

On Disturbance-Aware Minimum-Time Trajectory Planning: Evidence from Tests on a Dynamic Driving Simulator

Matteo Masoni¹, Vincenzo Palermo¹, Marco Gabiccini¹, Martino Gulisano¹
Giorgio Previati², Massimiliano Gobbi², Francesco Comolli²
Gianpiero Mastinu², and Massimo Guiggiani¹

Abstract

This work aims to investigate how disturbance-aware, robustness-embedding reference trajectories translate into actual driving performance when executed by professional drivers in a dynamic driving simulator. The study compares three planned reference trajectories against a free-driving baseline (NO-REF) to assess the trade-offs between lap time performance and steering effort: NOM, the nominal time-optimal trajectory; TLC, a track-limit-robust, time-optimal trajectory obtained by tightening margins to the track edges; and FLC, a friction-limit-robust, time-optimal trajectory obtained by tightening against axle/tire saturation. All reference trajectories share the same minimum lap-time objective with a small steering-smoothness regularizer, and are evaluated with two professional drivers driving a high-performance car on a virtual track.

The reference trajectories stem from a *disturbance-aware* minimum-lap-time framework recently proposed by some of the authors, where worst-case disturbance growth is propagated over a finite horizon and used to *tighten* tire-friction and track-limit constraints, preserving performance while delivering probabilistic safety margins.

Lap time (LT) and steering energy (SE) are evaluated as indicators of driving performance and steering effort, respectively, while RMS values of lateral deviation, speed error, and drift angle are used to characterize driving style. The results reveal a Pareto-like trade-off between LT and SE: NOM achieves the shortest LT but with the highest SE, TLC minimizes SE at the expense of longer LT, while FLC lies near the efficient frontier, markedly reducing SE relative to NOM with only a minor LT increase. Removing reference trajectories (NOREF) leads to both higher SE and longer LT, confirming that trajectory guidance improves pace and control efficiency. Overall, the findings highlight reference-based and disturbance-aware planning, particularly the FLC variant, as effective tools for training and for achieving fast yet stable trajectories.

1 Introduction

Vehicle stability is a central topic in automotive engineering, as it concerns a road vehicle's ability to maintain a desired trajectory and attitude when subjected to external disturbances, driver inputs, and variable road conditions. From a dynamics perspective, stability encompasses the control of lateral, longitudinal, and vertical motions, and is strongly influenced by vehicle parameters such as mass distribution, suspension geometry, tire characteristics, and steering system design.

Over the last decades, the growing demand for safety and performance has motivated extensive research on the mechanisms leading to loss of stability, including skidding, spin-out, and rollover, as well as on the development of active control systems to prevent them. Advanced systems such as Anti-lock Braking Systems (ABS), Traction Control Systems (TCS), and Electronic

¹Università di Pisa, Largo L. Lazzarino 1, 56122 Pisa, Italy.

²Politecnico di Milano, Via G. La Masa 1, 20156 Milano, Italy.

Stability Control (ESC) exploit real-time sensing and actuation to compensate for the intrinsic limitations of the vehicle-road-driver system, significantly improving handling robustness in critical maneuvers. In this context, both experimental testing and high-fidelity simulation play a key role in understanding stability limits and in validating new control strategies under reproducible and safety-critical scenarios.

Several strategies have been proposed in the scientific literature to assess vehicle stability. A common practice in control-oriented studies is to impose empirical bounds on selected state variables. Limits on the vehicle sideslip angle, for instance, are frequently employed in yaw-stability and ESC design [1, 2]. A more model-based research line analyzes simplified lateral-yaw dynamics—typically with single-track formulations—to characterize equilibrium configurations and their stability through bifurcation methods. Comprehensive continuation studies have revealed phenomena such as saddle-node and Hopf transitions in response to changing inputs and tire characteristics [3].

Beyond the identification of stable steady states, an important question in nonlinear dynamics concerns the estimation of the Regions of Attraction (ROAs), i.e., the sets of initial conditions that converge asymptotically to a desired equilibrium. Recent work has estimated ROAs by adopting enhanced tire models [4] or multi-point linearization techniques [5]. Another promising approach is based on Sum-of-Squares (SOS) programming, which enables systematic ROA computation through polynomial Lyapunov functions obtained via semidefinite programming. Several applications to vehicle lateral dynamics have already used SOS techniques to estimate stability regions [6, 7].

Alternative quantitative indicators include Lyapunov exponents to measure convergence rates to capture how disturbances decay during aggressive maneuvers [8]. However, asymptotic stability indicators offer limited operational guidance for planning: a point along the planned trajectory may lie inside the ROA of a stable equilibrium even if that equilibrium is reached far ahead in time or outside the admissible domain (e.g., off track), which makes asymptotic guarantees only loosely connected to the short-horizon constraint satisfaction required in practice.

