

Optimal Control of Autonomous Vehicles for Flow Smoothing in mixed-autonomy Traffic

Arwa Alanqary, Xiaoqian Gong, Alexander Keimer, Benjamin Seibold, Benedetto Piccoli, Alexandre Bayen

Abstract—This article studies the optimal control of autonomous vehicles over a given time horizon to smooth traffic. We model the dynamics of a mixed-autonomy platoon as a system of non-linear ODEs, where the acceleration of human-driven vehicles is governed by a car-following model, and the acceleration of autonomous vehicles is to be controlled. We formulate the car-following task as an optimal control problem and propose a computational method to solve it. Our approach uses an adjoint formulation to compute gradients of the optimization problem explicitly, resulting in more accurate and efficient numerical computations. The gradients are then used to solve the problem using gradient-based optimization solvers. We consider an instance of the problem with the objective of improving the fuel efficiency of the vehicles in the platoon. The effectiveness of the proposed approach is demonstrated through numerical experiments. We apply the proposed approach to different scenarios of lead vehicle trajectories and platoon sizes. The results suggest that introducing an AV can produce significant energy savings for the platoon. It also reveals that the solution is agnostic to the platoon size thus the fuel saving is mainly due to optimizing the trajectory of the AV.

I. INTRODUCTION

Car-following in mixed-autonomy traffic. Mixed-autonomy traffic is a system in which a fraction of the vehicles are automated (AVs) and they share the road with human-driven vehicles (HVs). The introduction of AVs in such predominantly human-driven traffic has the potential to improve traffic conditions locally and system-wide [1], [2]. This work focuses on the task of single-lane car-following, a fundamental challenge in autonomous driving.

It has been demonstrated, through simulations and field experiments, that human car-following behavior is sub-optimal. In the renowned ring experiment, Sugiyama et al. [3] showed the emergence of stop-and-go waves in a circular platoon due to human behavior. This was further investigated in [1], [4].

This material is based upon work supported by the National Science Foundation under Grant CNS-1837244. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. This material is also based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Vehicle Technologies Office award number CID DE-EE0008872. The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Arwa Alanqary and Alexandre Bayen are with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, USA. arwa@berkeley.edu, bayen@berkeley.edu

Xiaoqian Gong is with the School of Mathematical and Statistical Sciences, Arizona State University, Tempe, USA. xgong14@asu.edu

Alexander Keimer is with the Department of Mathematics, FAU Erlangen-Nuremberg, Germany. keimer@math.fau.de

Benjamin Seibold is with the Department of Mathematics, Temple University, Philadelphia, USA. seibold@temple.edu

Benedetto Piccoli is with the Department of Mathematical Sciences, Rutgers University, Camden, USA. piccoli@camden.rutgers.edu

Numerous studies have proposed control strategies to enhance car-following behavior in human-driven traffic by integrating AVs. We refer to some earlier work on intelligent cruise and platoon control in [5]–[9] and more recent developments in [1], [10]–[14], among many others. These controllers have shown remarkable success in attenuating stop-and-go waves and improving traffic conditions. For example, Stern et al. [1] experimentally demonstrated that a single AV equipped with a proportional-integral controller can dissipate stop-and-go waves in a circular platoon.

A key distinction between control strategies proposed in the literature is between local and non-local approaches. Local controllers rely on data from the AV's surroundings, while non-local controllers leverage information about upstream and downstream traffic. Non-local controllers can preemptively react to downstream oscillations, smoothing a broader range of stop-and-go patterns. These controllers can also incorporate the behavior of trailing vehicles and operate as actuators to improve the overall system performance.

In this work, we aim to create a baseline for performance benchmarking for the setting of non-local controllers. This baseline is set up with full knowledge of the downstream data (eliminating the effect of forecasting error) as well as the AV's ability to sense the trailing vehicles.

With this goal in mind, we formulate the task of car-following in mixed-autonomy traffic as an optimal control problem (OCP) over a finite time horizon and propose a method to solve it. We model the non-linear dynamics of a mixed-autonomy platoon where the AVs are controlled in their acceleration and the HVs' acceleration is governed by a car-following model (CFM), namely *Bando-Follow the leader* model. We formulate an objective functional that accounts for the AVs and all the HVs in the platoon.

Numerical solution of the optimal control problem. The literature on optimal control describes two primary solution approaches: optimize then discretize vs. discretize then optimize [15]. The first employs Pontryagin maximum principle (PMP) [16] to derive the necessary optimality conditions for the continuous-time problem, then discretizes and solves the resulting system of equations. The second class discretizes the OCP, rendering a finite-dimensional problem which is then solved by means of non-linear constrained optimization such as sequential quadratic programming (SQP).

In this work, we adopt the second approach (discretize then optimize) to solve the formulated OCP. While this approach can leverage existing non-linear optimization solvers, computing the objective functional's gradient proves to be non-trivial in many problems. For this, it is important to

provide the optimizer with accurate gradient information. We propose to use an adjoint-based approach to represent the objective gradient [17]. To this end, we derive the continuous-time adjoint equations corresponding to our OCP. Additionally, we parameterize the controls using piecewise constant functions and outline a suitable discretization scheme for both system dynamics and adjoint equations.

Related work. Optimal control is a prevalent tool for modeling and solving various tasks in the domain of automated vehicles, including car-following, trajectory optimization, and obstacle avoidance. The literature on this topic is extensive, but we will narrow our focus to advancements that employ comparable formulations and solution techniques to those introduced in this paper.

Wang et al. [18] studied the controllability and stabilizability of mixed-autonomy platoons using a single AV. They use the optimal velocity model (OVM) for both human-driven traffic and the AV's car following behavior. The AV has an additive acceleration that is being controlled using a state feedback controller. The controller design relies on the linearization of the system dynamics. Similar techniques were proposed for circular platoons [19], [20].

