

Reinforcement Learning Versus PDE Backstepping and PI Control for Congested Freeway Traffic

Huan Yu^{1b}, Saehong Park^{1b}, Alexandre Bayen^{1b}, *Senior Member, IEEE*,
Scott Moura^{1b}, and Miroslav Krstic^{1b}, *Fellow, IEEE*

Abstract—We develop reinforcement learning (RL) boundary controllers to mitigate stop-and-go traffic congestion on a freeway segment. The traffic dynamics of the freeway segment are governed by a macroscopic Aw–Rascle–Zhang (ARZ) model, consisting of 2×2 quasi-linear partial differential equations (PDEs) for traffic density and velocity. The boundary stabilization of the linearized ARZ PDE model has been solved by PDE backstepping, guaranteeing spatial L^2 norm regulation of the traffic state to uniform density and velocity and ensuring that traffic oscillations are suppressed. Collocated proportional (P) and proportional–integral (PI) controllers also provide stability guarantees for allowable control gains and are always applicable as model-free control options through gain tuning by trial and error, or by model-free optimization. Although these approaches are mathematically elegant, the stabilization result only holds locally and is usually affected by the change of model parameters. Therefore, we reformulate the PDE boundary control problem as an RL problem that pursues stabilization without knowing the system dynamics, simply by observing the state values. The proximal policy optimization (PPO), a neural network-based policy gradient algorithm, is employed to obtain RL controllers by interacting with a numerical simulator of the ARZ PDE. Being stabilization-inspired, the RL state-feedback boundary controllers are compared and evaluated against the rigorously stabilizing controllers in two cases: 1) in a system with perfect knowledge of the traffic flow dynamics and then 2) in one with only partial knowledge. We obtain RL controllers that nearly recover the performance of the backstepping, P, and PI controllers with perfect knowledge and outperform them in some cases with partial knowledge. It must be noted, however, that the RL controllers are obtained by conducting about one thousand episodes of iterative training on a simulation model. This training cannot be performed in a collision-free fashion in real traffic, nor convergence guaranteed when training. Thus, we demonstrate that the RL approach has learning (i.e., adaptation) potential for traffic PDE systems under uncertain and changing conditions, but RL is neither simple nor a fully safe substitute for model-based control in real traffic systems.

Manuscript received January 1, 2021; revised June 25, 2021; accepted September 20, 2021. Manuscript received in final form September 27, 2021. Recommended by Associate Editor A. Macchelli. (Huan Yu and Saehong Park contributed equally to this work.) (Corresponding author: Huan Yu.)

Huan Yu is with the Systems Hub, Intelligent Transportation Thrust, The Hong Kong University of Science and Technology, Hong Kong (e-mail: huanyu@ust.hk).

Saehong Park, Alexandre Bayen, and Scott Moura are with the Department of Civil and Environmental Engineering, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: sspark@berkeley; bayen@berkeley.edu; smoura@berkeley.edu).

Miroslav Krstic is with the Department of Mechanical and Aerospace Engineering, University of California at San Diego, La Jolla, CA 92093 USA (e-mail: krstic@ucsd.edu).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCST.2021.3116796>.

Digital Object Identifier 10.1109/TCST.2021.3116796

Index Terms—Feedback, Lyapunov methods, machine learning, traffic control.

I. INTRODUCTION

TRAFFIC congestion has been an inescapable problem in large and growing metropolitan areas across the world. Among various traffic congestion phenomena, stop-and-go traffic [7], [11], [15], [43] is a common variety that appears every day in congested freeways, causing unsafe driving conditions, more fuel consumption, and emissions. Traffic congestion is characterized by the propagation of oscillatory waves on roads and is caused by delayed driver response. Freeway traffic management usually relies on static road infrastructure to regulate traffic flow, such as ramp metering and varying speed limits (VSLs).

Depending on the employed traffic model and control objectives, numerous model-based traffic control strategies have been developed using Lyapunov stability analysis. Macroscopic traffic modeling is particularly well-suited to control design since it describes the overall spatial–temporal dynamics of traffic flow on a freeway segment. The state variables of the model are aggregated traffic values. They are continuous in time and space, easy to be measured and to be regulated in transportation management. The first-order hyperbolic partial differential equation (PDE) model by Lighthill, Whitham, and Richards (LWR) [36], [40] is a conservation law of traffic density, describing the formation and propagation of traffic density waves on the road. The oscillatory, unstable behaviors observed in stop-and-go traffic require a higher-order model rather than the static LWR model. The second-order Aw–Rascle–Zhang (ARZ) model [3], [57] addressed this issue by adding another PDE for the velocity state, leading to a 2×2 nonlinear coupled hyperbolic PDE system. Fan and Seibold [14] compared the data-fit LWR and ARZ model and found out that the second-order ARZ model is significantly more accurate in predicting traffic velocity. Therefore, we examine the state-of-art ARZ model to capture stop-and-go traffic oscillations in this article, and study the boundary control problem.

A. Lyapunov-Based Control Approach and Challenges

The control of traffic PDE models using VSL and ramp metering has been addressed by two categories of research. The first category reduces the infinite-dimensional traffic PDE model to discrete models, and then applies optimal and feedback control approaches. Gomes and Horowitz [22] proposed a cell-transmission model (CTM) derived from the LWR model.

Papageorgiou *et al.* [37] proposed the “METANET” model which is discretized from the second-order traffic model. Feedback control strategies have been proposed for ramp metering controllers such as Asservissement Linéaire d’Entrée Autoroutière (ALINEA), a local feedback ramp-metering strategy in [38] for PI-ALINEA in [50]. Optimal control [10], [22] and model predictive control [6] have also been widely studied for ramp metering control of freeway traffic systems. A more comprehensive survey is recently conducted in [44].

Many recent efforts [5], [7], [26], [52], [53]–[56], [58] focused on developing feedback control strategies directly for the macroscopic traffic PDE models, instead of their approximated versions. The Lyapunov approach was applied for stability analysis. In a nutshell, the Lyapunov-based control approach for traffic PDE models consists of two steps: 1) obtain the state-feedback controllers through various control design methods and 2) use Lyapunov analysis to study the closed-loop system dynamics.

Studies that examine the first-order LWR model include [5], [27], [55], [56], and all investigate the traffic bottleneck problem. Predictor feedback control laws were designed to compensate the delay of PDE, i.e., the linearized LWR model [5], [55]. Extremum seeking, a non-model based real-time adaptive control method was used in [56] to find the unknown optimal density of the bottleneck area. In [27], stabilization of a spatially uniform equilibrium profile for the LWR model was obtained under in-domain VSL control. All of these studies achieved closed-loop stability using Lyapunov analysis.

Studies that examine the second-order ARZ model include [26], [52], [58], and this article. In [26], a nonlinear boundary feedback law was designed that controls the inlet flow and achieves global stabilization for a modified second-order ARZ model. Proportional–integral (PI) controllers were designed in [58] for inlet traffic flow rate control through ramp metering and outlet velocity control through VSL. In [52], a backstepping controller was developed for ramp metering at the outlet, to regulate upstream traffic. A collocated P controller was also proposed for inlet flow control of the downstream traffic. The aforementioned control designs are representatives of boundary control of hyperbolic PDE systems, and guarantee Lyapunov stability of the closed-loop system.

Despite differences between the control design approaches, the key idea behind each method is similar. Namely, each method analyzes the particular structure of the PDE system to design a boundary feedback control law and then uses Lyapunov analysis to obtain closed-loop stability. Zhang *et al.* [58] analyzed PI boundary controllers for the linearized and homogeneous ARZ model and exponential stability of the closed-loop system was obtained using Lyapunov stability analysis. A boundary controller for the linearized inhomogeneous ARZ PDE is developed in [52] and [53]. A linear backstepping state-feedback controller and a static boundary feedback controller were proposed to achieve L^2 norm stabilization of traffic oscillations in finite time.

Although model uncertainty and nonlinearity have been studied for control of discrete traffic systems [17], [25], these challenges remain for the Lyapunov-based controllers for macroscopic traffic PDE models.

- 1) *Nonlinearity*: The backstepping, P and PI boundary control methods discussed above are developed for linearized PDEs. Thus only local stabilization is guaranteed for the nonlinear system. The linearization approach may fail for large perturbations from the linearization point.
- 2) *Model Knowledge*: Both model-based designs (e.g., PDE backstepping) and model-free designs (e.g., PI control) involve gain selection subject to certain theoretically determined bounds, which require some knowledge of the system such as real-time full state measurements as well as model parameters that characterize the overall system dynamics. The parameter values are usually obtained by calibrating the PDE model with traffic field data. This model calibration process can be laborious, and specific to certain traffic conditions. In addition, the deterministic traffic PDE model cannot address random events such as traffic accidents, changes in road attributes, and so on.
- 3) *Adaptive Optimization*: Besides traffic flow stabilization, there are other relevant performance metrics, including fuel consumption, travel time, and so on. When another objective is considered, then one needs to re-formulate and re-design the controller with a Lyapunov-based approach. Different design tools might be required. It is desirable to have an approach with modest tuning that can adapt to various problems and objectives arising in traffic systems.

The challenges mentioned above can be addressed by reinforcement learning (RL). This model-free approach relies on few assumptions about the system, including being agnostic about nonlinearities. From a control law perspective, RL produces a neural network-based feedback controller that maps observations (states) into actions (control inputs) through training the networks. An RL formulation to the traffic control problem has the potential to enable an *expert-free* traffic management system, in the sense that it does not require expert knowledge of the system dynamics or a control-theoretic training. Neural networks are used as parametric approximators for functions of states whose weights are updated automatically to learn the optimal control action. The model complexity and identification process faced by Lyapunov-based controllers are circumvented. RL controllers are trained to maximize the reward function, which in this article is chosen as the L^2 spatial norm of states to achieve stabilization. However, the proposed formulation can be easily adapted for different objectives by changing the reward formulation.

Extremum seeking, a real-time adaptive optimization technique, is an alternative to RL for tuning controller gains in a model-free fashion. The algorithms in [20], [29], and [30] can be applied for offline learning of feedback laws based on simulation models, as with RL and the ARZ model in this article. In contrast, the algorithm in [39] is designed for online tuning of feedback laws. In extremum seeking, a small excitation is used to perturb the feedback gains being tuned and to produce estimates of the gradient of a cost function. Convergence to a neighborhood of the optimal gain values is proven by means of averaging analysis and singular perturbation theory. Similarly, the adaptive PDE backstepping

control design in [52] achieves output feedback stabilization with gradient-based estimators for some unknown model parameters. In this article, we employ the RL boundary control approach to deal with the nonlinear ARZ model with uncertain steady-state conditions, which has not been studied before.

B. RL Boundary Control Approach

Recent developments in RL have enabled model-free control of high-dimensional continuous control systems. Model-free RL control does not require any prior knowledge of the model structure, nor model calibration. Instead, the control input is obtained by iterative interactions with a *simulator of the system dynamics*. Thus, the PDE boundary control problem is formulated as an RL problem, requiring minimal assumptions on the infinite-dimensional system yet in practice obtains reasonably good controllers through a completely data-driven process.

Learning-based methods have drawn attention from both the PDE modeling community and transportation researchers. Algorithms for solving high-dimensional parabolic PDE systems were derived using deep learning in [23] by reformulating the PDEs as stochastic differential equations and approximating the gradient of unknown solutions with deep neural networks (DNNs). More recent work in [49] proposed RL-learned numerical solutions for scalar conservation laws that autonomously generate accurate numerical schemes for various situations. For the PDE control problem, boundary control of time-varying 2-D convection-diffusion PDEs with an application for heating, ventilating, air conditioning (HVAC) control design was developed in [16] and [18].

