

# Incorporating Kinematic Wave Theory Into a Deep Learning Method for High-Resolution Traffic Speed Estimation

Bilal Thonnam Thodi<sup>1</sup>, Zaid Saeed Khan<sup>1</sup>, Saif Eddin Jabari<sup>2</sup>, *Senior Member, IEEE*,  
and Mónica Menéndez<sup>3</sup>, *Member, IEEE*

**Abstract**—We propose a kinematic wave-based Deep Convolutional Neural Network (Deep CNN) to estimate high-resolution traffic speed fields from sparse probe vehicle trajectories. We introduce two key approaches that allow us to incorporate kinematic wave theory principles to improve the robustness of existing learning-based estimation methods. First, we propose an anisotropic traffic kernel for the Deep CNN. The anisotropic kernel explicitly accounts for space-time correlations in macroscopic traffic and effectively reduces the number of trainable parameters in the Deep CNN model. Second, we propose to use simulated data for training the Deep CNN. Using a targeted simulated data for training provides an implicit way to impose desirable traffic physical features on the learning model. In the experiments, we highlight the benefits of using anisotropic kernels and evaluate the transferability of the trained model to real-world traffic using the Next Generation Simulation (NGSIM) and the German Highway Drone (HighD) datasets. The results demonstrate that anisotropic kernels significantly reduce model complexity and model over-fitting, and improve the physical correctness of the estimated speed fields. We find that model complexity scales linearly with problem size for anisotropic kernels compared to quadratic scaling for isotropic kernels. Furthermore, evaluation on real-world datasets shows acceptable performance, which establishes that simulation-based training is a viable surrogate to learning from real-world data. Finally, a comparison with standard estimation techniques shows the superior estimation accuracy of the proposed method.

**Index Terms**—Traffic state estimation, traffic anisotropy, kinematic wave theory, convolutional neural networks, deep learning.

## I. INTRODUCTION

**T**RAFFIC management agencies use a variety of monitoring and control tools to ensure the safe and efficient operation of network road traffic. To meet their operational goals, agencies employ tools that identify disturbances and deploy effective control strategies in real time [1]. However, this requires accurate and timely knowledge of traffic conditions over the entire network, which is currently not possible

Manuscript received 3 February 2021; revised 19 October 2021 and 11 February 2022; accepted 22 February 2022. Date of publication 23 March 2022; date of current version 11 October 2022. This work was supported by the New York University Abu Dhabi (NYUAD) Center for Interacting Urban Networks (CITIES) by Tamkeen through the NYUAD Research Institute under Award CG001. The Associate Editor for this article was X. Di. (*Corresponding author: Zaid Saeed Khan.*)

The authors are with the Division of Engineering, New York University Abu Dhabi, Abu Dhabi, United Arab Emirates, and also with the Department of Civil and Urban Engineering, New York University Tandon School of Engineering, Brooklyn, New York City, NY 11201 USA (e-mail: zaid.khan@nyu.edu).

Digital Object Identifier 10.1109/TITS.2022.3157439

given the limited sensory instrumentation in most (if not all) cities today. Fixed sensors are expensive and tend to be sparsely installed, offering limited spatial coverage. Data from mobile sensors are expected to become more widely available than data from point sensors, but remain extremely limited in practice; their sparsity is temporal [2]. To address such data sparsity (spatially or temporally), we need appropriate mechanisms that fill the gaps in the traffic observations. These are known as *traffic state estimation* (TSE) tools [3]. TSE is a critical precursor to a number of real-time traffic control strategies with either conventional vehicles or a mix with connected and autonomous vehicles [1], [4]. Such strategies include, but are not limited to, ramp metering, perimeter control, traffic signal control, and vehicle routing [5]–[7].

Existing TSE approaches can be broadly divided into two categories: model-based and data-driven [3]. The former approach adopts a mathematical model of traffic flow such as the first-order Lighthill-Whitham-Richards (LWR) model [8], [9] or one of its many higher-order extensions, like the Aw-Rascle-Zhang (ARZ) model [10], [11]. These methods assimilate flow model predictions with real-world observations using an exogenous filter (e.g., Ensemble Kalman filter) [12]–[19]. Traffic flow models ensure that estimates respect basic traffic principles. However, the models are based on simplifying assumptions of traffic physics that can lead to numerical bias when the assumptions are not met. Furthermore, approximation errors can arise from the data assimilation techniques used in TSE. For instance, it is common to linearize a non-linear flow model for the recursive estimation, and the approximations are poor around the capacity region [3], [16]. Lastly, model-based methods require additional inputs (e.g., boundary conditions) which are difficult to obtain in real-time.

The other category of TSE approaches include data-driven/learning techniques, which build statistical/machine learning models from large volumes of (historical) traffic data. Some commonly employed tools include (predominantly) deep neural networks [20], [21], support vector regression [22], principal component analysis [23], and matrix factorization methods [24]. The estimation results from data-driven methods are often reported to be more accurate than model-based approaches, but these methods also have shortcomings. Being purely data-driven, the models are agnostic to the physics of traffic flow and could lead to infeasible estimation results. These methods are also not often interpretable and

lack robustness. More importantly, the generalizability of the models is often weak and depends on the training data distribution.

We aim to develop a methodology that incorporates the desirable features of both categories, namely the combination of domain knowledge with representational power. Such structured learning methods can ensure robust and interpretable estimation results, parsimonious model complexity, and reduced data requirements. Some recent works along these lines include [25]–[30]. A common flavor in these approaches is to impose the physical constraints as cost function regularizers (i.e., as soft constraints) or derive the learning model architecture from physical principles. For instance, [28], [30] combine predictions of stochastic traffic flow model and limited probe vehicle data to infer vehicle trajectory distributions that are consistent with traffic physics. Reference [29] estimates queue lengths at signalized intersection as a solution to a convex optimization problem with queue propagation constraints guided by the kinematic wave theory of traffic flow. Reference [25] uses the dynamical equations of macroscopic flow models to regularize a Gaussian Process regression model, which is efficient in handling sparse and noisy data.

In the context of deep learning, which is a more attractive choice for non-linear modeling, [31]–[35] approximate the solution of a macroscopic traffic flow model using deep neural networks and use the governing physical dynamical equations (in the form of PDE/ODE) as a regularizer in the cost function. They demonstrate that these physics informed regularizers reduce the space of feasible solutions and learn solutions that are consistent with the chosen traffic flow models under limited real-world data. However, this requires training deep neural networks for every instance of initial/boundary conditions, which is computationally expensive for real-time implementation.

We propose an addition to this nascent literature on structured learning methods for TSE that incorporates traffic domain knowledge into learning models. Specifically, we propose a methodology to estimate high-resolution macroscopic traffic speed fields from limited probe vehicle measurements. We use a Deep Convolutional Neural Network (Deep CNN) as the learning model for estimation. The Deep CNN model takes as input sparse vehicle trajectory measurements and outputs a high-resolution speed field over a given space-time domain. The model is trained offline and can then be applied for real-time estimation. We incorporate traffic-specific features into the learning model in two ways, which are described below.

First, the naïve isotropic kernels in the Deep CNN model are modified to capture the wave propagation characteristics of free-flowing and congested traffic, in accordance with the kinematic wave theory (KWT) of traffic flow [36]. We develop a Deep CNN with *anisotropic kernels* designed to consider space-time inputs that are in the direction of feasible traffic waves, bounded by forward waves in free-flow and backward waves in congested traffic. As a result, we can significantly reduce the effective number of kernel parameters and hence the Deep CNN model complexity. Further, restricting the CNN to

consider only the relevant spatio-temporal input points results in feasible and robust estimation of traffic shockwaves.

Second, we train our Deep CNN model using simulated traffic data. Apart from resolving the data availability issue, this approach allows us to take an empirical distribution of any desirable traffic flow model and use it to train the Deep CNN. The empirical distribution is a surrogate representation of the traffic physics underlying the simulation model. This is a broader approach to incorporate the governing physics as it is easier to generate data corresponding to complex traffic behaviors rather than integrating them into the model architecture as in existing physics-informed learning methods. We demonstrate this by training the Deep CNN model with data generated from a microscopic traffic simulator, which consists of behavioral car-following, lane-changing and gap-acceptance models, and then test it with real-world data having similar traffic characteristics. A natural trade-off of this approach is that the learning model does not capture the exact physical traffic dynamics, but can incorporate a wide range of complex traffic behaviors. Similar methods have been explored in the context of automated systems such as robotic controls and object detection, whereby researchers use high-fidelity simulators or synthetic data instead of real-world data to train deep neural network models [37], [38].

To summarize, the contributions of this paper are:

- 1) We develop an anisotropic kernel design for CNNs following the wave propagation characteristics of traffic flow. This could be applied to traffic state estimation, prediction, and data imputation. We also suggest an optimization procedure to learn the optimal weights for the anisotropic kernels.
- 2) We propose to use simulated traffic data for fitting the anisotropic Deep CNN model and test its performance on real-world datasets.
- 3) We demonstrate the use of the anisotropic Deep CNN model for speed field estimation at fine space-time resolutions (10 meters  $\times$  1 second in our experiments) using limited input vehicle trajectories (5% probe vehicle penetration rates). We show sample estimations of real-world traffic data from multiple sources.
- 4) We extend our estimation methodology to handle unknown probe vehicle penetration rates by introducing an ensemble version of our Deep CNN model.

The rest of the paper is structured as follows. We present the estimation problem setting, the anisotropic kernel design, and the optimization procedure in Section II. We then describe the training data generation and the training experiments in Section III. In Section IV, we present estimation results, compare the anisotropic CNN with the naïve isotropic variant, discuss the transferability of the estimation model to real-world freeway traffic, and explore the sensitivity of the results to different probe vehicle penetration levels. Finally, we conclude the study in Section V.

