Reconstructing the traffic state by fusion of heterogeneous data," *Comput.-Aided Civil Infrastruct. Eng.*, vol. 26, no. 6, pp. 408–419, Aug. 2011.
- [24] X. Chen, S. Zhang, L. Li, and L. Li, "Adaptive rolling smoothing with heterogeneous data for traffic state estimation and prediction," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 4, pp. 1247–1258, Apr. 2019.
- [25] L. Ambühl and M. Menendez, "Data fusion algorithm for macroscopic fundamental diagram estimation," *Transp. Res. C, Emerg. Technol.*, vol. 71, pp. 184–197, Oct. 2016.
- [26] I. Dakic and M. Menendez, "On the use of Lagrangian observations from public transport and probe vehicles to estimate car space-mean speeds in bi-modal urban networks," *Transp. Res. C, Emerg. Technol.*, vol. 91, pp. 317–334, Jun. 2018.
- [27] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [28] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [29] J. Masci, U. Meier, D. Cireşan, and J. Schmidhuber, "Stacked convolutional auto-encoders for hierarchical feature extraction," in *Artificial Neural Networks and Machine Learning*, T. Honkela, W. Duch, M. Girolami, and S. Kaski, Eds. Berlin, Germany: Springer, 2011, pp. 52–59.
- [30] M. Menendez and C. Daganzo, "Assessment of the impact of incidents near bottlenecks: Strategies to reduce delay," *Transp. Res. Res. J. Transp. Res. Board*, vol. 1867, no. 1, pp. 53–59, Jan. 2004.
- [31] J. Mansoureh, J. Petronella, and S. Anthony, "Estimation of non-recurring post-incident traffic recovery time for different flow regimes: Comparing shock wave theory and simulation modeling," in *Proc. 90th Annu. Meeting Transp. Res. Board*, Washington, DC, USA, 2011, pp. 1–18, Paper 11-1756.
- [32] D. Scherer, A. Müller, and S. Behnke, "Evaluation of pooling operations in convolutional architectures for object recognition," in *Proc. Int. Conf. Artif. Neural Netw.*, 2010, pp. 92–101.
- [33] D. M. Dilip, N. M. Freris, and S. E. Jabari, "Sparse estimation of travel time distributions using Gamma kernels," in *Proc. 96th Annu. Meeting Transp. Res. Board*, Washington, DC, USA, 2017, pp. 1–13, Paper 02-17971.
- [34] S. E. G. Jabari, D. M. Dilip, D. Lin, and B. T. Thodi, "Learning traffic flow dynamics using random fields," *IEEE Access*, vol. 7, pp. 130566–130577, 2019.
- [35] United States Department of Transportation. (2006). *NGSIM—Next Generation Simulation*. [Online]. Available: <https://ops.fhwa.dot.gov/trafficanalysis/tools/ngsim.htm>
- [36] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in *Proc. Neural Inf. Process. Syst.*, 2012, pp. 2951–2959.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [38] X. Ban, R. Herring, P. Hao, and A. M. Bayen, "Delay pattern estimation for signalized intersections using sampled travel times," *Transp. Res. Res. J. Transp. Res. Board*, vol. 2130, no. 1, pp. 109–119, Jan. 2009.
- [39] X. Ban, P. Hao, and Z. Sun, "Real time queue length estimation for signalized intersections using travel times from mobile sensors," *Transp. Res. C, Emerg. Technol.*, vol. 19, no. 6, pp. 1133–1156, Dec. 2011.
- [40] P. Hao, X. Ban, K. P. Bennett, Q. Ji, and Z. Sun, "Signal timing estimation using sample intersection travel times," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 2, pp. 792–804, Jun. 2012.
- [41] P. Hao, X. Ban, D. Guo, and Q. Ji, "Cycle-by-cycle intersection queue length distribution estimation using sample travel times," *Transp. Res. B, Methodol.*, vol. 68, pp. 185–204, Oct. 2014.
- [42] P. Hao and X. Ban, "Long queue estimation for signalized intersections using mobile data," *Transp. Res. B, Methodol.*, vol. 82, pp. 54–73, Dec. 2015.
- [43] P. Izadpanah, B. Hellinga, and L. Fu, "Automatic traffic shockwave identification using vehicles' trajectories," in *Proc. 88th Annu. Meeting Transp. Res. Board*, 2009, pp. 1–14, Paper 09-0807.
- [44] Y. Cheng, X. Qin, J. Jin, and B. Ran, "An exploratory shockwave approach to estimating queue length using probe trajectories," *J. Intell. Transp. Syst.*, vol. 16, no. 1, pp. 12–23, Jan. 2012.