Building on theoretical studies that analyze the mechanisms of vehicle stability and loss of control, dynamic driving simulators offer a powerful experimental platform to test and further investigate these phenomena in a controlled, repeatable and risk-free environment. By allowing critical maneuvers and near-limit conditions to be safely reproduced, simulators enable systematic assessment of how drivers and vehicles behave close to the stability boundaries identified by theory, and how control strategies perform under such conditions. At the same time, their usefulness depends on how well simulator behavior reproduces real driving: a recent systematic review [9] shows that simulators can often achieve acceptable validity, but also highlights cases of mismatch and methodological limitations, underlining the need for rigorous validation and transparent reporting when using simulators for vehicle stability studies.

Driving simulators have been widely used to investigate vehicle stability in specific safety-related contexts. In [10], simulator tests with different vehicle speeds and crosswind profiles, including uniform and transient gusts, are carried out to study how crosswind-induced forces affect lateral deviation, yaw response and controllability on expressways. The resulting data are used to support safety-related decisions, such as defining expressway closure criteria and designing wind barriers. Driving simulators have also been employed to study sudden vehicle faults and their impact on controllability. In [11], drivers perform motorway-speed simulator tests in which single wheel hub motor failures are artificially triggered under straight-line and cornering conditions, creating strong yaw disturbances and longitudinal decelerations. The objective of these experiments is to quantify driver reaction times and resulting path deviations, and to assess the safety-criticality of such failures. Driving-simulator experiments on truck platooning are conducted in [12] to develop and validate a quantitative stability index for drivers' psychological stability under different time gaps and visualization conditions (with/without a see-through function). The tests relate vehicle-motion metrics and visual information to perceived comfort

and safety in closely spaced platoons, addressing driver acceptance and subjective stability.

Beyond safety-oriented applications, driving simulators have been instrumental in validating theoretical frameworks for vehicle stability and loss of control. A further line of research focuses on the straight-ahead stability of passenger cars. In [13], nonlinear bifurcation analysis is combined with experiments in a dynamic driving simulator to study how passenger cars lose straight-running stability as speed and disturbance levels increase, for both understeering and oversteering setups. In this case, the simulator experiments primarily serve to validate theoretical predictions about stability boundaries and loss-of-control mechanisms in straight-line driving. A broader theoretical perspective on stability with the driver in the loop is provided in [14], where a nonlinear vehicle-driver model is analyzed using bifurcation theory and Lyapunov methods, and then validated through experiments in a dynamic driving simulator. In this study, simulator runs are mainly used to confirm and illustrate the predicted domains of attraction and loss-of-control trajectories after disturbances, embedding them into a global stability framework for everyday road driving. In [15], dynamic models of the vehicle-driver system are combined with driving-simulator and track experiments to analyze how drivers lose control after strong disturbances such as severe lane changes, wind gusts or road irregularities. The test program concentrates on the onset and evolution of instability and loss-of-control trajectories near the stability boundary.

Understanding driver behavior and interaction with the vehicle is crucial for comprehensive stability analysis. In [16], the authors investigate, at a dynamic driving simulator, how the driver interacts with the steering wheel during cornering, in order to detect driving strategies. Such driving strategies allow to derive accurate holistic driver models for enhancing safety. Another important application of driving simulators is comparative vehicle dynamics testing. In [17], motion-based simulator experiments with heavy vehicles are used to evaluate whether drivers can reliably perceive handling differences induced by systematic variations of key parameters such as suspension and steering characteristics, combining subjective ratings with objective measurements. The main goal of these tests is to support early-stage vehicle dynamics development and parameter tuning in a cost-effective way. The handling performance of ultra-efficient lightweight vehicles (quadricycles) is studied in a dynamic driving simulator in [18]. The proper mechanical actions on the driver make the simulations trustworthy, enabling reliable assessment of handling characteristics for this class of vehicles.

Moving toward control-oriented applications, driving simulators enable validation of active systems designed to compensate for disturbances and enhance stability. In [19], driver-in-the-loop simulations at a dynamic driving simulator were performed for deriving a new electric power steering system able to give proper steering feedback in case of high torque vectoring at the front axle. Such a torque can be seen as a disturbance to the driver, and the simulator tests enable validation of the compensation strategy under realistic driving conditions. A more control-oriented use of driving simulators is presented in [20], where a real-time Hardware-in-the-Loop setup is used to compare advanced lateral stability controllers (LQR and sliding-mode) against a commercial ESP through standardized but stressed maneuvers such as step steer, sine steer and lane-change. These tests mainly serve to benchmark and refine stability control strategies near the limits of adhesion in demanding yet repeatable scenarios. In [21], future vehicles are studied aiming to exclude the human factor in limiting handling maneuvers. A dynamic driving simulator is used to compare robotic controllers with human drivers, showing that robotic controllers can outperform human drivers and posing interesting design challenges for future automated vehicles.

