

Optimal Weight Adaptation of Model Predictive Control for Connected and Automated Vehicles in Mixed Traffic with Bayesian Optimization

Viet-Anh Le, *IEEE Student Member*, Andreas A. Malikopoulos, *IEEE Senior Member*

Abstract—In this paper, we develop an optimal weight adaptation strategy of model predictive control (MPC) for connected and automated vehicles (CAVs) in mixed traffic. We model the interaction between a CAV and a human-driven vehicle (HDV) as a simultaneous game and formulate a game-theoretic MPC problem to find a Nash equilibrium of the game. In the MPC problem, the weights in the HDV’s objective function can be learned online using moving horizon inverse reinforcement learning. Using Bayesian optimization, we propose a strategy to optimally adapt the weights in the CAV’s objective function so that the expected true cost when using MPC in simulations can be minimized. We validate the effectiveness of the optimal strategy by numerical simulations of a vehicle crossing example at an unsignalized intersection.

I. INTRODUCTION

Recent advancements in connected and automated vehicles (CAVs) provide a promising chance in reducing both energy consumption and travel delay [1], [2]. In our previous work [3]–[5], we addressed coordination and routing problems for CAVs given full penetration of CAVs. However, CAVs will gradually penetrate the market and co-exist with human-driven vehicles (HDVs) in the next decades. Therefore, addressing safe and efficient motion planning and control for CAVs in mixed traffic given various human driving styles is highly important. Several control approaches have been proposed in the literature such as model predictive control [6], [7], learning-based control [8], [9], game-theoretic control [10], and socially-compatible control [11], [12].

Among those control approaches, model predictive control (MPC) has received significant attention since (1) it can be integrated into other methods such as learning-based control or socially-compatible control, and (2) it can handle multiple objectives and constraints concurrently. However, like in many MPC approaches for dynamical systems, some objectives, constraints, or system dynamics in motion planning and control for CAVs are usually simplified or approximated so that the resulting MPC problems can be solved in real-time. In addition, the objective function in MPC is generally formed by a linear combination of multiple features, in which the weights are chosen empirically. As a result, true cost optimization might not be achieved leading to performance degradation if the weights are chosen inappropriately. An efficient technique to overcome these difficulties in practice is automatic weight tuning [13] which aims to derive a strategy

to tune the weights of MPC so that the best true cost can be achieved. Marco *et al.* [14] used Bayesian optimization to optimize weights of a cost function to compensate for the discrepancy between the true dynamics and a linearized model. Gros and Zanon [15] utilized reinforcement learning for parameter adaptation in nonlinear MPC. Jain *et al.* [16] focused on finding an MPC rollout having a low true cost using covariance matrix adaptation evolution strategy.

Furthermore, in the control applications involving human decisions, e.g., CAVs interacting with HDVs in mixed traffic, the controller must address the stochasticity and diversity caused by human behavior. Generally, MPC with fixed weights cannot guarantee to work well in such applications. For example, overly weighting toward the safety objective in the MPC design while encountering a driving scenario with a conservative HDV may cause traffic delay. In contrast, if CAVs and HDVs behave aggressively then unsafe situations may occur. Therefore, the weights of the MPC problem need to be adapted online depending on the human driving model.

In the recent research effort [17], we developed a control framework to address the motion planning problem for CAVs in mixed traffic. We modeled the interaction between a CAV and an HDV as a simultaneous game and proposed an MPC objective function to find a Nash equilibrium of the game. The weights in the objective function are parameterized by social value orientation (SVO), and depending on the online estimate of the SVO for the HDV, the MPC weights are adapted heuristically. In this paper, we propose a method for *optimal weight adaptation* of MPC for CAVs in mixed traffic based on *Bayesian optimization*. Using the proposed method, we can derive offline the optimal weight adaptation strategy for the MPC with respect to the HDV’s objective weights so that the true desired performance can be achieved. Then by learning the objective weights that best describe human driving behavior online using real-time data and the moving horizon inverse reinforcement learning (IRL) technique [18], the MPC weights are adapted accordingly. We demonstrate the proposed method by a vehicle crossing example at an unsignalized intersection, and show the benefits by comparing with the heuristic method in [17].

The remainder of this paper is structured as follows. Section II presents the game-theoretic MPC formulation and the moving horizon IRL technique. In Section III, we develop the method to derive the optimal weight adaptation strategy with Bayesian optimization. In Section IV, we demonstrate the proposed framework by an intersection crossing example, while numerical simulation results are provided in Section V. Finally, we conclude the paper in Section VI.

This work was supported by NSF under Grants CNS-2149520 and CMMI-2219761.

The authors are with the Department of Mechanical Engineering, University of Delaware, Newark, DE 19716 USA. E-mail: vietale@udel.edu, andreas@udel.edu.

II. MOTION PLANNING FOR CAVS IN MIXED TRAFFIC WITH MODEL PREDICTIVE CONTROL

In this section, we present a game-theoretic MPC formulation for motion planning of a CAV while interacting with an HDV along with the moving horizon IRL technique to learn the objective weights of the HDV from real-time data.

