

LPV-MPC for Lateral Control in Full-Scale Autonomous Racing

Hassan Jardali, Ihab S. Mohamed, Durgakant Pushp and Lantao Liu

Abstract—Autonomous racing has attracted significant attention recently, presenting challenges in selecting an optimal controller that operates within the onboard system’s computational limits and meets operational constraints such as limited track time and high costs. This paper introduces a Linear Parameter-Varying Model Predictive Controller (LPV-MPC) for lateral control. Implemented on an IAC AV-24, the controller achieved stable performance at speeds exceeding 160 mph (71.5 m s^{-1}). We detail the controller design, the methodology for extracting model parameters, and key system-level and implementation considerations. Additionally, we report results from our final race run, providing a comprehensive analysis of both vehicle dynamics and controller performance. Moreover, a Python implementation of the framework can be accessed here: <https://tinyurl.com/LPV-MPC-acados>

Index Terms—Automated vehicles, autonomous racing, path tracking, MPC, lateral control, vehicle dynamics.

I. INTRODUCTION

CONTROLLING an autonomous vehicle safely and efficiently across diverse driving conditions requires addressing key challenges in vehicle dynamics, actuation limits, and environmental uncertainties. A control strategy must account for the intricate coupling between longitudinal and lateral dynamics, nonlinear tire behavior, and variations in tire-road friction and actuator response. These challenges become even more pronounced at high speeds, where modeling errors can lead to higher tracking errors. Autonomous racing serves as an ideal testbed for evaluating control strategies for vehicles operating at the limits of their physical capabilities. Racing competitions such as the Indy Autonomous Challenge (IAC) [1] and the Abu Dhabi Autonomous Racing League (A2RL) [2] enable direct benchmarking of control approaches, providing valuable insights for real-world autonomous driving.

In these competitions, teams participate in events such as time trials, passing, and multi-vehicle races, each requiring control strategies that balance precision, robustness, and adaptability. Our team, the IU-Luddy Autonomous Racing Team (IU-LART), competed in the Indy Autonomous Challenge (IAC) at the Indianapolis Motor Speedway (IMS) on September 6, 2024, and at the Las Vegas Motor Speedway (LVMS) on January 9, 2025 [3]. Despite limited practice track time, we achieved a maximum speed of 163 mph at LVMS using an IAC AV-24 (see Fig. 1). In this article, we focus on the motion control algorithm deployed during the final IAC race at LVMS and analyze its performance under high-speed racing conditions.

Authors are with the Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN 47408 USA (e-mail: {hjardali, mohamedi, dpushp, lantao}@iu.edu).



Fig. 1: IU-LART IAC AV-24 autonomous racing vehicle [3].

II. RELATED WORK

Although oval-shaped racetracks provide a structured environment similar to freeways, high-speed driving introduces complexities that challenge conventional control strategies. At these speeds, nonlinear vehicle dynamics, model uncertainties, such as empirically estimated tire parameters, and external disturbances, including wind and track inclination, significantly impact stability. Consequently, it is essential to develop the right control strategy to achieve reliable path tracking and safe handling of the car in high-performance autonomous racing [4]–[6].

Typically, the control module consists of two components: longitudinal control, which controls the car’s longitudinal motion, and lateral control, which controls its steering. Consequently, control frameworks for path trajectory tracking are generally categorized as *coupled* or *decoupled* approaches [6], [7]. Coupled strategies jointly optimize both lateral (steering) and longitudinal (acceleration/braking) inputs, capturing the intricate dependencies between vehicle dynamics to enhance tracking accuracy and stability [4], [8]–[11]. For instance, a learning-based MPC framework integrates Gaussian Process Regression (GPR) to take residual model uncertainty into account, allowing for improved lap time while maintaining the vehicle’s safety [8]. In [9], the authors used an LPV-MPC for longitudinal and lateral control to track an offline-generated racing path; however, its applicability to high-speed racing environments remains uncertain, as it has not been validated on a full-scale autonomous racing car [9]. Additionally, a tricycle model-based MPC explicitly accounts for the locked differential, improving stability and precision in high-speed maneuvers, especially on road-course tracks by coupling steering and acceleration dynamics [11]. However, these approaches introduce trade-offs, as the learning-based method incurs high computational costs and requires continuous data updates, and the tricycle model approach introduces modeling complexity and demands precise system identification.

Due to the challenges associated with coupled strategies, decoupled control remains widely adopted in both autonomous driving and racing scenarios, as it simplifies controller synthesis. In decoupled control frameworks, the system dynamics are typically split into longitudinal and lateral motion, with longitudinal behavior often modeled as a linear first-order system, while lateral dynamics require more sophisticated control strategies due to their inherent complexity [12], [13]. This complexity necessitates greater attention in racing applications, where precise steering control is crucial for maintaining stability and achieving optimal trajectory tracking at high speeds. Geometric-based algorithms for lateral motion, such as Pure Pursuit [14] and Stanley [15], are among the widely utilized path-tracking strategies due to their simplicity and computational efficiency. However, the performance of these kinematic models is limited by their core assumption of a no-slip condition. This premise, which assumes the tires roll perfectly without sliding, breaks down in high-speed racing where significant tire slip is essential for generating the large forces needed to maneuver at the limit of adhesion.

To overcome these limitations, optimization-based control strategies move beyond kinematic assumptions by using dynamic models that explicitly consider the relationship between tire slip and vehicle forces [16], [17]. In full-scale autonomous racing literature, [16] developed an LQR-based lateral tracking algorithm using a single-track lateral dynamics model and using a look-ahead search strategy to determine the target point along the path. Similarly, [17] employs an LQR controller but enhances it by incorporating the desired yaw rate and explicitly accounting for the track's banking angle, improving trajectory tracking and stability. Since both methods assume a constant longitudinal velocity, they incorporate a velocity-dependent weight scheduling. Additionally, due to the LQR framework, neither explicitly manages constraints. MPC offers a more robust alternative to LQR for lateral vehicle control by addressing key limitations, including the assumption of constant longitudinal velocity and the lack of explicit constraint handling. Outside the full-scale autonomous racing literature, [18] proposes an adaptive MPC with Laguerre functions for path tracking using an augmented single-track lateral dynamics model and optimization of the controller parameters and cost function weights with a Particle Swarm Optimization (PSO) algorithm.