In the domain of traffic management, researchers have been applying RL to various traffic problems. Different levels of traffic modeling are employed, depending on the problem. For example, Li *et al.* [33] examined traffic light signal timing at one intersection using deep Q-learning with a traffic simulator, Paramics, which is based on microscopic traffic models. Studies on the traffic signal scheduling problem were further extended to multi-agent control for five centrally connected traffic intersections [2], which was described with a traffic queuing model. The article [13] used a multi-agent RL algorithm to control the traffic light around a traffic junction. The authors proposed a framework where each agent was able to switch between independent and integrated modes in which the agent solved the multi-agent RL problem using modular Q-learning. Wu *et al.* [51] formulated a deep RL framework for mixed-autonomy traffic in some experimental scenarios using the traffic microsimulator Simulation of Urban MObility (SUMO). Under the same framework, [48] developed RL controllers for connected autonomous vehicles to de-congest traffic bottlenecks.

There are few RL works related to macroscopic traffic models in the literature, although the PDE model is particularly well-suited for modeling congested traffic flow patterns. The reason is that RL control of PDEs involves high-dimensional state spaces which makes the approximation of value functions challenging. Belletti *et al.* [8] considered the cell transmission model which is obtained from discretization of the LWR first-order PDE model. An RL-based controller was designed

using different policy gradient methods, such as REINFORCE, trust policy optimization (TRPO), and the truncated natural policy gradient (TNPG) algorithm. Incoming traffic flow is actuated by the RL controller such that the traffic flow is optimized for some target outflow. Our work on RL control differs on the traffic problem to be solved, the model being used, and the methodology being employed, which guides our study and analysis from a very different perspective. We focus on the stabilization problem and examine a second-order PDE model to describe stop-and-go traffic oscillations. The resulting RL controllers are tested, and we examine under what circumstances RL could be a better choice relative to a Lyapunov-based PDE controller.

C. Contribution

We formulate a state regulation control problem for the ARZ PDE model via boundary control. Then, we develop an RL control approach based on the proximal policy optimization (PPO) algorithm, which falls within the class of policy gradient methods. PPO ultimately yields a state-feedback boundary controller from iterative interactions with a simulation environment, as opposed to direct synthesis from a mathematical model. The performance of the RL controller is compared with a PDE backstepping controller, a P controller, and PI controllers.

Interestingly, RL controllers nearly recover the stabilization performance of the Lyapunov-based PDE control approaches for a system with perfect knowledge of the model. However, in a system with partial knowledge where the steady-state traffic is lighter or denser than what we assumed in constructing Lyapunov-based controllers, the RL controller obtained from a stochastic training process outperforms the Lyapunov-based controllers. However, it must be noted that the RL controllers are obtained by conducting about one thousand training episodes on a simulation model. Collision-free training is not possible in real-life traffic, nor is the iterative training guaranteed to converge. Although RL demonstrates learning (i.e., adaptation) potential under uncertain and changing conditions, it is neither simple nor a fully safe substitute for model-based control in real traffic systems. For real-world implementation, it is an active research area to develop safe RL strategies with some safety constraints guaranteed. Readers are referred to a more exhaustive review on this topic [21], and constrained policy optimization approach [1].

The contributions of this article are twofold: first, although several Lyapunov-based boundary controllers have been designed for the linearized ARZ model, there is no result for control of the original nonlinear ARZ model. This motivates the RL control in this article, which provides the first native controller for a nonlinear ARZ PDE model, to the best of our knowledge. Second, the performance of the RL controllers are compared with PDE backstepping, P, and PI controllers for stabilization of stop-and-go traffic, and are assessed for different scenarios. By evaluating performance with full knowledge of the dynamics and one with partial knowledge, we demonstrate via several numerical simulations under what circumstances RL could be a better choice from a traffic application perspective.

The outline of this article is as follows. Section II summarizes the ARZ traffic flow model and the Lyapunov-based boundary controllers. Section III details the RL approach for boundary control of the ARZ PDE model. Section IV includes the comparative analysis and numerical results, comparing the closed-loop performance of setpoint boundary inputs, backstepping, P, and PI controllers in different scenarios. The conclusion summarizes the main results and discusses future work.

II. PROBLEM STATEMENT

A. ARZ PDE Traffic Model

We consider the ARZ PDE model to describe the traffic dynamics on a freeway segment [3] [57]. The state variables are traffic density $\rho(x, t)$ and traffic speed $v(x, t)$, defined on the domains $x \in [0, L]$, $t \in [0, T]$

$$\rho_t + (\rho v)_x = 0 \quad (1)$$

$$(v - V(\rho))_t + v(v - V(\rho))_x = \frac{V(\rho) - v}{\tau} \quad (2)$$

where $(\cdot)_z$ is short-hand notation for the differential operator $\partial/\partial z$. Parameter τ is the relaxation time, and captures how quickly drivers adjust their velocity to the equilibrium. The equilibrium velocity-density relationship $V(\rho)$ is a decreasing function, e.g., the linear Greenshield's model $V(\rho) = v_m(1 - (\rho/\rho_m))$ representing the reduction of velocity with the increase of density where v_m is maximum velocity and ρ_m is maximum density on the freeway road. The traffic flow rate is defined with

$$q = \rho v. \quad (3)$$

Our control objective is to regulate the state around an equilibrium reference state (ρ^*, v^*) , where

$$v^* = V(\rho^*) \quad (4)$$

satisfies the equilibrium density and velocity relation. We choose the density ρ^* such that the reference system (ρ^*, v^*) is in the congested regime for dense traffic, which can be characterized by the two characteristics of the linearized PDE model [52]

$$\lambda_1 = v^* > 0 \quad (5)$$

$$\lambda_2 = v^* + \rho^* V'(\rho^*) < 0. \quad (6)$$

The first characteristic is always greater than 0. When traffic is light, $\lambda_2 > 0$ is satisfied. When traffic is dense, then $\lambda_2 < 0$ and in this regime there can be upstream propagation of oscillations in the states. This can also be characterized with the traffic Froude number (TFN) in [8]. Consequently, we have hetero-directional propagation of oscillations in congested traffic. The density oscillations are carried downstream by vehicles while the velocity oscillations are transported upstream. Intuitively, drivers are mostly affected by vehicles driving in front of them. The stop-and-go traffic is characterized by oscillations, caused by delayed driver reaction to vehicles in front of them. As shown in Fig. 1, the control inputs U_{in} and U_{out} are transported from the actuated boundaries to in-domain states with the two characteristic speeds, respectively.




Fig. 1. Stabilization of stop-and-go traffic on freeway with ramp metering located at boundaries of the segment.

B. Boundary Control Design

Our control objective is L^2 regularization of the density and velocity states to uniform steady-state values by ramp-metering on the boundary. The oscillations can occur due to delayed driver response without ramp-metering control. In order to reduce oscillations in congested traffic, we actuate traffic flow from the upstream inlet and downstream outlet of the freeway segment, using on-ramp metering as shown in Fig. 1. Alternatively, one can actuate velocity by installing VSLs.

Boundary control algorithms are designed to stabilize the traffic around the reference steady state. For the ARZ PDE model, the control design guarantees that the state variables $(\rho(x, t), v(x, t))$ are regulated to the reference system (ρ^*, v^*) in the spatial L^2 norm, i.e.,

$$\|\rho(x, t) - \rho^*\| \rightarrow 0 \quad (7)$$

$$\|v(x, t) - v^*\| \rightarrow 0 \quad (8)$$

where $\|\cdot\|$ is the spatial L^2 norm, i.e., $\|f(x, t)\| = (\int_0^L f^2(x, t) dx)^{1/2}$.

For the segment of freeway traffic, boundary actuation is implemented with ramp metering controlling the traffic flow rate entering from the on-ramp to the mainline. Both inlet and outlet flow rates can be controlled through ramp metering

$$\rho(0, t)v(0, t) = U_{in}(t) \quad (9)$$

$$\rho(L, t)v(L, t) = U_{out}(t). \quad (10)$$

We denote the boundary control inputs of different strategies with $U_{in,*}(t), U_{out,*}(t)$, where $*$ represents the name of the control design. In practice, there will be off-ramps on the freeway. For example, if we consider there is off-ramp at the outlet $x = L$ of the studied segment, boundary disturbances induced by some off-ramp traffic flow can be included in (10). Robust control designs of the hyperbolic PDE system are needed for the rejection of such boundary disturbances. Interested readers can refer to [4], [35], and [46]. As summarized in Table I, the Lyapunov-based boundary control algorithms consist of setpoint boundary inputs, backstepping PDE control [52], proportional [52], and PI control [58]. Despite different actuation locations and assumptions, three different state-feedback controllers achieve Lyapunov stabilization of the PDE system by collectively shifting the linearized PDE eigenvalues onto the open left half-plane. The control gains used for the backstepping method are computed numerically [52]. The static

TABLE I
LYAPUNOV-BASED BOUNDARY CONTROL DESIGN

Design method	U_{in}	U_{out}
Setpoint	q^*	q^*
Backstepping	q^*	$U_{\text{out,BKST}}$
Proportional	$U_{\text{in,L}}$	q^*
Proportional Integral	$U_{\text{in,PI}}$	$U_{\text{out,PI}}$

proportional gain in inlet control is derived by solving the linearized PDE system analytically [52]. The static PI gains in PI control are obtained through trial and error. It should be noted that the proposed controllers are continuous in time. For practical implementation of the ramp metering, data sampling strategies are needed to modulate the control signals that are continuous in time into binary ones for the traffic PDE model. A preliminary result can be found in [12].

1) *Setpoint Control*: We simply limit the outgoing flow rate to be the setpoint state value as the boundary control. The setpoint control is open-loop control which does not use feedback of the system. This simple method is considered as a baseline case to compare against more sophisticated control designs. The constant incoming and outgoing flow rate at the boundaries are implemented as

$$U_{\text{in,O}}(t) = q^* \quad (11)$$

$$U_{\text{out,O}}(t) = q^* \quad (12)$$

where $q^* = \rho^* v^*$ is the reference steady-state flow rate. The ARZ PDE system with this setpoint control can be unstable (in the linear sense) for certain combinations of q^* and model parameters $V(\rho)$ and τ [53, Sec. III].

2) *PDE Backstepping Control*: The backstepping controller produces a full-state-feedback control law located at the outlet boundary. The control design is developed in [52] for the linearized ARZ PDE model, and is applied to the original non-linear ARZ model. For initial conditions near the reference, the system is locally exponentially stable under outlet actuation. The backstepping method is used to construct the full-state-feedback control law [28]. The key idea of the backstepping method is to transform the original linearized ARZ model to a target system where the undesirable instabilities in the PDEs are transformed to the outlet boundary. The controller is designed such that the instability is canceled out and an exponentially stable target system with zero boundary input is obtained for the closed-loop system.

The boundary condition at the inlet, and the full-state-feedback outlet backstepping controller are given by

$$U_{\text{in,BKST}}(t) = q^* \quad (13)$$

$$U_{\text{out,BKST}}(t) = q^* + \rho^* \int_0^L c_v(\xi)(v(\xi, t) - v^*)d\xi + \int_0^L c_q(\xi)(q(\xi, t) - q^*)d\xi \quad (14)$$

where the control gains are given by $c_v(\xi) = M(L - \xi) + (\lambda_2/\lambda_1)K(L, \xi)\exp(\xi/\tau v^*)$, $c_q(\xi) = ((\lambda_1 - \lambda_2)/\lambda_1)K(L, \xi)\exp(\xi/\tau v^*)$ for $\xi \in [0, L]$. The control gain kernels $K(L, \xi)$,

$M(L - \xi)$ are obtained by solving hyperbolic equations that govern the kernel variables $K(x, \xi)$ in a triangular spatial domain $\mathcal{T} = \{(x, \xi) : 0 \leq \xi \leq x \leq L\}$. The numerical solution of the kernel equations is easily obtained, given the model parameters and steady states. The equilibrium (ρ^*, v^*) of the linearized system is exponentially stable in the L^2 sense and the equilibrium is reached in a finite time.

Theorem 1 [52, Th. 2]: Consider system (1) and (2) linearized around (4), initial conditions $(\rho(x, 0), v(x, 0)) \in L^2([0, L])$ and the control laws (13), (14). The equilibrium (ρ^*, v^*) of the linearized system is exponentially stable in the L^2 sense and the equilibrium is reached in a finite time $t = (L/|\lambda_1|) + (L/|\lambda_2|)$.

The proof of Theorem 1 is detailed in [52], using Lyapunov analysis.