Wang et al. [21] formulated the problem of smoothing single-lane mixed-autonomy traffic flow through optimal control of AVs. The dynamics combine the intelligent driver model (IDM) for HVs with a modified OVM with positive additive acceleration for the AVs. In contrast, our work uses the Bando-FtL model which enjoys well-posedness and inherent safety guarantees, eliminating the need for collision avoidance constraints for the HVs. Without these properties, handling such constraints becomes intricate, which was not discussed in [21]. For the AVs, we employ a more representative control parametrization compared to using CFMs with additive acceleration as suggested in [18], [21]. While the latter approach relies on the car-following dynamics for collision avoidance, it comes at the expense of restricting the AVs' behavior. Further, with the objective function proposed in [21], which only penalizes the perturbations of the AVs' speed, this modeling choice can lead to unrealistic vehicle dynamics. This can occur when the additive acceleration results in an arbitrarily small headway requiring an unrealistically high deceleration by the CFM to avoid collision. In our approach, we explicitly handle the safety constraints leading to natural bounds on the vehicle's acceleration independently of the objective functional. Given these modeling choices, we employ a solution approach similar to [21], leveraging gradient descent and the adjoint representation of the gradient.

A related line of research applies model predictive control (MPC) to car-following tasks [14], [22], [23]. In these controllers, the planning layer is often cast as a receding horizon OCP with a single AV. However, these formulations do not model human traffic and only respond to a leader trajectory. This enables the formulation of the OCP as a linearly constrained quadratic program which can be solved efficiently making it well-suited for real-time applications.

Outline. In section II, we formulate the continuous-time problem of mixed-autonomy car-following. In section III,

we state the adjoint equations associated with the problem and outline the discretization scheme. In section IV, we instantiate the problem with a specific objective functional and present numerical experiments to illustrate the use of the proposed approach. Lastly, we discuss the results and analyze the implications of our main design choices in section V.

II. PROBLEM STATEMENT

We study the problem of car-following in mixed-autonomy traffic on a single lane with $M \in \mathbb{N}_{\geq 1}$ autonomous vehicles (AVs), $N \in \mathbb{N}_{\geq 1}$ human-driven vehicles (HVs), and a lead vehicle. The AVs are controlled in their acceleration, the HVs' dynamics are governed by a car-following model (CFM) specified later, and the lead vehicle's trajectory is prespecified on a time horizon $T \in \mathbb{R}_{>0}$. Let $I = \{0, 1, \dots, M + N\}$ be the index set of the vehicles with $i = 0$ being the lead vehicle. We denote by $I_a \subset I$ and $I_h \subset I$ the sets of indices of AVs and HVs, respectively. For any $i \in I$, we denote the index of the vehicle ahead to be $i - 1$. Let $(\mathbf{x}, \mathbf{v}) \in \mathbb{R}^{|I|} \times \mathbb{R}^{|I|}$ be the position-velocity vector of the vehicles and denote by $(\mathbf{x}_o, \mathbf{v}_o) \in \mathbb{R}^{|I|} \times \mathbb{R}_{\geq 0}^{|I|}$ their initial positions and velocities.

Assumption 1: (Initial datum, lead vehicle's trajectory, and admissible set of controls): Given the vehicle length $l \in \mathbb{R}_{>0}$, we assume that

- **Minimal initial space headway:** there exists $d_o \in \mathbb{R}_{\geq 0}$, such that, the initial space headway between any two consecutive vehicles is at least d_o . That is, for any $i \in \{1, \dots, M + N\}$,

$$\mathbf{x}_{o,i-1} - \mathbf{x}_{o,i} - l \geq d_o.$$

- **Prespecified lead vehicle's trajectory:** The trajectory of the lead vehicle is given explicitly by its position $x_\ell \in W^{2,\infty}((0, T))$ and velocity $v_\ell := \dot{x}_\ell \geq 0$.
- **Admissible set of controls:** Given acceleration bounds $a_{\min} \in \mathbb{R}_{<0}$, $a_{\max} \in \mathbb{R}_{>0}$, time headway bounds $h_{\min} \in \mathbb{R}_{\geq 0}$, $h_{\max} \in \mathbb{R}_{>h_{\min}}$, and bounds on the headway at zero velocity $d_{\min} \in \mathbb{R}_{\geq d_o}$, $d_{\max} \in \mathbb{R}_{>d_{\min}}$, we consider the following set of controls

$$\begin{aligned} \mathcal{U} := & \left\{ \mathbf{u} \in L^\infty((0, T); \mathbb{R}^{|I|}) : \mathbf{u}_i \equiv 0 \ \forall i \in I \setminus I_a; \right. \\ & \wedge \forall (t, i) \in [0, T] \times I_a : \mathbf{u}_i(t) \in [a_{\min}, a_{\max}], \\ & \mathbf{v}_{o,i} + \int_0^t \mathbf{u}_i(s) ds \geq 0, \ d_{\max} + \mathbf{v}_i(t) h_{\max} \geq \dots \\ & \left. \geq \mathbf{x}_{i-1}(t) - \mathbf{x}_i(t) - l \geq h_{\min} \mathbf{v}_i(t) + d_{\min} \right\}. \end{aligned} \quad (1)$$

Specifically, only the AVs' accelerations are controlled and they satisfy the following

- 1) the controlled accelerations are bounded;
- 2) the AVs do not drive backward;
- 3) there is no collision between the autonomous vehicles and their lead vehicles.
- 4) The distance between the autonomous cars and their leaders is an envelope that increases the faster the autonomous vehicle's speed becomes.

The set of admissible controls does not encode collision avoidance conditions for the HVs. This is handled by the CFM which will be discussed later in the section.