A. Model Predictive Control for Motion Planning

We consider an interactive driving scenario including a CAV and an HDV whose indices are 1 and 2, respectively. The goal of the MPC motion planner is to generate the trajectory and control actions of CAV-1 while considering the real-time driving behavior of HDV-2. To guarantee that CAV-1 has data of HDV-2's real-time trajectories, we make the following assumption:

Assumption 1: A coordinator is available to collect trajectories of HDV-2 and transmit them to CAV-1 without any significant delay or error during the communication.

We formulate the problem in the discrete-time domain, in which the dynamic model of each vehicle i is given by

$$\mathbf{x}_{i,k+1} = \mathbf{f}_i(\mathbf{x}_{i,k}, \mathbf{u}_{i,k}), \quad (1)$$

where $\mathbf{x}_{i,k}$ and $\mathbf{u}_{i,k}$, $i = 1, 2$, are the vectors of states and control actions, respectively, at time step $k \in \mathbb{N}$. We utilize the control framework presented in [17], in which the interaction between CAV-1 and HDV-2 is modeled as a simultaneous game, i.e., the game without a leader-follower structure, in which the objective of each vehicle includes its individual objective and a shared objective. Let $l_1(\mathbf{x}_{1,k+1}, \mathbf{u}_{1,k})$ and $l_2(\mathbf{x}_{2,k+1}, \mathbf{u}_{2,k})$ be the individual objective functions of CAV-1 and HDV-2, respectively, and $l_{12}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k})$, where $\mathbf{x}_{12,k+1} = [\mathbf{x}_{1,k+1}^\top, \mathbf{x}_{2,k+1}^\top]^\top$ and $\mathbf{u}_{12,k} = [\mathbf{u}_{1,k}^\top, \mathbf{u}_{2,k}^\top]^\top$, be the cooperative term at time step k . We assume that CAV-1 and HDV-2 share the same cooperative objective, e.g., collision avoidance. Those objective functions are usually designed as weighted sums of some features as follows

$$l_i(\mathbf{x}_{i,k+1}, \mathbf{u}_{i,k}) = \boldsymbol{\omega}_i^\top \boldsymbol{\phi}_i(\mathbf{x}_{i,k+1}, \mathbf{u}_{i,k}), \quad i = 1, 2, \quad (2)$$

$$l_{12}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}) = \boldsymbol{\omega}_{12}^\top \boldsymbol{\phi}_{12}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}), \quad (3)$$

where $\boldsymbol{\phi}_i$, $\boldsymbol{\phi}_{12}$ are vectors of features and $\boldsymbol{\omega}_i \in \mathcal{W}_i$, $\boldsymbol{\omega}_{12} \in \mathcal{W}_{12}$ are corresponding vectors of weights, where \mathcal{W}_i and \mathcal{W}_{12} are the sets of feasible values. For ease of notation, we define $-i$ for each $i \in \{1, 2\}$ as the other vehicle than vehicle i . We consider that given any control actions $\mathbf{u}_{-i,k}$ of the other vehicle, each vehicle i applies the control actions $\mathbf{u}_{i,k}^*$ that minimizes a sum of its individual objective and the shared objective, i.e.,

$$\mathbf{u}_{i,k}^* = \arg \min_{\mathbf{u}_{i,k}} l_i(\mathbf{x}_{i,k+1}, \mathbf{u}_{i,k}) + l_{12}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}), \quad \forall \mathbf{u}_{-i,k}. \quad (4)$$

Next, we formulate an MPC problem with a control horizon of length $H \in \mathbb{N}$. Let t be the current time step and $\mathcal{I}_t = \{t, \dots, t+H-1\}$ be the set of all time steps in the control horizon at time step t . We can recast the simultaneous game between CAV-1 and HDV-2 presented above as a

potential game [19], the game in which all players minimize a single global function called the potential function. In the potential game, a Nash equilibrium can be found by minimizing the potential function. The potential function in this game at each time step k is

$$l_{\text{pot}}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}) = \sum_{i=1,2} l_i(\mathbf{x}_{i,k+1}, \mathbf{u}_{i,k}) + l_{12}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}) \quad (5)$$

Therefore, we propose utilizing the cumulative sum of the potential function over the control horizon as the objective function in the MPC problem, which can be given by

$$J_{\text{MPC}} = \sum_{k \in \mathcal{I}_t} l_{\text{pot},k}(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}). \quad (6)$$

Hence, the MPC problem for motion planning of CAV-1 is formulated as follows

$$\underset{\{\mathbf{u}_{12,k}\}_{k \in \mathcal{I}_t}}{\text{minimize}} \quad J_{\text{MPC}} \quad (7a)$$

subject to:

$$(1), \quad i = 1, 2, \quad (7b)$$

$$g_j(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}) \leq 0, \quad \forall j \in \mathcal{J}_{\text{ieq}}, \quad (7c)$$

$$h_j(\mathbf{x}_{12,k+1}, \mathbf{u}_{12,k}) = 0, \quad \forall j \in \mathcal{J}_{\text{eq}}, \quad (7d)$$

where (7b)–(7d) hold for all $k \in \mathcal{I}_t$. The constraints (7c) and (7d) are inequality and equality constraints with \mathcal{J}_{ieq} and \mathcal{J}_{eq} are sets of indices.