- [45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [46] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," 2013, *arXiv:1312.6034*. [Online]. Available: <http://arxiv.org/abs/1312.6034>
- [47] Q. Ge and M. Menendez, "Extending morris method for qualitative global sensitivity analysis of models with dependent inputs," *Rel. Eng. Syst. Saf.*, vol. 162, pp. 28–39, Jun. 2017.



OUFA BENKRAOUDA received the B.S. degree in architectural engineering from United Arab Emirates University, in 2013, and the M.S. degree in urban science and informatics from New York University, in 2015. She is currently pursuing the Ph.D. degree in urban and regional planning with the University of Illinois at Urbana–Champaign, USA. In 2015, she served as a Research Associate with the Center of Urban Science and Progress (CUSP) and NYC Parks. She was a Research Engineer with the Masdar Institute, United Arab Emirates and served as a Lead Researcher with National Pavilion U.A.E. La Biennale Di Venezia, from 2016 to 2018. Her research interests include artificial intelligence, urban science, and how connected and autonomous vehicles impact human behavior and infrastructure decisions.



BILAL THONNAM THODI received the B.Tech. degree in civil engineering from the Cochin University of Science and Technology, in 2014, and the M.Tech. degree in transportation engineering from the IIT Madras (IITM), in 2017. He is currently pursuing the Ph.D. degree in transportation planning and engineering with New York University, NY, USA. His research interests include the applications of artificial intelligence and machine learning techniques to problems in traffic operations.



HWASOO YEO (Member, IEEE) was born in Seoul, South Korea, in 1972. He received the B.S. degree in civil engineering from Seoul National University, Seoul, in 1996, and the M.S. and Ph.D. degrees in civil and environmental engineering from the University of California at Berkeley, Berkeley, CA, USA, in 2008. He is currently an Associate Professor with the Department of Civil and Environmental Engineering, KAIST. His current research interests include AI mobility, traffic flow and traffic operations, and intelligent transportation systems.



MÓNICA MENÉNDEZ (Member, IEEE) received the dual B.S. degree (*summa cum laude*) in civil and architectural engineering from the University of Miami, FL, USA, in 2002, and the M.S. and Ph.D. degrees (focusing on transportation) from the University of California at Berkeley, Berkeley, CA, USA, in 2003 and 2006, respectively.

Since January 2018, she has been an Associate Professor of civil and urban engineering with New York University, Abu Dhabi, and a Global Network Associate Professor of civil and urban engineering with the Tandon School of Engineering, New York University. She is also the Director of the Research Center for Interacting Urban Networks (CITIES). From 2010 to 2017, she was the Director of the Research Group Traffic Engineering, ETH Zurich and prior to that, a Management Consultant at Bain and Company. She has authored or coauthored over 60 journal articles and over 160 conference contributions and technical reports in the area of transportation. Her research interests include multimodal transportation systems paying special attention to new technologies and information sources. She is an active reviewer for over 20 journals, and a member of multiple editorial boards for top journals in transportation, including the IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS.



SAIF EDDIN JABARI received the B.Sc. degree in civil engineering from the University of Jordan and the M.S. and Ph.D. degrees in civil engineering from the University of Minnesota Twin Cities, in 2009 and 2012, respectively. He was a Postdoctoral Researcher with Math Department, IBM Watson Research Center, Yorktown Heights, NY, USA. He is an Assistant Professor of civil and urban engineering with New York University, Abu Dhabi, United Arab Emirates, and a Global Network Assistant Professor of civil and urban engineering with the Tandon School of Engineering, New York University, NY, USA. His research interests include the intersection of theoretical traffic flow modeling and statistical learning, with applications in traffic operations, including traffic state estimation, incident detection and localization, and real-time traffic control.

...