Given an acceleration function for the HVs, $A: \mathbb{R}^2 \times \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}$, the dynamics of the vehicles are governed by the following system of initial value problems for $t \in [0, T]$

$$\begin{aligned} \dot{\mathbf{x}}_0(t) &= \mathbf{x}_\ell(t) \\ \dot{\mathbf{v}}_0(t) &= \mathbf{v}_\ell(t) \\ \dot{\mathbf{x}}_i(t) &= \mathbf{v}_i(t), & i \in I \setminus \{0\} \\ \dot{\mathbf{v}}_i(t) &= A(\mathbf{x}_i(t), \mathbf{x}_{i-1}(t), \mathbf{v}_i(t), \mathbf{v}_{i-1}(t)) & i \in I_h \\ \dot{\mathbf{v}}_i(t) &= \mathbf{u}_i(t) & i \in I_a \\ \mathbf{x}(0) &= \mathbf{x}_o \\ \mathbf{v}(0) &= \mathbf{v}_o, \end{aligned} \quad (2)$$

where $\mathbf{x}_{o,0} = \mathbf{x}_\ell(0)$ and $\mathbf{v}_{o,0} = \mathbf{v}_\ell(0)$ are the lead vehicle's initial position and velocity.

Next, we instantiate the considered Bando-FtL CFM described in [10]. Given model parameters $(\alpha, \beta) \in \mathbb{R}_{>0}^2$, an optimal velocity function $V \in C^1(\mathbb{R}_{\geq 0}; \mathbb{R}_{\geq 0}) \cap L^\infty(\mathbb{R}_{\geq 0}; \mathbb{R}_{\geq 0})$ satisfying $V' > 0$, and a vehicle length $l \in \mathbb{R}_{>0}$, the Bando-FtL acceleration is

$$A(x_\ell, x, v_\ell, v) = \alpha(V(x_\ell - x - l) - v) + \beta \frac{v_\ell - v}{(x_\ell - x - l)^2}. \quad (3)$$

The ego vehicle's acceleration depends on the optimal velocity, space headway, and relative velocity. The Bando-FtL model is well-posed with a minimum safety distance and bounded velocity and acceleration. For more details, please refer to [24], where the well-posedness is shown for any leader's trajectory whenever there is no backward driving.

Now we define the cost functional as follows:

Definition 2.1 (Cost functional):

Let $L \in C^1(\mathbb{R}^{3|I|}; \mathbb{R}_{\geq 0})$ and $S \in C^1(\mathbb{R}^{2|I|}; \mathbb{R}_{\geq 0})$ be the running and terminal costs, respectively. The cost functional $J: L^\infty((0, T); \mathbb{R}^{|I|}) \times L^\infty((0, T); \mathbb{R}^{|I|}) \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ is defined as

$$J(\mathbf{x}[\mathbf{u}], \mathbf{v}[\mathbf{u}], \mathbf{u}) := \int_0^T L(\mathbf{x}[\mathbf{u}](t), \mathbf{v}[\mathbf{u}](t), \mathbf{u}(t)) dt + S(\mathbf{x}[\mathbf{u}](T), \mathbf{v}[\mathbf{u}](T)). \quad (4)$$

We aim to find a $\mathbf{u}^* \in \mathcal{U}$ that minimizes the cost functional in eq. (4). That is, we consider:

$$\mathbf{u}^* \in \arg\text{-inf}_{\mathbf{u} \in \mathcal{U}} J(\mathbf{x}[\mathbf{u}], \mathbf{v}[\mathbf{u}], \mathbf{u}), \quad (5)$$

where the vector $(\mathbf{x}[\mathbf{u}], \mathbf{v}[\mathbf{u}])$ satisfies the system in eq. (2).

Note that this notation emphasizes that \mathbf{x} and \mathbf{v} are functions of \mathbf{u} , which are guaranteed to exist by the well-posedness of the used CFM as well as the set of admissible controls. As we will be dealing with the ODE constraints by suitable Lagrange-multipliers (the adjoint method), we will skip the explicit dependency of \mathbf{x} and \mathbf{v} on the control \mathbf{u} .

Remark 1 (Existence of Minimizers): In contrast to what some literature [21, Theorem 3.1] claims, the set \mathcal{U} is not compact in any strong topology as we lack a uniform bound on its elements' derivatives (for example, the Arzela-Ascoli theorem). However, the lower and upper bounds on the admissible control result in the compactness of the set with

regard to the weak star topology. If the cost functional is weakly lower semi-continuous with regard to the control, one can indeed apply the typical existence of minimizer proof. This is because the velocities and positions are time integrations of the corresponding controls, and thus, the weak star convergence of the corresponding minimizing sequence of controls results in the strong (uniform) convergence of velocities and speed. As the car-following models only take these quantities as input, one can then rely on the continuity results of these. Another fact worth mentioning is that one needs the car-following models to be well-posed uniformly with regard to the chosen controls. Also, this is not a trivial fact, as the intelligent driver model (IDM) for instance does not possess this property [25] while the Bando-FtL model used here has been proven to satisfy this requirement [24].

Here we formulate and study the OCP on a single lane without lane-changing. While our proposed methods could potentially be extended to lane-change scenarios (by considering hybrid dynamical systems), this requires access to suitable lane-changing models that are beyond the scope of this work.

III. METHODOLOGY

In this section, we describe a numerical solution approach of the OCP in eq. (5). Firstly, we present the continuous-time adjoint equations and use them to write a representation of the gradient of the objective functional in eq. (4). Secondly, we convert the continuous-time OCP into a finite-dimensional optimization problem through control parameterization. We then solve the resulting optimization problem using a gradient-based solver.