In contrast to existing solutions, and given the constraints of limited physical access to the race car due to competition policies and restricted practice runs on the track, we present in this work a linear parameter-varying model predictive control (LPV-MPC) framework designed for full-scale high-speed autonomous racing while highlighting the implementation steps. The contributions of this work can be summarized as:

- 1) We present an LPV-MPC framework for lateral control and path tracking, based on a vehicle dynamic model that is formulated as an LPV system scheduled by longitudinal speed, trajectory curvature, and road banking.
- 2) We conduct a real-world evaluation on a full-scale autonomous race car, demonstrating stable control performance at speeds exceeding 71.5 m s^{-1} (160 mph), coupled with a system identification detailed technique, and an

in-depth analysis. This analysis includes a quantitative evaluation of performance and a discussion of the method's limitations.

- 3) We release an open-source, Python-based implementation¹, providing a reference framework for the research community to advance autonomous racing research.

The paper is organized as follows: Section III presents the single-track lateral dynamic model and identification of parameters. Section IV details the proposed LPV-MPC formulation. Section V covers the experimental setup, followed by performance analysis. Finally, Section VI summarizes key findings and outlines future research directions.

III. LATERAL VEHICLE DYNAMICS

This section derives the single-track lateral dynamic model and extends it by incorporating the steering angle as a state variable. It also details the system identification method used to estimate the dynamic parameters.

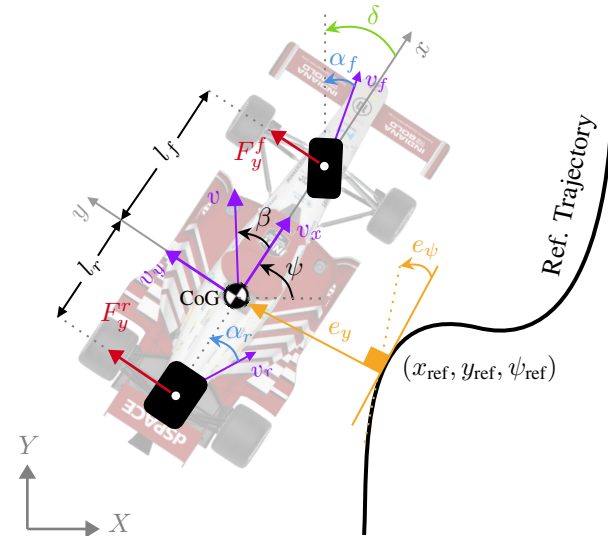


Fig. 2: Overview of the single-track bicycle model, highlighting relevant forces, angles, and lateral errors in their positive directions.

A. Single-Track Lateral Dynamic Model

Modeling vehicle dynamics accurately is challenging due to (i) varying road-tire friction and vertical loads; (ii) speed- and steering-dependent dynamics; and (iii) coupled lateral and longitudinal behavior. Owing to the typical symmetry of road vehicles, kinematic and dynamic bicycle (single-track) models are widely used to simplify vehicle dynamics, modeling the car as a rigid body with mass m and yaw inertia I_z , as depicted in Fig. 2. The single-track model assumes perfect lateral symmetry; however, although the race car considered in this study has an asymmetrical weight distribution due to its oval-track setup, this effect is neglected in the proposed model. To derive the governing equations of lateral dynamics, Newton's laws of motion are applied along the y - and z -axes, summing the lateral forces and yaw moment at the vehicle's center of gravity (CoG). The governing equations are given by

$$\begin{cases} \sum F_y = F_y^f + F_y^r = m(\ddot{y} + \dot{\psi}v_x), \\ \sum M_z = l_f F_y^f - l_r F_y^r = I_z \ddot{\psi}, \end{cases} \quad (1)$$

¹ <https://tinyurl.com/LPV-MPC-acados>

where F_y^f and F_y^r are the lateral forces at the front and rear tires, respectively, M_z is the resultant yaw moment, v_x is the longitudinal velocity, $\dot{\psi}$ and $\ddot{\psi}$ denote the yaw rate and yaw acceleration, l_f and l_r represent the distances from the CoG to the front and rear axles, and \ddot{y} is the lateral acceleration of the CoG. To derive such a linear model, several fundamental assumptions are made. The slip angles are considered small, the tires operate within the linear region of the slip-angle–lateral-force relationship, and the road–tire friction coefficient is assumed to be constant. Under these conditions, the lateral tire forces are approximated as

$$F_y^f \approx 2C_f\alpha_f, \quad F_y^r \approx 2C_r\alpha_r, \quad (2)$$

where C_f and C_r represent the cornering stiffnesses of the front and rear tires, respectively. To further describe the lateral dynamics, the front and rear slip angles, α_f and α_r , are introduced. These angles characterize the deviation between the tire direction and the actual velocity vector at each axle (see Fig. 2). Specifically, they are defined as $\alpha_f = \delta - \theta_f$ and $\alpha_r = -\theta_r$, where δ is the steering angle. Under small slip angles assumption, the front and rear velocity angles are approximated as $\theta_f \approx \beta + \frac{l_f\dot{\psi}}{v_x}$ and $\theta_r \approx \beta - \frac{l_r\dot{\psi}}{v_x}$, where $\beta = \arctan(\frac{v_y}{v_x}) \approx \frac{v_y}{v_x}$ is the sideslip angle, representing the lateral-to-longitudinal velocity ratio.