While most simulator-based studies focus on road-vehicle applications, recent work has begun exploring motorsport contexts. Exploiting a dynamic driving simulator, in [22] the authors show how to detect as quickly as possible whether the driver will lose control of a vehicle, after a disturbance has occurred. A degree of stability for the vehicle-and-driver combination is proposed which seems already applicable for motorsport applications.

Building on these contributions, while [22] touches on motorsport applications, most of the cited simulator-based studies are framed around road-vehicle applications, such as crosswind




Figure 1: Single-track model used as a basis for our stochastic vehicle dynamics model.

safety, driveline failures, global stability in everyday driving, platooning comfort, or early-stage vehicle dynamics development, and do not explicitly target sustained, near-limit motorsport operation with robust planning strategies. In contrast, the present work takes the disturbance-aware minimum-lap-time planner recently proposed in [23] and brings it to the DriSMi driving simulator [24], using professional drivers capable of operating at high performance on a race-like track to assess how robustness-embedding criteria translate into actual driving. Specifically, the open-loop, covariance-based framework in [23, Sec.3] is used to generate three reference trajectories – nominal time-optimal (NOM), track-limit-robust (TLC) and friction-limit-robust (FLC) – which are compared with a free-driving baseline (NO-REF). This enables quantification of how robustness-oriented constraint tightening influences lap time, steering effort and tracking/style indicators once human drivers attempt to execute these trajectories at high performance.

The rest of the paper is organized as follows. Section 2 revisits the open-loop disturbance-aware planner, detailing the vehicle model, uncertainty propagation, and the constraint-tightening mechanisms that yield the three reference trajectories (NOM, TLC, FLC). Section 3 documents how telemetry from the high-fidelity 14-DoF simulator was used to calibrate the single-track surrogate so that planning and execution remain consistent. Section 4 introduces the drivers, simulator setup, reference-generation protocol, and evaluation metrics, and outlines the randomized test campaign including the free-driving baseline NO-REF. In Section 5 lap-time, steer-energy, and tracking indicators are reported and discussed, highlighting the LT–SE trade-offs that emerge from the different references. Finally, Section 6 summarizes the main findings and discusses limitations and future research directions.

2 Open-loop disturbance-aware planning

This section recalls the essential notions of the disturbance-aware, robustness-embedding framework employed for trajectory planning. For a comprehensive treatment of the theoretical foundations, implementation details, and algorithmic aspects, the reader is referred to [23].

2.1 Vehicle and disturbance model

The planner relies on the nonlinear Single-Track model [25] with nonlinear tire characteristics and open differential schematically depicted in Figure 1. The state vector $\mathbf{x} = (u, v, r, x_G, y_G, \psi)^\top$ collects longitudinal and lateral velocities, yaw rate, and the pose of the vehicle with respect to

the center of mass. The control input $\mathbf{u} = (X_{2,a}, X_{2,b}, \delta)^\top$ includes the rear-axle acceleration force $X_{2,a}$, the braking force $X_{2,b}$ (distributed between axles via brake balance), and the steering angle δ . Let $f(\mathbf{x}, \mathbf{u})$ denote the resulting dynamics, where lateral forces $Y_j(\mathbf{x}, \mathbf{u})$ follow Pacejka's Magic Formula and vertical loads $Z_j(\mathbf{x}, \mathbf{u})$ capture load transfer. In open loop, stochastic perturbations related to the uncertainty on the vehicle state are modelled as additive Gaussian noise $\mathbf{w}(t)$,

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t)) + \mathbf{w}(t), \quad \mathbf{w}(t) \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}(t)), \quad (1)$$

where $\mathcal{N}(\mathbf{0}, \mathbf{Q}(t))$ denotes a Gaussian distribution with $\mathbf{0}$ mean and covariance matrix $\mathbf{Q}(t)$.

2.2 Mean and covariance propagation

Assuming a first-order disturbance propagation, the state becomes Gaussian with mean $\boldsymbol{\mu}(t)$ and covariance $\mathbf{P}(t)$. The mean trajectory satisfies

$$\dot{\boldsymbol{\mu}}(t) = f(\boldsymbol{\mu}(t), \mathbf{u}(t)), \quad \boldsymbol{\mu}(0) = \boldsymbol{\mu}_0, \quad (2)$$

while the covariance evolves through the Lyapunov differential equation

$$\dot{\mathbf{P}}(t) = \mathbf{A}(t) \mathbf{P}(t) + \mathbf{P}(t) \mathbf{A}^\top(t) + \mathbf{Q}(t), \quad \mathbf{P}(0) = \mathbf{P}_0 = \mathbf{P}_0^\top, \quad (3)$$

with Jacobian $\mathbf{A}(t) = \partial f / \partial \mathbf{x}|_{(\boldsymbol{\mu}(t), \mathbf{u}(t))}$. This approximation preserves tractability and provides closed-form expressions for the gradients used in the chance-constraint back-offs described in Sec. 2.5. A more detailed treatment is provided in [23].