## II. ESTIMATION METHODOLOGY

### A. Traffic Speed Field Estimation Problem

A space-time domain  $\mathcal{D} = \mathcal{X} \times \mathcal{T}$  representing a given road section is discretized into homogeneous segments  $x_i \subset \mathcal{X}$  and




Fig. 1. The architecture of the Deep CNN (speed reconstruction model).

time intervals  $t_i \subset \mathcal{T}$ , such that  $\cup_i x_i = \mathcal{X}$  and  $\cup_i t_i = \mathcal{T}$ . Let  $V(x, t)$  denote the value of the macroscopic speed field in a cell  $(x, t) \in \mathcal{D}$ . We use the cell size  $|x|$  closer to the length of a vehicle and  $|t|$  in the order of seconds (smaller than what existing estimation methods use [3]) to enable high-resolution speed field estimation. Probe vehicles (PVs) provide local speed measurements  $\{V^p(x_i^p, t_i^p)\}$  for some cells in  $\mathcal{D}$ ; we represent this partial information by the tensor  $\mathbf{z}^p$ . We assume sparse observation settings, where only a few cells (e.g., 5-10%) in  $\mathbf{z}^p$  have speed information. We denote by  $\mathbf{z}^f$  a tensor of estimates of the complete speed field  $V(x, t)$  for the entire space-time domain  $\mathcal{D}$ . The estimation problem can be formally stated as learning a mapping function  $g: \mathbf{z}^p \mapsto \mathbf{z}^f$ .

The speed field  $V(x, t)$  in each cell of the input tensor  $\mathbf{z}^p$  is encoded using a three-dimensional RGB array (with domain  $\{1, \dots, 255\}$ ) instead of a one-dimensional speed value. This is to differentiate cells occupied with a stopped vehicle (i.e., with  $V(x, t) = 0$ ) from empty cells. The output tensor  $\mathbf{z}^f$  represents the complete macroscopic speed field over the domain  $\mathcal{D}$  and can be encoded using the one-dimensional speed values. Thus, we have,  $\mathbf{z}^p \in \{1, \dots, 255\}^{|\mathcal{X}| \times |\mathcal{T}| \times 3}$  and  $\mathbf{z}^f \in \mathbb{R}_{\geq 0}^{|\mathcal{X}| \times |\mathcal{T}|}$ .

### B. Deep Convolutional Neural Network (Deep CNN) Model for Estimation

We use a Deep CNN model similar to the one in [20] to represent the mapping function  $g$ . The model architecture is shown in Fig. 1. It comprises an encoder  $g_{\text{enc}}$  and a decoder  $g_{\text{dec}}$ , each consisting of three CNN layers. Each CNN layer is composed of a 2D convolution operation, a non-linear activation operation called ReLU (Rectified Linear Unit), and a down-sampling operation called max-pooling (up-sampling operation called nearest neighbor in case of  $g_{\text{dec}}$ ). As shown in Fig. 1, the successive CNN layers of  $g_{\text{enc}}$  have reduced spatio-temporal widths and the successive CNN layers of  $g_{\text{dec}}$  have increased spatio-temporal widths. The Deep CNN model takes the input  $\mathbf{z}^p$ , passes it through the hierarchical convolution layers, and outputs the estimated speed field  $\mathbf{z}^f$ .

Unlike other neural network architectures, CNNs have proven to be effective in learning spatial data (e.g., images, video, etc.), which is useful for our application since the space-time diagram reflects spatial data. The CNN model has two properties favorable for learning macroscopic traffic features: local connectivity and parameter sharing. The former assumes the traffic speed fields are locally correlated, and the latter implies local traffic features can occur anywhere in




Fig. 2. Space-time correlations modeled by the isotropic kernel of the convolution operation, and that in the real traffic (free-flow and congested).

the space-time plane, i.e., they are space-time invariant. Furthermore, the specific encoder-decoder structure (bottleneck formation) of the model shown in Fig. 1 can efficiently handle the sparse nature of the model input [20].

The discrete convolution using local kernels in the CNN forms the basis of traffic speed field estimation. In a given CNN layer  $l$ , a convolution operation calculates the activation in a cell  $(x, t)$  as a weighted sum of cell activations observed in the previous layer  $(l - 1)$ :

$$\begin{aligned} \mathbf{z}^{(l)}(x, t, \chi) &= \mathbf{z}^{(l-1)}(\cdot, \cdot, \chi) * \Theta^{(l)}(\cdot, \cdot) \\ &= \sum_{(x_j, t_j) \in I_{\text{iso}}} \mathbf{z}^{(l-1)}(x_j, t_j, \chi) \Theta^{(l)}(x_j, t_j), \end{aligned} \quad (1)$$

where  $\mathbf{z}^{(l)}(x, t, \chi)$  is the feature map value in layer  $l$  associated with cell  $(x, t)$  and color channel  $\chi \in \{1, 2, 3\}$ ,  $\Theta^{(l)}(\cdot, \cdot) \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{T}|}$  is the kernel (matrix), which is identical for all cells.  $\Theta^{(l)}(x_j, t_j)$  on the right-hand side, an element of the kernel matrix, determines the extent to which neighboring cell  $(x_j, t_j) \in I_{\text{iso}}$  is correlated with the subject cell  $(x, t)$ . Hereafter, we simply write  $\Theta$  to represent the entire kernel, and drop the  $(\cdot, \cdot)$ .

The feature map value in cell  $(x, t)$  can be considered as equivalent to (or some function of) the speed field  $V(x, t)$  in that cell. Then, operation (1) simply says: the speed in cell  $(x, t)$  is a weighted interpolation of speeds observed in its immediate surrounding cells. The extent of local cell influence  $I_{\text{iso}}$  is depicted visually on the space-time plane in Fig. 2(a). Each kernel in a layer  $l$  represents a different weighting function; together, the kernels learn to identify different traffic features.

### C. Anisotropic Kernel Design for Deep CNN

The isotropic kernel shown in Fig. 2(a) says that the speed in cell  $(x, t)$  is correlated with the speeds observed *anywhere* in the shaded rectangular region  $I_{\text{iso}}$ . This assumes that a speed variation (such as that caused by slowdowns or speed-ups) at  $(x, t)$  can propagate at unbounded velocities in the space-time plane. However, this is not true in real traffic. In real traffic, (i) the speed/density variations propagate at finite velocities that are less than or equal to the free-flowing vehicle speed, and (ii) vehicles respond (predominantly) to frontal stimuli with a delay (approximately equal to the reaction time of driver). The former condition is called hyperbolicity and the latter is called anisotropy. Hyperbolicity is a necessary but not



sufficient condition for anisotropy in traffic flow models [39]–[41]. The use of local kernels (i.e., kernel dimensions  $\ll$  the dimension of space-time plane) captures the hyperbolicity property, whereas anisotropy can be captured by modifying the kernel shape as discussed below.

The actual propagation velocity of speed variations depends on the traffic state (i.e., speed or density). We assume that traffic at any point in the space-time plane is either in free-flow or is congested in relation to the Fundamental Diagram. There are different traffic conditions associated with free-flow and congestion, respectively. Then, a speed variation in cell  $(x, t)$  propagates downstream (i.e., in the direction of traffic) in free-flowing traffic and upstream (i.e., in the opposite direction of traffic) in congested traffic. This is an empirically and theoretically established feature of traffic [36], [42]–[45]. Thus, the extent of the space-time plane correlated with cell  $(x, t)$  depends on whether the traffic state is free-flow or congested. The respective correlated regions are shown in Fig. 2(b) and Fig. 2(c) as shaded areas  $I_{\text{free}}$  and  $I_{\text{cong}}$ .

The regions  $I_{\text{free}}$  and  $I_{\text{cong}}$  are bounded by the free flow traffic speed  $c_v$  and the backward shockwave speed  $c_w$  [39], [41], respectively. The speed in cell  $(x, t)$  influences the region  $I_{\text{free}}$  downstream, and the region  $I_{\text{cong}}$  upstream. Likewise, the regions  $I_{\text{free}}$  upstream and  $I_{\text{cong}}$  downstream influence the speed in cell  $(x, t)$ . In summary, the speed predicted in cell  $(x, t)$  is correlated with the speeds observed anywhere in  $I_{\text{free}} \cup I_{\text{cong}}$ . We use this knowledge of space-time correlations in designing an alternate and causally *correct* kernel (in the traffic sense) for the Deep CNN model in Fig. 1. We refer to this as the *anisotropic kernel*, and represent it by the tensor  $\Theta_{\text{ani}} = [\Theta_{\text{ani}}^{(l)}]$ . The corresponding convolution operation is slightly modified from (1) as,

$$\begin{aligned} \mathbf{z}^{(l)}(x, t, \chi) &= \mathbf{z}^{(l-1)}(\cdot, \cdot, \chi) * \Theta_{\text{ani}}^{(l)} \\ &= \sum_{(x_j, t_j) \in I_{\text{ani}}} \mathbf{z}^{(l-1)}(x_j, t_j, \chi) \Theta_{\text{ani}}^{(l)}(x_j, t_j), \quad (2) \end{aligned}$$

where the effective influence region is defined as  $I_{\text{ani}} := I_{\text{free}} \cup I_{\text{cong}}$ . This way, we direct the convolution operator to consider only that portion of the space-time plane which is relevant for the speed interpolation according to traffic physics.