Building on this formulation, we define the model in terms of reference path tracking errors, as this approach minimizes complexity in the cost function, reduces computational overhead, and simplifies both the design and tuning of the controller. This choice is particularly important in high-speed racing, where precise tracking is crucial to maintaining stability and performance, as even small deviations can have significant effects. The lateral position error e_y is defined as the signed perpendicular distance from the vehicle’s CoG to the desired raceline, commonly referred to as the cross-track error (CTE). Similarly, the heading error e_ψ is the difference between the vehicle’s yaw angle and the desired yaw angle, given by $e_\psi = \psi - \psi_{\text{ref}}$, where ψ and ψ_{ref} are the global yaw angle of the car and the desired yaw angle at the nearest point on the raceline, respectively. To incorporate the desired yaw rate, we define $\dot{\psi}_{\text{ref}} = v_x\kappa$, where κ is the instantaneous curvature of the local path.

Building on these tracking error definitions, the lateral vehicle dynamics in the curvilinear frame can be formulated as a state-space model in tracking error variables, leading to a linearized single-track model in matrix form given by

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} + \mathbf{C}\dot{\psi}_{\text{ref}}, \quad (3)$$

where $\mathbf{x} = [e_y, \dot{e}_y, e_\psi, \dot{e}_\psi]^\top \in \mathbb{R}^4$, is the state vector representing the system’s errors, $\mathbf{u} = \delta \in \mathbb{R}$ is the control input, $\dot{\psi}_{\text{ref}}$ is the external input derived from the reference trajectory, and $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are matrices defined as

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & a_{22} & a_{23} & a_{24} \\ 0 & 0 & 0 & 1 \\ 0 & a_{42} & a_{43} & a_{44} \end{bmatrix}, \mathbf{B} = \begin{bmatrix} 0 \\ b_{21} \\ 0 \\ b_{41} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} 0 \\ a_{24} - v_x \\ 0 \\ a_{44} \end{bmatrix},$$

with

$$a_{22} = -\frac{2(C_f+C_r)}{m v_x}, a_{23} = -v_x a_{22}, a_{24} = \frac{-2C_f l_f + 2C_r l_r}{m v_x},$$

$$a_{42} = -\frac{2(C_f l_f - C_r l_r)}{I_z v_x}, \quad a_{43} = -v_x a_{42},$$

$$a_{44} = -\frac{2(C_f l_f^2 + C_r l_r^2)}{I_z v_x}, \quad b_{21} = \frac{2C_f}{m}, \quad b_{41} = \frac{2C_f l_f}{I_z}.$$

Controlling a high-speed vehicle on an oval track requires smooth and precise steering to maintain stability. To better capture actuator dynamics while systematically handling steering constraints, we extend the model presented in (3) by incorporating the steering angle δ as an additional state, while treating the steering rate $\dot{\delta}$ as the control input, expressed as $\mathbf{u} = [\delta] \in \mathbb{R}$. This augmentation enables the explicit integration of steering constraints within the MPC framework, including lag and rate limitations. As a result, the augmented model is formulated as:

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ 0 & 0 \end{bmatrix}, \tilde{\mathbf{B}} = [0, 0, 0, 0, 1]^\top, \tilde{\mathbf{C}} = \begin{bmatrix} \mathbf{C} \\ 0 \end{bmatrix}.$$

Beyond steering dynamics, external factors such as road banking significantly influence lateral stability, particularly at high speeds. To account for these effects, we introduce the gravitational term $g \sin(\phi)$ into the lateral dynamics, where ϕ denotes the banking angle, reaching up to 20° at the testing track. This additional lateral gravitational component redistributes the cornering forces, reducing the load borne solely by the tires and enabling higher sustained cornering speeds. However, this effect also necessitates explicit compensation in the control model to prevent oversteer or instability. By incorporating ϕ , the model accounts for the gravitational assist in turns, leading to a more accurate prediction of the required lateral tire forces. Thus, the augmented state-space representation, including road banking effects, is formulated as:

$$\dot{\tilde{\mathbf{x}}} = \tilde{\mathbf{A}}\tilde{\mathbf{x}} + \tilde{\mathbf{B}}\mathbf{u} + \tilde{\mathbf{C}}\dot{\psi}_{\text{ref}} + D_\phi, \quad (4)$$

where the augmented state vector is defined as $\tilde{\mathbf{x}} = [e_y, \dot{e}_y, e_\psi, \dot{e}_\psi, \delta]^\top \in \mathbb{R}^5$, and the disturbance matrix due to road banking is given by $D_\phi = [0, g \sin(\phi), 0, 0, 0]^\top \in \mathbb{R}^5$. Since D_ϕ represents the effect of road banking, which is determined by track geometry rather than a controllable input, it is treated as a disturbance in the state-space model. A detailed derivation of the model and a discussion of its assumptions are provided in Chapter 2 of [19].

B. Dynamics Parameters Identification

To effectively implement the model given in (4) in a real-world racing scenario, key model parameters must be estimated, particularly under high-speed conditions where precise vehicle dynamics modeling is essential for stability and performance. Among these parameters, the cornering stiffness values C_f and C_r play a crucial role, as they directly influence the lateral force estimation and controller response. However, obtaining accurate estimates of these parameters is challenging due to measurement noise, costly track operations, and an increased risk of crashes during parameter tuning. In the autonomous racing literature, optimization-based methods [20] and advanced software solutions [21] have been utilized. Alternatively, learning-based approaches—such as hyperparameter optimization [22] and physics-informed neural networks [23]—offer data-driven strategies that improve model accuracy and reduce manual tuning. These techniques typically rely on

the nonlinear Magic Formula (Pacejka) tire model [24], which is commonly formulated as:

$$F_y^i = D_p^i \sin \left[C_p^i \arctan \left(B_p^i \alpha_i - E_p^i (B_p^i \alpha_i - \arctan(B_p^i \alpha_i)) \right) \right], \quad (5)$$

where F_y^i is the lateral tire force and α_i is the slip angle. The index i denotes the front (f) and rear (r) tires. The parameters are B_p^i , the stiffness factor that scales slip; C_p^i , the shape factor that influences curvature; D_p^i , the peak factor that defines the maximum force; and E_p^i , the curvature factor that adjusts peak behavior.

