Closed-Loop Neural Operator-Based Observer of Traffic Density

Alice Harting[*], Karl H. Johansson, Matthieu Barreau
Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden

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 microscopic-level simulations. 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 exhibit 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.

Introduction

Freeway congestion is a major issue in metropolitan areas, and a core component of addressing this is freeway traffic control [1]. A key state variable is the traffic density, which can be estimated from sparse data using stationary sensors. Traditional model-based methods such as Kalman filters may impose overly restrictive assumptions on the system. Purely data-driven methods, on the other hand, are often difficult to interpret and lightweight online prediction remains a challenge.

Thus, we introduce a data-driven, closed-loop neural observer of traffic density flow from sparse measurements. Our contributions are leveraging Fourier Neural Operators (FNOs) [2] 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. We train and test on ring-road data from the SUMO simulations.

Problem Statement

We model traffic density evolution as an n^{th}-order discrete-time dynamical system

$$\mathcal{S}[\bar{\zeta}_0] : \begin{cases} \rho(\cdot, t+\Delta t) = \mathcal{G}[\zeta_t], \quad t > 0, \\ \zeta_t = [\rho(\cdot, t), \rho(\cdot, t-\Delta t), \dots], \\ \zeta_0 = \bar{\zeta}_0, \end{cases} \tag{1}$$

where $\rho(\cdot, t)$ is the true density, ζ_t is the state at time t, $\bar{\zeta}_0$ is the initial state, and $\mathcal{G}$ represents the solution operator of the unknown dynamics.

[*]aharting@kth.se

The objective of this paper is to develop an algorithm that observes (1) in closed-loop by generating real-time predictions $\hat{\rho}$ of the system state ρ every Δt 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 ϵ at stationary, sparse sensor locations $x \in \mathcal{M}$, collectively denoted $\mathbf{y}_{\mathcal{M}}(t)$. We also assume access to offline data from physically identical systems, $\{\mathcal{S}[\zeta_0^i]\}_{i\in\mathcal{I}}$.

Methodology

To build a closed-loop observer of (1), the first step is to learn an approximate solution operator $\mathcal{G}_\theta \approx \mathcal{G}$ from offline data. We then compare two methods of iteratively applying $\mathcal{G}$ to observe (1) in an open-loop manner: $\hat{\mathcal{S}}^{ol}$, which runs autoregressively based on the estimated state $\hat{\zeta}_0$, and $\hat{\mathcal{S}}^{ol\text{-}r}$, which resets $\hat{\zeta}_t$ every time step based on recent measurements. To close the loop, we finally design an algorithm, $\hat{\mathcal{S}}^{cl}$, that integrates sparse online measurements with recursive predictions using $\mathcal{G}$ (Fig. 1a). The research problem is then addressed by using $\hat{\mathcal{S}}^{cl}$ with $\mathcal{G}_\theta$, resulting in a data-driven, closed-loop observer.

Results

A comparison between the observers applied to a single test example is shown in Fig. 1b. 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 observers $\hat{\mathcal{S}}^{ol}$ (top right) and $\hat{\mathcal{S}}^{ol\text{-}r}$ (bottom left), and the closed-loop observer $\hat{\mathcal{S}}^{cl}$ (bottom right). We observe that $\hat{\mathcal{S}}^{ol}$ quickly diverges, while $\hat{\mathcal{S}}^{ol\text{-}r}$ generates disconnected, volatile predictions. In contrast, $\hat{\mathcal{S}}^{cl}$ maintains stability while producing consistent predictions over the spatiotemporal domain.

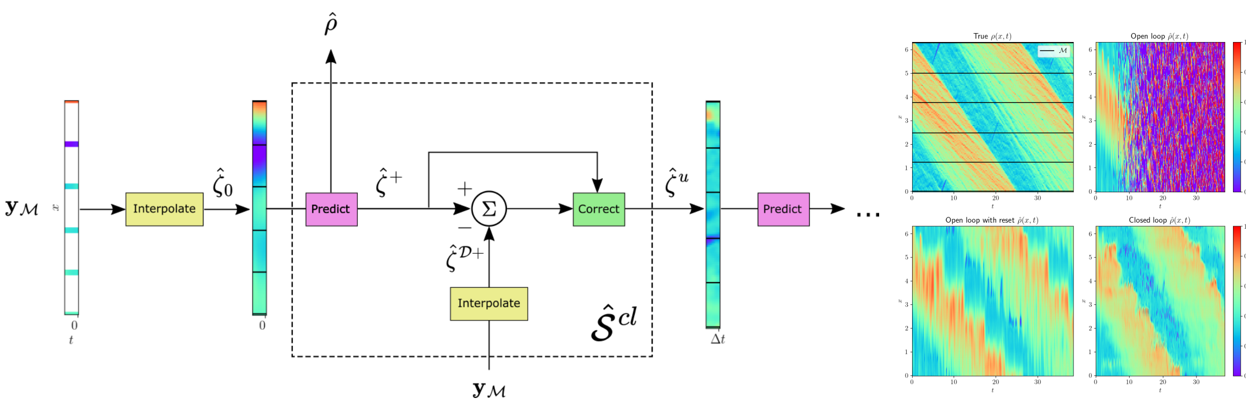

(a) Closed-loop observation with $\hat{\mathcal{S}}^{cl}$ (b) Observer variants

Figure 1: Traffic density estimation with learning-based observers

References

[1] A. Ferrara, S. Sacone, and S. Siri, *Freeway Traffic Modelling and Control*. Springer International Publishing, 2018.

[2] 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.