In this paper, we propose a specific anisotropic kernel design, whose influence region is further restricted, motivated by empirical observations: (i) congested traffic has a very narrow range of wave propagation velocities (such that they can be regarded as almost constant), and (ii) free-flow traffic wave propagation velocities are limited within the maximum and minimum desired vehicle speeds [36], [39], [41]–[43]. The anisotropic kernel design to replace the isotropic kernel (from Fig. 1) is illustrated in Fig. 3. We create two kernels, one each for free-flowing and congested traffic. The influence region  $I_{\text{free}}$  contains all the cells passing and bounded between the maximum ( $c_v^{\text{max}}$ ) and minimum ( $c_v^{\text{min}}$ ) desired vehicle speeds. This is relevant for heterogeneous traffic where the desired speed distribution has a wide range. The free-flow traffic kernel is shown in Fig. 3(a). The influence region for congested traffic,  $I_{\text{cong}}$ , contains only those cells passing through the backward propagating shockwave speed  $c_w$ ; see

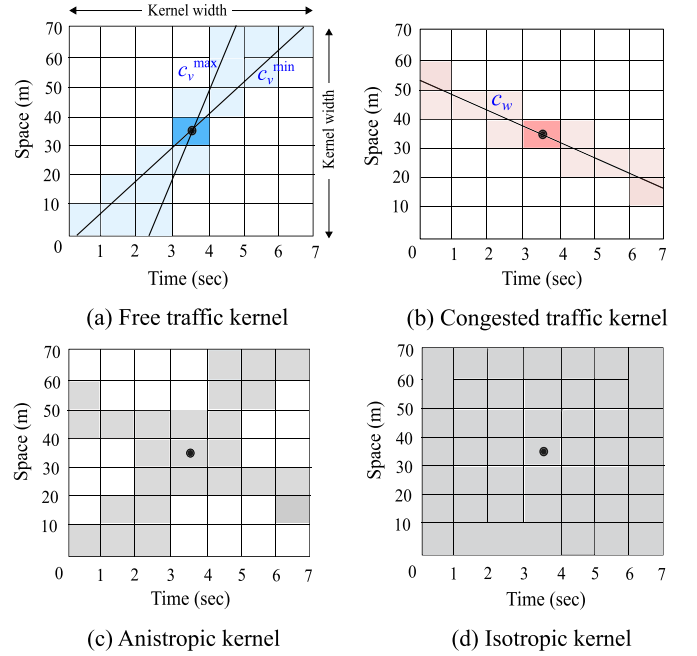


Fig. 3. The anisotropic kernel design for a  $7 \times 7$  CNN kernel (width = 7). The naïve isotropic kernel is also shown here for comparison. Parameters used:  $c_v^{\text{max}} = 100$  kmph,  $c_v^{\text{min}} = 60$  kmph and  $c_w = 18$  kmph.

Fig. 3(b). The proposed anisotropic kernel is a superposition of the free-flow and congested kernel. This is shown in Fig. 3(c). The corresponding isotropic kernel is shown in Fig. 3(d) for comparison. One can see that the anisotropic design requires 50% fewer parameters than its isotropic variant for a  $7 \times 7$  kernel.

In summary, our proposed anisotropic kernel design takes three input parameters  $\{c_v^{\text{max}}, c_v^{\text{min}}, c_w\}$ , whose values depend on the traffic characteristics of the road section. The proposed design aims to learn a broad range of forward propagation speeds and a narrow range of backward propagation speeds. Using a wide distribution for propagation speeds can simultaneously handle different road classes, e.g., highways with different speed limits and arterials. The variability in the free-flow speeds, in addition to capturing differences in speed limits, allows our kernels to capture a variety of kinematic wave speeds as combinations of free-flow waves and backward waves.

A practical benefit of the proposed anisotropic kernel design is the significant reduction in the model complexity of the Deep CNN model. Model complexity here refers to the total number of model parameters, and depends on the widths and depths of CNN kernels. We quantify the parameter requirements for isotropic and anisotropic kernels as a function of kernel widths in Fig. 4. The number of parameters scales *linearly* for anisotropic kernels as opposed to *quadratically* for isotropic kernels. This implies that anisotropic kernels result in a simpler, lower complexity CNN model which is easier to compute and optimize as compared to its isotropic counterpart. This scaling advantage is realized for higher kernel widths which naturally occur for larger problem sizes. We show these



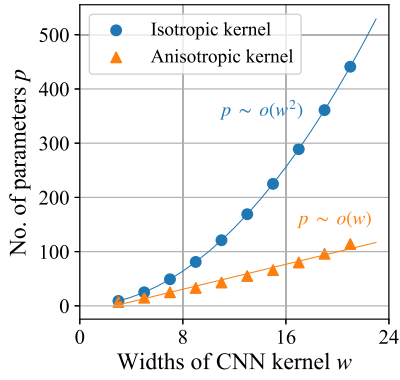


Fig. 4. Number of model parameters for different widths of isotropic and anisotropic kernel.

benefits experimentally in Section IV as we compare the model complexity requirements for different road network sizes.

We finally note that the proposed anisotropic kernel design is similar to Treiber and Helbing’s adaptive smoothing method for speed interpolation [42] except that: (a) we consider a range of wave propagation speeds in free-flowing traffic instead of a constant value, (b) the weights in the kernel are not set apriori as in [42] but learned from data, and (c) the actual speed predicted is a combination of several anisotropic kernels as opposed to a single anisotropic kernel.

#### D. Learning Anisotropic Kernels

We use anisotropic kernels in all layers of the Deep CNN model in Fig. 1. The optimal weights  $\Theta_{\text{ani}}^*$  for the anisotropic kernel are obtained from the following constrained optimization problem:

$$\Theta_{\text{ani}}^* := \arg \min_{\Theta \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{T}| \times L}} \{ \mathcal{L}(\mathbf{z}^f, g(\mathbf{z}^p, \Theta)) : \mathbb{1}_{\text{ani}} \odot \Theta = \mathbf{0} \}, \quad (3)$$

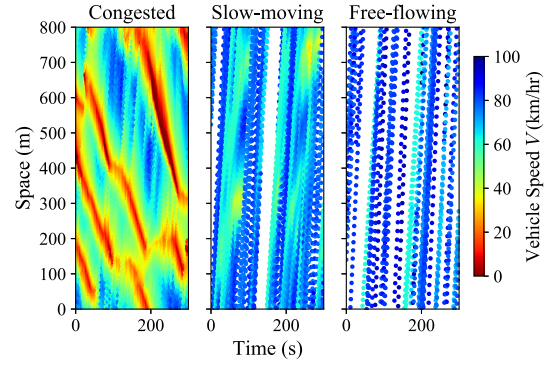
where  $g(\mathbf{z}^p, \Theta) : \{0, \dots, 255\}^{|\mathcal{X}| \times |\mathcal{T}| \times 3} \rightarrow \mathbb{R}_{\geq 0}^{|\mathcal{X}| \times |\mathcal{T}|}$  is the mapping function (i.e., the Deep CNN) with the kernel parameterization  $\Theta$  made explicit (i.e.,  $g$  performs the mapping  $\mathbf{z}^p \mapsto \mathbf{z}^f$ ),  $\mathbb{1}_{\text{ani}}$  is a binary tensor of the same dimension as  $\Theta$  with values of 0 for cells corresponding to the anisotropic influence cell region  $I_{\text{ani}}$ , e.g., the shaded cells in Fig. 3(c), and 1 elsewhere ( $\odot$  is the Hadamard product). The loss function  $\mathcal{L}$  captures any discrepancies between the estimated and true speed fields, e.g., the squared  $\ell_2$  distance (the squared error):

$$\mathcal{L}(\mathbf{z}^f, g(\mathbf{z}^p, \Theta)) = \left\| \mathbf{z}^f - g(\mathbf{z}^p, \Theta) \right\|_2^2. \quad (4)$$

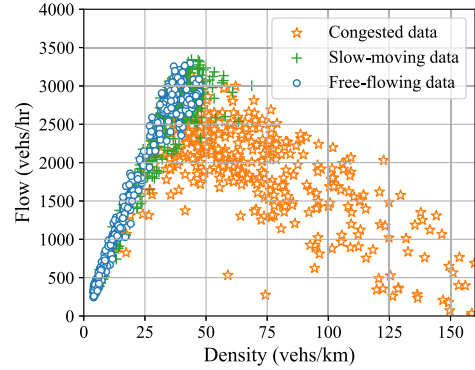
The constrained optimization problem (3) can be solved using iterative schemes which can handle feasibility constraints, such as the projected gradient descent. In each iteration  $i$ , the updates are calculated as follows:

$$\Theta_{\text{ani}}^{i+1} := \mathbf{P}_{I_{\text{ani}}}(\Theta_{\text{ani}}^i - \gamma^i G(\Theta_{\text{ani}}^i)), \quad (5)$$

where  $\gamma_i > 0$  is the step size (or learning rate) in iteration  $i$  and  $G(\Theta_{\text{ani}}^i)$  is a gradient tensor (descent direction) at  $\Theta_{\text{ani}}^i$ . The operator  $\mathbf{P}_{I_{\text{ani}}}$  assigns zeros to elements of  $\Theta_{\text{ani}}^i - \gamma^i G(\Theta_{\text{ani}}^i)$  corresponding to cells that lie outside of  $I_{\text{ani}}$ , thereby ensuring feasibility of the solutions.



(a) Space-time-speed contour plot.



(b) Flow-density scatter plot.

Fig. 5. Visualization of the richness or the traffic features contained in the simulated training dataset (300-second snapshot).

### III. DATA AND TRAINING

As mentioned earlier, we use simulated traffic data consisting of different traffic conditions for training the anisotropic Deep CNN model. In the following, we describe the data used for training and evaluating the model.

#### A. Training Data Generation

To generate data for training the CNN model, we simulate a freeway segment using the Vissim microscopic traffic simulator. The simulated segment corresponds to the *E-22 Abu Dhabi-Al Ain road, UAE* (2 miles in length and 3 lanes wide), and includes an entry and exit ramp to a nearby suburban region. The simulation model is calibrated with general traffic behavior, for instance, prioritizing through movements, appropriate yielding gaps for on-ramp vehicles, and minimum gap for lateral movements. A wide distribution of desired vehicle speeds (ranging from 60–100 kmph) is used to produce different free-flow wave propagation speeds as is the case for heterogeneous traffic.