We determined the linear cornering stiffness values C_f and C_r following the paradigm of [20]. In our work, Pacejka tire model parameters are estimated via an iterative expectation-maximization (EM) approach that employs a nonlinear least-squares optimization. The procedure iteratively refines the fit by excluding outliers identified using a Median Absolute Deviation (MAD) derived threshold until convergence is reached. The linear dynamics, used in (2), are then approximated using the relation $C_{\text{linear}} = B_p C_p D_p$. To estimate our initial dynamics parameters, we used data from our previous highest-speed lap, in which the vehicle, controlled via Pure Pursuit, completed a lap at 56 m s^{-1} . The lateral forces at the front and rear tires were then computed from the measured lateral acceleration $a_{y,\text{imu}}$ recorded by an Inertial Measurement Unit (IMU) using the relationships:

$$F_y^f = \frac{m l_r}{(l_r + l_f) \cos \delta} a_{y,\text{imu}}, \quad F_y^r = \frac{m l_f}{l_r + l_f} a_{y,\text{imu}}. \quad (6)$$

Using (5) and (6), we ran our optimization algorithm to determine the Pacejka parameters. These parameters were used to find the linear dynamics and served as initial inputs for tuning the MPC. Table I presents the optimized Pacejka parameters alongside the linear approximation. Fig. 3 illustrates the fitted Pacejka curve overlaid on the measured data, showing the lateral forces versus slip angles for both the front and rear tires. As an additional point of comparison, we implemented a

TABLE I: Initial Pacejka Parameters values using IMS data. D_p is in Newtons (N), and C_{linear} is in N/rad.

Tires	B_p	C_p	D_p	E_p	C_{linear}
Front	34.59	1.81	2100	-1.0	132 k
Rear	35.04	1.96	3036	-0.26	209 k




Fig. 3: Measured cornering forces versus slip angles with the corresponding fitted Pacejka curve, obtained from the final IMS run at a longitudinal velocity of 56 m s^{-1} . Note that the vertical offset in the Pacejka model (typically denoted as $s_{h,y}$) was not accounted for. Physics-Informed Neural Network (PINN) approach following

[23]. Among the tested architectures, the PINN employing Gated Recurrent Units (GRUs) produced parameter estimates most consistent with those obtained from our EM optimization method. This agreement provides complementary evidence of methodological consistency across independent estimation frameworks. Subsequently, during practice sessions, these parameters were further fine-tuned. In Section V-B, we present the estimated parameters derived from the data at 162 mph using our EM optimization method.

IV. PATH TRACKING WITH LPV-MPC

Building on the previously defined augmented lateral dynamics model, we now formulate the LPV-MPC strategy for high-speed autonomous racing.

A. Path Tracking Problem Formulation

Consider a smooth reference trajectory consisting of N_r discrete reference points, denoted as the reference state $\mathbf{x}_{\text{ref},k} = [x_{\text{ref},k}, y_{\text{ref},k}, \psi_{\text{ref},k}, v_{\text{ref},k}]^\top$, which specifies a desired position $(x_{\text{ref},k}, y_{\text{ref},k})$ in the global frame, yaw angle $\psi_{\text{ref},k}$, and velocity $v_{\text{ref},k}$ along the planned trajectory. Let us assume that the current vehicle state is given by $\mathbf{x}_k = [x_k, y_k, \psi_k, v_{x,k}, v_{y,k}, \dot{\psi}_k]^\top$, where x_k and y_k represent the vehicle's position in the global frame, ψ_k is the yaw angle, $v_{x,k}$ and $v_{y,k}$ are the longitudinal and lateral velocities, respectively, and $\dot{\psi}_k$ is the yaw rate. The reference trajectory is generated by the planning module, while the current vehicle state is provided by the state estimation module, as described in Section V-A. Since the system operates in the error space, the state vector $\tilde{\mathbf{x}}_k$, which is formally defined in Section III, is derived from the lateral components of the actual vehicle state \mathbf{x}_k and the reference state $\mathbf{x}_{\text{ref},k}$. Considering a finite time-horizon N , the sequence of control inputs is defined as $\mathbf{U} = [\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{N-1}]^\top \in \mathbb{R}^N$, while the corresponding state trajectory of the system is $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_0, \tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_N]^\top \in \mathbb{R}^{5(N+1)}$.

The objective of the lateral controller is to minimize the tracking errors in $\tilde{\mathbf{x}}_k$, ensuring accurate path following within the given system constraints, while accounting for variations in curvature, speed, and external disturbances to maintain stability. This is achieved by generating an optimal sequence of steering inputs \mathbf{U} over a finite prediction horizon N . To formulate the optimization problem, the continuous-time LPV model (4) is first discretized. Assuming the control inputs and scheduling parameters are held constant over a sampling interval T_s (a Zero-Order Hold assumption), the exact discretization of (4) yields the following discrete-time LPV model:

$$\tilde{\mathbf{x}}_{k+1} = \tilde{A}_d(\mathbf{p}_k) \tilde{\mathbf{x}}_k + \tilde{B}_d(\mathbf{p}_k) \mathbf{u}_k + \tilde{\mathbf{E}}_d(\mathbf{p}_k) \quad (7)$$

where the discrete-time matrices \tilde{A}_d , \tilde{B}_d , along with the affine term $\tilde{\mathbf{E}}_d$, which accounts for the effects of $\dot{\psi}_{\text{ref}}$ and ϕ , are derived from their continuous-time counterparts and depend on the scheduling parameters vector $\mathbf{p} = [v_x, \kappa, \phi]^\top \in \mathbb{R}^3$, which consists of the longitudinal velocity v_x , road curvature κ , and road banking angle ϕ . This parameter-dependent structure allows the proposed control strategy to be inherently adaptive to varying speed profiles and track geometry. Consequently, the matrices \tilde{A}_d , \tilde{B}_d and $\tilde{\mathbf{E}}_d$ are updated within the optimization

problem at each time step k . The scheduling parameter vector \mathbf{p} is required for the entire prediction horizon. At the current time step ($k = 0$), \mathbf{p} is known exactly from real-time sensor measurements. For future steps within the horizon ($k = 1, \dots, N$), a prediction strategy is employed: the future longitudinal velocity v_x and road curvature κ are extracted from the pre-computed reference trajectory, while the road banking angle ϕ is assumed to remain constant at its most recently measured value.

