Active Learning of Tire Parameters for High-Speed Autonomous Racing

Thomas A. Berrueta
Computing + Mathematical Sciences
California Institute of Technology
Pasadena, California 91125
Email: berrueta@caltech.edu

Nikhil Ranganathan
Graduate Aerospace Laboratories
California Institute of Technology
Pasadena, California 91125
Email: nrangana@caltech.edu

Soon-Jo Chung
Computing + Mathematical Sciences
California Institute of Technology
Pasadena, California 91125
Email: sjchung@caltech.edu

***Abstract*—High-performance autonomous racing presents a challenging proving ground for robot learning algorithms. In safety-critical learning scenarios such as these, typical exploration-exploitation tradeoffs carry real-world risk to both costly hardware and environmental harm. As a result, sample-efficient learning algorithms are essential to minimizing cumulative risk exposure during deployment. This paper presents preliminary results for a sample-efficient online algorithm for safe active learning. By leveraging contraction theory within an information-maximizing Model Predictive Contouring Control (MPCC) framework, our method generates online optimal perturbations to minimum time/curvature racelines, thereby enhancing sample efficiency while adhering to safety constraints. We plan to demonstrate this approach for learning predictive models of tire loads on a real-world hardware testbed—the Dallara AV-24 as part of the Indy Autonomous Challenge—showcasing its potential to improve both learning efficacy and operational safety in high-speed autonomous systems.**

I. Introduction

High-performance autonomous racing circuits such as the Indy Autonomous Challenge (IAC) push robot learning algorithms to their operational limits [1]. The IAC is a competition series featuring various race formats for autonomous vehicles, including time-trials and head-to-head racing at speeds exceeding 80 m/s (180 mph). Teams compete using the Dallara AV-24 (see Fig. 1), a modification of the Indy NXT racecar equipped with a suite of sensors including inertial measurement units (IMUs), light detection and ranging units (LiDARs), radars, and cameras. Due to their speed, the consequences of failure during deployment can be catastrophic, involving costly hardware damage and significant environmental risk [7]. In such extreme, safety-critical conditions, achieving competitive performance hinges on precise vehicle control, requiring accurate underlying system models. Among these, tire models are particularly crucial, as they govern the forces a vehicle can generate. However, accurately identifying empirical tire models, such as the widely used Pacejka model [4], sensitively depends on data from the limits of handling—a regime that is typically unsafe to explore during normal vehicle operation.

Given the difficulty of spontaneously acquiring the data needed for robust tire model identification, active learning strategies become indispensable [18]. Active learning presents a framework for intelligently selecting robot actions or designing vehicle inputs that maximally excite unmodeled system dynamics or reduce uncertainty in model parameters, thereby improving identification with fewer data points [3]. Common active learning approaches in robotics involve optimizing information-theoretic measures, such as those derived from the Fisher information matrix or from ergodicity, to guide exploration and data acquisition [19, 6]. However, control inputs derived purely from information maximization objectives can lead to aggressive or erratic behaviors. In the context of high-speed autonomous racing, information-maximizing actions pose a substantial safety risk precisely due to their optimality.

The tension between informative exploration and operational safety drives the need for safe active learning methodologies. In order to maintain vehicle safety without sacrificing performance, we propose to frame our approach within the Model Predictive Contouring Control (MPCC) framework [11]. MPCC problems represent a subset of model predictive control tasks in which the objective is to minimize the Euclidean distance between the agent's path and a reference path while simultaneously maximizing the speed at which the agent traverses the path. This framework has proven effective for time-optimal control in autonomous drone racing [16], and has recently been extended to provide safety guarantees that ensure persistent solution feasibility and collision avoidance [10]. By adapting MPCC to incorporate information-maximization, we aim to develop an optimal control methodology capable of actively learning critical system parameters that characterize tire loads in a manner that is both sample-efficient and provably safe.

To this end, we present preliminary work towards a safe active learning algorithm designed for the rigorous demands of high-speed autonomous racing. Our approach introduces a sample-efficient methodology for ensuring safety during robot learning. By applying contraction-theoretic analysis within an information-maximizing MPCC, our algorithm is capable of generating optimal, safety-aware perturbations to nominal racelines. We provide validation of our information-maximizing inference framework on data collected during testing at the Laguna Seca Raceway [2]. We plan to demonstrate our approach by learning tire load parameters in hardware experiments on the Dallara AV-24 platform, highlighting its potential for robot learning in safety-critical systems.