A. Adjoint System

We follow the general derivation of the adjoint representation of the gradient described in [17]. We first write the implicit form of the system dynamics from eq. (2). For $t \in [0, T]$,

$$\mathbf{F}[\dot{\mathbf{x}}, \dot{\mathbf{v}}, \mathbf{x}, \mathbf{v}, \mathbf{u}](t) = \left[\begin{array}{c} \dot{\mathbf{x}}_i(t) - \mathbf{v}_i(t) \\ \mathbf{f}_i[\mathbf{x}, \mathbf{v}, \mathbf{u}](t) \end{array} \right]_{i \in I \setminus \{0\}}, \quad (6)$$

where

$$\mathbf{f}_i[\mathbf{x}, \mathbf{v}, \mathbf{u}](t) = \begin{cases} \dot{\mathbf{v}}_i(t) - A(\mathbf{x}_i(t), \mathbf{x}_{i-1}(t), \mathbf{v}_i(t), \mathbf{v}_{i-1}(t)), & i \in I_h \\ \dot{\mathbf{v}}_i(t) - \mathbf{u}_i(t) & i \in I_a. \end{cases}$$

Then, we introduce a vector of adjoint variables $\zeta \in L^\infty((0, T); \mathbb{R}^{2(M+N)})$. We can now write the adjoint system of equations $\forall t \in [0, T]$ as

$$\begin{aligned} \dot{\zeta}(T) &= S_{(\mathbf{x}, \mathbf{v})}(\mathbf{x}(T), \mathbf{v}(T)) \\ \dot{\zeta}(t) &= \mathbf{F}_{(\mathbf{x}, \mathbf{v})}^\top[\dot{\mathbf{x}}, \dot{\mathbf{v}}, \mathbf{x}, \mathbf{v}, \mathbf{u}](t) \zeta(t) - L_{(\mathbf{x}, \mathbf{v})}(\mathbf{x}(t), \mathbf{v}(t), \mathbf{u}(t)) \end{aligned} \quad (7)$$

Note that the dynamics in eq. (2) are only coupled in one direction, that is the state of a vehicle $i \in I$ depends only on its preceding vehicle $i - 1$. This makes the time-dependent Jacobian matrix $\mathbf{F}_{(\mathbf{x}, \mathbf{v})}[\dot{\mathbf{x}}, \dot{\mathbf{v}}, \mathbf{x}, \mathbf{v}, \mathbf{u}](t)$ have a

sparse structure of the following form. For $t \in [0, T]$,

$$\mathbf{F}_{(\mathbf{x}, \mathbf{v})}[\dot{\mathbf{x}}, \dot{\mathbf{v}}, \mathbf{x}, \mathbf{v}, \mathbf{u}](t) = \begin{bmatrix} B(t) & C(t) \\ D(t) & E(t) \end{bmatrix}, \quad (8)$$

where $B(t), C(t), D(t), E(t) \in \mathbb{R}^{(M+N) \times (M+N)}$, $B(t) = \mathbf{0}$, $C(t) = -I$, and $D(t)$ and $E(t)$ are lower bidiagonal matrices. The entries of matrix D are given as follows:

$$D_{i,i} = \begin{cases} -A_{\mathbf{x}_i}(\mathbf{x}_i, \mathbf{x}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i-1}) & \forall i \in I_h \\ 0 & \forall i \in I_a, \end{cases} \quad (9)$$

$$D_{i+1,i} = \begin{cases} -A_{\mathbf{x}_i}(\mathbf{x}_{i+1}, \mathbf{x}_i, \mathbf{v}_{i+1}, \mathbf{v}_i) & \forall i \in I_h \setminus \{M+N\} \\ 0 & \forall i \in I_a \setminus \{M+N\} \end{cases} \quad (10)$$

The entries of matrix E are given as follows:

$$E_{i,i} = \begin{cases} -A_{\mathbf{v}_i}(\mathbf{x}_i, \mathbf{v}_{i-1}, \mathbf{v}_i, \mathbf{v}_{i-1}) & \forall i \in I_h \\ 0 & \forall i \in I_a, \end{cases} \quad (11)$$

$$E_{i+1,i} = \begin{cases} -A_{\mathbf{v}_i}(\mathbf{x}_{i+1}, \mathbf{x}_i, \mathbf{v}_{i+1}, \mathbf{v}_i) & \forall i \in I_h \setminus \{M+N\} \\ 0 & \forall i \in I_a \setminus \{M+N\} \end{cases} \quad (12)$$

The linear ODEs in eq. (7) describe a system of end value problems. The conditions at the end time T are defined by the terminal cost function $S(\mathbf{x}(T), \mathbf{v}(T))$. Thus, the adjoint system can be integrated backward from time T to zero.

The solution to the above system yields the following representation of the objective gradient in continuous-time. For $t \in [0, T]$,

$$\mathbf{G}_i(t) = \begin{cases} L_{\mathbf{u}_i}(\mathbf{x}(t), \mathbf{v}(t), \mathbf{u}(t)) - \zeta_{i+M+N}(t) & i \in I_a \\ 0 & i \in I_h. \end{cases} \quad (13)$$

We note that the OCP in eq. (5) has path inequality constraints encoded in the admissible control set \mathcal{U} . These constraints are not considered in the above adjoint equations and are to be discretized and handled by the optimizer. For a single autonomous vehicle, these constraints are affine in the control, so they are easy to handle by typical constrained optimization solvers.

B. Discretization scheme

We detail the discretization scheme used to convert the OCP in eq. (5) into a finite-dimensional problem.

Control and constraints discretization. The infinite-dimensional control space $L^\infty((0, T); \mathbb{R}^{|I|})$ is discretized into the space of piecewise constant functions over the grid

$$\mathcal{C}_K = \{\tau_k = kh_c, k \in \{0, 1, \dots, K\}\}, \quad (14)$$

with mesh size $h_c = \frac{T}{K}$ and $K+1 \in \mathbb{N}_{\geq 1}$ discretization points. This controller space is parameterized by the vectors $\mathbf{w}_i \in \mathbb{R}^{K+1}$ for each $i \in I_a$. Furthermore, for $\forall i \in I_a$, we denote the piecewise constant control function

$$\hat{\mathbf{u}}_i(t; \mathbf{w}) = \sum_{k=0}^{K-1} \chi_{[kh_c, (k+1)h_c]}(t) \mathbf{w}_{i,k},$$

where χ_Z is the characteristic function of a set $Z \subset \mathbb{R}$. This piecewise constant parametrization is a common choice due to its simplicity and convergence properties [26]. However, one could consider piecewise linear (and continuous)

parametrization or other more regular basis functions. We note that our parametrization is more representative than an alternative proposed in [21]. There, the control function consists of a car-following model and a small additive acceleration. This parametrization significantly restricts the behavior of the AVs and it fails to converge to the solution of the continuous-time OCP.

