

# Closed-Loop Neural Operator-Based Observer of Traffic Density

Alice Harting, Karl Henrik Johansson, and Matthieu Barreau

**Abstract**—We consider the problem of traffic density estimation with sparse measurements from stationary roadside sensors. Our approach uses Fourier neural operators to learn macroscopic traffic flow dynamics from high-fidelity data. During inference, the operator functions as an open-loop predictor of traffic evolution. To close the loop, we couple the open-loop operator with a correction operator that combines the predicted density with sparse measurements from the sensors. Simulations with the SUMO software indicate that, compared to open-loop observers, the proposed closed-loop observer exhibits classical closed-loop properties such as robustness to noise and ultimate boundedness of the error. This shows the advantages of combining learned physics with real-time corrections, and opens avenues for accurate, efficient, and interpretable data-driven observers.

## I. INTRODUCTION

Freeway congestion is a major issue in metropolitan areas, leading to increased travel times and excessive fuel consumption [1]. To address this, a core component of *Intelligent Transportation Systems* is freeway traffic control, which operates at the macroscopic level (ramp management and variable speed limits) or at the vehicle level (intelligent vehicle-based control). Traffic control relies on measurements of state variables such as traffic density, which can be collected from stationary sensors (inductive loops) or mobile sensors (connected vehicles). However, these measurements are often sparse and noisy. To complete and denoise the state observation, estimation approaches combine data and prior knowledge of traffic dynamics [2].

Model-based methods use physical models with calibrated parameters. Traffic state estimation then typically boils down to solving initial-boundary value problems (IBVPs) of partial differential equations (PDEs) with numerical methods such as the Godunov scheme [3], [4] or switching mode models [5]. To manage measurement noise and disturbance, it is common to use Kalman or particle filter-based approaches [6], [7]. For example, [8], [9] reformulate the continuity equation into a linear parameter-varying (LPV) system, where real-time flow measurements are treated as parameters. While this class of methods typically requires little data and offers high interpretability, the accuracy is highly dependent on model assumptions and calibration [2].

Data-driven methods, on the other hand, learn a system model directly from historical data, thereby relaxing ex-

PLICIT PDE assumptions and requirements of well-posedness [2]. Classical methods include ARIMA variants, Gaussian process (GP) models, and state-space Markov models. The rise of deep learning for nonlinear function approximation has enabled more expressive models of traffic flow prediction, with early examples such as the stacked autoencoder [10]. In particular, Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) are often used as building blocks to capture spatiotemporal correlation. While deep learning approaches are expressive and require few assumptions, they are limited by a high data demand and computational complexity, and low interpretability. Thus, lightweight real-time prediction with explainable models remains an open challenge [11].

In pursuit of data-efficient and interpretable models, hybrid models combine data-driven methods with prior physical knowledge. For example, [12] reconstructs the traffic density state using physics-informed neural networks (PINNs) trained on offline data from sparse mobile sensors while regularizing with a continuous traffic PDE model. In [13], traffic models are used in physics-regularized GPs (PRGPs), which rely on latent force models and regularized kernel learning. However, these methods are restricted by explicit model assumptions on global traffic behavior. Furthermore, online estimation can be computationally demanding, since it requires retraining the network (PINNs) or performing matrix inversion with cubic complexity in terms of the size of the dataset (PRGPs).

We identify the need for a data-driven observer of traffic density that encodes a learned governing model in an interpretable manner, and offers lightweight real-time prediction with the integration of online measurements. For finite-dimensional nonlinear systems, there are learning-based solutions in the form of closed-loop data-driven observers where training is done offline and low-cost inference is performed online [14]. The natural extension to PDE systems would be to consider Fourier neural operators (FNOs) [15], which have emerged as a method to learn operators between infinite-dimensional function spaces from finite-dimensional data.

Notably, [16] successfully applied FNOs to learn a PDE solution operator for IBVPs and inverse problems of a first-order macroscopic model of traffic density flow. With this method, inference with new input data is significantly cheaper than with PINNs, since it is performed by a single forward pass through the operator network. However, the solver in this work is designed to generate either an offline reconstruction or open-loop prediction, without closed-loop integration of online data.

Motivated by this, we introduce a data-driven, closed-

This work is supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

A. Harting, K.H. Johansson, and M. Barreau are with the Division of Decision and Control Systems, Digital Futures, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden {aharting,kallej,barreau}@kth.se

Method	Model approach			Observation mechanism		
	Model	Data	Learned closure	Online inference	Closed-loop	Grid-free
Numerical integration of IBVP [3]–[5]	✓	✗	✗	✓	✗	✗
Filtered state-space model [6], [7]	✓	✗	✗	✓	✓	✗
Filtered LPV-model with data-driven parameters [8], [9]	✓	✓	✗	✓	✓	✗
Stacked autoencoder [10]; CNN + RNN [11]	✗	✓	✓	✓	✗	✗
PINN [12]	✓	✓	✓	✗	✓	✓
PRGP [13]	✓	✓	✗	✓	✓	✓
KKL observer [14]	✓	✗	✗	✓	✓	✗
FNO [16]	✓ & ✗	✓ & ✗	✓ & ✗	✓	✗	✓
<b>FNO-based observer</b>	✓ & ✗	✓ & ✗	✓ & ✗	✓	✓	✓

TABLE I: Comparison of related methods for traffic state estimation.

loop neural observer of traffic density flow from sparse measurements. Our contributions are leveraging FNOs to learn a prediction operator of traffic density flow from high-fidelity data, and integrating Luenberger observer theory to get a robust density estimate using online measurements. We statistically evaluate the performance of the observer and compare it to open-loop variants. The relationship between the proposed observer and existing alternative methods is summarized in Table I.

The paper is organized as follows. Section II gives a brief introduction to traffic flow theory and neural operators, concluding with the problem statement. Section III describes the methodology for designing three types of observers: open-loop, open-loop with reset, and closed-loop. Section IV presents an extensive numerical evaluation of the proposed observers. Finally, Section V concludes the paper and outlines new research directions.

**Notation:** Let  $\mathcal{H}_D^d \triangleq H_{\text{per}}^2(D; \mathbb{R}^d)$  denote the Sobolev space of order two containing periodic functions with domain  $D$  and codomain dimension  $d$ . Further, we denote an operator with  $\mathcal{G}$  or  $\mathcal{N}$ . Finally, a GP is denoted  $\mathcal{GP}$ , and the posterior process associated with data  $\mathcal{D}$  is denoted  $\mathcal{GP}_{\mathcal{D}}$ . We assume that the mean and covariance functions are chosen such that sample paths belong to  $\mathcal{H}_D^d$  almost surely.