We simulate three traffic scenarios with different input vehicle demand profiles on the freeway segment: 800-1200 vehs/hr, 2400-3000 vehs/hr, and 4200-5400 vehs/hr. We used these demand profiles to replicate distinct traffic conditions on the simulated freeway, namely free-flowing, slow-moving (moderately congested), and (heavily) congested traffic. We used on-ramp inflows that constitute 15-20% of the total freeway flows.

Each traffic scenario is simulated for 2 hrs and the vehicle trajectory data for an 800 m homogeneous section on the freeway is recorded. The trajectory data corresponding to three traffic scenarios and their traffic dynamics are summarized in Fig. 5.

Fig. 5(a) shows a 300 second snapshot of vehicle trajectories for the three simulated traffic scenarios. One can note the backward and forward propagating waves due to the stop-and-go, slow-moving, and heterogeneous free-flowing traffic (respectively) in Fig. 5(a). The anisotropic kernel is designed based on the range of wave propagation speeds seen in these plots. Fig. 5(b) is a flow-density scatter-plot of the three scenarios. Together, these figures show the richness of traffic states contained in the training data.

### B. Definition of Macroscopic Speed Field

An important auxiliary task is to define the “true” speed field which the Deep CNN model uses as the “ground truth” for evaluating the quality of the estimation. This is achieved by *translating* the set of *all* vehicle trajectories (not just PVs) into a speed field  $V(x, t)$ . The commonly used generalized definition of macroscopic speeds [46] results in  $V(x, t) = 0$  for some cells due to the fine mesh size we use. Therefore, we propose a simple interpolation method for this purpose instead. Our method interpolates the speeds over the road cells at a fixed time according to:

$$V(x, t) = \begin{cases} V_{\text{up}} \left( \frac{d_{\text{dn}}}{d_{\text{up}} + d_{\text{dn}}} \right) + V_{\text{dn}} \left( \frac{d_{\text{up}}}{d_{\text{up}} + d_{\text{dn}}} \right), & \text{if } d_{\text{up}} < l_{\text{up}} \\ & \text{and } d_{\text{dn}} < l_{\text{dn}} \\ V_{\text{up}} \left( 1 - \frac{d_{\text{up}}}{l_{\text{up}}} \right) + V_{\text{max}} \left( \frac{d_{\text{up}}}{l_{\text{up}}} \right), & \text{if } d_{\text{up}} < l_{\text{up}} \\ & \text{and } d_{\text{dn}} > l_{\text{dn}} \\ V_{\text{dn}} \left( 1 - \frac{d_{\text{dn}}}{l_{\text{dn}}} \right) + V_{\text{max}} \left( \frac{d_{\text{dn}}}{l_{\text{dn}}} \right), & \text{if } d_{\text{up}} > l_{\text{up}} \\ & \text{and } d_{\text{dn}} < l_{\text{dn}} \\ V_{\text{max}}, & \text{otherwise,} \end{cases} \quad (6)$$

where  $V_{\text{max}}$  is the highest free-flow speed (or speed limit of the highway section),  $V_{\text{dn}}$  (resp.  $V_{\text{up}}$ ) is the speed of the downstream (resp. upstream) vehicle,  $d_{\text{dn}}$  (resp.  $d_{\text{up}}$ ) is the distance between the cell  $(x, t)$  and the cell containing the downstream (resp. upstream) vehicle, and  $l_{\text{dn}}$  (resp.  $l_{\text{up}}$ ) is the length of spatial interaction downstream (resp. upstream) of  $(x, t)$ .

Equation (6) can be understood as follows: the speed field  $V(x, t)$  in cell  $(x, t)$  is a weighted combination of the speeds upstream and downstream of the cell. The speed  $V_{\text{dn}}$  of the vehicle downstream of  $(x, t)$  has an effect only if it is within the downstream interaction range  $l_{\text{dn}}$  from  $(x, t)$ ; otherwise, its value is replaced by the maximum highway speed  $V_{\text{max}}$  (and analogously for the upstream vehicle). The weights of the upstream and downstream components are proportional to the proximity of the respective interactions. The spatial interaction lengths are chosen to satisfy  $l_{\text{up}} < l_{\text{dn}}$ , to reflect the asymmetrically greater influence of frontal interaction.

### C. Training Procedure

The simulation output for each scenario is 7200 seconds of trajectory data for each of the three lanes. We first map the trajectories from a single lane onto a space-time plane to form an input and output frame of dimension  $80 \times 7200$  (i.e., the mesh size is  $10 \text{ m} \times 1 \text{ s}$ ). The PV trajectories for the input frame are selected at random using a 5% sampling rate. The output frame that forms the ground truth speed field is generated using the interpolation procedure described in eq. (6). We then extract samples of the input ( $\mathbf{z}^{\text{p}}$ ) and output tensors ( $\mathbf{z}^{\text{f}}$ ) from the input and output frames respectively, using a  $80 \times 60$  sliding window. We generate 6000+ samples for each trajectory dataset using a 2 s spatial gap between sliding windows. We proceed similarly to generate more data with different sets of random input samples for each of the three traffic scenarios using a 5% sampling rate. The final augmented dataset has 64000+ input-output sample pairs for training the Deep CNN model. Note that the samples extracted from a specific trajectory record form a sequence, which violates the i.i.d assumption (independent and identically distributed) for the neural network training. However, this is rectified during the optimization stage, where only a random subset of the samples is used in each iteration of the CNN training (this is a common trick employed while training reinforcement learning models, for instance, the use of “replay memory” in [47]). We use the following additional parameters for training data generation:  $|x| = 10 \text{ m}$ ,  $|t| = 1 \text{ s}$ ,  $c_w = 18 \text{ kmph}$ ,  $c_v^{\text{max}} = 100 \text{ kmph}$ ,  $c_v^{\text{min}} = 60 \text{ kmph}$ ,  $V_{\text{max}} = 95 \text{ kmph}$ ,  $l_{\text{up}} = 80 \text{ m}$  and  $l_{\text{dn}} = 40 \text{ m}$ .

We train five instances of the anisotropic and isotropic CNN models, and report the average of their performance results. We use the *TensorFlow* framework [48] to train all the models. The two major hyper-parameters, namely the CNN kernel width and depth in each layer, are independently optimized using the Hyperband algorithm [49], which belongs to the class of bandit-based algorithms. Other hyper-parameter choices are: gradient descent batch size: 32 samples, total training epochs: 300, (fixed) learning rate:  $1e - 3$ , and optimizer: Adam [50]. We use a GPU cluster with NVIDIA Tesla V100 32GB for training the models. The run time for a single training experiment is between 120 and 150 min. Note that the training can be viewed as an offline procedure.

### D. Testing Data

We test our model using three datasets: (i) a hold-out set from the simulated data that is not used for training (from a different lane of the freeway section), (ii) the Next Generation Simulation Program (NGSIM) dataset [51], and (iii) the German Highway Drone (HighD) dataset [52]. We choose the US-101 highway trajectory data from NGSIM, which contains the locations and speeds of all vehicles crossing the observed area during a 45 min time period with a 0.1 s resolution. The HighD data consists of trajectory data from several German highways, each consisting of a frame-wise recording of all vehicles passing a 400 m section during a 20 min duration, with a resolution of 25 frames/second. The input-output test

TABLE I  
MODEL ARCHITECTURE AS OBTAINED FROM THE HYPER-PARAMETER OPTIMIZATION

Layer name	Kernel widths	Kernel depths
Conv-1	(5 × 5)	40
Conv-2	(7 × 7)	48
Conv-3	(7 × 7)	32
Conv-4	(5 × 5)	48
Conv-5	(5 × 5)	40
Conv-6	(9 × 9)	56
Output	(7 × 7)	1

samples are generated similarly to the training datasets with the respective space-time discretization parameters.

We emphasize that we do not use the *NGSIM* or *HighD* datasets for training the model. In other words, the model is trained with data from a simulation using a freeway in the United Arab Emirates; and then tested with additional data from that same simulation, as well as real data from a freeway in the United States and several freeways in Germany. This allows us to evaluate the model’s transferability to diverse traffic scenarios and dynamics not seen in the training set, and the viability of using simulated data instead of real data for training.

#### IV. RESULTS AND DISCUSSION

In this section, we present the anisotropic reconstruction results and compare the isotropic and anisotropic models. We also discuss the transferability of the trained models to real-world traffic conditions and extend the results to handle varying PV penetration rates.

The architecture of the CNN model obtained from the hyper-parameter optimization is shown in Table I. We use the same optimized architecture for both the anisotropic and isotropic CNN models.

##### A. Anisotropic CNN Model Reconstruction

Fig. 6 shows five sample estimated speed fields from the hold-out simulated test dataset using the anisotropic model. The reconstruction window is 800 m × 60 s with a 10 m × 1 s resolution. The true speed field, PV trajectories, and speed profiles at three time instants ( $t = 10, 30,$  and  $50$  s) are also shown for each sample. Three of the samples correspond to congested traffic conditions, one corresponds to slow-moving traffic conditions, and one corresponds to free-flowing traffic conditions.

There are several points of interest to note about the reconstruction: All the estimated speed fields are feasible in terms of traffic physics and capture the different traffic states well. The model reproduces the existence of free-flow, congested and transition traffic dynamics correctly despite having very limited input information from the PV trajectories. One can observe the accurate prediction of shockwave dynamics in the congested traffic samples (a)-(c). This is also evident from the speed profile comparison. The true speed profile is often noisy, and the reconstruction has a smoothing effect due to the local convolutional operations in the CNN layers.