The inequality constraints in the admissible control set in eq. (1) are also evaluated on the grid \mathcal{C}_K . We use a third-order explicit *Runge-Kutta* (RK-3) integration scheme with fixed mesh size to solve the dynamics in eq. (2). This results in the approximate state vector for each $i \in I \setminus \{0\}$

$$(\tilde{\mathbf{x}}_i(\tau_k), \tilde{\mathbf{v}}_i(\tau_k))_{\tau_k \in \mathcal{C}_K} \in \mathbb{R}^{2(K+1)} \quad (15)$$

Consequently, we write the discretized control set

$$\begin{aligned} \tilde{\mathcal{U}}_{\mathcal{C}_K} := & \left\{ \mathbf{w} \in \mathbb{R}^{K \times |I|} : \mathbf{w}_i \equiv \mathbf{0} \forall i \in I \setminus I_a; \right. \\ & \wedge \forall (k, i) \in \{0, \dots, K\} \times I_a : \mathbf{w}_{i,k} \in [a_{\min}, a_{\max}], \\ & \wedge \forall (\tau_k, i) \in \mathcal{C}_K \times I_a : h_{\max} \tilde{\mathbf{v}}_i(\tau_k) + d_{\max} \geq \dots \\ & \left. \tilde{\mathbf{x}}_{i-1}(\tau_k) - \tilde{\mathbf{x}}_i(\tau_k) - l \geq h_{\min} \tilde{\mathbf{v}}_i(\tau_k) + d_{\min} \right\}. \end{aligned}$$

State, objective, and adjoint system discretization. Let $R \in \mathbb{N}_{>0}$. We consider a different fixed grid for the states

$$\mathcal{S}_R = \{t_r = rh_s, r \in \{0, 1, \dots, R\}\} \quad (16)$$

with mesh size $h_s = \frac{T}{R}$. The states grid is selected as a refinement of the control grid ($R \gg K$ and $\mathcal{C}_K \subset \mathcal{S}_R$). This is done to avoid numerical instability and ensure that the objective functional is sensitive to every control parameter in the vectors \mathbf{w}_i .

We also use RK-3 (this time over the grid \mathcal{S}_R) to run both the forward simulation of the system dynamics in eq. (2) and the backward simulation of the adjoint system in eq. (7). We get for each $i \in I \setminus \{0\}$ the approximate states

$$(\hat{\mathbf{x}}_i(t_r), \hat{\mathbf{v}}_i(t_r))_{t_r \in \mathcal{S}_R} \in \mathbb{R}^{2(R+1)}, \quad (17)$$

and for each $i \in \{1, \dots, 2(M+N)\}$ the approximate adjoint variables

$$(\hat{\zeta}_i(t_r))_{t_r \in \mathcal{S}_R} \in \mathbb{R}^{2(R+1)}. \quad (18)$$

We note here that we use a fixed grid, rather than the common adaptive grid used for RK schemes, to ensure that both systems of ODEs are solved on the same time discretization. The use of adaptive mesh size creates a fundamental difference in the numerical scheme that is not considered in this study. Further, an adaptive mesh that potentially changes in each gradient evaluation might cause issues in the optimization routine as we would be optimizing over the discretization noise which changes each iteration.

We use the approximate states to evaluate the objective

$$\hat{J}_{\mathcal{S}_R} = \frac{1}{h_s} \sum_{t_r \in \mathcal{S}_R} L(\hat{\mathbf{x}}(t_r), \hat{\mathbf{v}}(t_r), \hat{\mathbf{u}}(t_r; \mathbf{w})) + S(\hat{\mathbf{x}}(t_R), \hat{\mathbf{v}}(t_R)). \quad (19)$$

Along with the approximate adjoint variables we evaluate its gradient at every point $\tau_k \in \mathcal{C}_K$

$$\hat{\mathbf{G}}_i(\tau_k) = \begin{cases} L_{\mathbf{u}_i}(\hat{\mathbf{x}}(\tau_k), \hat{\mathbf{v}}(\tau_k), (\mathbf{w}_i)_k) - \hat{\zeta}_{i+M+N}(\tau_k) & i \in I_a \\ 0 & i \in I_h, \end{cases} \quad (20)$$

which we can then supply to the optimization algorithm.

Note that we use the time discretization \mathcal{S}_R to evaluate the objective function and its gradient, while we use the coarser discretization \mathcal{C}_K to evaluate the constraints. For this, the RK-3 scheme is a suitable choice to ensure that the approximation error for both discretizations is of similar order. This is because RK-3 evaluates the controller at points $\hat{\mathbf{u}}(t_i; \mathbf{w})$, $\hat{\mathbf{u}}(t_i + 0.5t_i; \mathbf{w})$, and $\hat{\mathbf{u}}(t_i + 0.75t_i; \mathbf{w})$ which are all the same under our piecewise constant parameterization.

The discretized objective function in eq. (19) together with the constraints in the discrete control set in eq. (16) define a finite optimization problem with $K \times |I_a|$ variables. In our implementation, we use MATLAB's sequential quadratic programming (SQP) algorithm [27]. The optimization problem we formulate is, in general, non-convex (due to the non-linearity of the car-following dynamics), thus SQP is only guaranteed to converge locally. This makes the initialization of the algorithm crucial and we will discuss this further in section IV.

C. Numerical algorithm

We summarize the proposed numerical approach. We use the SQP algorithm to update the control parameters \mathbf{w}_i . In every iteration, we query the value of the objective functional and its gradient by the following steps:

- 1) Run a forward simulation of the system dynamics in eq. (2) (using RK-3 subroutine) to obtain the approximate states.
- 2) Evaluate the objective function as in eq. (19).
- 3) Run a backward simulation of the adjoint system in eq. (7) (using RK-3 subroutine) to obtain approximate adjoint variables.
- 4) Evaluate the gradient function as in eq. (20).