## II. BACKGROUND AND PROBLEM STATEMENT

We begin by briefly introducing basic traffic modeling theory, and then present neural operators in this context. This is followed by a formal statement of the research problem.

### A. Modeling of Traffic Density Flow

Traffic density is defined as the normalized number of vehicles per space unit at location  $x$  and time  $t$ ,  $\rho(x, t) \in [0, 1]$ . Macroscopic models describe the dynamics of  $\rho$  directly. While they are fast and simple, they typically rely on extensive assumptions. In contrast, microscopic models track each vehicle and offer greater precision, but are more computationally intensive and less convenient to use in freeway traffic control [1].

1) *Macroscopic scale:* A fundamental continuous first-order macroscopic model for traffic density flow is given by the *Lighthill–Whitham–Richards* framework. It expresses conservation of vehicles through the one-dimensional continuity equation, where the flow  $Q \in C^2([0, 1])$  is uniquely determined by the local density [17]:

$$\begin{aligned} \frac{\partial \rho(x, t)}{\partial t} + \frac{\partial Q(\rho(x, t))}{\partial x} &= 0, & x \in \mathbb{R}, \quad t \geq 0, \\ \rho(\cdot, 0) &= \rho_0, & x \in \mathbb{R}. \end{aligned} \quad (1)$$

In general, classical solutions to the initial value problem of the hyperbolic PDE (1) where  $Q$  is smooth may develop discontinuities in finite time even when the initial data  $\rho_0$  is smooth [18]. The weak solutions, on the other hand, are non-unique but can be filtered for physical relevancy using an entropy condition associated with the conservation law (1). In fact, whenever  $Q$  is smooth and  $\rho_0$  is well-behaved, there exists a unique admissible (entropic) solution  $\rho(\cdot, t)$  that is continuous in time and locally well-behaved [18, Theorem 6.2.1, 6.2.2]. In this view, traffic density estimation is a well-defined problem. However, a key challenge for macroscopic models is to properly model the flow  $Q$ .

2) *Microscopic scale:* To generate representations of higher fidelity without direct modeling of  $Q$ , traffic flow can instead be simulated using higher-order models at the vehicle level [1]. These models integrate interactions between vehicles and road infrastructure to produce realistic traffic flows. Vehicle interactions are commonly represented by car-following models, in which the behavior relative to the preceding vehicle or platoon is specified. For example, [19] proposed to let cars drive as fast as possible while maintaining a safe velocity to avoid collision. The model can be made increasingly complex by introducing parameters such as driver reaction time, estimation errors, and anticipation [20]. In this way, microscopic models may generate high-fidelity traffic flows with few restrictive assumptions. However, a key challenge is the lack of a compact representation of the dynamics that effectively integrates with macroscopic-level data and control strategies.

## B. Fourier Neural Operators (FNOs)

To address this challenge, neural operators provide a framework for learning the governing dynamics on the macroscopic level. Generally, neural operators learn a mapping between infinite-dimensional function spaces, that is, spaces containing functions with continuous domains, using finite-dimensional data. A special case is when the operator represents a forward solver for a hyperbolic PDE, such that the input function is the initial state and the output function is the solution at a final time  $T$ .

We introduce neural operators in the context of traffic density flow arising from higher-order dynamics on a ring road of length  $L$ , denoted  $\Omega \triangleq [0, L]$ . The closed course induces realistic traffic flows, including stop-and-go waves formed by tiny fluctuations [21], [22]. Consider pairs of functions

$$(P_0, \rho_T) \in \mathcal{H}_\Omega^n \times \mathcal{H}_\Omega^1,$$

where  $P_0 = [\rho(\cdot, 0), \rho(\cdot, -\Delta t), \rho(\cdot, -2\Delta t), \dots]$  represents the discrete  $n^{\text{th}}$ -order initial state of step length  $\Delta t$ , and  $\rho_T = \rho(\cdot, T)$  represents the resulting traffic density at time  $T$ . Assume there exists an operator  $\mathcal{G}$  such that

$$\begin{aligned} \mathcal{G} : \mathcal{H}_\Omega^n &\rightarrow \mathcal{H}_\Omega^1, \\ P_0 &\mapsto \rho_T, \end{aligned}$$

as can be motivated by the existence of a solution in the first-order case (1). Given observations of such pairs  $\{(P_0^i, \rho_T^i)\}_{i=1}^N$ , where the initial states are generated according to some distribution  $P_0^i \sim \mu$ , the goal is to learn an approximate solution operator  $\mathcal{G}_\theta \approx \mathcal{G}$  where  $\mathcal{G}_\theta$  is parametrized by a finite-dimensional vector  $\theta$ .

The approximate solution operator  $\mathcal{G}_\theta : \mathcal{H}_\Omega^n \rightarrow \mathcal{H}_\Omega^1$  can be modeled by a neural operator [15]:

$$\begin{aligned} v_1(x) &= \mathcal{P}[P_0](x), \\ v_M(x) &= \mathcal{L}_M \circ \dots \circ \mathcal{L}_1[v_1](x), \\ \rho_T(x) &= \mathcal{Q}[v_M](x), \quad \forall x \in \Omega, \end{aligned} \quad (2)$$

where  $\mathcal{P}$  is defined by a local transformation  $P : \mathbb{R}^n \times \Omega \rightarrow \mathbb{R}^{d_v}$ ,  $\mathcal{P}[P_0](x) \triangleq P(P_0(x), x)$  acting on function inputs and outputs directly to lift the input state to higher-dimensional space of size  $v$ . Similarly,  $\mathcal{Q}$  is defined by a local transformation  $Q : \mathbb{R}^{d_v} \times \Omega \rightarrow \mathbb{R}$ ,  $\mathcal{Q}[v](x) \triangleq Q(v(x), x)$  projecting back to the physical space. Both are typically shallow neural networks. The intermediate *spectral convolution layers*  $\mathcal{L}_l : \mathcal{V} \rightarrow \mathcal{V}$ , where  $\mathcal{V} = \mathcal{V}(\Omega; \mathbb{R}^{d_v})$ , are structured as

$$\mathcal{L}_l[v](x) = \sigma(W_l[v](x) + \mathcal{K}_{\phi_l}[v](x)), \quad \forall x \in \Omega,$$

where  $\sigma$  is a local nonlinear activation function, and  $W_l$  is a local pointwise affine transformation. The key component for the infinite-dimensional setting is the non-local bounded linear operator  $\mathcal{K}_{\phi_l} \in \mathcal{L}(\mathcal{V}, \mathcal{V})$ , which acts in function space.