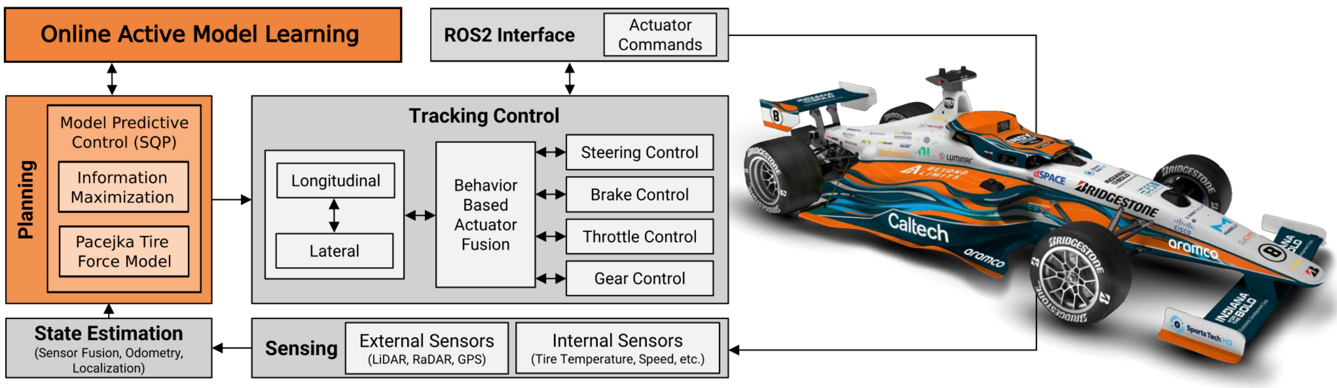

Fig. 1. **Caltech Racer active learning autonomy stack overview.** Our system architecture is comprised of modules for sensing, state estimation, tracking control, and model predictive planning with online model learning tailored for high-speed autonomous racing. In addition to standard state information, the planner estimates tire loads in real-time, which allows it to actively improve model learning.

II. Methods

A. Vehicle Dynamics & Tire Models

We model the autonomous racecar using a rear-wheel driven bicycle model (see Fig. 2), which lumps the front and rear axle dynamics [15]. We augment this model by considering the virtual dynamics of vehicle progress along its reference path, modeled through simple single integrator dynamics. The fictitious state, s, representing progress along the reference path and its rate, v_s, are central to the MPCC formulation. By treating s as a state and v_s as a virtual control input, we are capable of decoupling the spatial problem of following a geometric path from the temporal problem of how fast to traverse it. This contrasts with traditional trajectory tracking MPC, which typically follows a fixed reference.

Including the reference path progress variable, we define the augmented state of the vehicle as, $x = [X, Y, \varphi, v_x, v_y, \omega, s]^T \in \mathbb{R}^7$, where (X, Y) are the global Cartesian coordinates of the vehicle's center of gravity (CoG) in map frame, φ is the yaw angle, v_x and v_y are the longitudinal and lateral velocities in the vehicle body frame, respectively, and ω is the yaw rate. The control inputs to the vehicle are $u = [d, \delta, v_s]^T \in \mathbb{R}^3$ where d represents the commanded drive/brake force expressed as a slip ratio and δ is the steering angle of the front wheels. Equipped with our state representation, the continuous-time vehicle dynamics $\dot{x} = f(x, u)$ are given by:

$$
\begin{aligned}
\dot{X} &= v_x \cos(\varphi) - v_y \sin(\varphi) \\
\dot{Y} &= v_x \sin(\varphi) + v_y \cos(\varphi) \\
\dot{\varphi} &= \omega \\
\dot{v}_x &= \frac{1}{m}(F_{r,x} - F_{f,y} \sin\delta + m v_y \omega) \\
\dot{v}_y &= \frac{1}{m}(F_{r,y} + F_{f,y} \cos\delta - m v_x \omega) \\
\dot{\omega} &= \frac{1}{J}(F_{f,y} L_f \cos\delta - F_{r,y} L_r) \\
\dot{s} &= v_s
\end{aligned} \tag{1}
$$

where m is the vehicle mass, J is the moment of inertia about the vertical axis, and L_f and L_r are the distances from the CoG to the front and rear axles, respectively.