We note that evaluating the objective function requires solving the system dynamics. The cost of evaluating the gradient is equal to solving an additional (linear) system of ODEs (the adjoint system).

IV. NUMERICAL EXPERIMENTS

In this section, we consider an instance of the OCP in eq. (5). The goal is to smooth stop-and-go waves that appear in a lead vehicle trajectory and minimize the fuel consumption of a platoon of HVs by introducing an AV directly behind the lead vehicle (see fig. 1).



Fig. 1. The mixed-autonomy platoon considered in the experiments. The gray vehicle is the leader, followed by an AV (red), followed by several HVs (blue).

Remark 2 (Fuel consumption model): We evaluate the fuel consumption of a trajectory as the integral of the instantaneous energy function proposed in [28] over the time horizon $[0, T]$. We note that this energy function does not incur a cost on negative acceleration. This can be exploited if we optimize this function and result in a solution trajectory in which the acceleration drops heavily at the end of the

time horizon as all quantities in the objective functional will decrease. This behavior is an artifact of the energy model and the fact that we optimize over a fixed time horizon. For this reason, we refrain from using the model as an objective functional and only use it for evaluation.

For the reason discussed in remark 2, we consider the L^2 norm of the acceleration of a vehicle as a surrogate of the energy model. Optimizing for this objective functional may result in the AV driving at a speed much slower than its leader. This behavior negatively impacts other metrics such as throughput, density, and average velocity. There are many design choices to remedy this effect and *pull* the AV to exit the road. Here, we do so by imposing constraints on the maximum headway of the AV as shown in the admissible set of controls (eq. (1)). Here, we do not impose the box constraints on the acceleration as the objective functional favors small values. We will demonstrate through numerical examples that the solution remains bounded and is proportional to the acceleration of the leader.

Based on the above, we define the following OCP

$$\inf_{\mathbf{u}} \int_0^T \sum_{i \in I_a} \mathbf{u}_i^2(t) + \sum_{i \in I_h} (A(\mathbf{x}_{i-1}(t), \mathbf{x}_i(t), \mathbf{v}_{i-1}(t), \mathbf{v}_i(t)))^2 dt$$

where (\mathbf{x}, \mathbf{v}) satisfies eq. (2) and $\forall i \in I_a, \forall t \in [0, T]$

$$h_{\min} \mathbf{v}_i(t) + d_{\min} \leq \mathbf{x}_{i-1}(t) - \mathbf{x}_i(t) - l \quad (21)$$

$$h_{\max} \mathbf{v}_i(t) + d_{\max} \geq \mathbf{x}_{i-1}(t) - \mathbf{x}_i(t) - l$$

$$\mathbf{v}_i(t) \geq 0.$$

Remark 3 (Initial datum): If the initial condition vectors \mathbf{x}_o and \mathbf{v}_o are chosen such that the inequality constraints in eq. (21) are satisfied for $t = 0$, then the problem is feasible for all $t \in [0, T]$. This is obvious to see in the absence of constraints on the acceleration limit.

We use the solution of the problem in eq. (21) with a single AV following the given leader trajectory as an initial guess for other experiments with HVs in the platoon. With no HVs, the optimization problem becomes a convex quadratic linearly-constrained problem that can be solved efficiently.

A. Results

We demonstrate the use of the proposed method in section III on two sets of numerical experiments each with a specific lead vehicle trajectory and varying platoon size. We consider two cases for the lead vehicle: (1) a benchmark trajectory with a step acceleration profile and (2) an experimental trajectory collected on the I-24 in Nashville, TN [29]. Both leader trajectories resemble driving in congested traffic that exhibits stop-and-go waves as shown in fig. 2.

For each lead vehicle trajectory, we consider a platoon of the form shown in fig. 1 with 0, 10, and 20 HVs. We examine the effect of the platoon size on the solution. We model the human vehicles using Bando-FtL model described in eq. (3) with parameters $\alpha = 0.1/s$ and $\beta = 525m^2/s$, and the following optimal velocity function

$$V(h) = v_{\max} \frac{\tanh(kh-d) + \tanh(l+d)}{1 + \tanh(l+d)}. \quad (22)$$




Fig. 2. Acceleration and velocity profiles of the benchmark (top) and experimental (bottom) lead vehicle trajectories (in gray) followed by HVs (in blue) modeled by the Bando-FtL model.

The parameters are selected as $v_{\max} = 35m/s$, $k = 0.2$, and $d = 4m$. In fig. 2, we illustrate the behavior of the HV platoon consisting of 21 vehicles. Note that the model propagates and amplifies the stop-and-go waves – a behavior that indicates string instability. Such behavior is also observed in commercially deployed adaptive cruise control systems [30]. It is also believed that human driving exhibits such instability.

In our numerical experiments, we aim to examine the effect of introducing an AV in the platoon on smoothing the stop-and-go waves and to quantify its effect on fuel consumption. As a baseline, we consider a full HV platoon of the same length and measure the percentage of fuel saving for each experiment.

In all the experiments we fix the car length to $l = 5m$, and the mesh sizes $h_c = 1s$ and $h_s = 0.1s$. We fix the minimum and maximum time headways to $h_{\min} = 0.5s$ and $h_{\max} = 3s$ and the minimum and maximum allowable distances as zero velocity to $d_{\max} = d_{\min} = 2m$. The $0.5s$ minimum time headway is considered small for human drivers. However, it is an acceptable constraint for AVs as they are believed to have a faster reaction time. Later in our experiments we will vary the h_{\max} parameter and demonstrate how it affects the fuel saving.

Figure 3 and fig. 4 show segments of the mixed-autonomy trajectories with the different platoon sizes following the benchmark and the experimental lead vehicles, respectively. In table I, we summarize the percentage of fuel saving for each of the experiments relative to the baseline. We also report the results of the *greedy* solution, that is the solution to the OCP where only the acceleration of the AV is optimized in each platoon. We note that in all experiments, the AV respects the time headway constraints at every time step and there are no collisions between the HVs.