3) *P Control*: The proportional controller is an output-feedback control law that actuates traffic flow rate at the inlet boundary, as proposed in [52]. The control input only requires online measurements of velocity at the inlet boundary, collocated with the actuation location. The key intuition of the control design is to cancel the forward coupling in the system and then the closed-loop system can be directly solved. Local exponential stability and finite-time convergence is guaranteed. The boundary condition at the outlet is a constant flow rate and the inlet boundary is actuated by the collocated output-feedback controller

$$U_{\text{in,L}}(t) = q^* + g_P(v(0, t) - v^*) \quad (15)$$

$$U_{\text{out,L}}(t) = q^* \quad (16)$$

where $g_P = \rho^* + v^*/V'(\rho^*)$ is a constant control gain.

Theorem 2 [52, Th. 1]: Consider system (1) and (2) with initial conditions $(\rho(x, 0), v(x, 0)) \in L^2([0, L])$, linearized around (4) and the control law (15), (16). The equilibrium (ρ^*, v^*) of the linearized system is exponentially stable in the L^2 sense and the equilibrium is reached in a finite time $t = (L/|\lambda_1|) + (L/|\lambda_2|)$.

4) *PI Control*: In the PI controller, feedback control is applied to both the inlet and outlet boundary values. The PI boundary feedback controllers are proposed in [58] for the linearized ARZ model, which guarantees local exponential stability of the closed-loop system. The exit velocity at the outlet is controlled such that $v(L, t) = U_{\text{out,PI}}(t)$. Ramp metering controls the incoming flow rates, given by

$$U_{\text{in,PI}}(t) = q^* + k_P^r(\rho(L, t) - \rho^*) + k_I^r \int_0^t (\rho(L, t) - \rho^*)ds \quad (17)$$

$$U_{\text{out,PI}}(t) = v^* + k_P^v(v(0, t) - v^*) + k_I^v \int_0^t (v(0, t) - v^*)ds \quad (18)$$

where $k_P^v, k_I^v, k_P^r, k_I^r$ are tuning gains. For this anti-collocated output-feedback structure, a set of linear matrix inequalities are given for allowable control gains that guarantee Lyapunov stability. Within these conditions, the specific values are obtained through trial and error, as detailed in [58].

Theorem 3 [58, Th. 1]: Consider system (1) and (2) linearized around (4) and the control laws (17), (18).

Algorithm 1 Boundary Control Scheme

Result: Feedback boundary stabilization
initial conditions: $s_0 = \rho(x, 0), v(x, 0)$;
Input: parameters $V(\rho), \tau$, steady state $s^* = (\rho^*, v^*)$;
for $t = 1$ **to** $N - 1$ **do**
 Compute Lyapunov-based boundary controllers
 $U_{\text{in,out}} = u(s_t, V(\rho), \tau, s^*)$, where the functional
 of boundary control u depends on the choice of
 control design in Table.1
 Update the state variables s_{t+1} with actuated
 boundary control inputs $U_{\text{in,out}}$ and s_t ;
end

Algorithm 2 RL Control Procedure

Initialize parameters for actor θ_0 and critic ϕ_0
for $k = 1$ **to** E **do**
 Initialize the states s_0 from random distribution
 for $t = 1$ **to** N **do**
 Compute reward $r_t = \|s_t - s^*\|_2$;
 Draw an action a_t from stochastic policy, i.e.,
 $a_t \sim \mathcal{N}(\mu(s_t), \sigma^2(s_t))$;
 a_t goes to boundary conditions of PDE system;
 Update state s_{t+1} using a_t and s_t ;
 end
 Collect set of trajectories \mathcal{D} driven by policy π_{θ_k}
 Compute total discounted reward R_t .
 Compute advantage estimates, A_t from critic V_{ϕ_k} .
 Update the actor, θ_{k+1} by (38).
 Update the critic, ϕ_{k+1} by regression.
end

The equilibrium $\rho(x, t) \equiv \rho^*, v(x, t) \equiv v^*$ of the linearized system is exponentially stable in the L^2 sense.

All of the aforementioned controllers are designed from the linearized ARZ model, and therefore the stabilization guarantees are local. In addition, model uncertainty has not been taken into consideration, such as random events that frequently appear in traffic. This motivates us to explore an RL approach.

III. CONTROL OF ARZ MODEL WITH RL

In this section, we introduce an RL approach for boundary control of the nonlinear ARZ traffic flow model. Although explicit knowledge of the PDE model is not required, RL considers that the traffic dynamics are governed by a Markov decision process (MDP). In particular, we will use a policy gradient method since they are applicable to continuously valued control actions.

A. Boundary Control of PDE as an MDP

In the MDP setting, we seek the policy that maximizes the total reward received from the environment, i.e., the plant. At each time step t the environment conditions are described by a state vector, $s_t \in \mathcal{S}$, where \mathcal{S} is the state space while the control policy picks an action $a_t \in \mathcal{A}$, with \mathcal{A} being

the action space. The control policy is a full state feedback control law that selects an action a_t based on an observation of the state s_t . The action is applied to the environment, whose state evolves to $s_{t+1} \in \mathcal{S}$, according to the state-transition probability $\mathcal{P}(s_{t+1}|s_t, a_t)$, and the agent receives a scalar reward $r_{t+1} = r(s_t, a_t)$. The policy is represented by π which maps the state to the action and can be either deterministic or stochastic. The total discounted reward from time t onward can be expressed as

$$R_t = \sum_{k=0}^{\infty} \gamma^k r(s_{t+k}, a_{t+k}) \quad (19)$$

where $\gamma \in [0, 1]$ is the discount factor. Note, we have abused notation by allowing t to represent both continuous and discrete-valued times. Nevertheless, the meaning will be clear from context.

The implementation of the RL controller is summarized in Algorithm 1. We represent the state at time t with $s_t = (\rho(\cdot, t), v(\cdot, t))$. When the nonlinear ARZ model is boundary actuated, then it naturally forms a sequential decision-making problem, which is modeled as an MDP.

We consider a discretized approximation of the nonlinear ARZ PDE model using the Lax–Wendroff scheme [32] with conservative state variables. The discretized ARZ PDE model is a difference equation which numerically approximates the infinite-dimensional PDE model with second-order accuracy in space and time. The solution $\rho(x, t)$ and $v(x, t)$ to the ARZ PDE model is approximated by piecewise constant functions on discretized temporal and spatial domains. The solution domain is $[0, L] \times [0, T]$. The discretization resolution $\Delta t = T/(N - 1)$ and $\Delta x = L/(M - 1)$ are chosen such that the Courant–Friedrichs–Lewy (CFL) condition is met, i.e., $\Delta t \leq c \Delta x$, where M, N are the number of nodes for the spatial and temporal domains, respectively, and c is defined as the maximum characteristic speed of the nonlinear hyperbolic ARZ PDE model ($\Delta x / \Delta t \geq \max |\lambda_{1,2}|$). We can compactly write the discretized nonlinear ARZ PDE system (1) and (2) as the following difference equation:

$$s_{t+1} = f(s_t, u_t) \quad (20)$$

where s_t describes the density and velocity state at time t , and u_t represents the boundary control inputs, depending on the choice of control design in Table I. The function f represents the discretized deterministic dynamics for the temporal evolution of the PDE system. The PDE dynamics at the current time instant are fully described given the current state s_t and control inputs u_t .

Given the discretized ARZ PDE in (20), we may write the deterministic dynamics as

$$\mathcal{P}(s_{t+1}|s_t, a_t) = \delta(s_{t+1} - f(s_t, u_t)) \quad (21)$$

where $\delta(\cdot)$ is the Dirac delta function. In the MDP, $\mathcal{P}(s_{t+1}|s_t, a_t)$ represents the probability that action a in state s at time t will lead to the state at time $t+1$. Therefore, if there is stochasticity in the dynamics, such as some random model parameters, as introduced later in our article, the deterministic

temporal evolution of the PDE system can be generalized to stochastic dynamics by the MDP state transition probability

$$s_{t+1} \sim \mathcal{P}(s_{t+1} | s_t, \mathbf{a}_t). \quad (22)$$

The discretized states, s_t , and boundary control input, \mathbf{a}_t are written as

$$s_t = [\rho(0, t), \rho(\Delta x, t), \dots, \rho(L, t), \\ v(0, t), v(\Delta x, t), \dots, v(L, t)]^\top \quad (23)$$

$$\mathbf{a}_t = [q(0, t), q(L, t)]^\top \quad (24)$$

where $\rho(\cdot, t)$ and $v(\cdot, t)$ are the traffic density and velocity that are discretized in the spatial and temporal domains.

The reward $r(s, \mathbf{a})$ is defined by the L_2 norm of the states, namely

$$r_t(s_t, \mathbf{a}_t) = - \left\| \frac{\sum_i \rho(i \cdot \Delta x, t) - \rho^*}{\rho^*} \right\|_2 - \left\| \frac{\sum_i v(i \cdot \Delta x, t) - v^*}{v^*} \right\|_2. \quad (25)$$

The reward is equivalent to the control objective that achieves regulation of the traffic states to a spatially uniform density and velocity.

B. RL Background

In the following, we briefly review some essential RL concepts to aid the subsequent discussion. A thorough exposition can be found in [9] and [45].

The *state value function*, $V^\pi(s_t)$ is the expected total discounted reward starting from state s_t . In the controls community, this is sometimes called the cost-to-go, or reward-to-go. Importantly, note that the value function depends on the control policy. If the agent uses a given policy π to select actions starting from the state s_t , then the corresponding value function is given by

$$V^\pi(s_t) = \mathbb{E}[R_t | s_t]. \quad (26)$$

Then, the optimal policy π^* is the policy that corresponds to the maximum value $V^*(s_t)$ of the value function

$$\pi^* = \arg \max_{\pi} V^\pi(s_t). \quad (27)$$

The solution of (27) is pursued by dynamic programming (DP) methods. DP, however, requires knowledge of the model/environment.

The next definition, known as the “Q-function,” is fundamental since it enables the concept of model-free RL. Consider the *state-action value function*, $Q^\pi(s_t, \mathbf{a}_t)$, which is a function of the state-action pair that returns a real value. In other words, it corresponds to the expected total discounted reward when the action \mathbf{a}_t is taken in state s_t , and then the policy π is followed henceforth. Mathematically

$$Q^\pi(s_t, \mathbf{a}_t) = \mathbb{E}[R_t | s_t, \mathbf{a}_t]. \quad (28)$$

The optimal Q-function is given by

$$Q^*(s_t, \mathbf{a}_t) = \arg \max_{\pi} Q^\pi(s_t, \mathbf{a}_t) \quad (29)$$

and represents the expected total discounted reward received by an agent that starts in s_t , picks (possibly non-optimal) action \mathbf{a}_t , and then behaves optimally afterward. Since $V^*(s_t)$ is the maximum expected total discounted reward starting from state s_t , it will also be the maximum of $Q^*(s_t, \mathbf{a}_t)$ over all possible actions $\mathbf{a}_t \in \mathcal{A}$

$$V^*(s_t) = \max_{\mathbf{a}_t \in \mathcal{A}} Q^*(s_t, \mathbf{a}_t). \quad (30)$$

If the optimal Q-function is known, then the optimal action \mathbf{a}_t^* can be extracted by choosing the action \mathbf{a}_t that maximizes $Q^*(s_t, \mathbf{a}_t)$ for state s_t (i.e., the optimal policy π^* is retrieved)

$$\mathbf{a}_t^* = \arg \max_{\mathbf{a}_t \in \mathcal{A}} Q^*(s_t, \mathbf{a}_t) \quad (31)$$

without requiring knowledge of the environment dynamics. This last point is precisely why the Q-function enables model-free RL.

Lastly, the advantage function is defined to measure how advantageous the action is compared to the action drawn from policy

$$A^\pi(s_t, \mathbf{a}_t) = Q^\pi(s_t, \mathbf{a}_t) - V^\pi(s_t). \quad (32)$$

C. Actor-Critic

Actor-critic is an approximate DP (ADP) method that solves DP heuristically. Importantly, the actor-critic approach allows for continuous state/action spaces by using a function approximator, e.g., a neural network. In RL, as well as in DP, the action is taken by a *policy* to maximize the expected total discounted reward. By following a given policy and processing the rewards, one should estimate the expected return given states from the *value function*. In the actor-critic approach, the *actor* improves the policy based on the value function that is estimated by the *critic*. We specifically focus on the policy gradient-based actor-critic algorithm in this work, and, in particular, the PPO [42] is considered. The critic is the parameterized value function V_ϕ and the actor is the parameterized policy π_θ .