The terms $F_{f,y}$ and $F_{r,y}$ are the lateral tire forces at the front and rear axles, and $F_{r,x}$ is the longitudinal tire force at the rear axle. Lateral tire forces are modeled in accordance with a simplified Pacejka tire model [4, 5]—an empirical tire model that despite not being derived from first principles attains highly-accurate predictions of tire forces—while the longitudinal tire forces are modeled linearly for simplicity. Then, tire forces are given by

$$
\begin{aligned}
F_{\tau,x}(x) &= C_{\tau,x}\ d \\
F_{\tau,y}(x) &= D_{\tau,y} \sin(C_{\tau,y} \arctan(B_{\tau,y}\ \alpha_\tau(x)),
\end{aligned} \tag{2}
$$

where we defined $\theta_\tau = [C_{\tau,x}, B_{\tau,y}, C_{\tau,y}, D_{\tau,y}]^T$ for $\tau \in \{f, r\}$ as the vectors of parameters to be identified in order to characterize tire loads in terms of wheel slip ratios and angles, d and $\alpha_\tau(x)$, respectively. For the front and rear axles, the slip angles are

$$
\begin{aligned}
\alpha_f(x) &= \delta - \arctan\left(\frac{v_y + L_f \omega}{v_x}\right) \\
\alpha_r(x) &= -\arctan\left(\frac{v_y - L_r \omega}{v_x}\right).
\end{aligned} \tag{3}
$$

Since tire forces govern the impact that control inputs have on the dynamics in Eq. 1, correctly identifying θ_τ is essential to the design of effective motion plans.

B. Model Predictive Contouring Control

Equipped with a dynamical model, we review the standard MPCC formulation [12]. Given a reference path, the core objective of MPCC is to optimize a cost function that balances competing goals: minimizing contouring error, $\hat{e}^c_k$, minimizing lag error, $\hat{e}^l_k$, maximizing the rate of progress along the path, v_s, with regularization of control effort. This formulation allows the vehicle to autonomously adjust its speed along the path, slowing down for sharp corners and accelerating on straight sections, to achieve objectives such as minimum lap

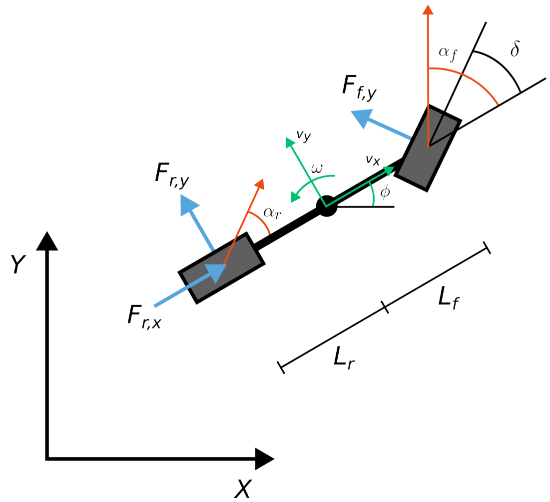

Fig. 2. **Bicycle model.** Diagram of rear-driven bicycle model that lumps front and rear axle dynamics, with forces acting on the car shown in blue.

time while respecting vehicle dynamics and operational constraints. The optimization problem is solved using sequential quadratic programming (SQP) [14] at each control step over a finite prediction horizon, yielding an optimal sequence of control inputs and desired setpoints that are then passed onto the tracking controller.

To define the tracking objectives, we consider a geometric reference path, $\mathcal{P}_d$, that is the result of an offline minimum time/curvature optimization [8]. Given vehicle states x_k at discretized time indices k, we use the path progress s_k to sample corresponding points on the reference path. Then, we can use the corresponding desired coordinates and orientations $[X^d_{s_k}, Y^d_{s_k}, \varphi^d_{s_k}]^T \in \mathcal{P}_d$ to formulate the MPCC optimization as follows:

$$\begin{aligned} \min_{u_{1:N}} \quad & \mathcal{J}_{\text{MPCC}}(x_{1:N}, u_{1:N}) \\ \text{s.t.} \quad & x_0 = x(0), \quad x_{k+1} = \hat{f}(x_k, u_k) \\ & x_k \in \mathcal{X}_{\text{Track}}, \quad \underline{x} \le x_k \le \overline{x} \\ & \Delta u_k = u_k - u_{k-1} \\ & \underline{u} \le u_k \le \overline{u}, \quad \underline{\Delta u} \le \Delta u_k \le \overline{\Delta u} \end{aligned} \tag{4}$$

with

$$\mathcal{J}_{\text{MPCC}} = \sum_{k=1}^{N} \begin{bmatrix} \hat{e}^c_k \\ \hat{e}^l_k \end{bmatrix}^T \begin{bmatrix} q_c & 0 \\ 0 & q_l \end{bmatrix} \begin{bmatrix} \hat{e}^c_k \\ \hat{e}^l_k \end{bmatrix} - q_v v_{s,k} + \Delta u_k^T R \Delta u_k,$$