In the objective function of the MPC problem (7), assume that we can pre-define the features $\boldsymbol{\phi}_i$, $i = 1, 2$ and $\boldsymbol{\phi}_{12}$, if we learn online $\boldsymbol{\omega}_2$ and $\boldsymbol{\omega}_{12}$ that best describe the human driving behavior, the CAV's objective weights $\boldsymbol{\omega}_1$ are adapted to achieve the desired performance. The optimal strategy for adapting $\boldsymbol{\omega}_1$ can be derived offline using Bayesian optimization as presented in Section III.

B. Moving Horizon Inverse Reinforcement Learning

To identify the weights $\boldsymbol{\omega}_2$ and $\boldsymbol{\omega}_{12}$ in the individual objective function of HDV-2 and the shared objective, we utilize the feature-based IRL approach [18], [20], a machine learning technique developed to learn the underlying objective or reward of an agent by observing its behavior. We define the vector of all features and the vector of all corresponding weights in HDV-2's objective function as $\mathbf{f} = [\boldsymbol{\phi}_2^\top, \boldsymbol{\phi}_{12}^\top]^\top$ and $\boldsymbol{\theta} = [\boldsymbol{\omega}_2^\top, \boldsymbol{\omega}_{12}^\top]^\top$, respectively. Let $\tilde{\mathbf{f}}$ be the vector of average observed feature values computed from data and $\mathbb{E}_p[\mathbf{f}]$ be the expected feature values with a given probability distribution p over trajectories. With feature-based IRL, the goal is to learn the weight vector $\boldsymbol{\theta} \in \Omega$, where $\Omega = \mathcal{W}_2 \times \mathcal{W}_{12}$ so that expected feature values can match observed feature values.

In moving horizon IRL, at each time step, we utilize the $L \in \mathbb{N}$ most recent trajectory segments to update the weight estimate, where L is the estimation horizon length. Let t be the current time step and $\mathcal{R}_t = \{\mathbf{r}_m\}_{m=1, \dots, L}$ be the set of L sample trajectory segments collected over the estimation horizon at time t , in which $\mathbf{r}_m = (\mathbf{x}_{12,t-m}, \mathbf{x}_{12,t-m+1}, \mathbf{u}_{12,t-m})$, for $m = 1, \dots, L$, is the

tuple representing the trajectory segment. We use the maximum entropy IRL approach [18] that utilizes an exponential family distribution for p and maximizes the entropy of the distribution, yielding the following optimization problem

$$\underset{\theta \in \Omega}{\text{maximize}} \sum_{\mathbf{r}_m \in \mathcal{R}} \log p(\mathbf{r}_m | \theta). \quad (8)$$

To solve (8), one can use gradient-based methods where the gradient can be approximated by the difference between the expected and the empirical feature values [18]

$$\nabla \mathcal{L}_\theta = \tilde{\mathbf{f}} - \mathbb{E}_p[\mathbf{f}]. \quad (9)$$

The average observed feature values $\tilde{\mathbf{f}}$ can be computed from an average of feature values for all training samples

$$\tilde{\mathbf{f}} = \frac{1}{L} \sum_{\mathbf{r}_m \in \mathcal{R}} \mathbf{f}(\mathbf{r}_m). \quad (10)$$

Meanwhile, $\mathbb{E}_p[\mathbf{f}]$ can be approximated by the expected feature values of the most likely trajectories as follows

$$\mathbb{E}_p[\mathbf{f}] \approx \mathbf{f}(\arg \max_{\mathbf{r}} \log p(\mathbf{r} | \theta)). \quad (11)$$

More specifically, for each sample trajectory \mathbf{r}_m , we fix θ , the trajectory $\{\mathbf{x}_{1,k}, \mathbf{x}_{1,k+1}, \mathbf{u}_{1,k}\}$ of CAV-1, and the initial condition $\mathbf{x}_{2,k}$, then find the optimized control actions of HDV-2 $\mathbf{u}_{2,k}$ that minimize $\theta^\top \mathbf{f}(\mathbf{r}_m)$. We denote the system trajectories resulted from the optimized HDV-2's actions as $\{\mathbf{r}_1^\theta, \dots, \mathbf{r}_L^\theta\}$. Next, we evaluate the features for all optimized trajectories and compute the approximated expected feature values $\tilde{\mathbb{E}}_p[\mathbf{f}]$ by

$$\tilde{\mathbb{E}}_p[\mathbf{f}] = \frac{1}{L} \sum_{\mathbf{r}_m \in \mathcal{R}} \mathbf{f}(\mathbf{r}_m^\theta). \quad (12)$$

Using (9), (10), and (12), the gradient of the objective function in (8) with respect to θ can be computed. Therefore, the estimate of θ can be updated by projected gradient ascent method as follows

$$\theta^{(j+1)} = \text{Proj}_{\Omega}(\theta^{(j)} + \eta \nabla \mathcal{L}_{\theta^{(j)}}), \quad (13)$$

where $\eta \in \mathbb{R}^+$ is the learning rate and $\theta^{(j)}$ denotes the estimate of θ at iteration $j \in \mathbb{N}$ of the algorithm.