A central aspect of FNOs [15] is that  $\mathcal{K}_{\phi_l}$  is defined by the convolution

$$\mathcal{K}_{\phi_l}[v] \triangleq \kappa_{\phi_l} * v, \quad (3)$$

where  $\kappa_{\phi_l} : \Omega \rightarrow \mathbb{R}^{d_v \times d_v}$  is a periodic, integrable function which has a Fourier series representation<sup>1</sup>. By the convolution theorem, (3) admits the frequency domain form

$$\mathcal{K}_{\phi_l}[v](x) = \mathcal{F}^{-1}(\mathcal{F}(\kappa_{\phi_l})(k) \cdot \mathcal{F}(v)(k))(x), \quad (4)$$

where  $\mathcal{F}$  computes the Fourier coefficients and  $k \in \mathbb{Z}$ . Parameterizing the convolution kernel directly in the frequency domain,  $R_{\phi_l} \triangleq \mathcal{F}(\kappa_{\phi_l})$ , and truncating the Fourier series to  $k_{max}$  frequency modes,  $R_{\phi_l} \in \mathbb{C}^{k_{max} \times d_v \times d_v}$  and  $\mathcal{F}(v) \in \mathbb{C}^{k_{max} \times d_v}$ , yields a finite-dimensional operator parametrization. Despite this truncation, the operator (3) remains defined on functions over the continuous domain  $\Omega$ , since  $\mathcal{F}^{-1}$  reconstructs from the basis functions  $e^{2\pi i \langle \cdot, k \rangle}$  for any  $x \in \Omega$ . In practice, input–output data consist of discretized rather than analytic functions, and the Fourier transform  $\mathcal{F}$  can be computed using the Fast Fourier Transform  $\hat{\mathcal{F}}$ . Crucially, truncating in the frequency domain – rather than the physical domain – allows for a parametrization that can be shared across resolutions. In other words, the resulting operator (2)–(4) is discretization convergent, and thus consistent in the infinite-dimensional setting.

Similar to neural networks, FNOs are capable of approximating a large class of operators between infinite-dimensional spaces up to arbitrary accuracy. Below, we present the universal approximation theorem from [23] as applied to the setting of this paper.

*Theorem 1:* Let  $\mathcal{G} : \mathcal{H}_\Omega^n \rightarrow \mathcal{H}_\Omega^1$  be a continuous operator, and  $K \subset \mathcal{H}_\Omega^n$  a compact subset. Let  $\mathcal{G}_\theta : \mathcal{H}_\Omega^n \rightarrow \mathcal{H}_\Omega^1$  be defined as in (2)–(4) with the coefficients  $\theta \triangleq \{\theta_P, \theta_Q, \theta_{W_l}, R_{\phi_l}\}$ . Then,  $\forall \epsilon > 0 \exists k_{max} \in \mathbb{N}$  and  $\theta \in \Theta$  such that  $\mathcal{G}_\theta$ , continuous as an operator, satisfies

$$\sup_{a \in K} \int_{\Omega} \|\mathcal{G}[a](x) - \mathcal{G}_\theta[a](x)\|^2 dx \leq \epsilon. \quad (5)$$

It is further suggested by [23] that, while the approximation error scales super-exponentially with the size of the FNO in the general case, it typically scales polynomially for PDEs of the form (1) under suitable assumptions. The expression (5) serves as the basis for a training loss function, and may be extended to include PINN-type regularization from a predefined PDE model [24].

## C. Problem Statement

Following this discussion, we model traffic density flow as an  $n^{\text{th}}$ -order discrete-time dynamical system

$$S[\bar{P}_0] : \begin{cases} \rho(\cdot, t + \Delta t) = \mathcal{G}[P_t], & t > 0, \\ P_t = [\rho(\cdot, t), \rho(\cdot, t - \Delta t), \dots], \\ P_0 = \bar{P}_0, \end{cases} \quad (6)$$

where  $\rho(\cdot, t) \in \mathcal{H}_\Omega^1$  is the true density,  $P_t \in \mathcal{H}_\Omega^n$  is the state at time  $t$ ,  $\bar{P}_0$  is the initial state, and  $\mathcal{G}$  represents the unknown solution operator of the dynamics.

The objective of this paper is to develop an algorithm in the form of a dynamical system that observes (6) in closed-loop by generating real-time predictions  $\hat{\rho}$  of the system state

<sup>1</sup>This aligns with our problem setting, as the ring road formulation entails that  $\rho(\cdot, t)$  is  $L$ -periodic on  $\Omega$ .

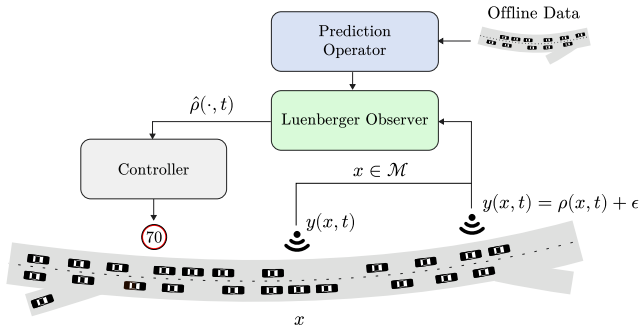


Fig. 1: Problem setting. Our goal is to design an observer using offline data to learn the dynamics, and online measurements from stationary sensors at  $x \in \mathcal{M}$  to calibrate predictions. The observer may then be used for macroscopic-level traffic control.

$\rho$  every  $\Delta t$  time unit such that  $\|\rho - \hat{\rho}\|$  is minimized. To achieve this goal, we assume access to online measurements  $y(x, t) = \rho(x, t) + \epsilon$  with i.i.d. noise  $\epsilon$  at stationary, sparse sensor locations  $x \in \mathcal{M} \subset \Omega$ , collectively denoted  $\mathbf{y}_{\mathcal{M}}(t)$ . We also assume access to offline observations  $\{\mathcal{S}[P_0^i]\}_{i \in \mathcal{I}}$ , where  $P_0^i \sim \mu$  for some distribution  $\mu$ . The problem setting is illustrated in Fig. 1.