As a result, this formulation yields an LPV-MPC path-tracking algorithm that conditions its optimal control problem (OCP) on the real-time measurement of the scheduling parameters, overcoming limitations imposed by fixed model assumptions, such as assuming constant velocity or neglecting banking and race-line curvature effects, that were assumed previously in full-scale autonomous racing (see [16], [17]).

B. Constrained Control Problem

To formulate the path-tracking task, we define a constrained nonlinear least-squares optimal control problem. The objective function incorporates penalties on state tracking errors, control effort, and the side-slip angle:

$$J = \sum_{k=0}^{N-1} (\|\tilde{\mathbf{x}}_k\|_Q^2 + \|\mathbf{u}_k\|_R^2 + Q_\beta \beta_k(\tilde{\mathbf{x}}_k)^2), \quad (8)$$

where the weighting matrix $Q \in \mathbb{R}^{5 \times 5}$ is diagonal and penalizes deviations in the state variables, while $R \in \mathbb{R}$ is a positive scalar that penalizes control effort.

Following insights from [11], we include a side-slip penalty to enhance vehicle stability by discouraging large slip angles. The side-slip angle is computed from the predicted states and is defined as $\beta_k(\tilde{\mathbf{x}}_k) = \arctan\left(\frac{\dot{e}_{y,k}}{v_{x,k}}\right)$, where $\dot{e}_{y,k}$ is the predicted lateral velocity error and $v_{x,k}$ is the longitudinal velocity treated as a scheduling parameter. Because β_k depends nonlinearly on the decision variable $\tilde{\mathbf{x}}_k$, its inclusion renders the overall objective function nonlinear. The optimal control problem can finally be formulated as:

$$\min_{\tilde{\mathbf{x}}, \mathbf{U}} J(\tilde{\mathbf{x}}, \mathbf{u}) \quad (9a)$$

$$\text{s.t. } \tilde{\mathbf{x}}_{k+1} = \tilde{A}_d(\mathbf{p}_k) \tilde{\mathbf{x}}_k + \tilde{B}_d(\mathbf{p}_k) \mathbf{u}_k + \tilde{\mathbf{E}}_d(\mathbf{p}_k), \quad k = 0, \dots, N-1, \quad (9b)$$

$$\delta_{\min} \leq \delta_k \leq \delta_{\max}, \quad k = 0, \dots, N-1, \quad (9c)$$

$$\mathbf{u}_{\min} \leq \mathbf{u}_k \leq \mathbf{u}_{\max}, \quad k = 0, \dots, N-1. \quad (9d)$$

Our proposed LPV-MPC strategy solves this optimization problem by minimizing the objective function J (9a), taking into account the system dynamics with time-varying parameters \mathbf{p} (9b), the steering constraints (9c), which are implicitly incorporated within $\tilde{\mathbf{x}}$, and the control input limits (9d), which impose rate constraints on $\mathbf{u}_k = \dot{\delta}_k$. This process generates an optimal control sequence \mathbf{U} , from which only the first control input \mathbf{u}_0 is applied to the autonomous racing vehicle. It is worth noting that the terminal cost in (8) is omitted since our empirical evaluations confirmed that its inclusion yielded negligible performance benefits for the chosen prediction horizon, while omitting it improved computational efficiency.




Fig. 4: High-level software architecture of our autonomous racing system for a single-car format [3].

V. EXPERIMENTAL VALIDATION

In this section, we validate our proposed LPV-MPC strategy through real-world experiments. We first outline the hardware and software architecture, followed by the experimental results from our final run at the racing competition, along with a detailed analysis and discussion of the system's performance.

A. Experimental Setup

To validate our proposed control strategy, we conducted experiments using an IAC AV-24, a full-scale autonomous racing vehicle (see Fig. 1). The vehicle is equipped with advanced hardware and a Drive-by-Wire (DBW) system for autonomous operation, including a dSPACE Autera AutoBox powered by an Intel Xeon 3GHz 12-core CPU and an NVIDIA A5000 GPU. It runs Ubuntu 22.04 with CycloneDDS for real-time communication and leverages ROS2 as the core development framework. For completeness, we briefly outline state estimation and planning, as they are essential for running our controller effectively, although these modules are beyond the scope of this work. The state estimation module fuses GNSS, IMU, wheel odometry, and steering angle data using an Unscented Kalman Filter (UKF) [25], thereby achieving accurate real-time localization. The planning module consists of both offline and online components: the offline planner generates a raceline using the minimum curvature method [26] as a global path, while the online planner continuously tracks the nearest path segment and publishes a local-frame trajectory.

As illustrated in Fig. 4, the vehicle's control architecture follows a decoupled design, comprising separate longitudinal and lateral control modules, with well-defined signal flow between components. These modules process the vehicle's primary control inputs, which include throttle (0–100%), brake pressure (in kPa), gear command (ranging from gear 1 to 6), and steering angle (in degrees). The longitudinal controller regulates the vehicle's speed by modulating throttle and brake inputs, while also executing gear shifts. In this work, a Proportional-Integral-Derivative (PID) controller is employed to track the desired velocity by modulating throttle and brake inputs, while gear shifting is regulated based on engine RPM. The lateral control module is implemented using our proposed LPV-MPC, designed to handle the vehicle's lateral dynamics in high-speed scenarios by generating the desired steering rate $\dot{\delta}$, which is then integrated to produce the steering angle δ required to track the reference trajectory provided by the planning module. For safety, a Pure Pursuit (PP) controller continuously runs as a backup in case the optimization fails due to constraint violations, numerical instability at low speeds, or exceeding real-time computation limits.