Therefore, given L sample trajectories over the estimation horizon, the moving horizon IRL procedure for learning HDV-2's objective weights is summarized as follows. At each time step, we start with initial weights $\theta^{(0)}$, and at each algorithmic iteration j , the gradient $\nabla \mathcal{L}_{\theta^{(j)}}$ of the objective function in (8) with respect to θ at $\theta = \theta^{(j)}$ is computed and used to update the estimate of θ by (13). For more details, the readers are referred to [18] on maximum entropy IRL and to [17] on moving horizon implementation.

III. OPTIMAL WEIGHT ADAPTATION WITH BAYESIAN OPTIMIZATION

In this section, we first introduce the optimal weight adaptation problem for MPC motion planning in mixed traffic, then propose using Bayesian optimization to solve the problem.

A. Optimal Weight Adaptation Problem

Let \mathbf{x}_{MPC} and \mathbf{u}_{MPC} be the state and control trajectories of the agents in the simulation using MPC to control CAV-1. We define the true cost in the simulation corresponding to using an MPC with a tuple of weights $\omega = (\omega_1, \omega_2, \omega_{12})$ as $J_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}})$. Note that generally the true cost function J_{true} can only be obtained after performing the simulations or experiments and evaluating the state and control trajectories of the agents. We aim to seek the optimal weights of CAV-1's individual objective $\omega_1^* \in \mathcal{W}_1$ corresponding to each (ω_2, ω_{12}) that minimize the expected true cost given prior distribution of initial conditions $\mathbf{x}_{\text{MPC}}(0)$. This can be achieved by solving the following optimization problem

$$\omega_1^* = \arg \min_{\omega_1 \in \mathcal{W}_1} \bar{J}_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}}) \quad (14)$$

where

$$\bar{J}_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}}) = \mathbb{E}_{\mathbf{x}_{\text{MPC}}(0)} [J_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}})], \quad (15)$$

in which the expected true cost can be computed approximately by the average true cost of $n_s \in \mathbb{N}$ independent and identically distributed (i.i.d.) simulations with the initial states sampled from a prior distribution.

Solving the problem in (14) can be computationally intractable since the objective is a black-box function of the optimization variable ω_1 . Moreover, it takes a significant amount of time to evaluate that objective function because it requires multiple simulations with different initial conditions to obtain the expected true cost. Those reasons motivate us to utilize Bayesian optimization to solve (14).

B. Bayesian Optimization

Bayesian optimization is a machine learning-based optimization technique commonly used for minimizing (or maximizing) a black-box objective function in which we can observe only the output of the function by sampling and no first- or second-order derivatives [21]. In Bayesian optimization, the objective function is learned by a surrogate model, e.g., Gaussian Process (GP), which can provide a posterior distribution of the function. The surrogate model is combined with an acquisition function to decide the next candidate of the optimal solution. As a result, at each algorithmic iteration, by optimizing the acquisition function over the current surrogate model, the next sampling candidate is found. The objective value at that sampling candidate is then evaluated and added to the training data set to re-train the surrogate model.

In our problem, let $f(\omega_1) = \bar{J}_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}})$ be the black-box objective function of the variable ω_1 which needs to be minimized with Bayesian optimization, i.e.,

$$\underset{\omega_1 \in \mathcal{W}_1}{\text{minimize}} f(\omega_1), \quad (16)$$

We use Gaussian process (GP) model [22] to learn the black-box objective function. The GP of $f(\omega_1)$ is denoted by $\mathcal{G}_f(\omega_1)$. The GP surrogate model is combined with an

Algorithm 1 Bayesian optimization for optimal weight adaptation

Require: $j_{\max}, j_{\text{init}} \in \mathbb{N}, \omega_2, \omega_{12}$

- 1: **procedure** INITIALIZATION
 - 2: **for** $j = 1, 2, \dots, j_{\text{init}}$ **do**
 - 3: Randomly sample $\omega_1^{(j)} \in \mathcal{W}_1$
 - 4: Compute average true cost $\bar{J}_{\text{true}}^{(j)}$
 - 5: Add $(\omega_1^{(j)}, \bar{J}_{\text{true}}^{(j)})$ to a training dataset \mathcal{D}
 - 6: Learn a GP model $\mathcal{G}(\omega_1)$ with \mathcal{D}
 - 7: **procedure** BAYESIAN OPTIMIZATION
 - 8: **for** $j = 1, \dots, j_{\max}$ **do**
 - 9: Find next candidate $\omega_1^{(j^*)}$ by optimizing the acquisition function
 - 10: Compute average true cost $\bar{J}_{\text{true}}^{(j^*)}$
 - 11: Add $(\omega_1^{(j^*)}, \bar{J}_{\text{true}}^{(j^*)})$ to \mathcal{D} and re-train $\mathcal{G}(\omega_1)$
 - 12: **return** ω_1^*
-