We have observed that the estimated speeds in slow-moving traffic have a higher root mean squared error (RMSE) than those in congested and free-flowing traffic; see Fig. 6 (d)-(e) and Table II. In slow-moving traffic, heterogeneity (caused by different vehicle characteristics, driving behaviors, etc.) is predominant, and one can see different forward wave propagation velocities in the speed field; see the example in Fig. 6 (d). Therefore, estimation is inherently a challenging problem unless we observe the actual travel speed. This is not the case for congested traffic, where the collective dynamics can be inferred from the trajectory of a single vehicle, or for free-flowing traffic, where the traffic heterogeneity is limited. In short, traffic speed fields with varied forward propagation wave velocities are still difficult to infer. Interestingly, in all the scenarios, the anisotropic model predicts the average desired vehicle speed in areas where there are no PV trajectories, which is a reasonable conclusion when no vehicles are observed.

We interpret the Deep CNN model as an interpolation function that locally propagates traffic characteristics (forward and backward waves) using the sparse information from input vehicle trajectories. The model ensures sound propagation of traffic information in space and time, resulting in speed field estimates with different traffic states - free-flowing, slow-moving, congested, and their transition states. Introducing anisotropic kernels further limits the propagation speeds of traffic information, in accordance with the Kinematic Wave Theory of traffic flow. This results in speed field estimates with a gradual and physically reasonable transition between the different traffic states. In contrast, traditional Kalman Filter based assimilation techniques only exploit state information from one (or a few) time step(s) when estimating the traffic speeds. This is inefficient in terms of data usage and fails to accurately reconstruct the dynamics.

In addition, we have tuned the Deep CNN model architecture to learn different traffic wave dynamics. Whether to produce a backward or forward wave depends on the traffic regime, which the model infers from the input trajectories. This is confirmed from the latent space projection of the data (i.e., the output from the encoder model), where three distinct clusters were generated, corresponding to free-flowing, slow-moving and congested traffic, respectively. Another way to put this is that a neural network model can solve an under-determined system - a major upside compared to other machine learning models. This is in contrast with traditional estimation methods which require additional information on initial/boundary conditions or traffic demands.

##### B. Comparison of Anisotropic and Isotropic Models

We next compare the performance and computational requirements of the anisotropic and isotropic models in Table II. The RMSE calculation shown in the table is the sample average for 4000+ simulated test samples. Overall, the anisotropic and isotropic models have similar performance in terms of accuracy, but the anisotropic model leads to more physically plausible shockwave dynamics (this is discussed below). In particular, the anisotropic model performs slightly



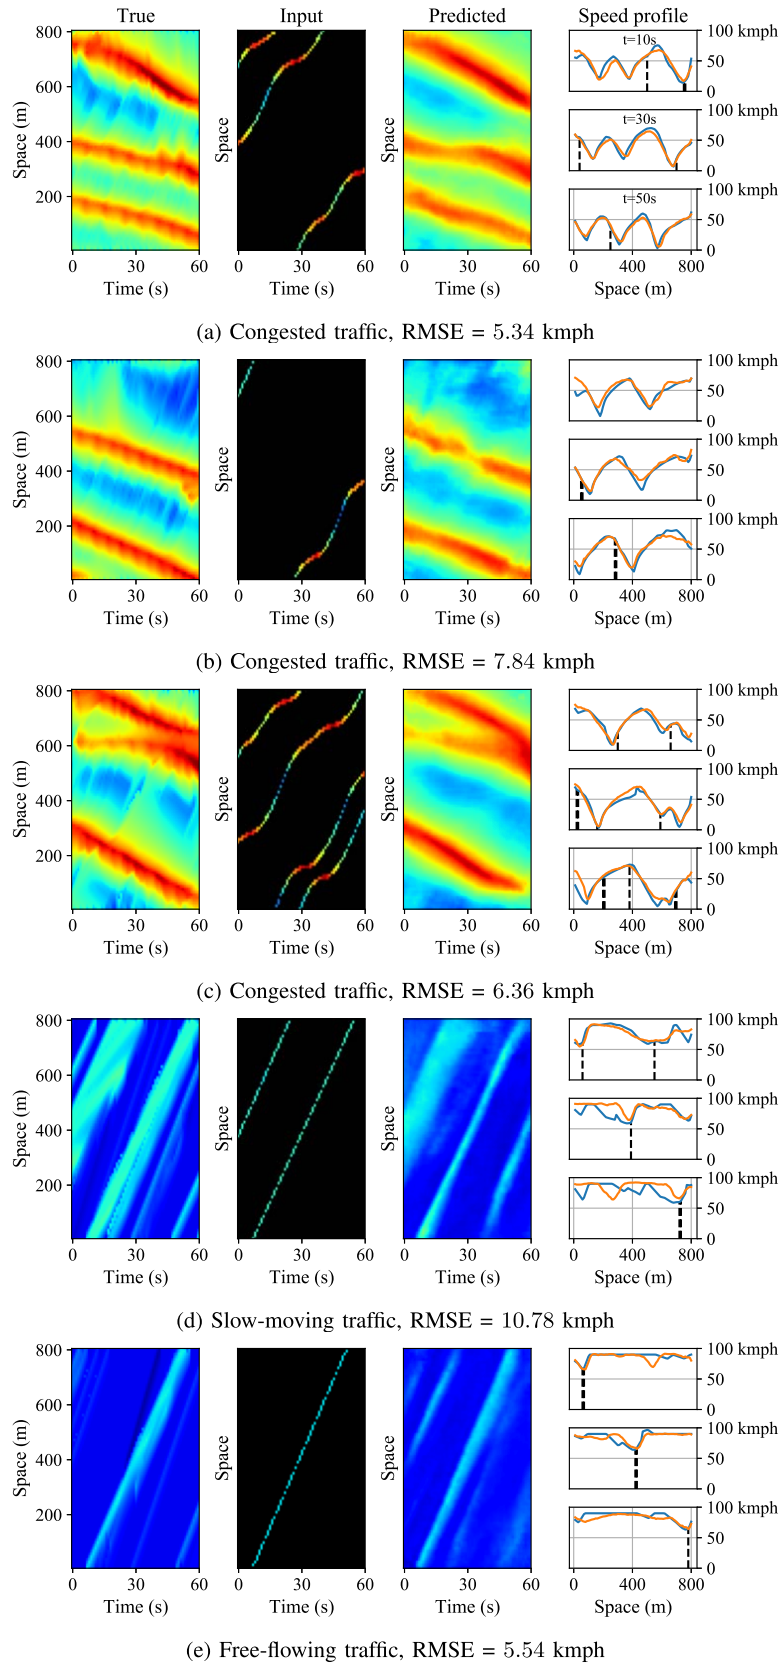


Fig. 6. Estimated speed field for selected samples in the simulated test data using the anisotropic CNN model. The fourth column shows the speed profile across the road section at  $t = 10, 30$ , and  $50$  secs (- true speeds, orange - estimated speeds, black vertical dashed line - input PV speed).

better in estimating the congested and free-flowing traffic in comparison to slow-moving traffic. This is because the

slow-moving data samples comprise heterogeneous traffic in the free-flow regime, which might be better observed by an

TABLE II

COMPARISON OF ANISOTROPIC AND ISOTROPIC MODELS. PERCENT CHANGE IS WITH RESPECT TO ISOTROPIC MODEL


Metric		Isotropic model	Anisotropic model	Percent change
Root mean squared error ( <i>kmph</i> )	Congested	8.60 (+/- 3.16)	8.50 (+/- 3.15)	-1.2%
	Slow-moving	10.37 (+/- 1.60)	10.53 (+/- 1.70)	+1.5%
	Free-flowing	7.42 (+/- 2.23)	7.40 (+/- 2.22)	-0.3%
	Total	8.71 (+/- 2.76)	8.76 (+/- 2.71)	+0.5%
Number of parameters		443193	215625	-51.4%

isotropic kernel than a restricted anisotropic kernel. Depending on the desired speed distribution, one can increase the extent of the anisotropic kernel (i.e., values of  $c_v^{\min}$  and  $c_w^{\max}$ ) and rectify this.


Table II also shows that the anisotropic model requires only half as many parameters as the isotropic model, which is a significant improvement in model complexity given that the performance of the two models is very similar. From a computational perspective, this is a substantial advantage, leading to faster model convergence (RMSE reduction per training epoch) and a potential reduction in the number of training samples required. This confirms that exploiting domain knowledge results in simpler and more interpretable learning models.

Although the isotropic and anisotropic models perform comparably in terms of the average error in estimating the speed, there are some examples where they differ in terms of the structure (speed and extent) of the shockwaves they produce. This is illustrated in Fig. 7, which shows certain examples where the anisotropic model clearly reconstructs more physically plausible shockwave dynamics, as mentioned below.


In the example in Fig. 7(a), the isotropic CNN underestimates the length of the shockwave at the top, whereas the anisotropic CNN correctly predicts that it existed some time prior to the two PV trajectories crossing it. This is because the anisotropic kernel gets more activation along the direction of the shockwave and hence reconstructs the stop-and-go region correctly, whereas the isotropic kernel considers all directions, which possibly results in averaging out all the nearby activations. Similar patterns have been observed in other test instances. The estimation in Fig. 7(b) is obtained using a single input trajectory. The anisotropic model gives a plausible reconstruction of the shockwave whereas the isotropic reconstruction shows large dispersion, which is also physically inconsistent with the input data. The design of anisotropic kernels can rule out such inconsistencies arising in the estimation. Fig. 7(c) shows a free-flowing traffic estimation. Again the forward wave produced by the isotropic kernel has more dispersion. In summary, one can see that the anisotropic model produces more accurate wave propagation



(a) Example 1 (Congested traffic, RMSE = 5.34 kmph)



(b) Example 2 (Congested traffic, RMSE = 7.84 kmph)



(c) Example 3 (Free-flowing traffic, RMSE = 6.36 kmph)

Fig. 7. Estimated speed fields for some selected samples in the simulated test data using the anisotropic and the isotropic CNN models. Black lines show the probe vehicle trajectories used for the reconstruction.

dynamics consistent with traffic physics, even though the RMSEs of the models are similar.