The control system operates in real-time, with the planning and control loop running at 50 Hz, while state estimation updates at 125 Hz providing precise localization. The optimal control problem in (9) is solved using *acados* [27], with code adapted from [28]. Because the side-slip term β_k depends nonlinearly on the predicted states, the formulation constitutes a constrained nonlinear least-squares optimal control problem. We therefore employ the real-time iteration sequential quadratic programming (SQP-RTI) scheme, in which a single quadratic programming (QP) subproblem is constructed and solved at each control step using HPIPM [29]. A Gauss-Newton approximation is used for the Hessian of the nonlinear least-squares objective. The problem is discretized over a prediction horizon of $T = 1.6$ s with $N = 45$ intervals and solved using a fixed-step ERK integrator ($\epsilon = 10^{-4}$).

Due to velocity-dependent terms in the model (Section III-A), numerical ill-conditioning arises as $v_x \rightarrow 0$. The controller therefore switches to a backup Pure Pursuit controller below 45 mph (20 m s^{-1}). The weighting matrices Q , R , and Q_β were tuned in simulation and further refined offline between track sessions alongside the vehicle’s dynamic parameters.

B. Results

In this section, we present results from our on-track practice sessions and the final racing competition, in which our team participated exclusively in the time-trials format. Leading up to the event, we completed 16 on-track sessions, each averaging about 25 minutes. Despite the considerable challenges of high logistical costs, limited track access, and expenses related to car repairs, we maximized every practice session to focus on parameter tuning, race-line validation, and other competition-specific preparations. It is important to note that, given the competitive nature of autonomous racing, teams do not share code or internal implementations. As a result, direct modular or quantitative comparisons with other teams’ approaches are not possible. Instead, in this subsection, we highlight the results from our final racing competition, which represent the most compelling outcome of our efforts—our fully autonomous run on an oval track during the challenge.

During the final competition, the run included 12 high-speed laps. On the final lap, we achieved our highest recorded speed of 163.1 mph (72 m s^{-1}), with a maximum absolute cross-track error of 1.6 m. The best lap time was 34.18 s with an average speed of 157.98 mph (70.6 m s^{-1}). It is worth noting that the maximum velocity reached during the practice runs was only 145 mph (64.8 m s^{-1}).

1) *Tracking Errors*: The tracking errors consist of the cross-track error e_y measured in m and the heading error e_ψ measured in deg. Additionally, we present the commanded and actual longitudinal velocities, measured in m s^{-1} . Fig. 5(a) shows the cross-track and heading errors during the high-speed lap. The mean cross-track error is approximately -1.02 m, with a standard deviation of 0.38 m. The heading error remains within $\pm 1^\circ$ throughout the lap. As shown in Fig. 5(a), the cross-track error e_y remains negative over the entire lap—indicating the vehicle consistently tracked to the left of the reference line—and this persistent behavior is attributable to either a steering offset in the hardware or the race car’s non-

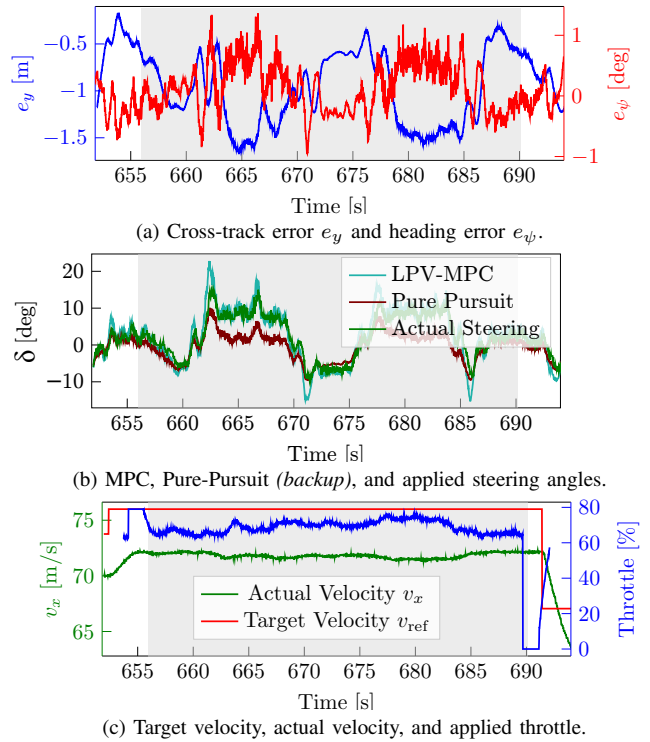


Fig. 5: Experimental results for path tracking during our highest speed lap. The highlighted areas mark the start and finish of the lap.




Fig. 6: Tracking errors vs. velocity during the complete run.

symmetrical weight configuration for oval racing (or both), neither of which is accounted for in our model. In Fig. 5(b), we compare the steering angle commands from the LPV-MPC controller and the backup PP controller with the actual steering applied by the vehicle. The LPV-MPC commands range from -20° to 20° , indicating no steering constraint violations. However, at $t = 663, 671, 678,$ and 686 s there is a large discrepancy between the applied steering and the LPV-MPC command, likely due to the latter’s rapid rate of change. This suggests the steering rate-limit constraints require further tuning. Notably, the MPC exhibits higher responsiveness, resulting in larger steering commands, whereas the PP controller takes a more gradual approach, demonstrating the MPC’s aggressive approach to minimizing lateral deviations. Finally, Fig. 5(c) shows the target velocity, actual velocity, and applied throttle. A persistent velocity error of approximately 4 m s^{-1} is observed, indicating that the longitudinal controller requires adjustment to improve speed tracking.

Fig. 6 illustrates the relationship between tracking errors and longitudinal velocity v_x across the 12 high-speed laps of the competition run with our proposed control strategy. The heading error e_ψ shows no clear dependence on velocity and remains within a relatively narrow band, whereas the cross-track error e_y increases approximately linearly with velocity.




Fig. 7: Measured cornering forces F_y versus slip angles, α_f and α_r , with the corresponding fitted Pacejka curve, obtained from the final run at a longitudinal velocity of 72 m s^{-1} . The *Linear Force* curve depicts the estimated force computed using the cornering stiffness parameters that were determined during the practice runs and employed during the final run.