acquisition function ξ leading to the following optimization problem for finding the next candidate of the optimal solution

$$\underset{\omega_1 \in \mathcal{W}_1}{\text{maximize}} \quad \xi(\mu(\omega_1), \sigma(\omega_1)). \quad (17)$$

where $\mu(\omega_1)$ and $\sigma(\omega_1)$ denote the mean and variance of the GP prediction, respectively. In this paper, we use the expected improvement acquisition function defined as follows

$$\text{EI}(\omega_1) = \mathbb{E} \left[\max\{\Delta(\omega_1), 0\} \right], \quad (18)$$

where $\Delta(\omega_1) = f(\omega_1^+) - \mu(\omega_1)$ is the difference between the previous best sample $f(\omega_1^+)$ at ω_1^+ and the predicted output at ω_1 . The expected improvement under the GP model can be derived analytically as follows [23]

$$\text{EI}(\omega_1) = \sigma(\omega_1) \varphi \left(\frac{\Delta(\omega_1)}{\sigma(\omega_1)} \right) + \Delta(\omega_1) \Phi \left(\frac{\Delta(\omega_1)}{\sigma(\omega_1)} \right), \quad (19)$$

where φ and Φ are the probability density function (PDF) and the cumulative distribution function (CDF) of the standard normal distribution, respectively.

The entire algorithm to determine the optimal value of ω_1 for each (ω_2, ω_{12}) is summarized in Algorithm 1. Note that we denote the candidate of the optimal solution obtained by optimizing the acquisition function at algorithmic iteration j as $\omega_1^{(j^*)}$, which is different to the global solution ω_1^* returned by Bayesian optimization that is the best candidate evaluated.

IV. ILLUSTRATIVE EXAMPLE

In this section, we demonstrate the control formulation presented in Section II and the optimal weight adaptation problem in Section III by a vehicle crossing example at an unsignalized intersection illustrated in Fig. 1. We define the surrounding area of the intersection inside of which the vehicles can communicate with the coordinator as a control zone, while the location where a lateral collision can occur is called a conflict point. The dynamics of each vehicle i are

described by the following double-integrator dynamics

$$\begin{aligned} p_{i,k+1} &= p_{i,k} + \Delta T v_{i,k} + \frac{1}{2} \Delta T^2 a_{i,k}, \\ v_{i,k+1} &= v_{i,k} + \Delta T a_{i,k}, \end{aligned} \quad (20)$$

where $\Delta T \in \mathbb{R}^+$ is the sampling time, $p_{i,k} \in \mathbb{R}$ is the longitudinal position of the vehicle with respect to the conflict point at time k , and $v_{i,k} \in \mathbb{R}$ and $a_{i,k} \in \mathbb{R}$ are the speed and acceleration of the vehicle i at time k , respectively. The state and control input of vehicle i are defined by $\mathbf{x}_{i,k} = [p_{i,k}, v_{i,k}]^\top$ and $u_{i,k} = a_{i,k}$, respectively.

The individual objective for each vehicle in the MPC problem includes: (1) minimizing the control input for smoother movement and energy saving, and (2) minimizing the deviation from the maximum allowed speed to reduce the time to cross the intersection, i.e.,

$$l_i(\mathbf{x}_{i,k+1}, u_{i,k}) = \begin{bmatrix} \omega_{i,1} \\ \omega_{i,2} \end{bmatrix}^\top \begin{bmatrix} a_{i,k}^2 \\ (v_{i,k+1} - v_{\max})^2 \end{bmatrix}, \quad (21)$$

for $i = 1, 2$, where $\omega_{i,1}, \omega_{i,2} \in \mathbb{R}^+$ are positive weights. The shared objective function takes the form of a logarithmic penalty function corresponding to a collision avoidance constraint as follows

$$l_{12}(\mathbf{x}_{1,k+1}, \mathbf{x}_{2,k+1}) = -\omega_{12} \log(\gamma(p_{1,k+1}^2 + p_{2,k+1}^2)), \quad (22)$$

where $\omega_{12} \in \mathbb{R}^+$ is a positive weight and $\gamma \in \mathbb{R}^+$ is a parameter of the logarithmic penalty function.

Next, we consider the following state and control constraints for CAV-1

$$v_{\min} \leq v_{1,k+1} \leq v_{\max}, \quad u_{\min} \leq a_{1,k} \leq u_{\max}, \quad \forall k \in \mathcal{I}_t, \quad (23)$$

where $u_{\min}, u_{\max} \in \mathbb{R}$ are the minimum deceleration and maximum acceleration, respectively, and $v_{\min}, v_{\max} \in \mathbb{R}$ are the minimum and maximum speed limits, respectively. Moreover, we impose the following safety constraint

$$r \leq \sqrt{p_{1,k+1}^2 + p_{2,k+1}^2}, \quad \forall k \in \mathcal{I}_t, \quad (24)$$

to guarantee that the predicted distances between CAV-1 and HDV-2 are greater than a safety threshold $r \in \mathbb{R}^+$.