We finally compare the anisotropic and isotropic models from the perspective of model complexity and over-fitting measures. In order to understand how the complexity of CNN models scales with the road network size, we optimize the anisotropic and isotropic CNN model architectures for different road lengths. The optimization is done using the Hyperband algorithm [49]. The results are shown in Fig. 8(a), where the optimal number of model parameters required for different road lengths are compared (see the scatter plot). We see that the model complexity scales *quadratically* for isotropic kernels, whereas for anisotropic models it scales *linearly* (see the curve plot). Thus, as the problem size becomes large (for e.g., for long road sections, multiple lanes, or network level settings), the optimal CNN model architecture required for learning traffic dynamics becomes significantly large with isotropic kernels. The anisotropic CNN model, on the other hand, scales well to large problem sizes, results in simpler and more manageable models, and is beneficial for

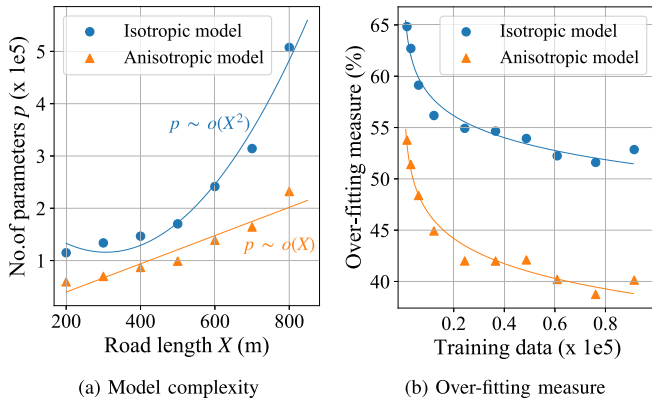


Fig. 8. Comparison of model-complexity and over-fitting measures for isotropic and anisotropic CNN models.

practical implementation. This observation is in-line with the scaling results obtained in Fig. 4.

In order to measure the over-fitting in CNN models, we define an over-fitting metric as,

$$O_{\text{fit}} = \left| \frac{\text{RMSE}_{\text{train}} - \text{RMSE}_{\text{test}}}{\text{RMSE}_{\text{test}}} \right| \times 100, \quad (7)$$

where  $\text{RMSE}_{\text{train}}$  and  $\text{RMSE}_{\text{test}}$  are the root mean squared error metrics for training and testing data respectively.  $O_{\text{fit}}$  measures the difference in the model's performance on the training and testing data; higher values for  $O_{\text{fit}}$  imply more over-fitting. Over-fitting is an undesirable feature, and indicates poor generalization to unseen testing data. Fig. 8(b) shows  $O_{\text{fit}}$  for the isotropic and anisotropic CNN models trained with different proportions of the total training data. The simulated training and simulated testing data are used to calculate  $O_{\text{fit}}$ . Note that the hyper-parameters of the CNN models are optimized independently to ensure that  $O_{\text{fit}}$  is compared for the optimal isotropic and anisotropic models. The trend line in Fig. 8(b) shows that the isotropic CNN model results in higher over-fitting. Since this observation is consistent at all data levels (and thus independent of model complexity), we conclude that the isotropic model has higher tendency to over-fit than the proposed anisotropic model. This is because the anisotropic CNN model reduces the number of parameters in a principled way, which lowers the model complexity without compromising test accuracy. In other words, the introduction of anisotropic kernels is a natural way to train CNN models that learn traffic speed dynamics with a lowered risk of over-fitting.

### C. Transferability to Real-World Traffic Dynamics

To understand how well the anisotropic CNN model performs in scenarios with different traffic characteristics than those observed in the training dataset, we test it on various real-world freeway sections. Figs. 9 to 11 show the estimation results for three sample freeway sections from the HighD and NGSIM datasets using data with a PV sampling rate of 5%.

A quick observation shows that all three example reconstructions are plausible, despite having different space-time dimensions from those used in the training dataset. This is possible because of the parameter sharing property of

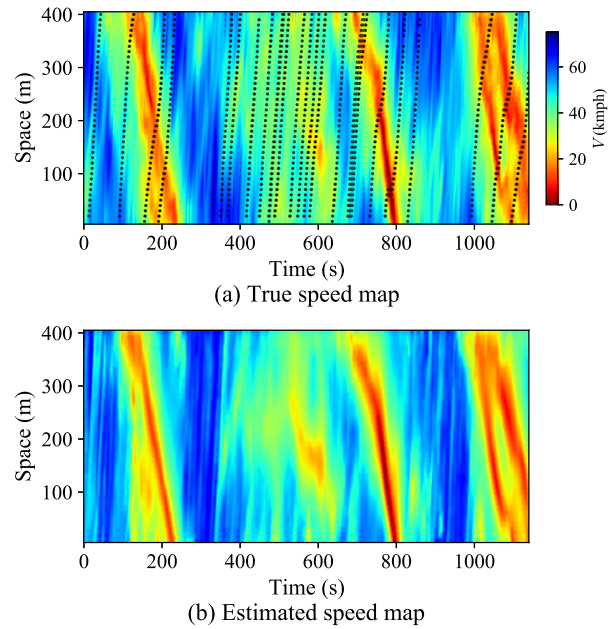


Fig. 9. Estimated speed field of lane 4 of highway No. 25 in the HighD dataset using 5% probe sampling rate. The road section is  $X = 400$  m long and the reconstruction period is  $T = 1140$  s. The RMSE is 6.80 kmph.

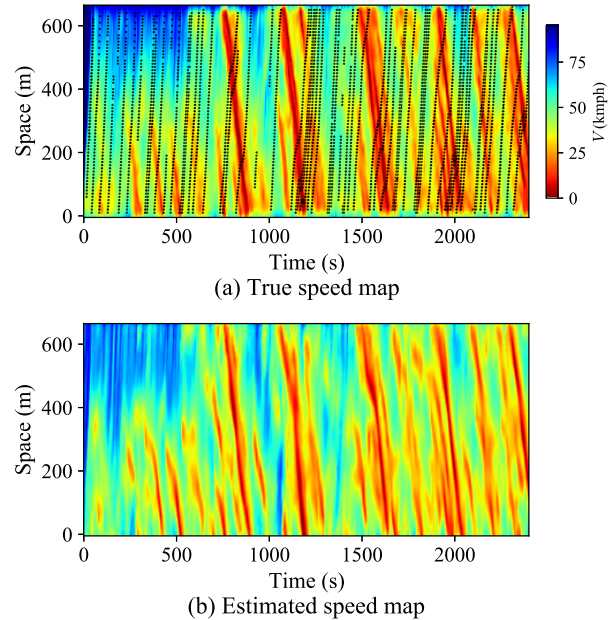


Fig. 10. Estimated speed field of lane 2 of U.S. Highway 101 in the NGSIM dataset using 5% probe sampling rate. The road section is  $X = 670$  m long and the reconstruction period is  $T = 2400$  s. The RMSE is 10.50 kmph.

CNNs, whereby the features learned during training (traffic characteristics in this case) are space-time invariant, and hence can be used with any spatio-temporal reconstruction window.

Closer observation reveals variation in the performance across the three examples. The estimated speed field in Fig. 9 has the lowest RMSE ( $\approx 6.80$  kmph), and the speed, width and duration of the predicted backward propagating shockwaves are accurate. In the estimation in Fig. 10, the shockwave reconstruction and speeds are reasonably correct, and the RMSE is moderate ( $\approx 10.50$  kmph). The model correctly predicts the onset of shockwaves 400 m upstream of the road






Fig. 11. Estimated speed field of lane 6 of Highway No. 44 in the HighD dataset using 5% probe sampling rate. The road section is  $X = 400$  m long and the reconstruction period is  $T = 1140$  s. The RMSE is 14.60 kmph.

section during the initial 600 seconds, though it underestimate the shockwave width. Thus, the model can accurately estimates speed fields for a road section that is simultaneously congested and free-flowing during the same period. The estimation result in Fig. 9 also supports this. The third freeway section, shown in Fig. 11, comprises free-flowing traffic and has the highest RMSE ( $\approx 14.60$  kmph). Apart from the inherent difficulties in the estimation of free-flowing traffic speeds, one can also observe that the speed of forward waves predicted by the model is slightly lower than that from the true waves; see the slow-moving band around 400 secs as an example. This is due to the difference in the traffic characteristics in the training and testing data, which we elaborate on below.

Recall that the CNN model trained on the simulated data encompasses the knowledge of traffic dynamics of a specific freeway section. How well this model transfers to other test scenarios depends on the traffic characteristics of the test segment. One can explain the difference in the RMSE errors in Figs. 9 to 11 by comparing the dynamics contained in the simulated data and the test data. One useful tool for this comparison is the flow-density scatter plot, which is shown in Fig. 12. The reason for the low RMSE in the first two examples is that the freeways are operating in the congested regime and the shockwave speeds in the simulated and test data are similar. Likewise, the reason for slightly lower prediction of free-flowing speed in the third freeway section (which operates in free-flowing traffic) is evident from Fig. 12(c). A similar reasoning is applicable when one discusses the transferability of the simulated section with its own real-world section, as the simulation doesn't capture the complete dynamics. Empirical FD comparisons can be further exploited to calibrate the trained deep learning model to match with the traffic dynamics of the test data. This is beyond the scope of the current work and represents a possible future extension.