1) *Critic*: The role of the critic is to evaluate the current policy prescribed by the actor. The action is drawn from a Gaussian distribution, namely

$$\mathbf{a}_t \sim \mathcal{N}(\mu, \sigma^2), \text{ where } [\mu, \sigma] = f_{\text{DNN}}(s_t; \theta) \quad (33)$$

where a mean μ and standard deviation σ computed from a DNN, $f_{\text{DNN}}(s_t; \theta) : \mathcal{S} \rightarrow \mathbb{R}^2$, and the DNN is parameterized by weight vector θ . After applying the action, we observe the reward r_{t+1} and the next state s_{t+1} . For each time step t , the tuple $(s_t, \mathbf{a}_t, r_{t+1}, s_{t+1})$ is stored in the buffer, \mathcal{D} . From a collected set of trajectories, the parameterized value function denoted by V_ϕ is updated to minimize the following loss function, \mathcal{L} :

$$\mathcal{L} = \frac{1}{N} \sum_{i \in \mathcal{D}} \sum_{t=0}^N (V_\phi(s_{i,t}) - R_{i,t})^2 \quad (34)$$

where $R_{i,t}$ is the total discounted reward at time t in the i th trajectory stored in the buffer. The critic network parameters, ϕ , are updated numerically via gradient descent.

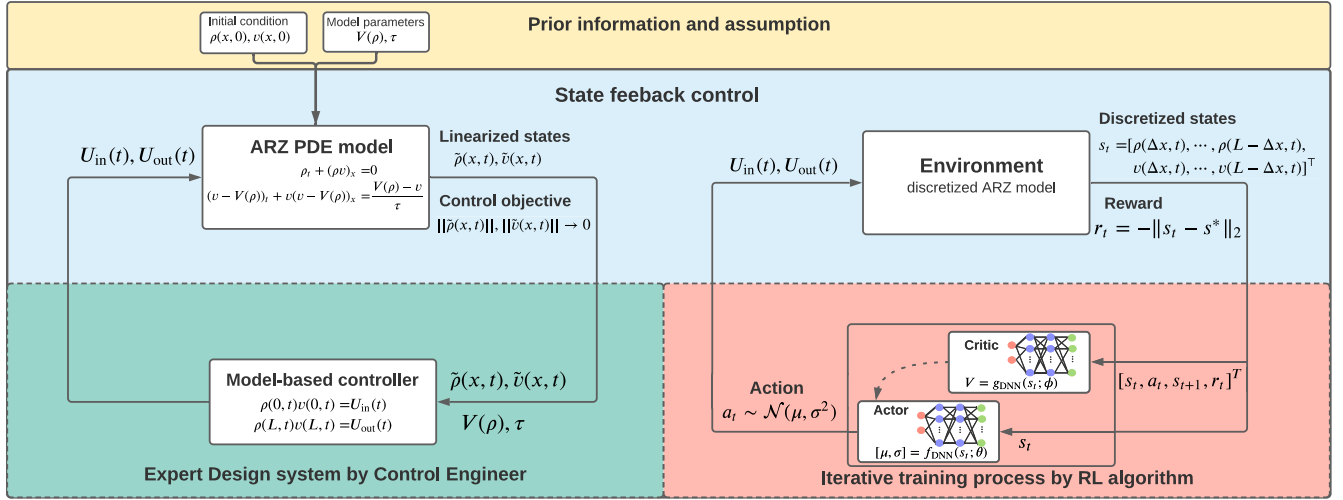


Fig. 2. State-feedback control loop of ARZ PDE Lyapunov-based control scheme and model-free RL boundary control scheme.

2) *Actor*: The actor is updated based on the value function estimates. The objective function for the actor is formulated in terms of the expected reward of policy π_θ and the advantage of $\pi_{\theta_{old}}$ [24], [41]

$$\max_{\theta} \hat{\mathbb{E}}_t \left[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \hat{A}_t \right] \quad (35)$$

where the hats on $\hat{\mathbb{E}}_t$ signify a sample mean, and \hat{A}_t indicates an estimated advantage function obtained from the critic. Schulman *et al.* [41] prove the expected reward corresponding to π_θ increases relative to $\pi_{\theta_{old}}$, if a distance measure between π_θ and $\pi_{\theta_{old}}$ is sufficiently bounded. This motivates the following trust region policy optimization (TRPO) algorithm:

$$\max_{\theta} \hat{\mathbb{E}}_t \left[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \hat{A}_t \right] \quad (36)$$

$$\text{s.t. } \hat{\mathbb{E}}_t [\text{KL}[\pi_{\theta_{old}}(\cdot|s_t), \pi_\theta(\cdot|s_t)]] \leq \delta \quad (37)$$

where θ_{old} is the vector of policy parameters before the update. Kullback–Leibler (KL)-divergence measures the difference between the old policy and current policy. The constraint ensures that the new policy does not deviate from the old policy by δ .

In this work, we adopt the PPO RL algorithm [42], which is based on TRPO [41]. The PPO algorithm similarly limits the new policy from being excessively far from the previous one. However, it does so with a modified objective that penalizes changes to the policy that move $r_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{old}}(a_t|s_t)$ away from 1. The key idea is to use probability clipping, as follows:

$$\max_{\theta} \hat{\mathbb{E}}_t [\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1-\varepsilon, 1+\varepsilon)\hat{A}_t)]. \quad (38)$$

The main idea of PPO is to modify the objective by clipping the probability ratio. This removes the incentive for moving r_t outside of the interval $[1-\varepsilon, 1+\varepsilon]$. With this clipping method, the lower bound of objective function is maximized. Readers are referred to [42] for more details.

3) *Proximal Policy Optimization*: We adopt a policy gradient-based approach to obtain a continuous-valued stochastic control policy. Mathematically, the goal is to find

$$\theta^* = \arg \max_{\theta} \mathbb{E}[\mathcal{R}_t] \quad (39)$$

where the expectation is taken with respect to $\mathcal{P}(s_{t+1}|s_t, a_t)$ and $\pi_\theta(a_t|s_t)$, and θ parameterizes the control policy distribution. Policy gradient methods essentially solve (39) via gradient ascent. The key challenge is estimating the gradient since it is computationally intractable to compute it exactly.

One can reformulate this optimization problem (39) in terms of the expected reward of policy π_θ and the advantage of $\pi_{\theta_{old}}$ [24], [41]

$$\max_{\theta} \hat{\mathbb{E}}_t \left[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \hat{A}_t \right] \quad (40)$$

where the hats on $\hat{\mathbb{E}}_t$ signify a sample mean, and \hat{A}_t indicates an estimated advantage function from simulations. Schulman *et al.* [41] prove the expected reward corresponding to π_θ increases relative to $\pi_{\theta_{old}}$, if a distance measure between π_θ and $\pi_{\theta_{old}}$ is sufficiently bounded. This motivates the following TRPO algorithm:

$$\max_{\theta} \hat{\mathbb{E}}_t \left[\frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{old}}(a_t|s_t)} \hat{A}_t \right] \quad (41)$$

$$\text{s.t. } \hat{\mathbb{E}}_t [\text{KL}[\pi_{\theta_{old}}(\cdot|s_t), \pi_\theta(\cdot|s_t)]] \leq \delta \quad (42)$$

where θ_{old} is the vector of policy parameters before the update. KL-divergence measures the difference between the old policy and current policy. The constraint ensures that the new policy does not deviate from the old policy by δ .

Both the RL and Lyapunov-based approaches provide state-feedback controllers for the ARZ PDE model. Fig. 2 compares the RL and Lyapunov-based approaches as signal flow diagrams. The expert design system developed by a control engineer is replaced with an iterative learning process by an RL algorithm. In addition, the prior assumption of the initial condition and model parameters are not required in the RL

approach. In Section IV, we will conduct several numerical simulations to compare these two methodologies.

IV. COMPARATIVE STUDY AND SIMULATION

In this section, we numerically test the control designs of the Lyapunov-based controllers and the RL controller. Two settings are considered: 1) full knowledge of the system dynamics and 2) partial knowledge of the system. In the full knowledge setting, both the Lyapunov-based and model-free RL controllers have perfect knowledge of the model with known model parameters and steady-state conditions. For the partial knowledge setting, both controllers have partial knowledge of the model, i.e., the true steady state has deviated from the ones used in the model. For each setting, we test the model-free RL controller with the following cases: 1) outlet control; 2) inlet control; and 3) outlet and inlet control, as defined in (9) and (10). For each case, one Lyapunov-based controller is used to evaluate the performance of the RL controller.

A. Simulation Configuration

We simulate traffic on a freeway segment with length $L = 500$ m for a time period $T = 240$ s = 4 min for system with full knowledge of dynamics and 8 min for system with partial knowledge. Steady-state density $\rho^* = 120$ veh/km, $v^* = 36$ km/h are chosen given the maximum density $\rho_m = 160$ veh/km, and maximum velocity $v_m = 40$ m/s. The steady-state traffic is lightly congested such that $\lambda_1 = 10$ m/s, $\lambda_2 = -20$ m/s, which satisfy the conditions in (5) and (6). This configuration represents heterogeneous propagation of oscillatory waves of density and velocity in congested traffic. The 500 m long road segment is considered and the simulation is run for 4 min. The spatial grid size is chosen as $\Delta x = 10$ m and the temporal grid size is chosen as $\Delta t = 0.25$ s. The discretization in spatial-temporal domain guarantees the convergence of the numerical scheme. Besides, the computational time for training an RL controller increases with the number of grid size but the closed-loop performance is not affected significantly for the order of discretization we choose. The following simulations are performed on a Python running an Intel core i9-9900K central processing unit (CPU) with a clock rate of 3.60 GHz, and graphics processing unit (GPU) device, GeForce RTX 2080 Ti.

B. Comparative Study With Full Knowledge of System Dynamics

We assume sinusoidal initial conditions

$$\rho(x, 0) = 0.1 \sin\left(\frac{3\pi x}{L}\right) \rho^* + \rho^* \quad (43)$$

$$v(x, 0) = -0.1 \sin\left(\frac{3\pi x}{L}\right) v^* + v^*. \quad (44)$$

1) *Learning Process of RL Controllers:* The learning process for RL controllers is illustrated in Fig. 3. The learning curve for RL corresponding to each control scheme is presented in Fig. 3(a). The evolution of the cumulative reward R_t , defined in (19), reflects the overall learning performance.

In addition, several episodes of the simulation in the learning process are compared more closely in Fig. 3(b). We plot the reward values r_t , defined in (25), over several hundreds of iterations of testing candidate RL outlet feedback laws on ARZ model simulations over a time window of 4 min for each simulation.

In Fig. 3(a), the shaded area represents min/max of different actor-critic parameterizations while the bold line measures the average of the performance during training. Note that the training process converges over different numbers of episodes across the controllers. This motivates us to set a large enough number of episodes in order to make sure that RL converges. We obtain the actor-critic parameters by the end of training. It is also observed that there is a spike in cumulative reward for the inlet and outlet RL controller at the episode around 500 which indicates that the closed-loop system at the iteration becomes unstable and states diverge. For example, if density state is greater than the maximum density, there will be collision in the ARZ model. Therefore, this training process cannot perform in a collision-free fashion in real traffic.

In Fig. 3(b), we choose the outlet control case to illustrate the RL learning process. We plot the reward evolution for the closed-loop system with RL outlet controllers that are obtained from training episodes 1, 100, 200, 400, and 800, respectively. In addition, the reward of the setpoint control is also plotted as a baseline comparison. It turns out that the initial reward is worse than the baseline performance. After 100 iterations of training, the outlet RL controller plotted with dotted blue line performs even worse with reward decreasing with time and the lowest cumulative reward $R = -615.5$. However, the RL performance improves after the iteration number increases from 200 to 400 and reaches its best at episode 800. As demonstrated in this figure, the learning process does not guarantee a monotonically improving performance. Certain episodes can be even worse than the baseline case due to some iterations of unsuccessful training.