2.3 Collocation discretization and H -step predictions

The track centerline is parameterized by the curvilinear variable $\alpha \in [0, 1]$ and sampled at $N + 1$ nodes $\alpha_0, \dots, \alpha_N$. At each interval $[\alpha_k, \alpha_{k+1}]$ we approximate $\boldsymbol{\mu}(\cdot)$ and $\mathbf{P}(\cdot)$ with degree- d polynomials $\pi_k(\tau)$ defined on the normalized interval $\tau \in [0, 1]$ and collocated at Gauss–Legendre points τ_ℓ . The associated collocation states and covariances are denoted $\boldsymbol{\xi}_{k,\ell}$ and $\boldsymbol{\Sigma}_{k,\ell}$. To emulate the absence of closed-loop correction (i.e., the delay in driver corrections), we propagate $H + 1$ covariance replicas: $\mathbf{P}_k^0$ is re-initialized at each node with $\bar{\mathbf{P}}_0 \succeq 0$, whereas $\mathbf{P}_k^j$ for $j \geq 1$ captures the covariance evolved j steps earlier. The residual maps $\boldsymbol{\Psi}_k^\mu(\cdot)$ and $\boldsymbol{\Psi}_k^P(\cdot)$ enforce, respectively, the *collocation constraints* for $(\boldsymbol{\mu}, \boldsymbol{\xi})$ and the *continuity* of every replica $\mathbf{P}_k^j$ across the grid. Algebraic quantities such as tire forces or track offsets are grouped in $\mathbf{z}_k$ and constrained through $\boldsymbol{\Omega}_k(\cdot)$.

2.4 Open-loop disturbance-aware OCP

The disturbance-aware reference generator solves the nonlinear program

$$\underset{\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{P}_k, \mathbf{z}_k}{\text{minimize}} \quad J_k(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k) \quad (4a)$$

$$\text{s.t.} \quad \mathbf{0} = \boldsymbol{\Psi}_k^\mu(\boldsymbol{\mu}_{k-1}, \boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{z}_k), \quad k = 1, \dots, N \quad (4b)$$

$$\boldsymbol{\mu}_0 = \bar{\boldsymbol{\mu}}_0 \quad (4c)$$

$$\mathbf{0} = \boldsymbol{\Psi}_k^P(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{P}_{k-1}^{j-1}, \mathbf{P}_k^j, \boldsymbol{\Sigma}_k^j, \mathbf{z}_k), \quad \begin{matrix} k=1, \dots, N; \\ j=1, \dots, H \end{matrix} \quad (4d)$$

$$\mathbf{P}_k^0 = \bar{\mathbf{P}}_0 \succeq 0 \quad k = 1, \dots, N \quad (4e)$$

$$\mathbf{0} = \boldsymbol{\Omega}_k(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{z}_k), \quad k = 0, \dots, N \quad (4f)$$

$$0 \geq h_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{z}_k) + \beta_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{P}_k^H, \mathbf{z}_k), \quad k = 1, \dots, N; \quad i \in \mathcal{I} \quad (4g)$$

where J_k is the standard minimum-lap-time cost (i.e. our planning objective). Equations (4b)–(4f) collect the collocation residuals, initial conditions, and algebraic constraints. Constraint (4g) evaluates the safety back-off using the most propagated covariance $\mathbf{P}_k^H$, hence embedding worst-case open-loop uncertainty growth over the horizon H . A more detailed description of the constraints formulation and the associated back-off terms is provided in Sec. 2.5.

2.5 Constraint back-offs and planner variants

For each inequality $h_i(\mathbf{x}, \mathbf{u}, \mathbf{z}) \leq 0$, the deterministic back-off β_i is computed from the linearized chance constraint $\Pr\{h_i(\mathbf{x}, \mathbf{u}, \mathbf{z}) \leq 0\} \geq p_i$ as

$$\beta_i = \Phi^{-1}(p_i) \sigma_i, \quad \sigma_i = [\nabla_{\mathbf{x}} h_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{z}_k)^\top \mathbf{P}_k^H \nabla_{\mathbf{x}} h_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{z}_k)]^{1/2}, \quad (5)$$

with Φ^{-1} the standard-normal quantile.

Track boundaries (TLC) are enforced by shifting the allowable lateral offset e as $e_{\min} + \beta^{\text{TLC}} \leq e \leq e_{\max} - \beta^{\text{TLC}}$, where $\nabla_{\mathbf{x}} h^{\text{TLC}}$ depends on the track normal.