TABLE III  
COMPARISON RESULTS TO EXISTING ESTIMATION TECHNIQUES


Estimation techniques	Root mean squared error ( <i>kmph</i> ) at different PV penetration rate		
	3%	5%	10%
Aniso CNN model (this paper)	11.60 (+/- 1.46)	10.70 (+/- 0.59)	8.88 (+/- 0.24)
GASM method (Treiber et al. 2011 [43])	13.49 (+/- 2.50)	11.80 (+/- 1.90)	9.50 (+/- 1.06)
LWR-v EnKF method (Daniel et al. 2008 [53])	14.93 (+/- 0.05)	14.64 (+/- 0.08)	13.63 (+/- 0.20)

In short, we observe that the Deep CNN models trained with simulated data from minimally calibrated traffic flow models transfer well to real-world traffic scenarios. Better results are to be expected with simulation models that are calibrated to the testing scenarios. Since the data required for calibration is lower than that needed for training deep neural networks, we can take advantage of well-developed microscopic traffic simulations to fit data-hungry models like CNNs.

We finally compare the anisotropic Deep CNN model performance with two other existing traffic speed estimation techniques in the literature: (a) General Adaptive Smoothing Method (GASM) from [43], and (b) Velocity-based LWR Ensemble Kalman Filtering technique (LWR-v EnKF) from [53]. While both techniques directly estimate the speed field over a given time-space plane, the former is a data assimilation technique using a macroscopic traffic flow model and the latter is an informed traffic interpolation procedure. The LWR-v EnKF method additionally requires initial and boundary conditions as inputs. The estimation results for the NGSIM US-101 highway lane 2 using the anisotropic CNN model, GASM and LWR-v EnKF are compared in Table III. The RMSE metric is evaluated for different input PV penetration rates to understand how well the techniques perform in sparse observation setting. Overall, we see that the anisotropic Deep CNN model results in the least estimation error for all PV penetration rates.

We found that the GASM method provides reasonable estimates at higher PV penetration rates (for e.g.,  $\geq 10\%$ ). However, at lower PV penetration rates (i.e., when the input only consist of one or two PV trajectories), the GASM method fails to reproduce correct traffic speed waves and results in higher estimation error. We have also noticed that the GASM method produces large dispersion in their estimates, which implies it poorly captures short-term traffic variations and is not suitable for high-resolution estimation. This is because the GASM method only uses two (pre-defined) kernels for interpolation, while our anisotropic CNN model uses an ensemble of (learned) kernels and hence interpolates low-level traffic features well.

Traffic speed estimated using the LWR-v EnKF method results in highest RMSE as shown in Table III. This poor performance could be due to its myopic character – only uses traffic speed inputs at the current time steps whereas the anisotropic CNN model considers inputs from multiple time steps. We also notice that the performance gets worse for



(a) HighD Highway 25 lane 4 (b) NGSIM US Highway 101 lane 2 (c) HighD Highway 44 lane 6

Fig. 12. Flow-density scatter plot comparing the traffic characteristics of the real-world datasets with the simulated training data.

longer estimation intervals, since the LWR model predictions deviate significantly from the actual data. Similar to GASM, the LWR-v EnKF also captures macroscopic traffic features and gets better at higher PV penetration rates.

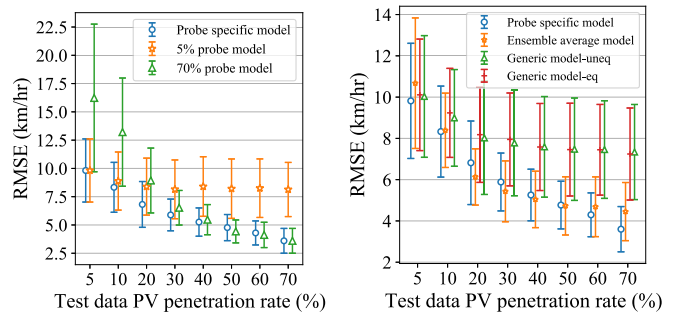
In short, the anisotropic CNN model outperforms the existing traffic speed estimation techniques, especially at lower PV penetration rates.

#### D. Variable Probe Vehicle Penetration Rates

To conclude our evaluation of the CNN model’s performance, we investigate the effect of changing the PV penetration rate. We train six separate models using data consisting of specific PV penetration rates 10%, 20%, . . . , 70% respectively (in addition to the 5% *probe model* discussed so far in this paper). The input-output pairs for training are generated in the same way as explained in Sec. III B. Each of these *probe specific models* is evaluated using testing data which has a corresponding PV penetration rate to the respective model. The average test RMSE results are shown in Fig. 13(a) (labeled “probe specific model”). As expected, the RMSE decreases with higher PV penetration rates.

However, we find that these probe specific models are not trivially generalizable to handle penetration rates other than what they were trained on, i.e., the models are penetration rate dependent. We demonstrate this by evaluating the performance of the two extreme probe specific models (i.e. the 5% probe model and the 70% probe model) on testing data across the whole range of PV penetration rates. It is clear that these probe specific models perform well only in/near their training domain. This could be due to the unconstrained latent space representation while training the CNN models, and is inevitable in any data driven models unless physical constraints are imposed.

The actual PV penetration rate depends on the prevailing traffic demand on the freeway, which is hard to measure in practice. We aim for an estimation model that performs well irrespective of the PV penetration rate. In other words, we want an estimation model that doesn’t require prior knowledge of the PV penetration rate. Therefore, we test three methods to handle varying PV penetration rates. The first two methods are brute force approaches, which consist of training the CNN model on a dataset containing the whole range of PV penetration rates 5%, 10%, 20%, . . . , 70%. The third method



(a) Probe specific models

(b) Generic models

Fig. 13. Probe vehicle (PV) penetration rate analysis.

is to use an ensemble CNN model. The RMSE results from these models are compared in Fig. 13(b).

The first model (labeled *generic model-eq*) is trained on a dataset consisting of all PV penetration rates sampled in equal proportion. The second model (labeled *generic model-uneq*) is similar except that we give more importance to lower PV penetration rates which are more difficult to learn. This is achieved by including more data samples for lower PV penetration rates, i.e., training data  $\propto 1/(\text{PV penetration rate})$ . Both these models, however, perform sub-optimally compared to the probe specific models. The best method is the third model (labeled *ensemble average model*), which takes the average of the predictions of all the probe specific models. This is referred to as “ensemble bagging” in the machine learning literature, and performs better than a single model trained on a wide range of penetration rates. As seen in Fig. 13(b), the ensemble CNN performs consistently well across all the PV penetration rates, even outperforming the respective probe specific models in certain cases. In addition to the performance, the individual models in the ensemble CNN (also called weak learners) can be trained in parallel, resulting in significantly lower training time than the other two generic models.

## V. CONCLUSION

Deep learning models have shown success in solving several inverse problems in traffic flow, but they are limited by their lack of robustness and poor model interpretability. In this paper, we overcome these limitations by proposing an anisotropic Deep Convolutional Neural Network (CNN)

model for estimating high-resolution traffic speed field using measurements from probe vehicles. The model employs anisotropic traffic kernels which are designed to explicitly capture a broad range of forward and backward propagation speeds in macroscopic traffic. Additionally, the Deep CNN model is trained using simulated traffic data. Since the generalization of Deep CNN performance depends on the distribution of training data, we note that using a targeted simulated data is an alternate method of imposing desirable traffic physics on the estimation model. For instance, we generate data corresponding to different traffic conditions (congested, slow-moving, free-flowing, etc.) so that the Deep CNN can learn traffic wave propagation speeds originating in heterogeneous traffic.

We present estimation results with input PV penetration rates as low as 5% and output resolution as high as  $10\text{m} \times 1\text{s}$ . In the experiments, we primarily focused on the benefits of using anisotropic kernels in the Deep CNN model over the naïve isotropic kernels. We found that anisotropic kernels result in parsimonious model complexity and are less prone to model over-fitting, although the estimation error is similar to their isotropic counterparts. The model complexity grows linearly with problem size for anisotropic kernels whereas it grows as quadratic for isotropic kernels. Specific examples are provided to demonstrate that the anisotropic kernels better produce physically correct traffic shockwaves. We further evaluated the anisotropic Deep CNN on real-world traffic datasets and found acceptable transferability performance. This suggests that simulated data is a viable surrogate to real-world data for training Deep CNNs. We also found that the Deep CNN model performance is PV penetration rate dependent and proposed an ensemble model to handle unknown PV penetration rates.

We believe that the optimal way to apply learning techniques to a specific domain such as traffic state estimation is to integrate the fundamental principles of the domain into the framework of the learning model. This paper represents only one possible example of this general approach. In future work, we aim to explore other methods to incorporate traffic flow theory into learning models such as CNNs.

#### ACKNOWLEDGMENT

The views expressed in this article are those of the authors and do not reflect the opinions of CITIES or its funding agencies.