The MPC problem for CAV-1 in this example is thus formulated as follows

$$\underset{\{u_{1,k}, u_{2,k}\}_{k \in \mathcal{I}_t}}{\text{minimize}} \quad \sum_{k \in \mathcal{I}_t} \left(\sum_{i=1,2} l_i(\mathbf{x}_{i,k+1}, u_{i,k}) + l_{12}(\mathbf{x}_{1,k+1}, \mathbf{x}_{2,k+1}) \right), \quad (25a)$$

subject to:

$$(20), \quad \forall k \in \mathcal{I}_t, \quad i = 1, 2, \quad (25b)$$

$$(23), (24), \quad \forall k \in \mathcal{I}_t, \quad (25c)$$

We define a true cost function called *time-energy efficiency with safety* that is computed by

$$J_{\text{true}}^\omega(\mathbf{x}_{\text{MPC}}, \mathbf{u}_{\text{MPC}}) = \alpha t_{1,f} + \beta E_t + \lambda \mathbb{I}(g(\mathbf{x}_{\text{MPC}})), \quad (26)$$

where α, β , and $\lambda \in \mathbb{R}^+$ are constant weights and λ is sufficiently large compared to α and β to prioritize safety rather than time and energy efficiency, $t_{1,f}$ is the time that

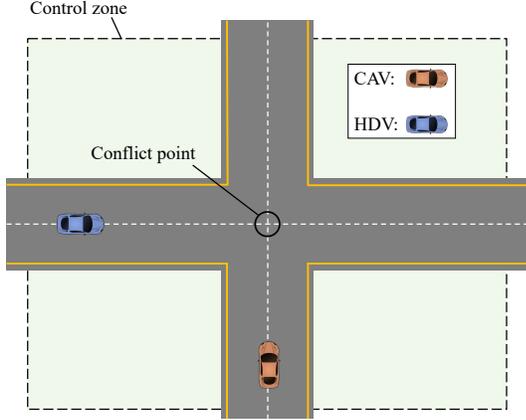


Fig. 1: An unsignalized intersection scenario with a CAV and an HDV.

CAV-1 exits the control zone, E is the total amount of energy consumption of CAV-1 while traveling in the control zone, and $\mathbb{I}(g(\mathbf{x}_{\text{MPC}}))$ is the indicator function of the safety constraint g defined as

$$\mathbb{I}(g(\mathbf{x}_{\text{MPC}})) = \begin{cases} 0, & \text{if } g(\mathbf{x}_{\text{MPC}}) \leq 0 \\ 1, & \text{otherwise} \end{cases}. \quad (27)$$

The safety constraint is $g(\mathbf{x}_{\text{MPC}}) = r - d_{12,\min} \leq 0$ where $d_{12,\min}$ is the minimum distance between two vehicles. Within Bayesian optimization framework that requires a continuous objective function, we approximate the indicator function by a sigmoid function. To evaluate the total fuel consumption E_t of CAV-1, we consider the polynomial meta-model and coefficients from an engine torque-speed-efficiency map of a typical car presented in [24].

V. SIMULATION RESULTS

To demonstrate the effectiveness of the proposed method, we conduct numerical simulations for the intersection crossing example described in Section IV.

A. Simulation Setup

For the implementation, since the solution of the MPC problem does not change if all the weights are scaled by a positive factor, we fix the shared objective weight $\omega_{12} = 10^3$ to reduce the dimension of the problem. We consider $\mathcal{W}_i = \{\omega_{i,1}, \omega_{i,2} \mid 10^{-2} \leq \omega_{i,1}, \omega_{i,2} \leq 10^2\}$, for $i = 1, 2$ and create a grid of size 9×9 linearly spaced in log scale for ω_2 . For each ω_2 in the grid, we employ Bayesian optimization to find the optimal value for ω_1 . The average true cost of MPC is computed by averaging the true cost values in $n_s = 100$ i.i.d. simulations with uniformly distributed initial positions and velocities. The parameters in the Bayesian optimization algorithm and in the true cost are chosen as $j_{\max} = 25$, $j_{\text{init}} = 5$, $\alpha = 1.0$, $\beta = 1.0$, $\lambda = 10^3$, $r = 10.0$ (m), $\xi = 10.0$. The grid and corresponding solutions are then used as training data for GP regression to learn the weight adaptation strategy. The derived optimal weight adaptation strategy can be illustrated by heat maps in Fig. 2.