where $q_c, q_l, q_v \in \mathbb{R}^+$ assign weights to the contouring, lag and progress terms of the objective function, respectively, and $R \succ 0$ is a positive definite matrix. The contouring and lag error terms are defined in the following way,

$$\begin{aligned} \hat{e}^c_k &= \sin(\varphi^d_{s_k})(X_k - X^d_{s_k}) - \cos(\varphi^d_{s,k})(Y_k - Y^d_{s_k}) \\ \hat{e}^l_k &= -\cos(\varphi^d_{s_k})(X_k - X^d_{s_k}) - \sin(\varphi^d_{s_k})(Y_k - Y^d_{s_k}), \end{aligned} \tag{5}$$

which specify notions of lateral and longitudinal error in the frame of the desired trajectory. The discrete-time system, $\hat{f}$, results from time integration of the dynamics in Eq. 1. Lastly, we note that the constraints specified in Eq. 4 capture track bounds (e.g., geometric constraints), state limits (e.g. maximum velocities), actuator limits (e.g., steering range), and actuator bandwidths (e.g., brake pressure rate), respectively.

C. Fisher Information & Belief Updates

Prior to retooling the MPCC formulation for information-maximization, we must introduce a measure for quantifying information amenable to real-time optimization, as well as our procedure for updating our parameters online. To this end, in this work we make use of the Fisher information [9], which has a long history in control for the design of probing signals for optimal parameter estimation [13]. The Fisher information provides a general means of quantifying the amount of information that a random variable Z contains about the estimate of an unknown vector of parameters θ. More formally, the Fisher information is defined as

$$\mathcal{F}(\theta) = E_{z\sim Z}\left[\left(\frac{\partial}{\partial\theta}\log p(z|\theta)\right)\left(\frac{\partial}{\partial\theta}\log p(z|\theta)\right)^T \middle| \theta \right] \tag{6}$$

where $p(z|\theta)$ represents the conditional probability density of observing sample z given the choice of parameters θ. In other words, $\mathcal{F}(\theta)$ captures the local sensitivity of our observations to changes in our parameters.

Within our problem domain, we consider the vector $z = [\hat{F}_{f,x}, \hat{F}_{r,x}, \hat{F}_{f,y}, \hat{F}_{r,y}]^T$ to be comprised of noisy observations of the vehicle tire forces. Then, let $p(z_k|\theta) = \mathcal{N}(h(x_k, u_k, \theta), \Sigma_z)$, where the measurement model $h(x_k, u_k, \theta)$ represents the functional form of the tire loads in Eq. 2, $\theta = [\theta_f, \theta_r]$ represents a stacked vector of parameters, and $\Sigma_z \succ 0$ is a constant. Under this set of modeling assumptions, for a given dataset $\mathcal{D} = \{x_1, u_1, z_1, \cdots, x_M, u_M, z_M\}$ we may write the following a sample-based simplified form of the Fisher information:

$$\hat{\mathcal{F}}(\theta, x_{1:M}, u_{1:M}) = \sum_{k=1}^{M} \frac{\partial h(x_k, u_k, \theta)}{\partial\theta}^T \Sigma_z^{-1} \frac{\partial h(x_k, u_k, \theta)}{\partial\theta}$$

In turn, we may now define the following scalar-valued information metric

$$\mathcal{I}(\theta, x_{1:M}, u_{1:M}) = Tr(\hat{\mathcal{F}}(\theta, x_{1:M}, u_{1:M})) \tag{7}$$

which we may readily incorporate into our MPCC formulation to perform information maximization.

Lastly, for active learning we require an online procedure for updating model parameters. To this end, we maintain a probabilistic belief over the unknown tire parameters θ, represented by $p(\theta_k) = \mathcal{N}(\hat{\theta}_k, P_k)$, where $\hat{\theta}_k$ is the mean estimate and P_k is the covariance matrix at time step k. Upon acquiring a new measurement z_k, we update our belief using Bayes' rule. Given the previous estimate $p(\theta_{k-1}) = \mathcal{N}(\hat{\theta}_{k-1}, P_{k-1})$, and the likelihood $p(z_k|\theta_k)$, the posterior distribution $p(\theta_k|z_k)$ can be approximated by a Gaussian. In conjunction with a linearized measurement model, this procedure results in the