#### REFERENCES


- [1] M. Papageorgiou, "Overview of road traffic control strategies," *IFAC Proc. Volumes*, vol. 37, no. 19, pp. 29–40, Oct. 2004.
- [2] 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.
- [3] 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.
- [4] L. Li, V. Okoth, and S. E. Jabari, "Backpressure control with estimated queue lengths for urban network traffic," *IET Intell. Transp. Syst.*, vol. 15, no. 2, pp. 320–330, Feb. 2021, doi: [10.1049/itr2.12027](https://doi.org/10.1049/itr2.12027).
- [5] 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.
- [6] K. Yang, N. Zheng, and M. Menendez, "Multi-scale perimeter control approach in a connected-vehicle environment," *Transp. Res. C, Emerg. Technol.*, vol. 94, pp. 32–49, Sep. 2018.
- [7] D. Lin and S. E. Jabari, "Pay for intersection priority: A free market mechanism for connected vehicles," *IEEE Trans. Intell. Transp. Syst.*, early access, Jan. 15, 2021, doi: [10.1109/TITS.2020.3048475](https://doi.org/10.1109/TITS.2020.3048475).
- [8] M. J. Lighthill and G. B. Whitham, "On kinematic waves II. A theory of traffic flow on long crowded roads," *Proc. Roy. Soc. A, Math. Phys. Eng. Sci.*, vol. 229, no. 1178, pp. 317–345, 1955.
- [9] P. I. Richards, "Shock waves on the highway," *Oper. Res.*, vol. 4, no. 1, pp. 42–51, 1956.
- [10] A. Aw and M. Rascle, "Resurrection of 'second order' models of traffic flow," *SIAM J. Appl. Math.*, vol. 60, no. 3, pp. 916–938, 2000.
- [11] H. M. Zhang, "A non-equilibrium traffic model devoid of gas-like behavior," *Transp. Res. B, Methodol.*, vol. 36, no. 3, pp. 275–290, 2002.
- [12] P. B. C. van Erp, V. L. Knoop, and S. P. Hoogendoorn, "On the value of relative flow data," *Transp. Res. C, Emerg. Technol.*, vol. 113, pp. 74–90, Apr. 2020.
- [13] 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.
- [14] M. Fountoulakis, N. Bekiaris-Liberis, C. Roncoli, I. Papamichail, and M. Papageorgiou, "Highway traffic state estimation with mixed connected and conventional vehicles: Microscopic simulation-based testing," *Transp. Res. C, Emerg. Technol.*, vol. 78, pp. 13–33, May 2017.
- [15] N. Bekiaris-Liberis, C. Roncoli, and M. Papageorgiou, "Highway traffic state estimation with mixed connected and conventional vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 12, pp. 3484–3497, Dec. 2016.
- [16] Y. Yuan, J. W. C. van Lint, R. E. Wilson, F. van Wageningen-Kessels, and S. P. Hoogendoorn, "Real-time Lagrangian traffic state estimator for freeways," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 1, pp. 59–70, Mar. 2012.
- [17] A. Nantes, D. Ngoduy, A. Bhaskar, M. Miska, and E. Chung, "Real-time traffic state estimation in urban corridors from heterogeneous data," *Transp. Res. C, Emerg. Technol.*, vol. 66, pp. 99–118, May 2016.
- [18] S. 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.
- [19] S. E. Jabari and H. X. Liu, "A stochastic model of traffic flow: Gaussian approximation and estimation," *Transp. Res. B, Methodol.*, vol. 47, pp. 15–41, Jan. 2013.
- [20] O. Benkraouda, B. T. Thodi, H. Yeo, M. Menendez, and S. E. Jabari, "Traffic data imputation using deep convolutional neural networks," *IEEE Access*, vol. 8, pp. 104740–104752, 2020.
- [21] 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.
- [22] J. Xiao, C. Wei, and Y. Liu, "Speed estimation of traffic flow using multiple kernel support vector regression," *Phys. A, Stat. Mech. Appl.*, vol. 509, pp. 989–997, Nov. 2018.
- [23] L. Li, Y. Li, and Z. Li, "Efficient missing data imputing for traffic flow by considering temporal and spatial dependence," *Transp. Res. C, Emerg. Technol.*, vol. 34, pp. 108–120, Sep. 2013.
- [24] W. Li, C. Yang, and S. E. Jabari, "Nonlinear traffic prediction as a matrix completion problem with ensemble learning," *Transp. Sci.*, vol. 56, no. 1, pp. 52–78, Jan. 2022.
- [25] Y. Yuan, Z. Zhang, X. T. Yang, and S. Zhe, "Macroscopic traffic flow modeling with physics regularized Gaussian process: A new insight into machine learning applications in transportation," *Transp. Res. B, Methodol.*, vol. 146, pp. 88–110, Apr. 2021.
- [26] Z. Zhang, Y. Yuan, and X. Yang, "A hybrid machine learning approach for freeway traffic speed estimation," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2674, no. 10, pp. 68–78, Oct. 2020.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] S. E. 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.



- [31] R. Shi, Z. Mo, K. Huang, X. Di, and Q. Du, "A physics-informed deep learning paradigm for traffic state and fundamental diagram estimation," *IEEE Trans. Intell. Transp. Syst.*, early access, Sep. 8, 2021, doi: [10.1109/TITS.2021.3106259](https://doi.org/10.1109/TITS.2021.3106259).
- [32] M. Barreau, J. Liu, and K. H. Johansson, "Learning-based state reconstruction for a scalar hyperbolic PDE under noisy Lagrangian sensing," in *Proc. 3rd Conf. Learn. Dyn. Control*, vol. 144, A. Jadbabaie *et al.*, Eds. Jun. 2021, pp. 34–46.
- [33] R. Shi, Z. Mo, and X. Di, "Physics-informed deep learning for traffic state estimation: A hybrid paradigm informed by second-order traffic models," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, vol. 35, no. 1, May 2021, pp. 540–547.
- [34] M. Barreau, M. Aguiar, J. Liu, and K. H. Johansson, "Physics-informed learning for identification and state reconstruction of traffic density," 2021, *arXiv:2103.13852*.
- [35] J. Huang and S. Agarwal, "Physics informed deep learning for traffic state estimation," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–6.
- [36] G. F. Newell, "A simplified theory of kinematic waves in highway traffic—Part I: General theory," *Transp. Res. B, Methodol.*, vol. 27, no. 4, pp. 281–287, Aug. 1993.
- [37] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2017, pp. 23–30.
- [38] J. Tremblay *et al.*, "Training deep networks with synthetic data: Bridging the reality gap by domain randomization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 969–977.
- [39] F. van Wageningen-Kessels, B. van't Hof, S. P. Hoogendoorn, H. van Lint, and K. Vuijk, "Anisotropy in generic multi-class traffic flow models," *Transportmetrica A, Transp. Sci.*, vol. 9, no. 5, pp. 451–472, May 2013.
- [40] C. F. Daganzo, "Requiem for second-order fluid approximations of traffic flow," *Transp. Res. B, Methodol.*, vol. 29, no. 4, pp. 277–286, Aug. 1995.
- [41] M. Treiber and A. Kesting, *Traffic Flow Dynamics: Data, Models, and Simulation*. Berlin, Germany: Springer, 2013.
- [42] 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.
- [43] 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.
- [44] C. F. Daganzo, "A variational formulation of kinematic waves: Basic theory and complex boundary conditions," *Transp. Res. B, Methodol.*, vol. 39, no. 2, pp. 187–196, Feb. 2005.
- [45] M. Menendez, and C. F. Daganzo, and H. S. Mahmassani, "A variational formulation of kinematic waves: Bottleneck properties and examples: Proceedings of the 16th international symposium on transportation and traffic theory (ISTTT)," in *Transportation and Traffic Theory: Flow, Dynamics and Human Interaction*. USA: Elsevier, 2005, pp. 345–364.
- [46] L. C. Edie, "Discussion of traffic stream measurements and definitions," in *Proc. Port New York Authority Int. Symp. Theory Road Traffic Flow*, New York, NY, USA, 1963, pp. 139–154.
- [47] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015.
- [48] M. Abadi *et al.* (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. [Online]. Available: <https://www.tensorflow.org/>
- [49] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *J. Mach. Learn. Res.*, vol. 18, no. 185, pp. 1–52, 2018.
- [50] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [51] United States Department of Transportation. (2006). *NGSIM-Next Generation Simulation*. [Online]. Available: <https://ops.fhwa.dot.gov/trafficanalysisistools/ngsim.htm>
- [52] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "The highD dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 2118–2125.
- [53] D. B. Work, O.-P. Tossavainen, S. Blandin, A. M. Bayen, T. Iwuchukwu, and K. Tracton, "An ensemble Kalman filtering approach to highway traffic estimation using GPS enabled mobile devices," in *Proc. 47th IEEE Conf. Decis. Control*, 2008, pp. 5062–5068.



learning techniques to problems in traffic operations.



Emirates. Previously, he was a Teaching Fellow at the Lahore University of Management Sciences from 2017 to 2019. His research interests include traveler information, traveler behavior, connected and autonomous vehicles, and congestion management.



University, he was a Post-Doctoral Researcher with the Math Department, IBM Watson Research Center, Yorktown Heights, New York. His research interests lie at 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.



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. Prior to that, she was a Management Consultant at Bain and Company. She has authored or coauthored over 85 journal articles and over 180 conference contributions and technical reports in the area of transportation. Her research interest includes 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*.

**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 Indian Institute of Technology Madras (IITM) in 2017. He is currently pursuing the Ph.D. degree in transportation planning and engineering with New York University, NY, USA. He is also a Global Ph.D. Student Fellow with New York University Abu Dhabi, United Arab Emirates. His research interests include applications of artificial intelligence and machine

**Zaid Saeed Khan** received the B.S. degree in computer science from the Lahore University of Management Sciences, Lahore, Pakistan, in 2015, and the S.M. degree in computation for design and optimization from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2017. He is currently pursuing the Ph.D. degree in transportation planning and engineering from New York University, New York, NY, USA. He is currently a Global Ph.D. Student Fellow and a Graduate Research Assistant with New York University Abu Dhabi, United Arab Emirates. Previously, he was a Teaching Fellow at the Lahore University of Management Sciences from 2017 to 2019. His research interests include traveler information, traveler behavior, connected and autonomous vehicles, and congestion management.

**Saif Eddin Jabari** (Senior Member, IEEE) 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 is currently an Associate Professor of civil and urban engineering with New York University Abu Dhabi, United Arab Emirates, and a Global Network Associate Professor of civil and urban engineering with the Tandon School of Engineering, New York University, NY, USA. Prior to joining New York University, he was a Post-Doctoral Researcher with the Math Department, IBM Watson Research Center, Yorktown Heights, New York. His research interests lie at 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.

**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, Berkeley, CA, USA, in 2003 and 2006, respectively. She is currently an Associate Dean of engineering for graduate affairs and a Professor of civil and urban engineering with New York University Abu Dhabi, and a Global Network 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. Prior to that, she was a Management Consultant at Bain and Company. She has authored or coauthored over 85 journal articles and over 180 conference contributions and technical reports in the area of transportation. Her research interest includes 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*.