2) *State Evolution, Reward, and Control Inputs:* In Figs. 4–7, the initial condition is highlighted in blue. The actuated inlet and/or outlet boundaries are highlighted in red. In other words, the red curves visualize the control inputs for the setpoint control, backstepping control, P control, PI control, and the RL controllers in each case. Specifically, the closed-loop results are compared between the outlet backstepping controller and the outlet RL controller, as shown in Fig. 5, between inlet P controller and inlet RL controller as shown in Fig. 6, between the inlet and outlet PI controller and inlet and outlet RL controller as shown in Fig. 7.

Fig. 4 plots the evolution of density and velocity under setpoint control. The setpoint control is open-loop that does not use feedback of the system. We observe the persisting oscillations, although they appear lightly damped. We regard the setpoint control as a baseline case against which the subsequent feedback control designs should outperform. In setpoint control, the outlet and inlet boundary flow rates are controlled to be constant (11), (12). As highlighted with red in Fig. 4, the product of density and velocity values is constant at inlet and outlet boundary whereas the boundary states of density and velocity oscillate over time.




Fig. 3. Learning curves for RL and reward evolution in the learning process. (a) Learning curve of RL controllers with different initializations of the actor-critic network. (b) Reward evolution for the closed-loop system with outlet RL controller that is obtained for different episode numbers k .




Fig. 4. Density and velocity evolution and reward evolution with setpoint inputs. (a) Density and velocity of ARZ model under setpoint control. (b) Reward evolution for setpoint control.

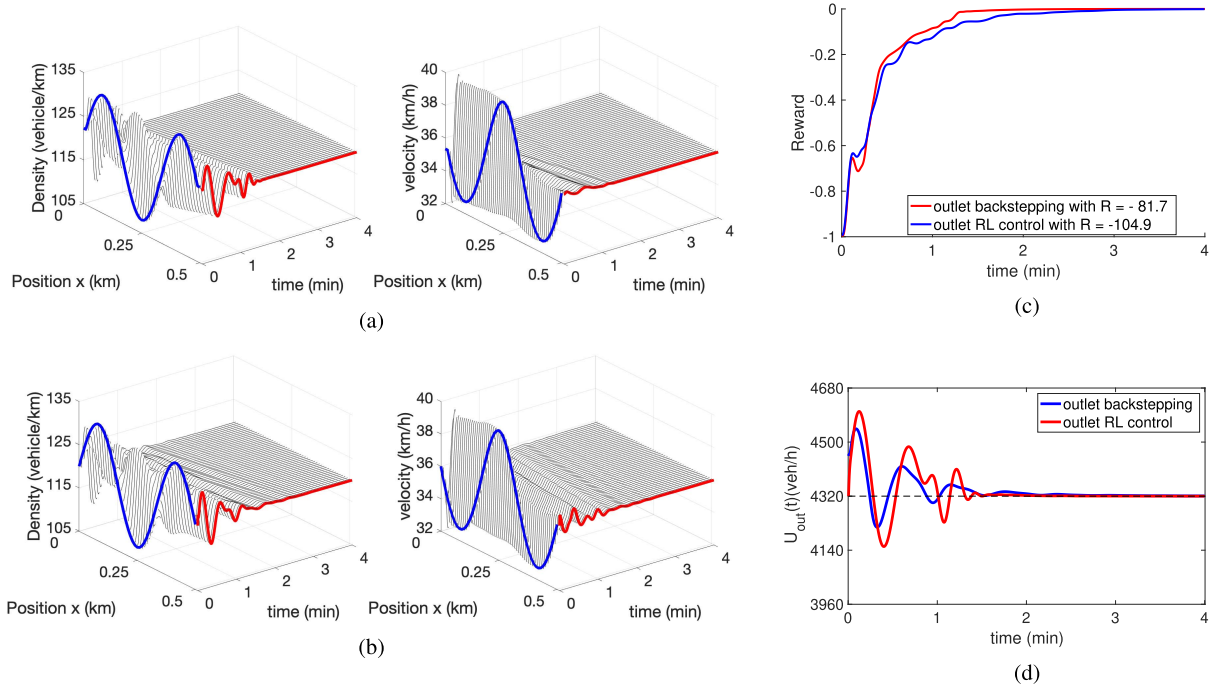


Fig. 5. Density and velocity evolution, reward evolution and control inputs of closed-loop with outlet controller. (a) Closed-loop with **outlet backstepping controller**. (b) Closed-loop with **outlet RL controller**. (c) Reward evolution for closed-loop with backstepping and RL controller at outlet. (d) Backstepping and RL control inputs at outlet.

Fig. 5 presents the closed-loop result of the outlet backstepping controller and outlet RL controller. In Fig. 5(a) and (b), the states are stabilized to spatially uniform steady-state values

by the outlet backstepping controller and outlet RL controller, respectively. An interesting finding from comparing between the two controllers is that RL learns a policy that produces

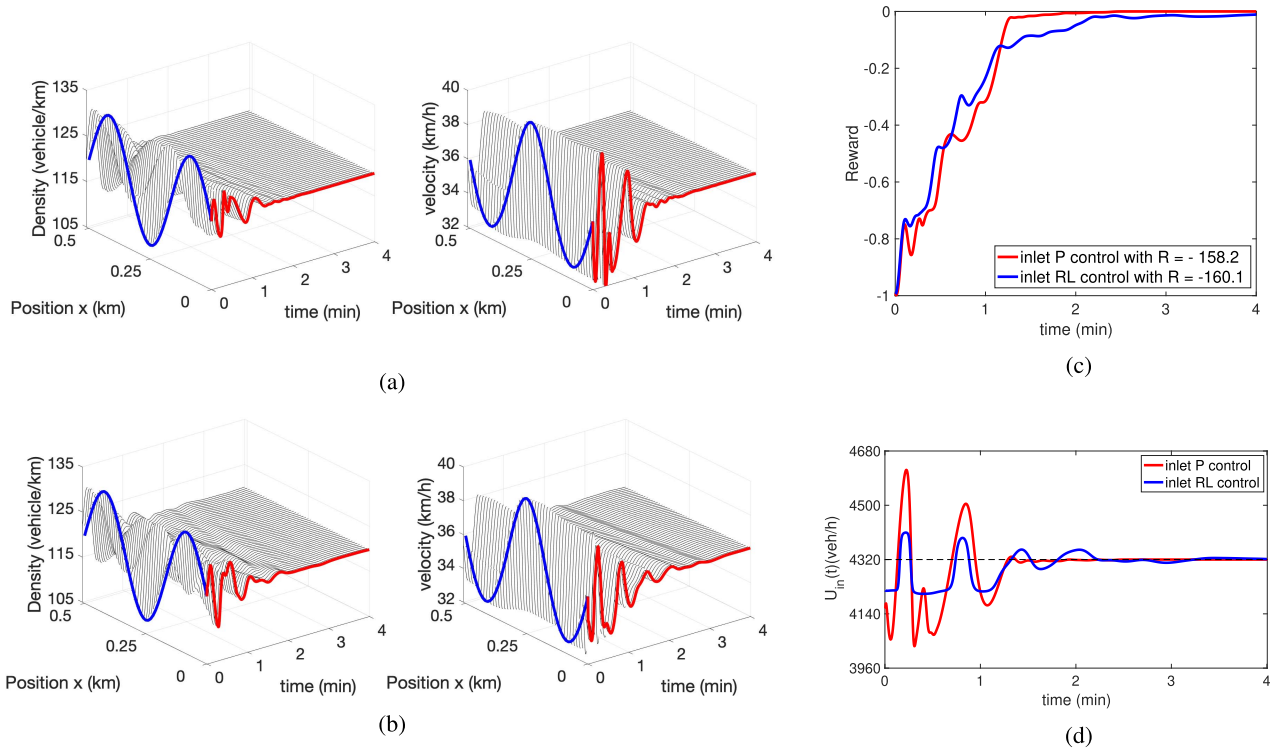


Fig. 6. Density and velocity evolution, reward evolution and control inputs of closed-loop with inlet controller. (a) Closed-loop with **inlet P controller**. (b) Closed-loop with **inlet RL controller**. (c) Reward evolution for closed-loop with P and RL controller at inlet. (d) P and RL control inputs at inlet.

a control input [red line in Fig. 5(b)] that closely replicates the backstepping control input [red line in Fig. 5(a)]. The RL policy is developed without explicit knowledge of the differential equations and parameters. It is trained iteratively on the nonlinear simulation model. As a state-feedback controller, the RL policy produces an action for boundary input given the current state. In contrast, the PDE backstepping state-feedback control law is obtained by rigorous theoretical control design assuming perfect knowledge of the model. As shown in Fig. 5(d), both methods yield similar control input trajectories. However, RL underperforms relative to backstepping in terms of the instantaneous reward over time. The convergence of the L^2 spatial norm represents the stabilization of the closed-loop system. Unlike the finite-time convergence for the backstepping controller shown with red line, the RL controller's reward in blue takes a longer time about 3 min for convergence. The cumulative reward for the RL outlet controller is $R = -104.9$ while the cumulative reward of backstepping is $R = -81.7$.

In the same fashion, we compare closed-loop results for inlet P controller and inlet RL controller in Fig. 6. The incoming flow rate is actuated either with the P controller or RL controller. The control inputs at the inlet are highlighted with red in Fig. 6(a) and (b). As shown in Fig. 6(d), the inlet RL controller needs more than 4 min to converge, longer than the P control. Although the inlet RL controller almost recovers the performance of the inlet P controller, we can see from Fig. 6(c) that the RL control input is quite different than the P control input, requiring longer convergence time and is oscillated.

Fig. 7 describes the closed-loop results for PI controller and RL controller, actuating the incoming and outgoing flow rates. We observe that the RL controller achieves a better reward evolution in Fig. 7(c). The convergence of the closed-loop system with RL is faster than that of the PI controller, and with a slightly smaller cumulative reward. It could also be noted that control efforts are smaller for the RL controller in terms of magnitude, as shown in Fig. 7(c). These results will vary, however, with the PI controller gain selection.

The comparison of closed-loop results between the Lyapunov-based and RL approaches is summarized in Table II. The stabilization result is mainly evaluated by convergence time S_T and control effort S_E . The convergence time is defined as the time instance such that the instant reward converges to zero $r_{t=S_T} = 0$. We mainly focus on the control efforts that varies around the nominal values q^* . Therefore, the control effort is defined as $S_E = \int_0^{S_T} |U_{in,out}(t) - q^*| dt$.