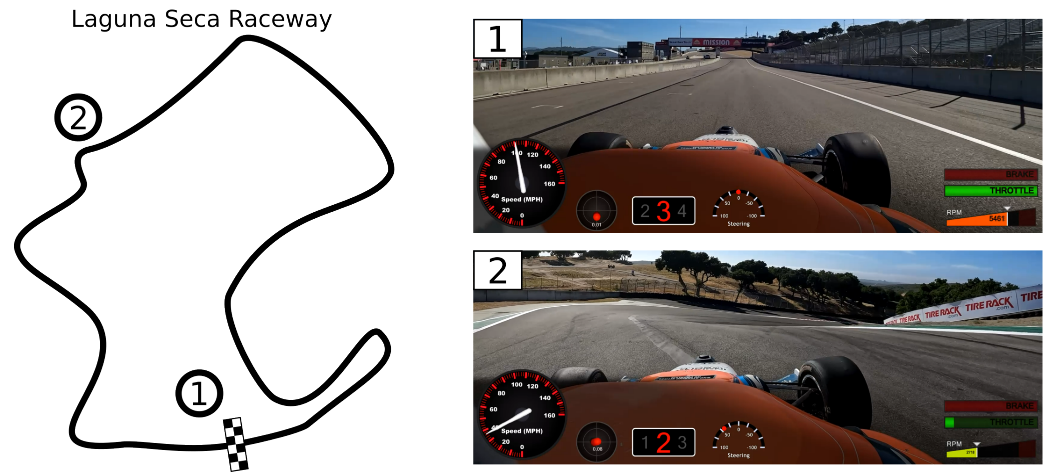

Fig. 3. **Caltech Racer deployment at Laguna Seca Raceway.** As part of the Indy Autonomous Challenge, we deployed our Dallara AV-24 autonomous racecar at the WeatherTech Laguna Seca Raceway. The track features high-speed straights where we achieved speeds of over 100 mph (see frame 1), as well as challenging corners with elevation changes (see frame 2).

extended Kalman filter update equations [17]. The posterior mean $\hat{\theta}_k$ and covariance P_k are updated as follows:

$$\begin{aligned} K_k &= P_{k-1} H_k^T (H_k P_{k-1} H_k^T + \Sigma_z)^{-1} \\ \hat{\theta}_k &= \hat{\theta}_{k-1} + K_k (z_k - h(x_k, u_k, \hat{\theta}_{k-1})) \\ P_k &= (I - K_k H_k) P_{k-1} \end{aligned} \tag{8}$$

where $H_k = \frac{\partial h(x_k, u_k, \theta)}{\partial \theta}\Big|_{\theta=\hat{\theta}_{k-1}}$ is the measurement model Jacobian evaluated at the prior mean, and K_k is the Kalman gain. This iterative update allows the system to continuously refine its estimate of the tire parameters as new data becomes available.

D. Information-Maximizing MPCC

With these preliminaries out of the way, the information-maximizing (Infomax) MPCC problem statement is

$$\begin{aligned} \min_{u_{1:N}} \quad & \mathcal{J}_{\text{MPCC}}(x_{1:N}, u_{1:N}) - q_I \mathcal{I}(\theta, x_{1:N}, u_{1:N}) \\ \text{s.t.} \quad & x_0 = x(0), \quad x_{k+1} = \hat{f}(x_k, u_k) \\ & x_k \in \mathcal{X}_{\text{Track}}, \quad \underline{x} \le x_k \le \overline{x} \\ & \Delta u_k = u_k - u_{k-1} \\ & \underline{u} \le u_k \le \overline{u}, \quad \underline{\Delta u} \le \Delta u_k \le \overline{\Delta u} \\ & \Delta \mathcal{I}_k = \mathcal{I}(\theta, x_{1:k}, u_{1:k}) - \mathcal{I}(\theta, x_{1:k-1}, u_{1:k-1}) \\ & |\Delta \mathcal{I}_k| \le \overline{\Delta \mathcal{I}}. \end{aligned} \tag{9}$$

where $q_I \in \mathbb{R}^+$ assigns a weight to the cost incurred by the application of the information measure, and $\overline{\Delta \mathcal{I}}$ bounds the amount of information gain per horizon. We note that since the Fisher information matrix is positive semi-definite, its trace is always non-negative and so we introduce the negation of the information measure into the objective function.