In the testing simulations, we generate the actions of the human drivers by using the solution of (4) in which

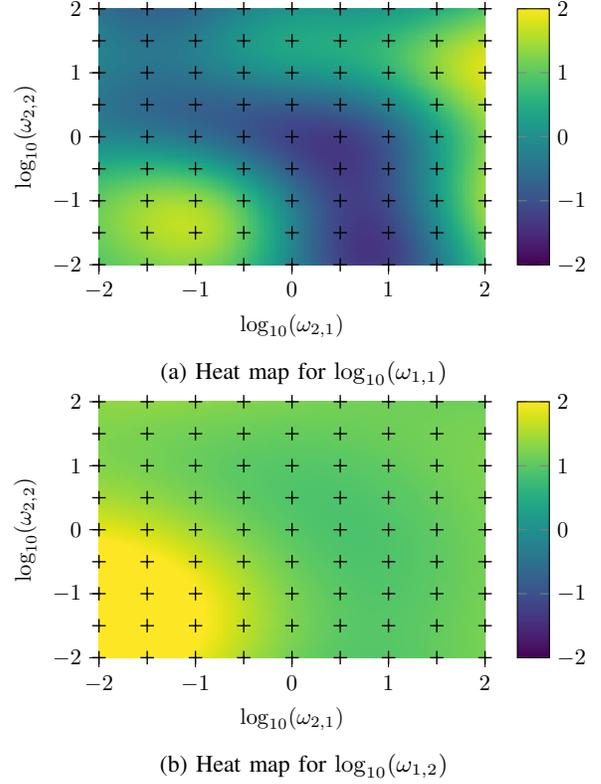


Fig. 2: Heat maps for the optimal weight adaptation strategy (in log scale). The black crosses represent the sampled values of $\log_{10}(\omega_2)$.

the weights are varied to imitate different driving behavior. Note that in all the simulations the HDV-2's objective weights are unknown to CAV-1 and must be learned online by moving horizon IRL. The parameters of MPC and moving horizon IRL are chosen as: $\Delta T = 0.2$ s, $H = 10$, $\gamma = 1.0$, $v_{\min} = 0.0$ (m/s), $v_{\max} = 12.0$ (m/s), $u_{\min} = -5.0$ (m/s²), $u_{\max} = 3.0$ (m/s²), $L = 20$, $\eta = 0.01$. The simulation is implemented in Julia programming language, and KNITRO solver [25] is used for solving MPC problems. The code for simulations is available online at <https://github.com/vietanhle0101/MPC-BayesOpt-Mixed-Traffic>.

B. Results and Discussion

Using the obtained strategy for MPC weight adaptation, we first evaluate the control framework in two specific simulations with an altruistic driver and with an egoistic driver to demonstrate that CAV-1 behaves differently depending on human driving behavior. The video for those simulations can be found in <https://sites.google.com/view/ud-ids-lab/mpc-bayesopt>.

Comparison with a baseline strategy: We compare the performance of the optimal weight adaptation strategy with a baseline strategy using SVO [17]. To extensively assess the benefits of the proposed method, we conduct 5000 simulations with different initial conditions of the vehicles and heterogeneous driving styles of the human drivers. First, we compare by two metrics: (1) the number of simulations without unsafe situations, and (2) the number of simulations

TABLE I: Comparison between weight adaptation strategies using Bayesian optimization (BayOpt) and SVO.

| Comparison metrics | BayOpt | SVO |
|-----------------------------------------------------------------|--------------|--------------|
| Number of simulations with safety | 4983 (99.7%) | 4981 (99.6%) |
| Number of simulations with time-energy improvement ¹ | 3921 (79.0%) | 1043 (21.0%) |

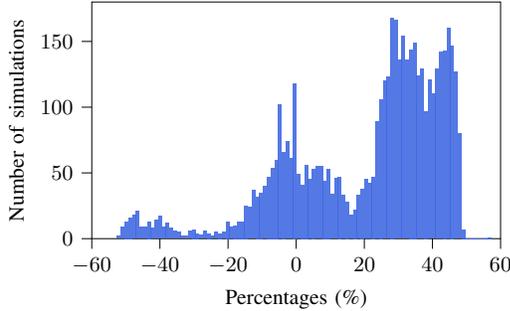


Fig. 3: A histogram for percentages of improvement in 5000 simulations.

with time-energy improvement among all the configurations in which using both strategies do not cause unsafe situations, as indicated in Table I. It can be observed that with a roughly similar level of safety (higher than 99%), MPC weight adaptation with the optimal strategy performs better than with the socially cooperative strategy in approximately 80% of the simulations. Furthermore, we also compute the percentages of improvement in time-energy costs and show the results in a histogram form in Fig. 3. We have been able to improve the average performance by 20.7%.

VI. CONCLUSIONS

In this paper, we presented a method to derive an optimal weight adaptation strategy of MPC for CAVs in mixed traffic with Bayesian optimization. By numerical simulations of a vehicle crossing example at an unsignalized intersection, we showed that the proposed optimal weight adaptation strategy has approximately 20% improvement on average over a baseline strategy using social value orientation. As a future research direction, we plan to focus on (1) enhancing the framework with a safety-guarantee mechanism, and (2) validating it in an experimental testbed [26].