*Remark 1:* The problem is inherently ill-posed due to the incomplete observation of the initial state. Thus, our goal is to learn the most probable solution that remains physically consistent with the dynamics in the offline observations.  $\blacklozenge$

### III. METHODOLOGY

To build a closed-loop observer of (6), the first step is to obtain a representation of the open-loop prediction,  $\mathcal{G}$ . Thus, our approach starts by learning an approximate solution operator  $\mathcal{G}_\theta \approx \mathcal{G}$  from offline data. We then introduce two methods of iteratively applying  $\mathcal{G}$  to observe (6) in an open-loop manner,  $\hat{\mathcal{S}}^{ol}$  and  $\hat{\mathcal{S}}^{ol-r}$ . To close the loop, we finally design an algorithm,  $\hat{\mathcal{S}}^{cl}$ , that integrates sparse online measurements with recursive predictions using  $\mathcal{G}$ . The research problem is then addressed by using  $\hat{\mathcal{S}}^{cl}$  together with  $\mathcal{G}_\theta$ , resulting in a data-driven, closed-loop observer.

#### A. Solution Operator: Model Identification

Given the offline scenarios  $\{\mathcal{S}[P_0^i]\}_{i \in \mathcal{I}}$ , the goal is to learn a representation of the solution operator

$$\mathcal{G}_k : \mathcal{H}_\Omega^n \rightarrow \mathcal{H}_\Omega^1, \\ P_0 \mapsto \rho_{k\Delta t},$$

which maps the input state  $P_0$  to the solution function at  $k$  time steps later,  $\rho_{k\Delta t}$ . While both  $P_0$  and  $\rho_{k\Delta t}$  are infinite-dimensional functions, we assume that the available data is represented in finite-dimensional form, obtained by evaluating them on a uniformly discretized grid  $\mathcal{X} \subset \Omega$ . This yields  $P_0|_{\mathcal{X}} \in \mathbb{R}^{n \times |\mathcal{X}|}$  and  $\rho_{k\Delta t}|_{\mathcal{X}} \in \mathbb{R}^{|\mathcal{X}|}$ . To bridge the gap between finite-dimensional data and function space learning, the approximate operator  $\mathcal{G}_\theta$  is modeled by an FNO (2)–(4).




Fig. 2: Open-loop prediction with solution operator  $\mathcal{G}_\theta$ . It predicts  $\hat{\rho}_{\Delta t}, \hat{\rho}_{2\Delta t}, \dots, \hat{\rho}_{n_{out}\Delta t}$  given an initial state  $P_0 = [\rho_0, \rho_{-\Delta t}, \dots, \rho_{-(n-1)\Delta t}]$ . Here,  $n = 10$  and  $n_{out} = 100$ .

With a target offset  $n_{out}$  for the observer predictions,  $\mathcal{G}_\theta$  is trained to predict the density over the entire horizon,

$$\mathcal{G}_\theta : \mathcal{H}_\Omega^n \rightarrow \mathcal{H}_\Omega^{n_{out}}, \\ P_0 \mapsto [\hat{\rho}_{\Delta t}, \hat{\rho}_{2\Delta t}, \dots, \hat{\rho}_{n_{out}\Delta t}]. \quad (7)$$

Including the horizon up to the target estimate  $\hat{\rho}_{n_{out}\Delta t}$  was empirically found to improve performance, and it reflects the setup in [25]. Thus, the corresponding data-driven learning objective<sup>2</sup> is given by

$$\min_{\theta} \mathbb{E}_{a \sim \mu} \left[ \sum_{k=1}^{n_{out}} \int_{\Omega} \|\mathcal{G}_k[a](x) - \mathcal{G}_\theta^{(k)}[a](x)\|^2 dx \right]. \quad (8)$$

Fig. 2 illustrates an example forward pass of a learned operator  $\mathcal{G}_\theta$  for  $n_{out} = 100$ . We refer to this as open-loop prediction in the observer framework.

#### B. Open-Loop Observers

Given a representation of  $\mathcal{G}$ , the goal is now to develop the observer algorithm. We begin by considering standard autoregression, wherein past estimates serve as input to succeeding predictions through the *predicted state estimate*

$$\hat{P}_t = [\hat{\rho}(\cdot, t), \hat{\rho}(\cdot, t - \Delta t), \dots] \in \mathcal{H}_\Omega^n.$$

Thus, triggering an autoregressive rollout only requires an estimate of the initial state  $\hat{P}_0$ . However, this is only partially observed in our problem, since measurements are restricted to the sparse sensor locations  $\mathcal{M} \subsetneq \Omega$ . To complete the initial state estimate, we interpolate between the sparse measurements  $\mathbf{y}_{\mathcal{M}}(t)$  with GP regression,

$$\hat{P}_0 \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(0), \mathbf{y}_{\mathcal{M}}(-\Delta t), \dots]} \in \mathcal{H}_\Omega^n.$$

Thereafter, the estimates  $\hat{\rho}(\cdot, t)$  are generated by shifting the autoregressive prediction window in increments of  $\Delta t$ .

We found that using multi-step ahead prediction  $\mathcal{G}_{n_d+1}$  with a delayed state estimate  $\hat{P}_{t-n_d\Delta t}$  offset by  $n_d > n - 1$

<sup>2</sup>While we focus on a data-driven objective in this paper, we highlight that it is possible to regularize using lower-order models such as (1) with a predefined or jointly learned flow function  $Q$  following [16], [24].

improved stability. This creates the *open-loop observer*  $\hat{S}^{ol}$ :

$$\hat{S}^{ol}[\mathcal{G}, \mathbf{y}_{\mathcal{M}}] : \begin{cases} \hat{\rho}(\cdot, t + \Delta t) = \mathcal{G}_{n_d+1}[\hat{\mathbf{P}}_{t-n_d\Delta t}], \\ \quad \quad \quad t = 0, \Delta t, \dots \\ \hat{\mathbf{P}}_{t-n_d\Delta t} = [\hat{\rho}(\cdot, t - n_d\Delta t), \dots], \\ \hat{\mathbf{P}}_{-n_d\Delta t} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(-n_d\Delta t), \dots]}. \end{cases} \quad (9)$$

The gap between the initial state and the first estimate,  $\hat{\rho}(\cdot, t)$ ,  $t \in [-n_d\Delta t + \Delta t, \dots, 0]$ , is filled by interpolating the corresponding measurements with GP regression. We illustrate the observer in Fig. 3.