In addition to the bound on information gain, $\overline{\Delta \mathcal{I}}$, in order to guarantee vehicle safety in our problem domain, we introduce an additional set of constraints to ensure that the vehicle does not lose traction over the course of an information-maximizing maneuver. In particular, we require the following constraint is satisfied in addition to the constraints listed in Eq. 9:

$$|F_{\tau,x}(x_k, u_k, \theta_\tau) + F_{\tau,y}(x_k, u_k, \theta_\tau)| \le \overline{F}_\tau \tag{10}$$

for each k and for each choice of θ. The bounds $\overline{F}_f$ and $\overline{F}_r$ represent the best known empirical estimates of the loads that the vehicle's tires are capable of bearing. As identification of θ improves through our information-maximizing estimation procedure, our approach enables safe, iterative envelope

---

Algorithm 1 Active Learning of Tire Parameters

Require: Initial parameter $\hat{\theta}_0$, initial covariance P_0, reference path $\mathcal{P}_d$, measurement covariance Σ_z, MPCC weights q_c, q_l, q_v, R, information weight q_I, information rate bound $\overline{\Delta \mathcal{I}}$, tire load limits $\overline{F}_f, \overline{F}_r$, as well as track $\mathcal{X}_{\text{Track}}$, state $\overline{x}, \underline{x}$, actuation $\overline{u}, \underline{u}$, bandwidth bounds $\overline{\Delta u}, \underline{\Delta u}$, and planning horizon N.

Init: $t \leftarrow 1$, $u_{\text{init}} \leftarrow$ InitialControlSequence();
while NotDone() **do**
  $x_t \leftarrow$ GetState();
  $z_t \leftarrow$ GetMeasurement();
  BeliefUpdate():
    $H_t \leftarrow \frac{\partial h(x_t, u_t, \theta)}{\partial \theta}\Big|_{\theta=\hat{\theta}_{t-1}}$
    $K_t \leftarrow P_{t-1} H_t^T (H_t P_{t-1} H_t^T + \Sigma_z)^{-1}$
    $\hat{\theta}_t \leftarrow \hat{\theta}_{t-1} + K_t (z_t - h(x_t, u_t, \hat{\theta}_t))$
    $P_t \leftarrow (I - K_t H_t) P_{t-1}$
  SolveInfomaxMPCC():
    $\{u_k^*\}_{k=0}^N \leftarrow$ SolveSQP$(x_t, u_{\text{init}}, \hat{\theta}_t, \mathcal{P}_d)$
    $u_{\text{init}} \leftarrow \{u_k^*\}_{k=1}^N$
    $u_{t-1} \leftarrow u_1^*$
  ApplyAction(u_{t-1})
  $t \leftarrow t + 1$
end while

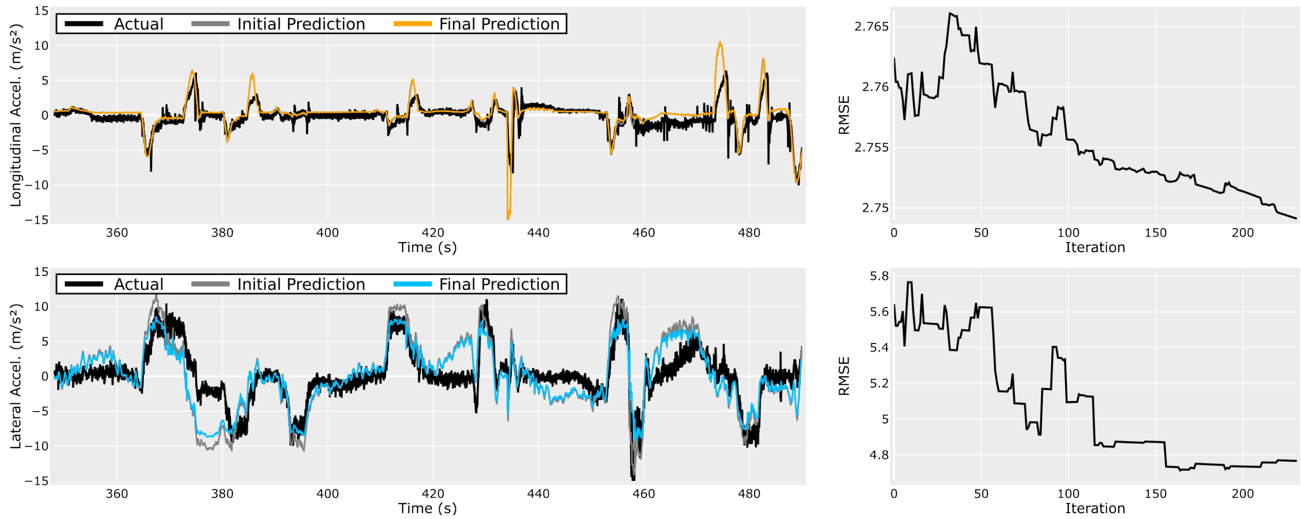

Fig. 4. **Tire parameter identification and predicted vehicle body accelerations.** We show 2.5 min of body acceleration data (left column) from a 30 min run collected during field tests at the Laguna Seca Raceway we fit the parameters of our tire load models. Parameter updates are done in accordance with our specified procedure in batches of 100 samples, resulting the shown RMSE curves (right column) for the predicted longitudinal and lateral body accelerations.