Fig. 8 includes all the closed-loop results with the Lyapunov-based controllers and their corresponding RL controllers. All of the Lyapunov-based feedback controllers and RL controllers outperform the setpoint controllers. The RL agent achieves suboptimal performance compared to Lyapunov-based approaches. In addition, backstepping outperforms the P and PI controllers for the stabilization of the traffic flow. RL controllers almost recover the closed-loop results of the Lyapunov-based design in the outlet control or inlet control cases. RL outperforms the PI controllers when both the inlet and outlet boundaries are controlled. The aforementioned results are uncertain model parameters, which we will discuss in Section IV-C. Note that the reward in (25) is defined as the

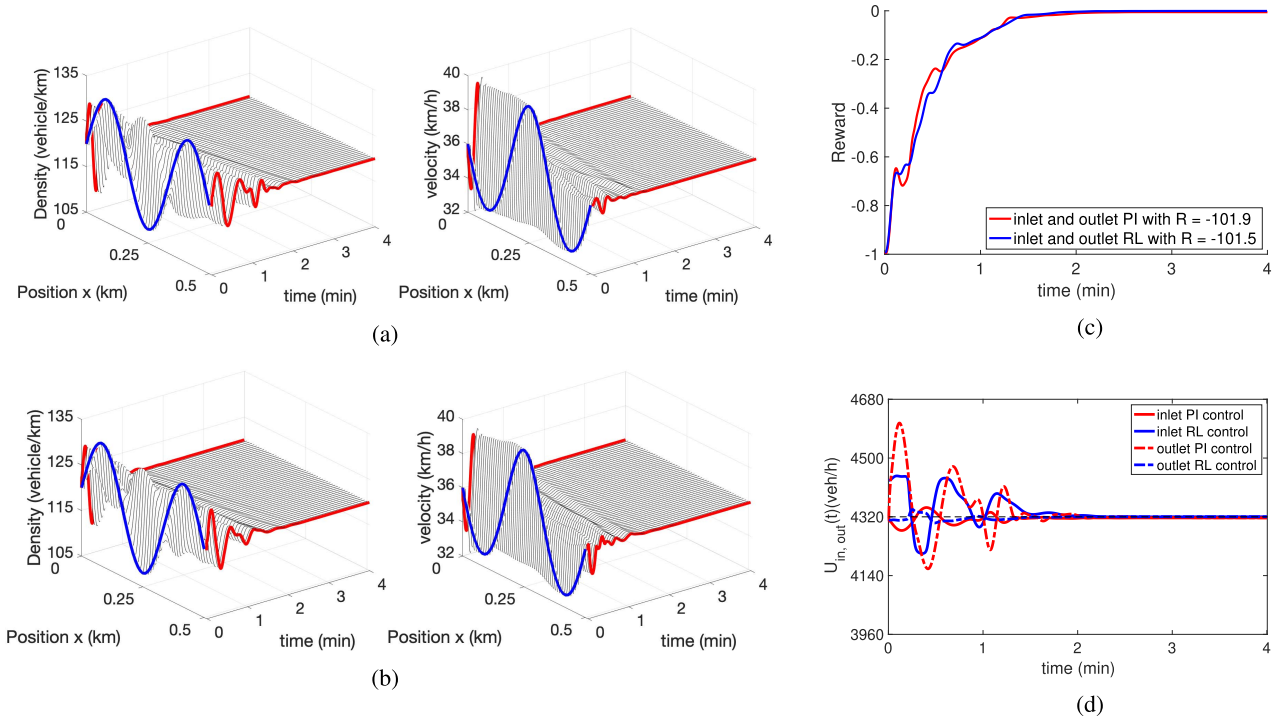


Fig. 7. Density and velocity evolution, reward evolution and control inputs of closed-loop with inlet and outlet controllers. (a) Closed-loop with **PI controllers** at inlet and outlet. (b) Closed-loop with **RL controllers** at inlet and outlet. (c) Reward evolution for closed-loop with PI and RL controller at inlet and outlet. (d) Inlet and outlet PI and RL control inputs.

TABLE II
CLOSED-LOOP RESULTS COMPARISON

	Lyapunov-based			RL		
Stability	backstepping	P	PI	outlet RL	inlet RL	outlet&inlet RL
S_T	79.5	84	140	121	204	139
S_E	7.9	10.5	12.6	5.5	9.4	9.5




Fig. 8. Reward evolution of the Lyapunov-based approaches (blue), the corresponding RL controllers after 2000 episodes of iterative training (red).

L^2 spatial norm of the state values without considering control effort. The RL controllers were designed for comparison with the stabilizing Lyapunov-based controllers, and thus control effort was not penalized. However, RL controllers can be trained to consider control effort by modifying (25) with a control effort penalty.

3) *Other Performance Measures:* Common traffic performance indices include total travel time (TTT), fuel consumption, and travel comfort. These indices are calculated for the closed-loop system with either the Lyapunov-based controllers or the RL controllers. The performance indices are given by

$$J_{TTT} = \int_0^T \int_0^L \rho(x, t) dx dt \quad (45)$$

$$J_{fuel} = \int_0^T \int_0^L \max\{0, b_0 + b_1 v(x, t) + b_3 v(x, t) + b_4 v(x, t) a(x, t)\} \rho(x, t) dx dt \quad (46)$$

$$J_{comfort} = \int_0^T \int_0^L (a(x, t)^2 + a_t(x, t)^2) \rho(x, t) dx dt \quad (47)$$

where $a(x, t)$ is defined as the local acceleration $a(x, t) = v_t(x, t) + v(x, t)v_x(x, t)$ and b_i are constant coefficients chosen as $b_0 = 25 \cdot 10^{-3} [1/s]$, $b_1 = 24.5 \cdot 10^{-6} [1/m]$, $b_3 = 32.5 \cdot 10^{-9} [1/s^3/m^2]$, $b_4 = 125 \cdot 10^{-6} [1/s^2/m^2]$, based on [47]. As shown in Table III, the performance indices of each controller are compared with their improvement percentage over the baseline setpoint controller. Among the three performance indices, drivers' comfort is the most significantly improved for all Lyapunov-based and RL controls, since the stop-and-go

oscillations are suppressed in the closed-loop system. For fuel consumption of total traffic, RL actually performs better than its corresponding Lyapunov-based controller. This may be coincidental, or could be related to RL's control effort and the induced acceleration. In any case, RL provides the flexibility to define the reward as fuel consumption, which can be directly optimized. Backstepping does not provide this capability. We also see that TTT is only marginally improved, since the traffic dynamics are stabilized to steady state and the average speed remains relatively similar.

C. Comparison Study With Partial Knowledge of System

In the previous comparative study, we assume perfect knowledge of the traffic system for the Lyapunov-based controllers. In practice, model parameters are obtained by calibrating the nonlinear ARZ PDE model with field data obtained from loop detectors measuring the traffic flow rate or by high-speed cameras recording vehicle trajectories. The model calibration process can be laborious. More importantly, it is hard to determine some model parameters such as the steady-state density, fundamental diagram, and relaxation time. These macroscopic models are just that—model idealizations of reality. In [52], an adaptive output feedback controller is designed to stabilize the linearized ARZ model with a gradient-based estimator for unknown relaxation time. The adaptive stabilization problem has not been studied when the steady state is uncertain. In traffic field data, it is observed that a certain steady-state density in the congested regime possesses a significant spread of flow rates. Therefore, it is hard to accurately determine the steady state for a real-world traffic system that is often periodically evolving, as exhibited in freeway traffic data [19]. On the other hand, the steady states in the reward function encode our belief in the current averaged and aggregated traffic condition which could deviate from reality. Moreover, there are measurement errors in raw data. The data need to be pre-processed leading to more approximation and processing errors. In [14] and [54], the ARZ model was calibrated with the next generation simulation (NGSIM) traffic data which records the trajectories of vehicles on a 500 m freeway segment over a 45 min rush-hour period. Data reconstruction is conducted to obtain the aggregated state values. All in all, the Lyapunov-based control approaches cannot overcome the limitation of model uncertainty when they are implemented in practice. In other words, their performance is limited by the best prediction the PDE model can make.

Therefore, it is of practical relevance to evaluate the closed-loop performance of the Lyapunov-based and RL controllers given partial or inaccurate models. Here, we investigate the performance of the Lyapunov-based controllers that employ incorrect steady-state density, and then compare with RL controllers obtained from a stochastic training process.

We choose two representative scenarios characterized by the steady-state density. We define ρ_r as the real steady-state density while $\rho^* = 120$ veh/km is assumed by the Lyapunov-based feedback controllers. We consider $\rho_r = 115$ veh/km for the outlet backstepping control case, representing that the actual traffic is lighter than the parameters used by the PDE

backstepping controller. In the second scenario, we consider $\rho_r = 125$ veh/km for the inlet control such that the actual traffic is denser than the value employed by the P controller. For both scenarios, RL controllers are trained in a stochastic environment where a steady-state density value of the traffic system ρ^* is randomly chosen from a uniform distribution of the integer values $\{115, 120, 125\}$ in each episode. Then the RL controllers are validated for each specific scenario.

The RL learning curve is plotted in Fig. 9(a). Note that the steady-state density used to calculate the cumulative reward is randomly chosen at each episode, which affects the initial condition and the associated steady-state velocity in (4) during training. The shaded area represents min/max of different actor-critic parameterizations, while the bold line measures the average performance. Similar to Fig. 3, it should be noted that the performance of RL controllers does not monotonically improve in the learning process as the number of training episodes increases. Additionally, cumulative reward convergence for some initiations of the actor-critic parameters is not guaranteed.

We observe that the RL controllers are more adaptable to the stochastic environment. Fig. 9(b) shows all the closed-loop controller results in Scenario 1 where the Lyapunov-based controllers assume a greater steady-state density value than the actual traffic environment. It is interesting to find out that all of the RL controllers in blue outperform the setpoint control in green, with larger reward values. Among them, the inlet RL controller ultimately performs the best. In addition, all the Lyapunov-based feedback controllers in red have smaller rewards, indicating the traffic states do not converge to the actual steady-state value after applying the controllers. Assuming an incorrect steady-state density value deteriorates the closed-loop stabilization results. The backstepping outlet controller, which performs the best in a system with perfect knowledge from Section IV-B, turns out to be the worst controller in terms of cumulative reward in Scenario 1, as depicted in Fig. 9(b). This observation demonstrates that Lyapunov-based control is sensitive to full knowledge of the system dynamics, and even perturbing the steady-state values can result in worse performance than a baseline control method.

1) *Scenario 1 of Lighter in-Domain Traffic:* As shown in Fig. 10, we apply the outlet backstepping controller to a nonlinear ARZ model with lighter steady-state traffic density, $\rho_r = 115$ veh/km than the controller assumes, i.e., $\rho_r < \rho^*$. The closed-loop simulation is run for 8 min, since the stabilization requires a longer time. When we apply the PDE backstepping controller in (14), which is constructed using $\rho^* = 120$ veh/km, the traffic is actually slowed down, leading to a reduction of velocity and an increase of density in Fig. 10(a). The closed-loop traffic is more congested, due to the controller's incorrect model assumptions. In contrast, since the RL controller is trained in an environment with stochastic conditions, where some episodes use a lighter traffic scenario, it successfully stabilizes the traffic state to a uniform value. We find out in Fig. 10(b) that the closed-loop system with the RL outlet controller converges close to the actual steady state. This leads to a larger cumulative reward $R = -534.3$ for RL

TABLE III
PERFORMANCE IMPROVEMENT

Percent	Lyapunov-based/baseline			RL/baseline		
	backstepping	P	PI	outlet RL	inlet RL	outlet&inlet RL
J_{TTT}	1.6%	1.5%	1.5%	1.4%	1.6%	1.4%
J_{fuel}	3.4%	2.5%	3.3%	3.9%	3.6%	4.0%
$J_{comfort}$	30.6%	37.3%	30.1%	48.3%	47.8%	47.9%

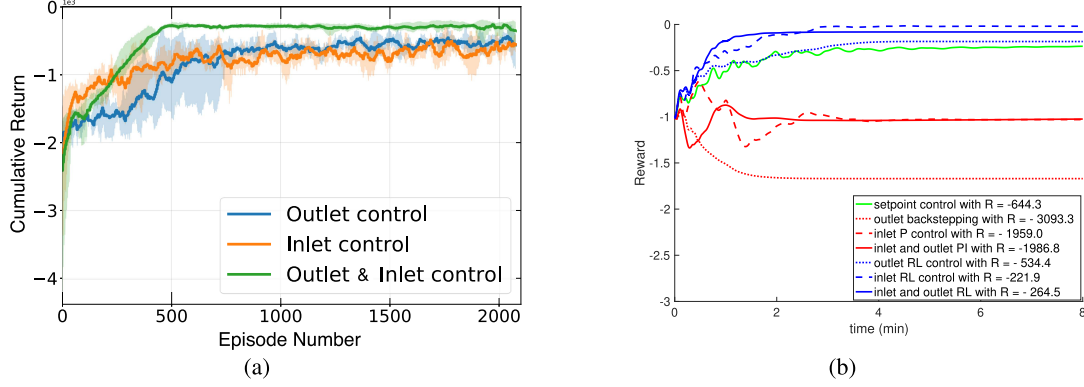


Fig. 9. Learning curves for RL and reward comparison with the Lyapunov controllers in Scenario 1. (a) RL learning curve with different initializations of actor-critic network, for randomly generated steady-state values. (b) In the system with partial knowledge, the reward evolution of the Lyapunov-based approaches (red), corresponding RL approaches (blue), and setpoint control (green).