Friction limits (FLC) rely on the axle saturation ratio

$$S_j(\mathbf{x}, \mathbf{u}) = \frac{\left(\frac{X_j(\mathbf{x}, \mathbf{u})}{\mu_{x,j}}\right)^2 + \left(\frac{Y_j(\mathbf{x}, \mathbf{u})}{\mu_{y,j}}\right)^2}{Z_j^2(\mathbf{x}, \mathbf{u})}, \quad j \in \{1, 2\}, \quad (6)$$

constrained by $S_j - 1 + \beta_j^{\text{FLC}} \leq 0$. Choosing $\beta_i \equiv 0$ yields the nominal planner (NOM). Emphasizing the track limit ($\beta^{\text{TLC}} > 0$) produces the TLC references, while focusing on the friction limit ($\beta^{\text{FLC}} > 0$) produces the FLC trajectories discussed in Sec. 5. A thorough theoretical treatment can be found in our previous work [23].

Remark. The horizon H controls conservatism and computational effort; in practice it is tuned to cover the time span before significant driver corrections intervene.

2.6 Parameters selection

The initial covariance $\mathbf{P}_0$ and the disturbance covariance $\mathbf{Q}(t)$ are key parameters that govern the magnitude of the back-off terms β_i in (5). These terms effectively “inflate” the constraints, providing a safety margin that ensures constraint satisfaction under uncertainty. The selection of the above parameters must be based on engineering criteria to balance robustness against excessive conservatism. Specifically, $\mathbf{P}_0$ and $\mathbf{Q}(t)$ are tuned to ensure that the resulting back-off terms do not become excessively large, which would unnecessarily reduce the potential of the axles by forcing the planner to operate well below the friction limits.

Since the effect of these parameters on the back-off terms is not clearly visible a priori, a steering maneuver with high lateral acceleration that saturates the axles above 90% was simulated. During this process, the covariance matrix was propagated for H steps from each discretization step, and the back-off terms were computed using the most propagated covariance matrix. Based on this analysis, the parameters were tuned via trial and error.

3 Vehicle model identification procedure

Planning under uncertainty is computationally demanding: at each time step the covariance matrix is reset and propagated over a horizon of H steps ($H = 4$ in our case) to capture the growth of uncertainty in the dynamics. This propagated covariance is then used to determine the amount of constraint tightening required at each H -th step. To keep runtime manageable we currently adopted a single-track vehicle model. The driving simulator, however, implements a high-fidelity 14-DoF dynamic model with detailed suspension kinematics; consequently, an identification procedure was required to align the two.

3.1 DriSMi layout and vehicle dynamics model

The dynamic simulator employed is located at Politecnico di Milano (DriSMi lab [24]). It is a new-generation, mid-size facility. A six-actuator electric hexalift provides full six-DoF motion with 20 Hz closed-loop bandwidth and is mounted on three aerostatic pads that enable frictionless

planar travel. Longitudinal and lateral motion on a machined 6×6 m deck is produced by four independent cable drives, which decouple x/y translations and yaw; yaw rotations exceed 60° , with 3 Hz bandwidth. The enlarged workspace supports peak accelerations up to 1.5 g (longitudinal/lateral) and 2.5 g (vertical), with maximum end-to-end latency of 20 ms. To offset moving-mass inertia and protect the cable system, three counter-masses are employed. Immersion is enhanced by five projectors on a 270° screen, an active seat and belt tensioners for sustained acceleration cues, electric steering torque feedback, an active hydraulic brake (including ABS effects), and a five-speaker audio system. For comfort/NVH studies, eight shakers (bandwidth up to 200 Hz) reproduce engine- and road-induced vibrations.

The high-fidelity vehicle representation is a 14-DoF multibody model implemented in the VI-CarRealTime environment [26] (hereafter, the 14-DoF model). The chassis has six degrees of freedom (three translations plus roll, pitch, and yaw). Each of the four wheels contributes two additional DoF (rolling rotation and vertical displacement). Additional implementation details on the 14-DoF model can be found in [13, 26, 27].

3.2 Identification procedure

The vehicle employed in this study is a high performance car whose main parameters are reported in Table 1. Geometric and inertial parameters and tire characteristics have been measured from the actual vehicle and introduced into the simulator model [13–15]. Tire forces are modeled via a Pacejka Magic Formula 6.1 [26]. These data provided the initial baseline for the single-track model used in planning. The parametrization corresponds to a decidedly oversteer-biased setup—deliberately chosen to make the car more demanding to drive and to accentuate any differences in the designed reference trajectory the driver would follow. The geometric and inertial parameters employed in the single-track surrogate were extracted directly from the simulator model; the adopted values are summarized in Table 1.