expansion through successive modifications of $\overline{F}_f$ and $\overline{F}_r$ across experimental deployments. Lastly, we note that in addition to these safety measures, the bound on $\Delta\mathcal{I}_k$ enables a contraction-theoretic analysis of the Infomax MPCC optimization, the proof of which will be presented in a subsequent archival submission of this work.

The combined active learning methodology is summarized in Algorithm 1. Given properly initialized model and optimization parameters, Infomax MPCC begins by refining its model parameter estimates based on data using Eq. 8 before instantiating an optimal control problem instance. Then, equipped with the best current model estimates we use SQP to solve Eq. 9. Lastly, we take the first action in the solution horizon and save the rest of the solution as a warm-start to the next optimization iteration. This process repeats successively as the model parameter estimate converges, or until the test concludes.

III. Results

Prior to deploying our active learning framework in a hardware-in-the-loop fashion, we validate our identification framework on data collected passively during field experiments conducted at the WeatherTech Laguna Seca Raceway (see Fig. 3). We use data collected over the course of a 30 minute test slot consisting of 8 laps with progressively faster lap times in which speeds of up to 115 mph were reached (see frame 1 in Fig. 3). The collected data set consists of complete state, x_k, and control, u_k, information in alignment with our dynamics model in Eq. 1. Since directly observing longitudinal and lateral tire forces is challenging, we instead make use of $z_k = [a_{x,k}, a_{y,k}]$ measurements collected from the vehicle's IMU. This is feasible because we are able to plug in the tire force equations from Eq. 2 directly into the vehicle dynamics model in Eq. 1 to predict body accelerations. As a result of the linear relationship between tire forces and body accelerations, the rest of our parameter identification pipeline remains unchanged. As a final simplification we assume that the front and rear axle parameters are the same. That is, our goal is to identify $\theta = [C_x, B_y, C_y, D_y]^T$.

Given an initial guess of parameters listed in Table I, we regressed over the 30 min of data in batches of 100 samples at a time. For each batch, we linearized the measurement model once and averaged the loss and descent directions to perform one parameter update per batch. Our results are illustrated in Fig. 4, where the right column depicts the root mean square error (RMSE) in the prediction of longitudinal and lateral body accelerations over the identification procedure. The final parameters are listed in Table I and produce more accurate predictions of body accelerations. This is shown in the left column of Fig. 4, where initial predicted, final predicted, and actual accelerations are depicted over the course of a 2.5 min window of the test dataset. While the overall quality of predictions improves as a result of our identification procedure, there are still mismatches between predicted body accelerations and actual as a result of model inadequacy. For one, the linear tire force model tends to overshoot its predicted longitudinal accelerations (e.g., between 470s and 480s), and the lateral tire force model seems to not fully capture all relevant effects, leading to some mismatch (e.g., between 370s and 380s). These results suggest that there is an opportunity

TABLE I
Comparison Between Identified Model Parameters

Parameter	Initial	Final	Change
C_x	7483.17	7480.77	−0.03%
B_y	19.79	17.07	−13.76%
C_y	1.50	2.67	78.23%
D_y	1.50	0.91	−39.12%

for the use of machine learning based techniques in order to augment our tire models and optimize their learning within the context of our safe active learning framework, which we plan to explore in future work.

IV. Conclusion

In this work, we presented preliminary results in support of a novel framework for safe, sample-efficient active learning of tire model parameters tailored to high-performance autonomous racing. By integrating an information-maximization objective based on the trace of the Fisher information matrix directly into a safety-constrained MPCC formulation, our approach generates informative control actions while ensuring operational safety. The online estimation of tire parameters is achieved through an iterative estimation procedure, which refines the parameter beliefs using real-time measurements. Preliminary validation on real-world data from Laguna Seca demonstrates the potential of key components central to our approach, and future work will focus on a full-stack hardware deployment on the Dallara AV-24 platform to showcase the complete closed-loop active learning cycle and realize its potential to enhance model accuracy and vehicle safety at the limits of handling. This research paves the way for more robust and adaptive autonomous systems capable of safely learning in complex, safety-critical environments.