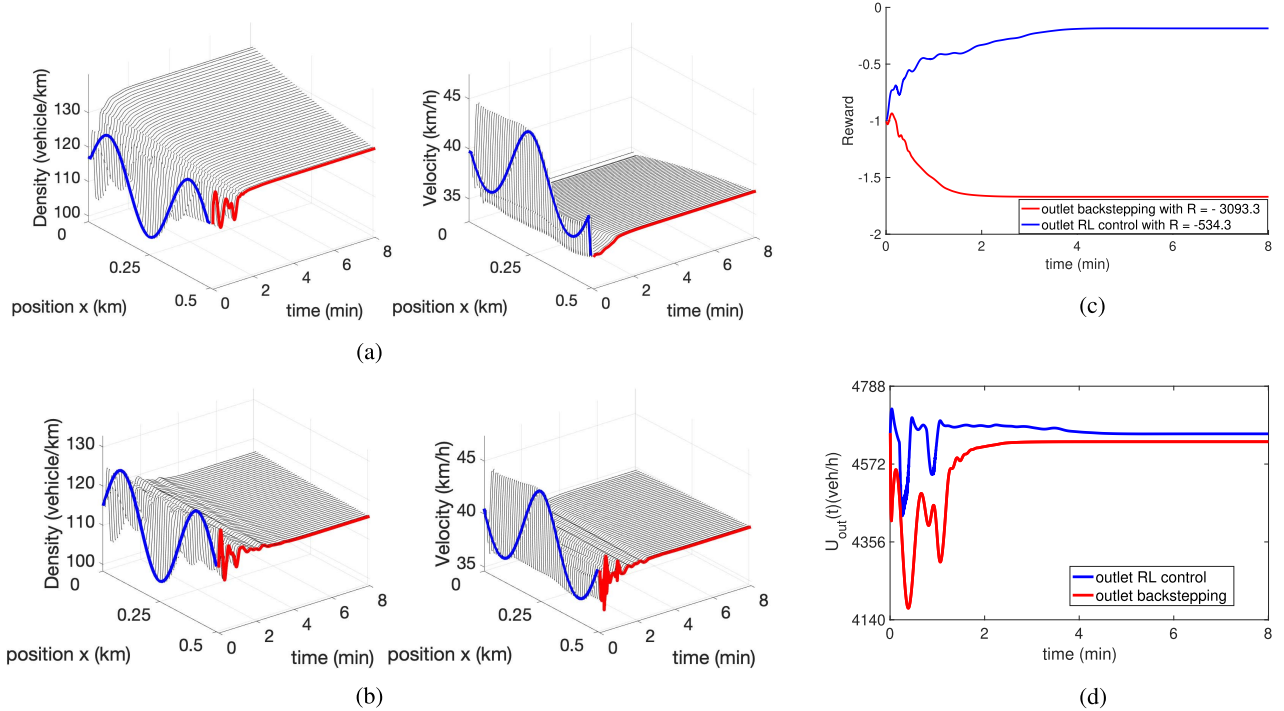


Fig. 10. Density and velocity evolution, reward evolution and outlet control inputs in Scenario 1. (a) Closed-loop with outlet backstepping controller in a lighter traffic. (b) Closed-loop with outlet RL controller in a lighter traffic. (c) Reward evolution for closed-loop with backstepping and RL controller at outlet in a lighter traffic. (d) Outlet control inputs for backstepping and RL controller.

than the backstepping controller $R = -3093.3$ in Fig. 10(c). Comparing the backstepping control input and RL outlet control input in Fig. 10(d), we find out that backstepping applies less control effort than RL. Backstepping's discharging flow

rate is less than the actual steady-state flow rate in the segment, thus unnecessarily holding vehicles back from leaving the segment and thus exacerbating traffic congestion. As shown in Fig. 10(a), traffic congestion forms and propagates backward

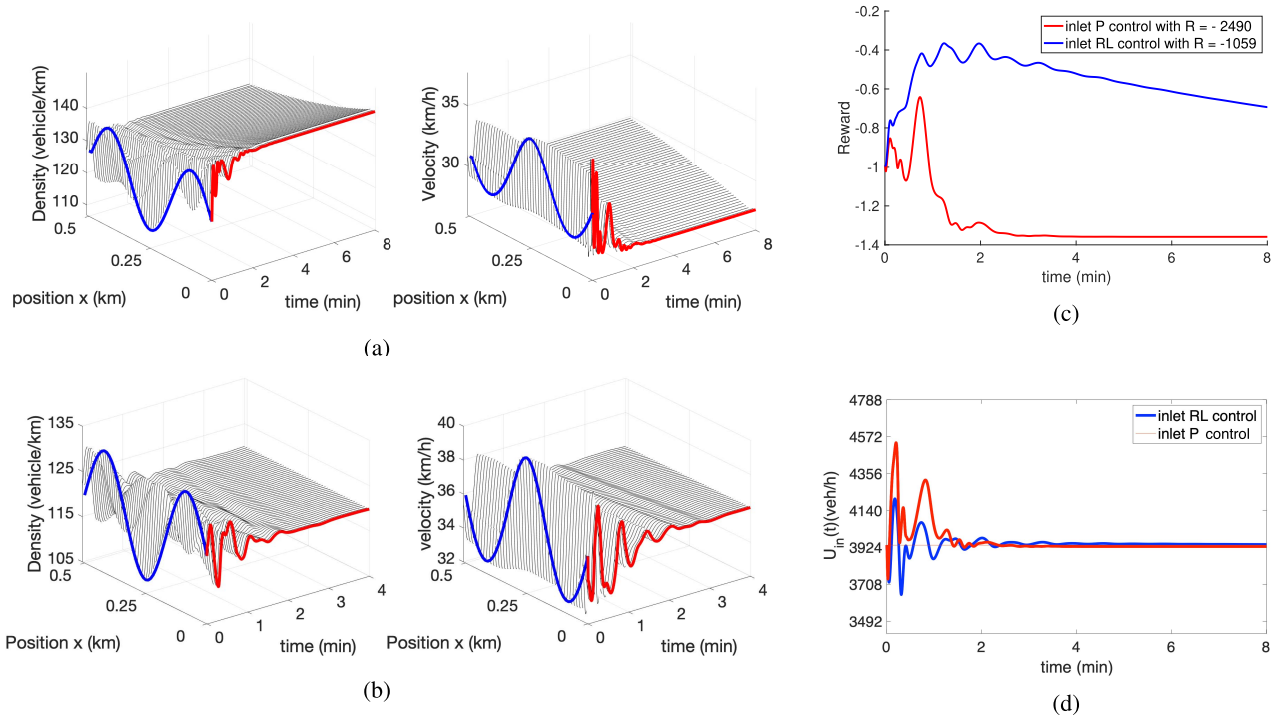


Fig. 11. Density and velocity evolution, reward evolution and inlet control inputs in Scenario 2. (a) Closed-loop with inlet P controller **in a denser traffic**. (b) Closed-loop with inlet RL controller **in a denser traffic**. (c) Reward evolution for closed-loop with P and RL controller at inlet **in a denser traffic**. (d) P and RL control inputs at inlet **in a denser traffic**.

to the segment inlet, whereas there is no such congestion in Fig. 10(b).

2) *Scenario 2 of Denser in-Domain Traffic*: In Fig. 11, we apply the inlet P controller to a nonlinear ARZ model with denser traffic $\rho_r = 125$ veh/km than the controller assumes, i.e., $\rho_r > \rho^*$. The inlet controller in (15) assumes $\rho^* = 120$ veh/km to design the control gain. As a result, the P controller applies a larger incoming traffic flow rate than it should, since it receives the feedback of slower velocity than the incorrect reference value. This creates congestion in the downstream traffic. As shown at the inlet of Fig. 11(a), the traffic density increases while velocity reduces. In contrast, the RL controller trained from the stochastic environment alleviates this situation in Fig. 11(c), and also exhibits an improved cumulative reward. However, the RL controller cannot stabilize the system to the true steady-state ρ_r , as shown by the blue reward initially increasing and then decreasing. Fig. 11(d) describes the evolution of the P and RL inlet control inputs. The incoming flow rate is eventually regulated by the inlet controllers to the real steady-state value, at the inlet. However, neither the RL controller nor the P controller successfully stabilizes the states to their true steady-state values across the segment, as highlighted in red Fig. 11(a) and (c). This is also demonstrated by the reward evolution in Fig. 11(c). However, the RL controller does not diverge as quickly, and in this sense is more robust.

In both scenarios, the performance of the Lyapunov-based controllers deteriorates with small errors to the assumed steady-state conditions. Comparing the state and reward evolution with outlet control and perfect system knowledge in

Figs. 5(a) and (c) and 10(a) and (c) in Scenario 1, with the same comparison between Figs. 6(a) and (c) and 11(a) and (c) in Scenario 2, the Lyapunov-based control systems do not converge to the actual steady states and the corresponding rewards do not converge to 0. On the other hand, the RL controllers trained in a stochastic environment perform worse in the deterministic setting, but are more robust to steady-state errors compared to the Lyapunov-based controllers, as shown in Fig. 10 for Scenario 1 and in Fig. 11. It should also be noted in Fig. 11(b) and (c) that RL controllers do not guarantee convergence of the closed-loop system, despite some improvements of the reward.

V. CONCLUSION

In this work, we address the freeway traffic control problem with an RL paradigm. We are motivated by two challenges currently unaddressed in the literature: 1) controllers natively designed for nonlinear traffic dynamics do not exist, and existing stabilization results are local and derived with linearized systems and 2) model-based controllers depend on accurate models. However, even the state-of-the-art ARZ second-order PDE model does not capture the complex realities of true real-world traffic flow, making model calibration a significant challenge. Even if one does successfully calibrate an ARZ model on a dataset, the parameters can shift over time due to weather, evolving commute patterns, discrete events, and more. RL methods provide an opportunity to address this practical challenge.

Ramp metering boundary controllers are developed to stabilize the freeway traffic oscillations. The second-order ARZ

PDE model is employed to describe the traffic dynamics. We formulate the control of the hyperbolic PDE model as an RL problem, where the control objective is encoded into the reward function. RL controllers are designed for various boundary control schemes, along with corresponding Lyapunov-based feedback control strategies. The nonlinear ARZ PDE model is used as the environment to simulate and generate data for RL. Among the RL algorithms, the actor-critic scheme is considered, and specifically the PPO algorithm has been adopted. The closed-loop performance of RL controllers is compared against Lyapunov-based state-feedback controllers, including PDE backstepping, P, and PI controllers.

Simulation results show that RL controllers achieve comparable performance with the Lyapunov-based controllers in a traffic system with perfect knowledge of model parameters, with slightly longer convergence times. The analysis becomes more interesting when we consider the reality of model uncertainty and imperfect knowledge. Specifically, the RL controllers outperform the Lyapunov-based controllers when the assumed steady-state conditions contain errors. Importantly, we find the Lyapunov-based controller loses its stabilizing properties, whereas the RL controllers obtained from a stochastic training process maintain stability or at least slow down divergence. Nevertheless, the RL controllers are obtained from thousands of iterative training episodes, and reward convergence is not guaranteed. For practical considerations, RL controllers should be developed in a simulation environment where trial and error are allowed.

One of the main advantages of the proposed RL formulation is that it does not require any prior knowledge of the traffic model, besides a simulator, but can achieve comparative and possibly better closed-loop performance compared to a Lyapunov-based approach. Future work may validate the RL-based controller on a freeway network comprised of several segments, merging and diverging lanes. The traffic performance on the ramp such as total waiting time should be incorporated into the evaluation of the mainline traffic. Another direction is to utilize the Lyapunov-based approach with stability guarantees to improve the RL controllers to obtain faster learning rates and more robust performance.

CODE AVAILABILITY

All the data and experiment codes are made available online. <https://github.com/saehong/RL-ARZ-PPO>

ACKNOWLEDGMENT

The authors would like to thank Eugene Vinitsky and Aboudy Kreidieh at the University of California at Berkeley for their insights and help at the beginning of this work.