Table 1: Baseline geometric and inertial parameters extracted from the simulator model and used in the single-track surrogate.

Parameter	Value
Mass	1875 kg
Yaw inertia	3341 kg m ²
Wheelbase	2.97 m
Weight distribution (front/rear)	53/47
CoG height	0.5 m

However, when driven with identical inputs, the two models produced noticeably different responses at high lateral acceleration levels. The mismatch stems both from neglected body dynamics in the single-track and from the fact that axle characteristics represented with an MF structure are not directly mappable from the per-tire parameters of the 14-DoF simulator model.

To model the lateral forces at the front and rear axles of the single-track vehicle we adopted a simplified version of the pure side slip Magic Formula [28]. In particular, for each axle we assume

$$F_y(\alpha, F_z; \mathbf{p}) = D_y \sin(C_y \arctan(B_y \alpha) - E_y(B_y \alpha - \arctan(B_y \alpha))), \quad (7)$$

with $\mathbf{p} = (F_{z0}, p_{Cy1}, p_{Dy1}, p_{Dy2}, p_{Ey1}, p_{Ey2}, p_{Ky1}, p_{Ky2})$ and where

$$\begin{aligned} df_z &= \frac{F_z - F_{z0}}{F_{z0}}, & D_y &= (p_{Dy1} + p_{Dy2} df_z) F_z, \\ C_y &= p_{Cy1}, & B_y &= p_{Ky1} F_{z0} \frac{\sin(2 \arctan(F_z / (p_{Ky2} F_{z0})))}{C_y D_y}, \\ E_y &= p_{Ey1} + p_{Ey2} df_z. \end{aligned}$$

Here, α denotes the axle slip angle, F_z is the corresponding vertical load, and, for each axle, the coefficient vector $\mathbf{p}$ collects the Magic Formula parameters that characterize the lateral-force curve and is treated as the set of calibration variables to ensure consistency between the lateral forces of the simulator model and those of the single-track surrogate.

To reconcile the two-axle model with the 14-DoF one, we adopted the following identification workflow:

(i) Data acquisition. For tire calibration, preliminary simulator sessions were run by driving the 14-DoF vehicle model on the same track used for the planning studies, the Siena kart circuit scaled $\times 2$ to better match the dynamic characteristics of the vehicle considered here. These sessions were used to collect the *training data* for tire identification: from the recorded telemetry, axle slip angles were computed as the mean of the left/right tire slip angles, $\bar{\alpha} = \frac{\alpha_{\text{left}} + \alpha_{\text{right}}}{2}$, while axle vertical and lateral forces were obtained as $\bar{F}_z = F_z^{\text{left}} + F_z^{\text{right}}$, $\bar{F}_y = F_y^{\text{left}} + F_y^{\text{right}}$.

(ii) Parameter identification. Once the training data had been assembled, the single-track tire model was calibrated to reproduce the simulator forces by solving a nonlinear least-squares problem. Denoting by $\mathbf{p}$ the Magic Formula parameter vector associated with a given axle, the calibration problem for that axle was posed as

$$\bar{\mathbf{p}} = \arg \min_{\mathbf{p}} \sum_{k=1}^{N_{\text{train}}} (\bar{F}_{y,k} - F_y(\bar{\alpha}_k, \bar{F}_{z,k}; \mathbf{p}))^2, \quad (8)$$

where $\bar{F}_{y,k}$ denotes the lateral axle force extracted from the 14-DoF simulator training data at sample k , and $F_y(\bar{\alpha}_k, \bar{F}_{z,k}; \mathbf{p})$ is the corresponding force predicted by the single-track tire model, as defined in (7), for the same axle slip angle $\bar{\alpha}_k$ and vertical load $\bar{F}_{z,k}$. The optimization (8) is solved independently for the front and rear axles, starting from an initial guess given by the parameter values in the `.tir` files of the simulator model, yielding distinct parameter vectors for each and thereby aligning the single-track lateral-force response with the 14-DoF telemetry over the training set.

The numerical values of the calibrated Magic Formula parameters collected in $\bar{\mathbf{p}}$ for each axle, together with the starting values used to initialize the optimization, are summarized in Table 2.

Table 2: Calibrated Magic Formula parameter vectors $\bar{\mathbf{p}}$ for the front and rear axles.

	F_{z0} (N)	p_{Cy1}	p_{Dy1}	p_{Dy2}	p_{Ey1}	p_{Ey2}	p_{Ky1}	p_{Ky2}
Starting Values	6500	1.45	1.09	-0.20	0.81	0.34	26.94	3.20
Front axle	11050	2.47	1.85	-0.34	0.57	0.59	28.29	3.04
Rear axle	9210	1.92	1.03	-0.34	0.57	0.59	28.29	3.04

Figure 2 compares the axle lateral forces extracted from the simulator telemetry with those predicted by the calibrated Magic Formula tire model at the front and rear axles, showing close agreement over the operating range of interest.