Acknowledgments

We thank Matthew Anderson, Joshua Cho, and the Indy Autonomous Challenge organization for their support in maintaining the Dallara AV-24 hardware platform. Additionally, we acknowledge funding and technical support from Aramco and BeyondLimits. N. R. was partially supported by the Office of Naval Research as well. The conclusions reached by this work only reflect the opinions of the authors.

References

[1] Indy Autonomous Challenge - Official Website. https://www.indyautonomouschallenge.com/, Accessed: May 2025.

[2] WeatherTech Raceway Laguna Seca. https://weathertechraceway.com/, Accessed: May 2025.

[3] Ian Abraham and Todd D. Murphey. Active Learning of Dynamics for Data-Driven Control Using Koopman Operators. *IEEE Transactions on Robotics*, 35(5):1071–1083, 2019.

[4] E. Bakker, L. Nyborg, and H. B. Pacejka. Tyre Modelling for Use in Vehicle Dynamics Studies. *SAE Technical Paper*, (870421), 1987.

[5] E. Bakker, H. B. Pacejka, and L. Lidner. A new tire model with an application in vehicle dynamics studies. *SAE Technical Papers*, 98:101–113, 1989.

[6] Thomas A. Berrueta, Ian Abraham, and Todd Murphey. *Experimental Applications of the Koopman Operator in Active Learning for Control*, pages 421–450. Springer International Publishing, Cham, 2020. ISBN 978-3-030-35713-9.

[7] Johannes Betz, Tobias Betz, Felix Fent, Maximilian Geisslinger, Alexander Heilmeier, Leonhard Hermansdorfer, Thomas Herrmann, Sebastian Huch, Phillip Karle, Markus Lienkamp, Boris Lohmann, Felix Nobis, Levent Ögretmen, Matthias Rowold, Florian Sauerbeck, Tim Stahl, Rainer Trauth, Frederik Werner, and Alexander Wischnewski. TUM autonomous motorsport: An autonomous racing software for the Indy Autonomous Challenge. *Journal of Field Robotics*, 40(4):783–809, 2023.

[8] Fabian Christ, Alexander Wischnewski, Alexander Heilmeier, and Boris Lohmann and. Time-optimal trajectory planning for a race car considering variable tyre-road friction coefficients. *Vehicle System Dynamics*, 59 (4):588–612, 2021.

[9] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. EBSCO Computers & Applied Sciences. Wiley, 1991. ISBN 9780471062592.

[10] Maria Krinner, Angel Romero, Leonard Bauersfeld, Melanie Zeilinger, Andrea Carron, and Davide Scaramuzza. MPCC++: Model Predictive Contouring Control for Time-Optimal Flight with Safety Constraints. In *Robotics: Science and Systems*, 2024.

[11] Denise Lam, Chris Manzie, and Malcolm Good. Model predictive contouring control. In *49th IEEE Conference on Decision and Control (CDC)*, pages 6137–6142, 2010.

[12] Alexander Liniger, Alexander Domahidi, and Manfred Morari. Optimization-based autonomous racing of 1:43 scale RC cars. *Optimal Control Applications and Methods*, 36(5):628–647, 2015.

[13] N. E. Nahi and G. A. Napjus. Design of optimal probing signals for vector parameter estimation. In *1st IEEE Conference on Decision and Control (CDC)*, pages 162–168, 1971.

[14] J. Nocedal and S. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer New York, 2006. ISBN 9780387227429.

[15] Rajesh Rajamani. *Vehicle Dynamics and Control*. Mechanical Engineering Series. Springer US, 2011. ISBN 9781461414339.

[16] Angel Romero, Sihao Sun, Philipp Foehn, and Davide Scaramuzza. Model Predictive Contouring Control for Time-Optimal Quadrotor Flight. *IEEE Transactions on Robotics*, 38(6):3340–3356, 2022.

[17] D. Simon. *Optimal State Estimation: Kalman, H-Infinity, and Nonlinear Approaches*. Wiley, 2006. ISBN 9780470045336.

[18] Annalisa T. Taylor, Thomas A. Berrueta, and Todd D. Murphey. Active learning in robotics: A review of control principles. *Mechatronics*, 77:102576, 2021.

[19] Andrew D. Wilson, Jarvis A. Schultz, and Todd D. Murphey. Trajectory Synthesis for Fisher Information Maximization. *IEEE Transactions on Robotics*, 30(6): 1358–1370, 2014.