Fig. 3: Open-loop observer  $\hat{S}^{ol}$

While the purely autoregressive approach benefits from consistent predictions, it is sensitive to disturbance in the initial state. As reported in [26], autoregressive rollouts of neural operators suffer from unbounded out-of-distribution (OOD) error growth for long-term forecasts. Since we have access to online measurements, we compare  $\hat{S}^{ol}$  with another open-loop observer that resets the input state for every new prediction based on recent measurements. Specifically, the *data-based state estimate*  $\hat{\mathbf{P}}_t^{\mathcal{D}}$  is obtained by interpolating the corresponding measurements  $[\mathbf{y}_{\mathcal{M}}(t), \mathbf{y}_{\mathcal{M}}(t - \Delta t), \dots]$  using GP regression,

$$\hat{\mathbf{P}}_t^{\mathcal{D}} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(t), \mathbf{y}_{\mathcal{M}}(t - \Delta t), \dots]} \in \mathcal{H}_{\Omega}^n. \quad (10)$$

For comparability with (9), we consider a delayed state estimate  $\hat{\mathbf{P}}_{t-n_d\Delta t}^{\mathcal{D}}$ . The resulting dynamical system is denoted *open-loop observer with reset*  $\hat{S}^{ol-r}$ , and defined as

$$\hat{S}^{ol-r}[\mathcal{G}, \mathbf{y}_{\mathcal{M}}] : \begin{cases} \hat{\rho}(\cdot, t + \Delta t) = \mathcal{G}_{n_d+1}[\hat{\mathbf{P}}_{t-n_d\Delta t}^{\mathcal{D}}], \\ \quad \quad \quad t = 0, \Delta t, \dots \\ \hat{\mathbf{P}}_{t-n_d\Delta t}^{\mathcal{D}} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(t-n_d\Delta t), \dots]}, \\ \hat{\mathbf{P}}_{-n_d\Delta t}^{\mathcal{D}} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(-n_d\Delta t), \dots]}. \end{cases} \quad (11)$$

The observer is illustrated in Fig. 4.




Fig. 4: Open-loop observer with reset  $\hat{S}^{ol-r}$

A limitation of  $\hat{S}^{ol-r}$  is that each prediction relies on a few data points. By the property of the disconnected predictions, the observer does not leverage the full historical dataset.

### C. Closed-Loop Observer

Combining the consistent predictions of  $\hat{S}^{ol}$  with the online updates of  $\hat{S}^{ol-r}$  leads to our main contribution: a *closed-loop observer*  $\hat{S}^{cl}$  that integrates an autoregressive rollout of  $\mathcal{G}$  with online measurements using a correction operator  $\mathcal{N}$ . Analogous to a Kalman filter,  $\mathcal{N}$  adjusts the predicted state estimate using recent measurements before it is passed as input to the next prediction. Measurements are introduced via a Luenberger-type error operator  $\mathcal{E}$ , which computes the error estimate  $\hat{e}$  as the difference between the predicted state estimate  $\hat{\mathbf{P}}$  and the data-based estimate  $\hat{\mathbf{P}}^{\mathcal{D}}$ . The correction operator  $\mathcal{N}$  then processes  $\hat{\mathbf{P}}$  and  $\hat{e}$  to generate an *updated state estimate*  $\hat{\mathbf{P}}^u \in \mathcal{H}_{\Omega}^n$ , which is passed as input to the next prediction. We summarize  $\hat{S}^{cl}$  as

$$\hat{S}^{cl}[\mathcal{G}, \mathcal{N}, \mathbf{y}_{\mathcal{M}}] : \begin{cases} \hat{\rho}(\cdot, t + \Delta t) = \mathcal{G}_{n_d+1}[\hat{\mathbf{P}}_{t-n_d\Delta t}^u], \\ \quad \quad \quad t = 0, \Delta t, \dots \\ \hat{\mathbf{P}}_{t-n_d\Delta t}^u = \mathcal{N}[\hat{\mathbf{P}}, \hat{e}], \\ \hat{e} = \mathcal{E}[\hat{\mathbf{P}}, \hat{\mathbf{P}}^{\mathcal{D}}], \\ \hat{\mathbf{P}}_{t-n_d\Delta t} = [\hat{\rho}(\cdot, t - n_d\Delta t), \dots], \\ \hat{\mathbf{P}}_{t-n_d\Delta t}^{\mathcal{D}} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(t-n_d\Delta t), \dots]}, \\ \hat{\mathbf{P}}_{-n_d\Delta t}^u \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(-n_d\Delta t), \dots]}, \end{cases} \quad (12)$$

and illustrate the system in Fig. 5.

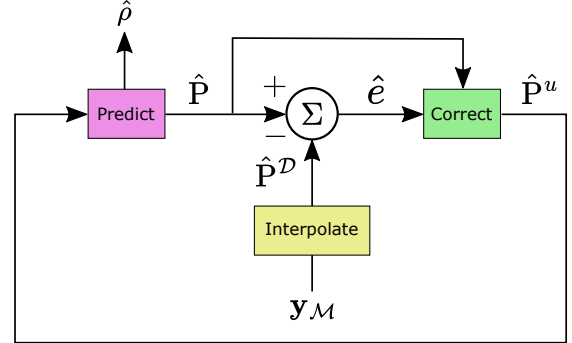


Fig. 5: Closed-loop observer  $\hat{S}^{cl}$

In implementing (12), corrections are applied to a moving window of historical predictions,

$$\begin{aligned} \hat{\mathbf{P}}_{t-n_d\Delta t}^+ &= [\hat{\rho}(\cdot, t - (n-1)\Delta t), \dots, \\ &\quad \hat{\rho}(\cdot, t - n_d\Delta t), \dots, \\ &\quad \hat{\rho}(\cdot, t - n_d\Delta t - (n-1)\Delta t)], \end{aligned} \quad (13)$$

assuming  $n_d > n - 1$  as in the previous section. Thus, the *extended* state estimate  $\hat{\mathbf{P}}^+$  contains  $\hat{\mathbf{P}}$  and extends forward in time such that the window length matches the prediction horizon (7) where  $n_{out} = n_d + 1$  similar to (9) and (11). After applying the correction to this window, the updated state estimate  $\hat{\mathbf{P}}^u$  is taken as the subset  $\hat{\mathbf{P}}_{-n:-1}^{u+}$ . The procedure is illustrated in Fig. 6.