REFERENCES

- [1] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 22–31.
- [2] I. Arel, C. Liu, T. Urbanik, and A. G. Kohls, "Reinforcement learning-based multi-agent system for network traffic signal control," *IET Intell. Transp. Syst.*, vol. 4, no. 2, pp. 128–135, 2010.
- [3] 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.
- [4] J. Aurial and F. Di Meglio, "Robust output feedback stabilization for two heterodirectional linear coupled hyperbolic PDEs," *Automatica*, vol. 115, May 2020, Art. no. 108896.
- [5] N. Bekiaris-Liberis and M. Krstic, "Compensation of actuator dynamics governed by quasilinear hyperbolic PDEs," *Automatica*, vol. 92, pp. 29–40, Jun. 2018.
- [6] T. Bellemans, B. D. Schutter, and B. D. Moor, "Model predictive control for ramp metering of motorway traffic: A case study," *Control Eng. Pract.*, vol. 14, no. 7, pp. 757–767, 2006.
- [7] F. Belletti, M. Huo, X. Litrico, and A. M. Bayen, "Prediction of traffic convective instability with spectral analysis of the Aw–Rascle–Zhang model," *Phys. Lett. A*, vol. 379, no. 38, pp. 2319–2330, Oct. 2015.
- [8] F. Belletti, D. Haziza, G. Gomes, and A. M. Bayen, "Expert level control of ramp metering based on multi-task deep reinforcement learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 4, pp. 1198–1207, Apr. 2018.
- [9] D. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA, USA: Athena Scientific, 2005.
- [10] R. C. Carlson, I. P. M. Papamichail, and A. Messmer, "Optimal main-stream traffic flow control of large-scale motorway networks," *Transp. Res. C, Emerg. Technol.*, vol. 18, no. 2, pp. 193–212, 2010.
- [11] C. F. Daganzo, "On the macroscopic stability of freeway traffic," *Transp. Res. B, Methodol.*, vol. 45, no. 5, pp. 782–788, Jun. 2011.
- [12] N. Espitia, H. Yu, and M. Krstic, "Event-triggered varying speed limit control of stop-and-go traffic," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 7509–7514, 2020. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02931407>
- [13] S. El-Tantawy, B. Abdulhai, and H. Abdelgawad, "Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC): Methodology and large-scale application on downtown Toronto," *IEEE Trans. Intell. Transp. Syst.*, vol. 14, no. 3, pp. 1140–1150, Sep. 2013.
- [14] S. Fan and B. Seibold, "Data-fitted first-order traffic models and their second-order generalizations: Comparison by trajectory and sensor data," *Transp. Res. Rec., J. Transp. Res. Board*, vol. 2391, no. 1, pp. 32–43, Jan. 2013.
- [15] M. R. Flynn, A. R. Kasimov, J.-C. Nave, R. R. Rosales, and B. Seibold, "Self-sustained nonlinear waves in traffic flow," *Phys. Rev. E, Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top.*, vol. 79, no. 5, May 2009, Art. no. 056113.
- [16] A.-M. Farahmand, S. Nabi, and D. N. Nikovski, "Deep reinforcement learning for partial differential equation control," in *Proc. Amer. Control Conf. (ACC)*, May 2017, pp. 3120–3127.
- [17] A. Ferrara, G. P. Incremona, and G. Piacentini, "A hierarchical MPC and sliding mode based two-level control for freeway traffic systems with partial demand information," *Eur. J. Control*, vol. 59, pp. 152–164, May 2021.
- [18] Y. Pan, A.-M. Farahmand, M. White, S. Nabi, P. Grover, and D. Nikovski, "Reinforcement learning with function-valued action spaces for partial differential equation control," 2018, *arXiv:1806.06931*. [Online]. Available: <http://arxiv.org/abs/1806.06931>
- [19] FHWA, U.S. Department of Transportation. *Next Generation Simulation (NGSIM)*. Accessed: Jul. 30, 2018. [Online]. Available: <http://ops.fhwa.dot.gov/trafficanalysis/tools/ngsim.html>
- [20] P. Frihauf, M. Krstić, and T. Başar, "Finite-horizon LQ control for unknown discrete-time linear systems via extremum seeking," *Eur. J. Control*, vol. 19, pp. 399–407, Sep. 2013.
- [21] J. García and F. Fernández, "A comprehensive survey on safe reinforcement learning," *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 1437–1480, 2015.
- [22] G. Gomes and R. Horowitz, "Optimal freeway ramp metering using the asymmetric cell transmission model," *Transp. Res. C, Emerg. Technol.*, vol. 14, no. 4, pp. 244–262, 2006.
- [23] J. Han, E. J. Weinan, and A. Jentzen, "Solving high-dimensional partial differential equations using deep learning," *Proc. Nat. Acad. Sci. USA*, vol. 115, no. 34, pp. 8505–8510, 2018.
- [24] S. Kakade and J. Langford, "Approximately optimal approximate reinforcement learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2002, pp. 267–274.
- [25] I. Karafyllis, M. Kontorinaki, and M. Papageorgiou, "Robust global adaptive exponential stabilization of discrete-time systems with application to freeway traffic control," *IEEE Trans. Autom. Control*, vol. 62, no. 12, pp. 6195–6208, Dec. 2017.

- [26] I. Karafyllis, N. Bekiaris-Liberis, and M. Papageorgiou, "Feedback control of nonlinear hyperbolic PDE systems inspired by traffic flow models," *IEEE Trans. Autom. Control*, vol. 64, no. 9, pp. 3647–3662, Sep. 2019.
- [27] I. Karafyllis and M. Papageorgiou, "Feedback control of scalar conservation laws with application to density control in freeways by means of variable speed limits," *Automatica*, vol. 105, pp. 228–236, Jul. 2019.
- [28] M. Krstic and A. Smyshlyaev, *Boundary Control of PDEs: A Course on Backstepping Designs*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2008.
- [29] S. Z. Khong, D. Nesic, and M. Krstic, "A non-model based extremum seeking approach to iterative learning control," *Automatica*, vol. 66, pp. 238–245, Apr. 2016.
- [30] N. Killingsworth and M. Krstic, "PID tuning using extremum seeking," *Control Syst. Mag.*, vol. 26, pp. 70–79, Feb. 2006.
- [31] I. Kostrikov, *PyTorch Implementations of Reinforcement Learning Algorithms*. Github Repository. Accessed: Jan. 1, 2021. [Online]. Available: <https://github.com/p-christ/Deep-Reinforcement-Learning-Algorithms-with-PyTorch>
- [32] P. Lax and B. Wendroff, "Systems of conservation laws," *Commun. Pure Appl. Math.*, vol. 13, no. 2, pp. 217–237, 1960.
- [33] L. Li, Y. Lv, and F.-Y. Wang, "Traffic signal timing via deep reinforcement learning," *IEEE/CAA J. Automat. Sinica*, vol. 3, no. 3, pp. 247–254, Apr. 2016.
- [34] E. Liang *et al.*, "RLlib: Abstractions for distributed reinforcement learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Jul. 2018, pp. 3053–3062. [Online]. Available: [arXiv:1712.09381](https://arxiv.org/abs/1712.09381)
- [35] T. Liard, I. Balogoun, S. Marx, and F. Plestan, "Boundary sliding mode control of a system of linear hyperbolic equations: A Lyapunov approach," Tech. Rep. hal-03084995, 2020.
- [36] M. J. Lighthill and G. B. Whitham, "On kinematic waves: II. A theory of traffic flow on long crowded roads," *Proc. Roy. Soc. London. Ser. A, Math. Phys. Sci.*, vol. 229, no. 1178, pp. 317–345, 1955.
- [37] M. Papageorgiou, J.-M. Blosseville, and H. Hadj-Salem, "Modelling and real-time control of traffic flow on the southern part of boulevard peripherique in Paris: Part I: Modelling," *Transp. Res. A, Gen.*, vol. 24, no. 5, pp. 345–359, Sep. 1990.
- [38] M. Papageorgiou, H. Hadj-Salem, and J. M. Blosseville, "ALINEA: A local feedback control law for on-ramp metering," *Transp. Res. Rec.*, vol. 1320, no. 1, pp. 58–67, 1991.
- [39] M. S. Radenkovic and M. Krstic, "Adaptive control via extremum seeking: Global stabilization and consistency of parameter estimates," *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2350–2359, May 2017.
- [40] P. I. Richards, "Shock waves on the highway," *Oper. Res.*, vol. 4, no. 1, pp. 42–51, 1956.
- [41] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Jun. 2015, pp. 1889–1897.
- [42] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*. [Online]. Available: [http://arxiv.org/abs/1707.06347](https://arxiv.org/abs/1707.06347)
- [43] B. Seibold, M. R. Flynn, A. R. Kasimov, and R. R. Rosales, "Constructing set-valued fundamental diagrams from jamiton solutions in second order traffic models," 2012, *arXiv:1204.5510*. [Online]. Available: <http://arxiv.org/abs/1204.5510>
- [44] S. Siri, C. Pasquale, S. Saccone, and A. Ferrara, "Freeway traffic control: A survey," *Automatica*, vol. 130, Aug. 2021, Art. no. 109655.
- [45] R. Sutton, D. McAllester, S. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2000, pp. 1057–1063.
- [46] S. Tang and M. Krstic, "Sliding mode control to the stabilization of a linear 2×2 hyperbolic system with boundary input disturbance," in *Proc. Amer. Control Conf.*, Jun. 2014, pp. 1027–1032.
- [47] M. Treiber and A. Kesting, *Traffic Flow Dynamics: Data, Models and Simulation*. Berlin, Germany: Springer, 2013.
- [48] E. Vinitzky, K. Parvate, A. Kreidieh, C. Wu, and A. Bayen, "Lagrangian control through deep-RL: Applications to bottleneck decongestion," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 759–765.
- [49] Y. Wang, Z. Shen, Z. Long, and B. Dong, "Learning to discretize: Solving 1D scalar conservation laws via deep reinforcement learning," 2019, *arXiv:1905.11079*. [Online]. Available: [http://arxiv.org/abs/1905.11079](https://arxiv.org/abs/1905.11079)
- [50] Y. Wang, E. B. Kosmatopoulos, M. Papageorgiou, and I. Papamichail, "Local ramp metering in the presence of a distant downstream bottleneck: Theoretical analysis and simulation study," *IEEE Trans. Intell. Transp. Syst.*, vol. 15, no. 5, pp. 2024–2039, Oct. 2014.
- [51] C. Wu, A. Kreidieh, E. Vinitzky, and A. M. Bayen, "Emergent behaviors in mixed-autonomy traffic," in *Proc. Conf. Robot Learn.*, 2017, pp. 398–407.
- [52] H. Yu and M. Krstic, "Traffic congestion control for Aw-Rascle-Zhang model," *Automatica*, vol. 100, pp. 38–51, Feb. 2019.
- [53] H. Yu and M. Krstic, "Varying speed limit control of Aw-Rascle-Zhang traffic model," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 1846–1851.
- [54] H. Yu, Q. Gan, A. Bayen, and M. Krstic, "PDE traffic observer validated on freeway data," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 3, pp. 1048–1060, May 2021, doi: [10.1109/TCST.2020.2989101](https://doi.org/10.1109/TCST.2020.2989101).
- [55] H. Yu, M. Diagne, L. Zhang, and M. Krstic, "Bilateral boundary control of moving shockwave in LWR model of congested traffic," *IEEE Trans. Autom. Control*, vol. 66, no. 3, pp. 1429–1436, Mar. 2021.
- [56] H. Yu, S. Koga, T. R. Oliveira, and M. Krstic, "Extremum seeking for traffic congestion control with a downstream bottleneck," *J. Dyn. Syst., Meas., Control*, vol. 143, no. 3, Mar. 2021, Art. no. 031007, doi: [10.1115/1.4048781](https://doi.org/10.1115/1.4048781).
- [57] 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.
- [58] L. Zhang, C. Prieur, and J. Qiao, "PI boundary control of linear hyperbolic balance laws with stabilization of ARZ traffic flow models," *Syst. Control Lett.*, vol. 123, pp. 85–91, Jan. 2019.