TABLE I

PERCENTAGE OF FUEL SAVING RELATIVE TO THE HVs BASELINE.

Platoon Size	Benchmark leader		Experimental leader	
	Platoon optimization	Greedy optimization	Platoon optimization	Greedy optimization
0 HV	1.42	-	7.99	-
10 HV	1.42	1.42	8.84	8.47
20 HV	1.43	1.42	9.92	9.12

As expected, introducing the AV in the platoon reduces fuel consumption by up to almost 10% as stated in table I. The qualitative improvement can also be observed in fig. 3

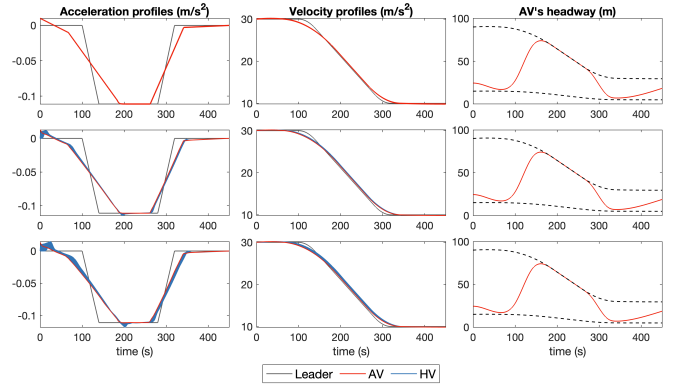


Fig. 3. Segments of the acceleration, velocity and AV headway profiles of the platoons following the *benchmark lead vehicle trajectory*. Platoon size: 0 HV (top row), 10 HV (middle row), and 20 HV (bottom row). The dashed lines in the headway plots represent the feasible headway envelope.

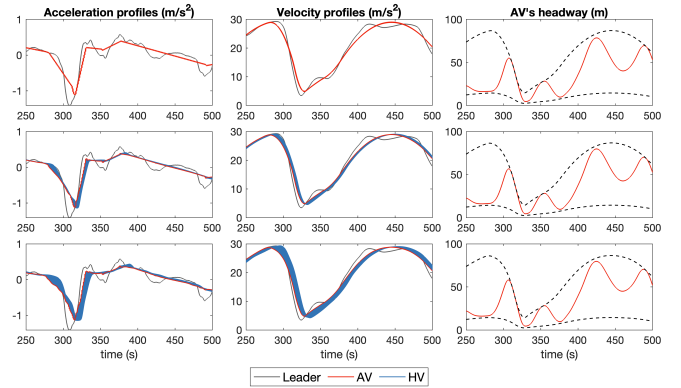


Fig. 4. Segments of the acceleration, velocity and headway profiles of the platoons following the *experimental lead vehicle trajectory*. Platoon size: 0 HV (top row), 10 HV (middle row), and 20 HV (bottom row). The dashed lines in the headway plots represent the feasible headway envelope.

and 4 in comparison to the baseline in fig. 2. Introducing the AV smooths the acceleration and velocity profiles of the leader by opening large gaps in anticipation of large drops in the velocity. This allows the vehicle to maintain a more steady speed profile and reduce the fluctuations in its acceleration. This also improves the driving behavior of the HVs behind the AV. However, we observe that the platoon optimization does not offer significant improvement over the greedy optimization. This suggests that the optimal car-following behavior of the platoon is agnostic to our choice of CFM and platoon size.

V. DISCUSSION AND FUTURE WORK

In this section, we discuss three important design choices and their effect on the optimal solution of the problem. **First** is the choice of the car following model. In our study, we used Bando-FtL as it enjoys desirable theoretical guarantees. Our results suggest that the solution does not depend heavily on the CFM. We note that this might not be true for other classes of CFMs and further investigation of this is needed.

The **second** design choice is the feasible headway envelope. It is obvious that with larger h_{\max} values, less fuel consumption can be achieved. As noted in table II, the

choice of this parameter has a big impact on the solution. The appropriate choice of this parameter depends on the practical setting. Opening large gaps is not a human-like driving behavior and it can cause the vehicles to drive much slower than the speed limit which is considered a traffic violation. On the other hand, if the AV is viewed as an actuator that is not expected to follow traffic rules, larger values of h_{\max} would be appropriate. The key takeaway we want to emphasize here is that even with complete knowledge of the downstream information, the optimal solution and the performance gain are bound by local constraints.

Lastly, we comment on the choice of the platoon. Many other interesting configurations can be studied in which multiple AVs are distributed in the platoon. However, deviating from the platoon configuration examined here makes the problem more complicated. In the present analysis, the constraints in eq. (21) are all linear. However, when an AV is preceded by an HV, these constraints become non-linear and significantly more difficult to handle. One possible approach to work with these constraints involves employing a penalty method, as described in [26]. Investigating such approaches and conducting experiments with a more general platoon configuration is left for future work.

TABLE II
PERCENTAGE OF FUEL SAVING RELATIVE TO THE HVs BASELINE AT
DIFFERENT VALUES h_{\max} .