With  $\hat{\mathbf{P}}^+$  and the corresponding extended data-based estimate  $\hat{\mathbf{P}}^{\mathcal{D}+}$ , the Luenberger error operator  $\mathcal{E}$  is defined as

$$\mathcal{E}[\hat{\mathbf{P}}^+, \hat{\mathbf{P}}^{\mathcal{D}+}] \triangleq \hat{\mathbf{P}}^+ - \hat{\mathbf{P}}^{\mathcal{D}+} \in \mathcal{H}_{\Omega}^{n_d+1}.$$

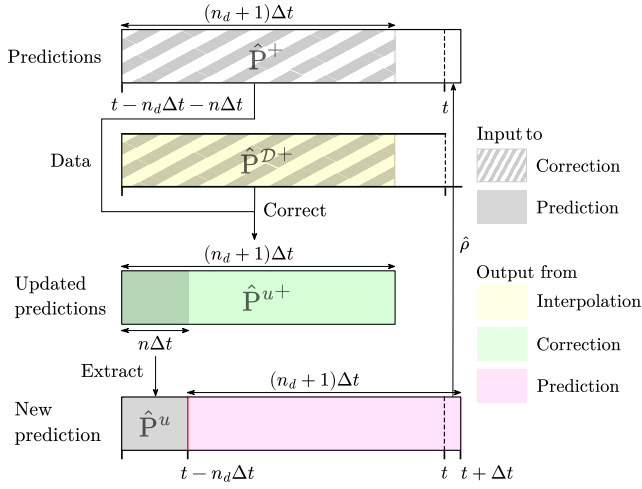


Fig. 6: Implementation of  $\hat{S}^{cl}$  with extended state estimates

The correction operator  $\mathcal{N}$  processes  $\hat{P}^+$  and the error  $\hat{e}^+ = \mathcal{E}[\hat{P}^+, \hat{P}^{\mathcal{D}+}]$  to generate the updated state estimate  $\hat{P}^{u+}$ ,

$$\mathcal{N} : \mathcal{H}_{\Omega}^{n_d+1} \times \mathcal{H}_{\Omega}^{n_d+1} \rightarrow \mathcal{H}_{\Omega}^{n_d+1}, \quad (14)$$

$$\hat{P}^+, \hat{e}^+ \mapsto \hat{P}^{u+}.$$

The correction update window  $\hat{P}_{t-n_d\Delta t}^{u+}$  is shifted in increments of  $\Delta t$  in parallel with the prediction input window  $\hat{P}_{t-n_d\Delta t}$ , cf. (9). Since  $\hat{P}_{t-n_d\Delta t}^u \subset \hat{P}_{t-n_d\Delta t}^{u+}$ , the state estimates are updated before being passed as input to the next prediction. The implementation details are outlined in Algorithm 1.

---

**Algorithm 1:** Closed-Loop Observer  $\hat{S}^{cl}$

---

**input :**  $\mathcal{G}, \mathcal{N}, \mathbf{y}_{\mathcal{M}}$   
**output:**  $\hat{\rho}$   
 $t = 0;$   
 $\hat{\rho}(\cdot, t) \sim \mathcal{GP}_{\mathbf{y}_{\mathcal{M}}(t)}, \quad -n_d\Delta t < t \leq 0, \quad x \in \mathcal{X};$   
 $\hat{P}_{-n_d\Delta t}^u \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(-n_d\Delta t), \dots]}, \quad x \in \mathcal{X};$   
**repeat**  
  // Predict  
   $\hat{\rho}(\cdot, t + \Delta t) \leftarrow \mathcal{G}_{n_d+1}[\hat{P}_{t-n_d\Delta t}^u], \quad x \in \mathcal{X};$   
   $\hat{P}^+ \leftarrow [\hat{\rho}(\cdot, t - (n-1)\Delta t + \Delta t), \dots];$   
  // Collect measurements  
   $\hat{P}^{\mathcal{D}+} \sim \mathcal{GP}_{[\mathbf{y}_{\mathcal{M}}(t-(n-1)\Delta t + \Delta t), \dots]};$   
  // Correct  
   $\hat{e}^+ = \mathcal{E}[\hat{P}^+, \hat{P}^{\mathcal{D}+}];$   
   $\hat{P}^{u+} \leftarrow \mathcal{N}[\hat{P}^+, \hat{e}^+];$   
  // Extract next input  
   $\hat{P}_{t-n_d\Delta t + \Delta t}^u \leftarrow \hat{P}_{-n:-1}^{u+};$   
   $t \leftarrow t + \Delta t;$   
**until end;**  
**return**  $\hat{\rho}$

---

The correction operator  $\mathcal{N}$  is approximated using an FNO defined on a two-dimensional domain:

$$\mathcal{N}_{\psi} : \mathcal{H}_{\Omega \times \Gamma}^1 \times \mathcal{H}_{\Omega \times \Gamma}^1 \rightarrow \mathcal{H}_{\Omega \times \Gamma}^1, \quad (15)$$

where  $\Gamma = [\Delta t, (n_d + 1)\Delta t]$  and the processed functions are translated in time to take values in this domain. Note that  $\Gamma$  reflects the choice of window size in (13). This setting enables an extension of the learning objective for open-loop prediction (8), such that  $\mathcal{N}_{\psi}$  is obtained as the minimizer of

$$\min_{\psi} \mathbb{E}_{a \sim \mu} \left[ \sum_{k=1}^{n_d+1} \int_{\Omega} \|J_k[a](\psi; x)\|^2 dx \right], \quad (16)$$

$$J_k[a](\psi; \cdot) \triangleq \mathcal{G}_k[a](\cdot) - \mathcal{N}_{\psi}[\mathcal{G}_{\theta}[a], \mathcal{E}[\mathcal{G}_{\theta}[a], \rho_{\mathcal{D}}[a]]](\cdot, k\Delta t),$$

where  $\rho_{\mathcal{D}}[a]$  is a data-based estimate of the true density  $\mathcal{G}_{1:n_d+1}[a]$  using samples at the sensor locations,  $\{\mathcal{G}_{1:n_d+1}[a](\cdot)\}_{x \in \mathcal{M}}$ . By comparing the objective function of  $\mathcal{N}_{\psi}$  in (16) with that of the open-loop predictor  $\mathcal{G}_{\theta}$  in (8), we observe that  $\mathcal{N}_{\psi}$  is trained to refine the prediction of  $\mathcal{G}_{\theta}$ . This refinement is obtained by taking in measurements in a Luenberger manner through  $\mathcal{E}$ .

#### IV. NUMERICAL SIMULATIONS AND DISCUSSION

Numerical simulations are conducted using the open-source microscopic traffic simulator SUMO [27]. This framework integrates road network data, infrastructure like traffic lights, and stochastic demand to produce realistic vehicle flows. Following the methodological section, we focus on ring road dynamics [28]: vehicles are gradually introduced to the road during an initialization phase to reach the initial state, after which the higher-order microscopic dynamics are run in a loop. To convert the discrete distribution of vehicles to a continuous density function, the grid of vehicle-populated cells is convolved with a Gaussian filter [29].