(iii) Validation. To validate the calibrated single-track model, we compared data from a Nominal Minimum-lap-time (NOM) optimization, i.e., without robustness criteria, against simulator data collected with a human driver attempting to track the NOM reference trajectory generated by the lap-time optimization. Figure 3 illustrates representative comparisons between the longitudinal, lateral, and vertical axle forces obtained from the optimization and those measured in the simulator while the human driver was trying to track the optimized trajectory, confirming the consistency of the single-track model with the high-fidelity simulator dynamics.

Figure 4 shows a comparison between the steering-wheel angle recorded at the driving simulator (red) and the steering input generated by the single-track minimum-lap-time optimization (black) along the same reference state trajectory.




Figure 2: Comparison between axle lateral forces obtained from the simulator telemetry (red) and from the calibrated Magic Formula tire model (black) for the front and rear axles.

4 Experimental Setting and Test Protocol

4.1 Drivers, vehicle and track

Two human drivers participated: VP (Formula SAE and sim-racing experience) and MM (rally background, more expert overall). Both were selected for demonstrated motorsport skills. Trials with passenger-car drivers showed that they could not consistently follow racing-style references: they either deviated from the line or could not sustain the target speed profile, and often incurred catastrophic failures (spins or off-track events requiring a simulator reset), yielding too few valid laps for statistical analysis.

The simulator rendered the high-performance car on a closed circuit (the Kartodromo di Siena, scaled $\times 2$ to match the vehicle’s dimensions). This track was chosen because it is familiar and offers a balanced mix of fast corners, slow corners, and transitional sections. We conveyed the racing reference as a 1 m-wide *racing ribbon* overlaid on the roadway (Figure 5). Drivers were instructed to keep their reference point centered within the ribbon while driving. The ribbon’s centerline was obtained by translating the planned center-of-mass (CoM) path to the driver reference frame (fixed rigid-body offset), so that the visual cue corresponded to the driver’s position rather than to the CoM trajectory. The finite width provided a tolerance band consistent with human execution.

The speed profile was encoded by the ribbon’s color map, transitioning from green (full throttle/acceleration) through intermediate hues to red (hard braking). This dual-cue design (lateral ribbon + longitudinal color) offered continuous guidance without numeric overlays or ghost vehicles.

All tests were completed in a single day. For each driver, we randomized and counterbalanced (i) the sequence of trajectory types to be followed (details provided below) and (ii) the order of driver sessions. Rest periods were inserted between blocks. Crucially, the driver was kept unaware of the trajectory category being provided (single-blind), to reduce strategy or effort biases. At the beginning of the day, each driver completed 30 to 40 laps without guidance to build confidence and familiarize themselves with the track layout.

4.2 Reference trajectories and baseline

All references were generated with our *lapsim* (optimal-planning) model, well aligned with the simulator dynamics after identification. The track was discretized into 2000 spatial intervals, each 1.3 m long. We consider four scenarios:




Figure 3: Comparison between longitudinal, lateral, and vertical axle forces obtained from the simulator (Sim data, red) and from the single-track minimum-lap-time optimization (ST data, black) while attempting to follow the optimized trajectory, plotted against the normalized distances along the track.




Figure 4: Comparison between steering-wheel angle recorded at the driving simulator (Sim data, red) and steering input from the single-track minimum-lap-time optimization (ST data, black), plotted against the normalized distance along the track while attempting to follow the optimized trajectory.

- NOM: fastest feasible line. The planner minimizes lap time with a small secondary penalty that promotes smooth steering actuation so a human can execute it. The *cost structure is identical* across variants; introducing robustness does not change the planning objective, i.e., lap-time minimization.
- TLC: robust planning with constraint tightening focused on *track limits*, i.e., maintaining explicit margins to the road edges.
- FLC: robust planning with tightening focused on *friction limits*, i.e., maintaining margins to axle/tire saturation and biasing states toward more stable tire utilization.
- NO-REF: free driving without a reference, used to assess the impact of following a planned trajectory versus raw driver skill.

Figure 6 shows the two robust trajectories (TLC in the left panel and FLC in the right panel) overlaid on the NOM trajectory in light grey, for a sector of the track. The trajectories are colored to highlight the longitudinal speed difference with respect to the nominal. Notably, TLC follows a noticeably different trajectory that allows, in certain portions of the track, a higher speed than NOM, while FLC follows a trajectory that is closer to NOM with a lower longitudinal



Figure 5: Dynamic Simulator view from the driver’s perspective with the *racing ribbon* overlaid on the scene, indicating the desired trajectory to follow.