Fig. 8: Lateral vs. longitudinal acceleration (g-g diagram) during the complete run. Accelerations were measured using the on-board IMU.

By characterizing e_y as a function of longitudinal velocity v_x , we anticipated tracking error at race speeds and tailored our strategy.

TABLE II: Pacejka parameter values obtained via EM optimization using our final run data. D_p is in Newtons [N], and C_{linear} is in [N/rad]. Note that these parameters have not been tested with the proposed controller.

Tires	B_p	C_p	D_p	E_p	C_{linear}
Front	22.30	2.00	3885.85	-1.00	173 k
Rear	26.08	2.00	5342.89	-1.00	278 k

2) *Dynamics and Model Errors*: During the final run, the vehicle achieved speeds exceeding those attained in practice. As shown in Fig. 7, slip angles reached up to 2 deg, briefly entering the nonlinear region. The figure shows that, at higher slip angles, the linear force model (orange; calibrated on practice data) overestimates the lateral forces estimated from the on-board IMU. Additionally, Table II lists the Pacejka parameters identified from the same IMU-derived final-run data using the method in Section III-B.

In Fig. 8, the vehicle's lateral and longitudinal accelerations (g-g diagram) over the entire run are shown, colored by longitudinal velocity v_x . At the highest longitudinal speeds, the car achieved a lateral acceleration a_y of up to 27 m s^{-2} (yellow points, $v_x \approx 70 \text{ m s}^{-1}$). It is worth noting that we adopted a conservative raceline with large boundary margins to ensure safety, resulting in higher curvature. Despite these constraints, the controller maintained stable performance throughout.




Fig. 9: Lateral model error vs. velocity during the final run.

Fig. 9 presents the lateral model error, computed offline as the one-step-ahead difference between the MPC-predicted cross-track error \hat{e}_y and the actual e_y from the state-estimation output, thus, the lateral model error can be formulated as $e_y^{\text{mdl}} = \hat{e}_y - e_y$. Beyond approximately 55 m s^{-1} , the maximum model error stabilizes around 0.06 m. However, during high-acceleration events (e.g., at 42 m s^{-1} and 57 m s^{-1}), the error spikes, reaching a peak of 0.14 m. This occurs when aggressive throttle inputs induce longitudinal slip in the rear tires, reducing their lateral force capacity—a dynamic coupling not fully captured by our model.




Fig. 10: Computation time t_{comp} of the control loop during the high-speed lap.

3) *Computation Time*: Fig. 10 demonstrates the real-time performance of the controller, with a mean computation time of 5.78 ms. Some spikes are attributed to other processes such as data logging, which intermittently increase CPU contention. Despite these transient delays, the control loop consistently operates within a 10 ms cycle time. These low computation times confirm the feasibility of operating at 100 Hz, and future iterations of this work will target this frequency to enhance command smoothness and reactivity.

C. Discussion

Despite the encouraging results, several limitations of this work warrant mention:

1) *Cross-Track Error*: The observed cross-track error was moderate to high. The primary causes were an unmodeled steering offset in the hardware, which we cannot calibrate or access; the car's asymmetric lateral weight distribution; inaccuracies in extracting vehicle-dynamic parameters; and suboptimal tuning of the weighting matrices. Due to limited track time and competition requirements, we prioritized controller stability over minimizing tracking error. While this trade-off was acceptable for single-car racing, it may be unsuitable in multi-vehicle contexts requiring precise trajectory tracking. With more time, improved parameter identification and better tuning of the weighting matrices and constraints could have reduced the cross-track error.

2) *Terminal Cost*: Although removing the terminal cost had no noticeable effect in our tests, retaining it is theoretically advisable. Future work should include a robust terminal cost

to improve closed-loop stability.

3) *Longitudinal Control*: For single-car oval racing with minimal speed variation, our simple velocity-based controller was adequate. However, for road-course tracks or multi-vehicle racing, a more advanced longitudinal control approach is required.

4) *Limitations of the Linear Model*: The LPV formulation treats v_x , κ , and ϕ as exogenous scheduling parameters, neglecting cross-sensitivities and preventing the controller from exploiting their variation within the prediction horizon. While track banking kept operation largely within the tire model's linear region, a flat track would induce nonlinear behavior at lower speeds. Future work should consider a fully nonlinear MPC formulation to better capture these coupled effects.

VI. CONCLUSION

In this study, we introduced the Linear Parameter-Varying Model Predictive Controller (LPV-MPC) for lateral control, as implemented by our team. The proposed controller enabled the racing vehicle to maintain bounded tracking errors at speeds exceeding 160 mph (71.5 m s^{-1}). Furthermore, we detailed a methodology for extracting dynamic parameters, supported by a comprehensive analysis of both vehicle dynamics and controller performance, while also addressing the limitations inherent in our approach. To promote reproducibility and stimulate further research, the Python-based source code for the framework is openly available. Future research efforts will focus on mitigating the identified limitations to further enhance performance and broaden the framework's applicability to increasingly complex autonomous racing environments, such as multi-vehicle racing on road-course tracks. In addition, we will conduct a systematic sensitivity analysis to rigorously evaluate the controller's robustness to model uncertainties and varying environmental conditions.

VII. ACKNOWLEDGMENTS

We thank all the students and researchers who contributed to the development of the software stack, especially Youwei Yu and Mahmoud Ali for their work on key modules and architectures. We are also grateful to the IAC organization for supporting this work. ChatGPT was used to help improve the presentation of the manuscript.