The remainder of this section is organized as follows: we first outline implementation details, then demonstrate the observers on a single test case, evaluate their performance in noiseless, noisy, and out-of-distribution (OOD) scenarios, and finally analyze input-to-state stability.

##### A. Implementation Details

The SUMO simulations are generated on a circular road of length  $L = 6.2$  km, discretized into 123 cells  $\mathcal{X}$  with six equidistant sensors  $\mathcal{M}$ . The simulations are run for a total time of  $T = 40$  minutes divided into time steps of length  $\Delta t = 1$  s, following [28]. The training dataset contains 20 noiseless, independent simulations per average  $\rho \in [0.1, 0.2, \dots, 0.8]$  split into a total of 3360 non-overlapping input-output pairs  $(P_0^j, [\rho_{k\Delta t}^j]_{k=1}^{n_d+1})$ . We found that  $n = 10$  and  $n_d + 1 = 100$  works well in this setting.

The test dataset consists of 10 simulations per average  $\rho \in [0.3, \dots, 0.8]$ , focusing on the settings where congestion occurs. In addition, we test the observers on a noisy test set and a noiseless OOD test set. The noisy data is produced by adding i.i.d. Gaussian noise with standard deviation  $\sigma = 0.1$  to the measurements. The OOD data is generated by varying parameters in SUMO to induce more traffic jams.

The approximate solution operator  $\mathcal{G}_{\theta}$  and correction operator  $\mathcal{N}_{\psi}$  share a similar FNO structure: the lifting function

is a linear layer  $P$  of output dimension 16, and the projection function  $Q$  is a fully-connected neural network  $Q$  with one hidden layer of size 128. In between,  $\mathcal{G}_\theta$  has four spectral convolution layers  $\mathcal{L}_l$  where  $W_l$  are one-dimensional convolution layers of sizes [24, 24, 32, 32] and the Fourier expansions are truncated to  $k_{\max}$  modes [15, 12, 9, 9]. The correction operator  $\mathcal{N}_\psi$  has two spectral convolution layers  $\mathcal{L}_l$  where  $W_l$  are two-dimensional convolution layers of sizes [24, 32] and the Fourier expansions are truncated to  $k_{\max}$  modes [15, 9]. For both operators, we use GELU activation functions and warp the output with a sigmoid to impose  $\hat{\rho} \in [0, 1]$ . GP regression is implemented using a prior with a zero mean function and a squared exponential kernel with a length scale of 1, and we sample once during training while taking the posterior mean function during inference.

We use the Adam optimizer and train for 500 epochs. The training and testing were run on a laptop with an Intel(R) Core(TM) i7-1370p CPU @ 1.90 GHz and 14 processing cores. The code and experiments are available on GitHub for reproduction<sup>3</sup>.

### B. Numerical Simulations and Discussion

A comparison between the open-loop observer  $\hat{S}^{ol}$ , open-loop observer with reset  $\hat{S}^{ol-r}$ , and closed-loop observer  $\hat{S}^{cl}$  applied to a single test example is shown in Fig. 7. We observe that  $\hat{S}^{ol}$  quickly diverges, while  $\hat{S}^{ol-r}$  produces disconnected predictions, with the vertical dispersion of the measurements resulting from the vertical GP interpolation between observation points. In contrast,  $\hat{S}^{cl}$  maintains stability while producing consistent predictions over the spatiotemporal domain. This highlights the benefits of  $\hat{S}^{cl}$ : the continuous calibration of the autoregressive inputs both improves the prediction estimates and counteracts instability. Hence,  $\hat{S}^{cl}$  unites the consistent predictions of  $\hat{S}^{ol}$  with the online updates of  $\hat{S}^{ol-r}$ , achieving an improved observer. We continue by exploring the performance across the test set in terms of accuracy, robustness, and input-to-state stability.

The prediction accuracy of the observers across the test set in the noiseless, noisy, and OOD settings is shown in Fig. 8. In the noiseless setting, the closed-loop observer  $\hat{S}^{cl}$  generally performs better than the open-loop observer  $\hat{S}^{ol}$  and the open-loop observer with reset  $\hat{S}^{ol-r}$ . The poor performance of  $\hat{S}^{ol}$  is explained by the instability caused by disturbance in the initial state. The improved performance of  $\hat{S}^{cl}$  over  $\hat{S}^{ol-r}$  can be explained by the fact that  $\hat{S}^{cl}$  benefits from the full historical dataset, as previous corrections influence all subsequent predictions through the continuous chain of autoregression.

Moreover,  $\hat{S}^{cl}$  is generally robust to noise, as the median error remains around the same level when the observer is applied to noisy data. This is not the case of  $\hat{S}^{ol-r}$ , which sees a significant downgrade in performance. This can be explained by the high dependence on a few measurements in  $\hat{S}^{ol-r}$ , rendering it vulnerable to noise. Instead,  $\hat{S}^{cl}$  weighs

<sup>3</sup><https://github.com/aharting/closed-loop-neural-operator-based-observer-of-traffic-density>




Fig. 7: Traffic density estimation with observer variants. The true density is shown in the top left figure, where measurement locations are indicated in black. The density estimates are shown for the open-loop observer (top right), the open-loop observer with reset (bottom left), and the closed-loop observer (bottom right). Note that the solution periodicity is a result of the ring road implementation.

the recent measurements against the prediction suggested by past measurements through the correction operator, resulting in a more accurate estimate. Furthermore, the accuracy of  $\hat{S}^{cl}$  generally remains the same in the OOD dataset, suggesting that it is robust to distribution shifts. Overall, our simulations indicate that  $\hat{S}^{cl}$  is both more accurate and more resilient to noise compared to  $\hat{S}^{ol}$  and  $\hat{S}^{ol-r}$ , as well as robust to distributional changes.




Fig. 8: Robustness of the observers under noise and disturbance.

Finally, the input-to-state stability for the observers is analyzed. In Fig. 9, the evolution of the prediction accuracy is compared across the observers as they are rolled out. The error grows unboundedly for the open-loop observer  $\hat{S}^{ol}$ , while it remains stable for the open-loop observer with reset  $\hat{S}^{ol-r}$ , and gradually decreases for the closed-loop observer  $\hat{S}^{cl}$ . By construction,  $\hat{S}^{ol}$  is completely dependent

on an imperfect estimate of the initial state, which gives rise to the unbounded growth of the error. Meanwhile,  $\hat{S}^{ol-r}$  performs the same operation iteratively, so the error naturally remains stable. In contrast,  $\hat{S}^{cl}$  cumulatively integrates data, resulting in an increasingly informed prediction manifested by a decreasing error. Our simulations thus indicate that  $\hat{S}^{cl}$  exhibits the classical closed-loop property of ultimate-boundedness of the error.