REFERENCES

- [1] J. Guanetti, Y. Kim, and F. Borrelli, “Control of connected and automated vehicles: State of the art and future challenges,” *Annual reviews in control*, vol. 45, pp. 18–40, 2018.
- [2] T. Ersal, I. Kolmanovsky, N. Masoud, N. Ozay, J. Scruggs, R. Vasudevan, and G. Orosz, “Connected and automated road vehicles: state of the art and future challenges,” *Vehicle system dynamics*, vol. 58, no. 5, pp. 672–704, 2020.
- [3] A. A. Malikopoulos, L. E. Beaver, and I. V. Chremos, “Optimal time trajectory and coordination for connected and automated vehicles,” *Automatica*, vol. 125, no. 109469, 2021.
- [4] B. Chalaki and A. A. Malikopoulos, “Optimal control of connected and automated vehicles at multiple adjacent intersections,” *IEEE Transactions on Control Systems Technology*, vol. 30, no. 3, pp. 972–984, 2022.

¹Among all the configurations in which using both strategies can avoid unsafe situations.

- [5] H. Bang, B. Chalaki, and A. A. Malikopoulos, “Combined Optimal Routing and Coordination of Connected and Automated Vehicles,” *IEEE Control Systems Letters*, vol. 6, pp. 2749–2754, 2022.
- [6] A. M. I. Mahbub, V.-A. Le, and A. A. Malikopoulos, “Safety-aware and data-driven predictive control for connected automated vehicles at a mixed traffic signalized intersection,” in *10th IFAC International Symposium on Advances in Automotive Control*. IFAC, 2022, pp. 51–56.
- [7] J. Wang, Y. Zheng, Q. Xu, and K. Li, “Data-driven predictive control for connected and autonomous vehicles in mixed traffic,” in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 4739–4745.
- [8] B. Chalaki, L. E. Beaver, B. Remer, K. Jang, E. Vinitzky, A. Bayen, and A. A. Malikopoulos, “Zero-shot autonomous vehicle policy transfer: From simulation to real-world via adversarial learning,” in *IEEE 16th International Conference on Control & Automation (ICCA)*, 2020, pp. 35–40.
- [9] R. Valiente, B. Toghi, R. Pedarsani, and Y. P. Fallah, “Robustness and adaptability of reinforcement learning-based cooperative autonomous driving in mixed-autonomy traffic,” *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 397–410, 2022.
- [10] R. Chandra and D. Manocha, “Gameplan: Game-theoretic multi-agent planning with human drivers at intersections, roundabouts, and merging,” *IEEE Robotics and Automation Letters*, 2022.
- [11] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, “Social behavior for autonomous vehicles,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, 2019.
- [12] L. Wang, L. Sun, M. Tomizuka, and W. Zhan, “Socially-compatible behavior design of autonomous vehicles with verification on real human data,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3421–3428, 2021.
- [13] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, “Learning-based model predictive control: Toward safe learning in control,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020.
- [14] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, “Automatic lqr tuning based on gaussian process global optimization,” in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 270–277.
- [15] S. Gros and M. Zanon, “Data-driven economic nmPC using reinforcement learning,” *IEEE Transactions on Automatic Control*, vol. 65, no. 2, pp. 636–648, 2019.
- [16] A. Jain, L. Chan, D. S. Brown, and A. D. Dragan, “Optimal cost design for model predictive control,” in *Learning for Dynamics and Control*. PMLR, 2021, pp. 1205–1217.
- [17] V.-A. Le and A. A. Malikopoulos, “A Cooperative Optimal Control Framework for Connected and Automated Vehicles in Mixed Traffic Using Social Value Orientation,” in *2022 61th IEEE Conference on Decision and Control (CDC)*, 2022, pp. 6272–6277.
- [18] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, et al., “Maximum entropy inverse reinforcement learning,” in *Aaai*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.
- [19] J. R. Marden, G. Arslan, and J. S. Shamma, “Cooperative control and potential games,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 6, pp. 1393–1407, 2009.
- [20] M. Kuderer, S. Gulati, and W. Burgard, “Learning driving styles for autonomous vehicles from demonstration,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2641–2646.
- [21] P. I. Frazier, “A tutorial on bayesian optimization,” *arXiv preprint arXiv:1807.02811*, 2018.
- [22] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [23] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient global optimization of expensive black-box functions,” *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, 1998.
- [24] M. A. S. Kamal, M. Mukai, J. Murata, and T. Kawabe, “Model predictive control of vehicles on urban roads for improved fuel economy,” *IEEE Transactions on control systems technology*, vol. 21, no. 3, pp. 831–841, 2012.
- [25] R. H. Byrd, J. Nocedal, and R. A. Waltz, “Knitro: An integrated package for nonlinear optimization,” in *Large-scale nonlinear optimization*. Springer, 2006, pp. 35–59.
- [26] B. Chalaki, L. E. Beaver, A. M. I. Mahbub, H. Bang, and A. A. Malikopoulos, “A research and educational robotic testbed for real-time control of emerging mobility systems: From theory to scaled experiments,” *IEEE Control Systems*, vol. 42, no. 6, pp. 20–34, 2022.