h_{\max}	3 s	5 s	7 s	10 s	12 s
Fuel saving (%)	7.99	11.98	15.38	19.54	21.75

REFERENCES

- [1] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, H. Pohlmann, F. Wu, B. Piccoli *et al.*, "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transportation Research Part C: Emerging Technologies*, vol. 89, pp. 205–221, 2018.
- [2] C. Wu, A. M. Bayen, and A. Mehta, "Stabilizing traffic with autonomous vehicles," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 6012–6018.
- [3] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, "Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam," *New journal of physics*, vol. 10, no. 3, p. 033001, 2008.
- [4] F.-C. Chou, A. R. Bagabaldo, and A. M. Bayen, "The lord of the ring road: a review and evaluation of autonomous control policies for traffic in a ring road," *ACM Transactions on Cyber-Physical Systems (TCPS)*, vol. 6, no. 1, pp. 1–25, 2022.
- [5] P. Ioannou, Z. Xu, S. Eckert, D. Clemons, and T. Sieja, "Intelligent cruise control: theory and experiment," in *Proceedings of 32nd IEEE Conference on Decision and Control*. IEEE, 1993, pp. 1885–1890.
- [6] S. Darbha and K. Rajagopal, "Intelligent cruise control systems and traffic flow stability," *Transportation Research Part C: Emerging Technologies*, vol. 7, no. 6, pp. 329–352, 1999.
- [7] L. Davis, "Effect of adaptive cruise control systems on traffic flow," *Physical Review E*, vol. 69, no. 6, p. 066110, 2004.
- [8] R. Rajamani, S. B. Choi, J. K. Hedrick, B. Law, R. Prohaska, and P. Kretz, "Design and experimental implementation of control for a platoon of automated vehicles," in *ASME International Mechanical Engineering Congress and Exposition*, vol. 15861. American Society of Mechanical Engineers, 1998, pp. 681–689.
- [9] J.-J. Martinez and C. Canudas-de Wit, "A safe longitudinal control for adaptive cruise control and stop-and-go scenarios," *IEEE Transactions on control systems technology*, vol. 15, no. 2, pp. 246–258, 2007.
- [10] M. L. Delle Monache, T. Liard, A. Rat, R. Stern, R. Bhadani, B. Seibold, J. Sprinkle, D. B. Work, and B. Piccoli, "Feedback control algorithms for the dissipation of traffic waves with autonomous vehicles," *Computational Intelligence and Optimization Methods for Control Engineering*, pp. 275–299, 2019.
- [11] B. S. Kerner, "Physics of automated driving in framework of three-phase traffic theory," *Physical Review E*, vol. 97, no. 4, p. 042303, 2018.
- [12] Y. Jia, C. Cai, and D. Gorges, "An LSTM-based speed predictor based on traffic simulation data for improving the performance of energy-optimal adaptive cruise control," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–7.
- [13] E. Vinitzky, N. Lichtlé, K. Parvate, and A. Bayen, "Optimizing mixed autonomy traffic flow with decentralized autonomous vehicles and multi-agent reinforcement learning," *ACM Transactions on Cyber-Physical Systems*, 2023.
- [14] R. Schmied, H. Waschl, R. Quirynen, M. Diehl, and L. del Re, "Nonlinear mpc for emission efficient cooperative adaptive cruise control," *IFAC-papersonline*, vol. 48, no. 23, pp. 160–165, 2015.
- [15] F. Biral, E. Bertolazzi, and P. Bosetti, "Notes on numerical methods for solving optimal control problems," *IEEJ Journal of Industry Applications*, vol. 5, no. 2, pp. 154–166, 2016.
- [16] L. S. Pontryagin, *Mathematical theory of optimal processes*. CRC press, 1987.
- [17] A. E. Bryson, *Applied optimal control: optimization, estimation and control*. CRC Press, 1975.
- [18] J. Wang, Y. Zheng, C. Chen, Q. Xu, and K. Li, "Leading cruise control in mixed traffic flow: System modeling, controllability, and string stability," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 8, pp. 12 861–12 876, 2021.
- [19] J. Wang, Y. Zheng, Q. Xu, J. Wang, and K. Li, "Controllability analysis and optimal control of mixed traffic flow with human-driven and autonomous vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 12, pp. 7445–7459, 2020.
- [20] Y. Zheng, J. Wang, and K. Li, "Smoothing traffic flow via control of autonomous vehicles," *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 3882–3896, 2020.
- [21] S. Wang, R. Stern, and M. W. Levin, "Optimal control of autonomous vehicles for traffic smoothing," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3842–3852, 2022.
- [22] X. Lin, D. Gorges, and A. Weißmann, "Simplified energy-efficient adaptive cruise control based on model predictive control," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 4794–4799, 2017.
- [23] F. Wu and A. M. Bayen, "A hierarchical MPC approach to car-following via linearly constrained quadratic programming," *IEEE Control Systems Letters*, vol. 7, pp. 532–537, 2022.
- [24] X. Gong and A. Keimer, "On the well-posedness of the "Bando-follow the leader" car following model and a time-delayed version," *Networks and Heterogeneous Media*, vol. 18, no. 2, pp. 775–798, 2023.
- [25] S. Albeaik, A. Bayen, M. T. Chiri, X. Gong, A. Hayat, N. Kardous, A. Keimer, S. T. McQuade, B. Piccoli, and Y. You, "Limitations and improvements of the intelligent driver model (IDM)," *SIAM Journal on Applied Dynamical Systems*, vol. 21, no. 3, pp. 1862–1892, 2022.
- [26] Q. Lin, R. Loxton, and K. L. Teo, "The control parameterization method for nonlinear optimal control: a survey," *Journal of Industrial and management optimization*, vol. 10, no. 1, pp. 275–309, 2014.
- [27] N. Jorge and J. W. Stephen, "Numerical optimization," 2006.
- [28] J. W. Lee, G. Gunter, R. Ramadan, S. Almatrudi, P. Arnold, J. Aquino, W. Barbour, R. Bhadani, J. Carpio, F.-C. Chou *et al.*, "Integrated framework of vehicle dynamics, instabilities, energy models, and sparse flow smoothing controllers," *arXiv preprint arXiv:2104.11267*, 2021.
- [29] M. Nice, N. Lichtle, G. Gumm, M. Roman, E. Vinitzky, S. Elmadani, M. Bunting, R. Bhadani, K. Jang, G. Gunter, M. Kumar, S. McQuade, C. Denaro, R. Delorenzo, B. Piccoli, D. Work, A. Bayen, J. Lee, J. Sprinkle, and B. Seibold. (2021) The I-24 trajectory dataset. [Online]. Available: <https://zenodo.org/record/6456348>
- [30] G. Gunter, D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle *et al.*, "Are commercially implemented adaptive cruise control systems string stable?" *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 11, pp. 6992–7003, 2020.