REFERENCES

- [1] "Indy Autonomous Challenge — indyautonomouschallenge.com," [Accessed 02-03-2025].
- [2] "Abu Dhabi Autonomous Racing League in UAE | A2RL — a2rl.io," [Accessed 02-03-2025].
- [3] H. Jardali, D. Pushp, Y. Yu, M. Ali, I. S. Mohamed, A. Murillo-Gonzalez, P. D. Coen, M. A.-M. Khan, R. C. Pulivendula, S. Park *et al.*, "From zero to high-speed racing: An autonomous racing stack," in *Proceedings of The 19th International Symposium on Experimental Robotics (ISER)*, 2025.
- [4] A. Wischnewski, T. Herrmann, F. Werner, and B. Lohmann, "A tube-MPC approach to autonomous multi-vehicle racing on high-speed ovals," *IEEE Trans. Intell. Veh.*, vol. 8, no. 1, pp. 368–378, 2022.
- [5] M. Reda, A. Onsy, A. Y. Haikal, and A. Ghanbari, "Path planning algorithms in the autonomous driving system: A comprehensive review," *Robotics and Autonomous Systems*, vol. 174, p. 104630, 2024.
- [6] J. Betz, H. Zheng, A. Liniger, U. Rosolia, P. Karle, M. Behl, V. Krovi, and R. Mangharam, "Autonomous vehicles on the edge: A survey on autonomous vehicle racing," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 458–488, 2022.

- [7] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Trans. Intell. Veh.*, vol. 1, no. 1, pp. 33–55, 2016.
- [8] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019.
- [9] E. Alcalá, V. Puig, J. Quevedo, and U. Rosolia, "Autonomous racing using linear parameter varying-model predictive control (LPV-MPC)," *Control Engineering Practice*, vol. 95, p. 104270, 2020.
- [10] A. Raji, A. Liniger, A. Giove, A. Toschi, N. Musiu, D. Morra, M. Verucchi, D. Caporale, and M. Bertogna, "Motion planning and control for multi vehicle autonomous racing at high speeds," in *IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, 2022, pp. 2782–2782.
- [11] A. Raji, N. Musiu, A. Toschi, F. Prignoli, E. Mascaro, P. Musso, F. Amerotti, A. Liniger, S. Sorrentino, and M. Bertogna, "A tricycle model to accurately control an autonomous racecar with locked differential," in *IEEE Int. Conf. Syst. Control (ICSC)*, 2023, pp. 782–789.
- [12] P. Stano, U. Montanaro, D. Tavernini, M. Tufo, G. Fiengo, L. Novella, and A. Sorniotti, "Model predictive path tracking control for automated road vehicles: A review," *Annual reviews in control*, vol. 55, pp. 194–236, 2023.
- [13] A. Artuñedo, M. Moreno-Gonzalez, and J. Villagra, "Lateral control for autonomous vehicles: A comparative evaluation," *Annual Reviews in Control*, vol. 57, p. 100910, 2024.
- [14] R. C. Coulter, "Implementation of the pure pursuit tracking algorithm," *Robotics Institute, Carnegie Mellon University*, 1992.
- [15] G. M. Hoffmann, C. J. Tomlin, M. Montemerlo, and S. Thrun, "Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing," in *American Control Conference (ACC)*, 2007, pp. 2296–2301.
- [16] A. Saba, A. Adetunji, A. Johnson, A. Kothari, M. Sivaprakasam, J. Spisak, P. Bharatia, A. Chauhan, B. Duff Jr, N. Gasparro *et al.*, "Fast and modular autonomy software for autonomous racing vehicles," *arXiv preprint arXiv:2408.15425*, 2024.
- [17] C. C. Chung, A. Finazzi, H. Seong, D. Lee, S. Lee, B. Kim, G. Gang, and D. H. Shim, "Autonomous system for head-to-head race: Design, implementation and analysis; team KAIST at the Indy autonomous challenge," *IEEE Transactions on Field Robotics*, 2024.
- [18] Y. Kebbaty, V. Puig, N. Ait-Oufroukh, V. Vigneron, and D. Ichlal, "Optimized adaptive MPC for lateral control of autonomous vehicles," in *Int. Conf. Control Mechatron. Autom. (ICCA)*, 2021, pp. 95–103.
- [19] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [20] J. Becker, N. Imholz, L. Schwarzenbach, E. Ghignone, N. Baumann, and M. Magno, "Model-and acceleration-based pursuit controller for high-performance autonomous racing," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 5276–5283.
- [21] A. Raji, D. Caporale, F. Gatti, A. Toschi, N. Musiu, M. Verucchi, F. Prignoli, D. Malatesta, A. F. Jesus, A. Finazzi *et al.*, "er. autopilot 1.1: A software stack for autonomous racing on oval and road course tracks," *IEEE Transactions on Field Robotics*, 2024.
- [22] H. Seong, C. Chung, and D. H. Shim, "Model parameter identification via a hyperparameter optimization scheme for autonomous racing systems," *IEEE Control Systems Letters*, vol. 7, pp. 1652–1657, 2023.
- [23] J. Chrosniak, J. Ning, and M. Behl, "Deep dynamics: Vehicle dynamics modeling with a physics-constrained neural network for autonomous racing," *IEEE Robotics and Automation Letters*, 2024.
- [24] H. B. Pacejka and E. Bakker, "The magic formula tyre model," *Vehicle system dynamics*, vol. 21, no. S1, pp. 1–18, 1992.
- [25] Y. Yu, Y. Liu, F. Fu, S. He, D. Zhu, L. Wang, X. Zhang, and J. Li, "Fast extrinsic calibration for multiple inertial measurement units in visual-inertial system," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 01–07.
- [26] A. Heilmeier, A. Wischnewski, L. Hermansdorfer, J. Betz, M. Lienkamp, and B. Lohmann, "Minimum curvature trajectory planning and control for an autonomous race car," *Vehicle System Dynamics*, 2020.
- [27] R. Verschuere, G. Frison, D. Kouzoupis, J. Frey, N. v. Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados—a modular open-source framework for fast embedded optimal control," *Math. Program. Comput.*, vol. 14, no. 1, pp. 147–183, 2022.
- [28] D. Kloeser, T. Schoels, T. Sartor, A. Zanelli, G. Prison, and M. Diehl, "NMPC for racing using a singularity-free path-parametric model with obstacle avoidance," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 14324–14329, 2020.
- [29] G. Frison and M. Diehl, "HPIPM: a high-performance quadratic programming framework for model predictive control," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 6563–6569, 2020.