Fig. 9: Evolution of prediction accuracy over the observation horizon.

## V. CONCLUSION

In this work, we introduced a data-driven, closed-loop observer for traffic density estimation using sparse stationary sensor measurements. Our method trains a Fourier neural operator to learn macroscopic traffic flow dynamics from high-fidelity data and integrates it with a correction operator to update predictions with real-time data. Numerical simulations with SUMO show that the closed-loop observer outperforms open-loop alternatives in terms of accuracy, robustness to noise, and stability over time, and indicate that the closed-loop observer is robust to distribution shifts.

This study highlights the potential of this kind of learning-based, closed-loop observer and opens several interesting research directions. Future work includes leveraging real traffic data – possibly with varying resolution – to learn the dynamics, analyzing sensor placement for observability, and conducting formal analyses of convergence and stability.

## REFERENCES

- [1] A. Ferrara, S. Sacone, and S. Siri, *Freeway Traffic Modelling and Control*. Springer International Publishing, 2018.
- [2] T. Seo, A. M. Bayen, T. Kusakabe, and Y. Asakura, “Traffic state estimation on highway: A comprehensive survey,” *Annual Reviews in Control*, 2017.
- [3] J. Lebacque, “The Godunov scheme and what it means for first order traffic flow models,” in *Proceedings of the 13th International Symposium on Transportation and Traffic Theory, Lyon, France*, 1996.
- [4] M. Barreau, A. Selivanov, and K. H. Johansson, “Dynamic traffic reconstruction using probe vehicles,” in *2020 59th IEEE Conference on Decision and Control (CDC)*, 2020.
- [5] L. Muñoz, X. Sun, R. Horowitz, and L. Alvarez, “Traffic density estimation with the cell transmission model,” *Proceedings of the American Control Conference*, 2003.
- [6] X. Sun, L. Muñoz, and R. Horowitz, “Highway traffic state estimation using improved mixture Kalman filters for effective ramp metering control,” in *42nd IEEE International Conference on Decision and Control*, 2003.
- [7] L. Mihaylova and R. Boel, “A particle filter for freeway traffic estimation,” in *2004 43rd IEEE Conference on Decision and Control (CDC)*, 2004.

- [8] N. Bekiaris-Liberis, C. Roncoli, and M. Papageorgiou, “Highway traffic state estimation with mixed connected and conventional vehicles,” *IEEE Transactions on Intelligent Transportation Systems*, 2016.
- [9] —, “Highway traffic state estimation per lane in the presence of connected vehicles,” *Transportation research part B: methodological*, 2017.
- [10] Y. Lv, Y. Duan, W. Kang, Z. Li, and F.-Y. Wang, “Traffic flow prediction with big data: A deep learning approach,” *IEEE Transactions on Intelligent Transportation Systems*, 2015.
- [11] X. Yin, G. Wu, J. Wei, Y. Shen, H. Qi, and B. Yin, “Deep learning on traffic prediction: Methods, analysis, and future directions,” *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [12] M. Barreau, M. Aguiar, J. Liu, and K. H. Johansson, “Physics-informed learning for identification and state reconstruction of traffic density,” in *60th IEEE Conference on Decision and Control*, 2021.
- [13] 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,” *Transportation Research Part B: Methodological*, 2021.
- [14] M. U. B. Niazi, J. Cao, M. Barreau, and K. H. Johansson, “KKL observer synthesis for nonlinear systems via physics-informed learning,” *arXiv preprint arXiv:2501.11655*, 2025.
- [15] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, “Fourier neural operator for parametric partial differential equations,” *arXiv preprint arXiv:2010.08895*, 2020.
- [16] B. T. Thodi, S. V. R. Ambadipudi, and S. E. Jabari, “Fourier neural operator for learning solutions to macroscopic traffic flow models: Application to the forward and inverse problems,” *Transportation research part C: emerging technologies*, 2024.
- [17] M. J. Lighthill and G. B. Whitham, “On kinematic waves ii. a theory of traffic flow on long crowded roads,” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 1955.
- [18] C. M. Dafermos, *Hyperbolic Conservation Laws in Continuum Physics*. Springer Berlin Heidelberg, 2000.
- [19] S. Krauß, “Microscopic modeling of traffic flow: Investigation of collision free vehicle dynamics,” *Hauptabteilung Mobilität und Systemtechnik des DLR Köln*, 1998.
- [20] M. Treiber, A. Kesting, and D. Helbing, “Delays, inaccuracies and anticipation in microscopic traffic models,” *Physica A: Statistical Mechanics and its Applications*, 2006.
- [21] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, “Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam,” *New journal of physics*, 2008.
- [22] S.-i. Tadaki, M. Kikuchi, M. Fukui, A. Nakayama, K. Nishinari, A. Shibata, Y. Sugiyama, T. Yosida, and S. Yukawa, “Phase transition in traffic jam experiment on a circuit,” *New Journal of Physics*, 2013.
- [23] N. Kovachki, S. Lanthaler, and S. Mishra, “On universal approximation and error bounds for Fourier neural operators,” *Journal of Machine Learning Research*, 2021.
- [24] Z. Li, H. Zheng, N. Kovachki, D. Jin, H. Chen, B. Liu, K. Azizzadenesheli, and A. Anandkumar, “Physics-informed neural operator for learning partial differential equations,” *ACM/IMS Journal of Data Science*, 2024.
- [25] V. Gopakumar, A. Gray, J. Oskarsson, L. Zanisi, S. Pamela, D. Giles, M. Kusner, and M. P. Deisenroth, “Uncertainty quantification of surrogate models using conformal prediction,” *arXiv preprint arXiv:2408.09881*, 2024.
- [26] M. McCabe, P. Harrington, S. Subramanian, and J. Brown, “Towards stability of autoregressive neural operators,” *Transactions on Machine Learning Research*, 2023.
- [27] P. A. Lopez *et al.*, “Microscopic traffic simulation using sumo,” in *2018 21st international conference on intelligent transportation systems (ITSC)*, 2018.
- [28] M. Barreau, “SUMOsimulator,” <https://github.com/mBarreau/SUMOsimulator>, 2025.
- [29] R. Scheepens *et al.*, “Composite density maps for multivariate trajectories,” *IEEE Transactions on Visualization and Computer Graphics*, 2011.