Figure 6: Comparison of robust planning trajectories (TLC and FLC) overlaid on the nominal trajectory (NOM) for a representative track sector. Trajectory color indicates longitudinal speed difference relative to the nominal. The TLC trajectory (left) deviates more from NOM and achieves higher speeds in some sections, while FLC (right) closely follows NOM at reduced speeds.

speed.

Robust back-offs (TLC, FLC) are computed by resetting the state covariance at each curvilinear grid point and propagating it over a horizon of H steps. In our experiments, setting $H=4$ results in a maximum delay of 5.2m in the driver’s corrective actions. The grid spacing and the reset covariance assumptions were calibrated from pilot tests so that the induced tightenings correspond to a 99% constraint-satisfaction confidence, implemented via a safety factor $\gamma=3$ (three-sigma). In practice, these settings yield a tangible increase in constraint margins while preserving executability.

4.3 Metrics

We evaluate performance and tracking using the following quantities (with respect to the *active* reference, i.e., NOM, TLC, or FLC):

Lap Time (s). Total lap time LT, in seconds, over valid laps.

Steer energy (rad²/s). As a proxy for *steering activity*, we use

$$E_s = \int_0^T \dot{\delta}(t)^2 dt, \quad (9)$$

so lower values indicate smoother, more executable actuation. We introduce the root-mean-square (RMS) over a time horizon $[0, T]$. For a signal $f(\cdot)$:

$$\text{RMS}(f) \triangleq \sqrt{\frac{1}{T} \int_0^T f(t)^2 dt}. \quad (10)$$

and use it to define both *tracking fidelity* and *style indicators*.

Tracking fidelity. RMS lateral position and speed errors, measuring adherence to the planned line and speed profile:

$$\text{RMS}(e_y) \text{ (m)}, \quad \text{RMS}(e_v) \text{ (m/s)}. \quad (11)$$

Statistics are computed over all valid laps per driver and condition.

Style indicators. RMS side-slip usage:

$$\text{RMS } \beta_{\text{drv}} \quad \text{and} \quad \text{RMS } \beta_{\text{ref}},$$

to compare the driver's typical sideslip β_{drv} with that implied by the planned trajectory β_{ref} .

5 Results




Figure 7: Lap times by test and driver (box plots). Grey lines indicate the reference laptime for each trajectory type. For driver VP, NOM is fastest and FLC is moderately slower; for driver MM, NOM and FLC yield comparable medians; TLC is slower for both; NO-REF for driver MM is slower and more variable while driver VP shows a more consistent performance with higher medians.

5.1 Lap-time distributions

Figure 7 juxtaposes the lap-time (LT) box plots for each trajectory–driver pair and provides the main visual reference for this subsection.

For VP, the leftmost boxes (NOM) anchor the fastest medians, confirming the benchmark role of the nominal planner for this driver. For MM, however, the FLC and NOM boxes are




Figure 8: Steer Energy by test and driver (log scale). NOM requires the highest effort among reference-following tests; FLC and TLC markedly reduce effort. For driver MM, NO-REF shows *high* SE coupled with slow LTs.

practically superposed: their medians lie within each other’s IQR and any LT gap is visually negligible, indicating comparable lap times under these two references. Notably, for the robust trajectories (FLC and TLC), driver MM achieves laptimes faster than the planner’s reference (indicated by the grey lines), suggesting that the calibration phase can be improved: MM operates with lower margins than those assumed by the planner, effectively exploiting the vehicle envelope more aggressively. This observation connects to the RMS tracking errors discussed in Section 5.4: MM’s lower trajectory-following fidelity (higher RMS errors, as visible in Figure 11) enables these faster laptimes by deviating from the planned path to exploit tighter margins. The FLC column for VP shifts vertically by roughly 1–2s relative to NOM yet preserves a compact Interquartile Range (IQR), signalling a modest time penalty in exchange for tighter dispersion. TLC occupies the clearly slower columns in Figure 7: both drivers show medians several seconds above NOM, consistent with the more conservative curvature shaping. Finally, the NO-REF column for MM exhibits the tallest whiskers and a median higher than NOM and FLC, highlighting the variability that accompanies unguided laps. For NO-REF runs, both drivers show medians above their NOM and FLC benchmarks, but with distinct dispersion patterns: VP displays relatively short whiskers with a few isolated outliers, whereas MM shows taller whiskers without clear outliers, indicating more run-to-run variability in typical laps but no extreme excursions.

5.2 Steering effort

Figure 8 mirrors the structure of Figure 7 but maps the steer energy SE distributions on a logarithmic axis, making the effort differences visually salient. The NOM boxes and the NO-REF box for MM occupy the upper part of the axis, meaning that both drivers pay the highest steering cost whenever the nominal reference is enforced or free driving is performed. Moving to the FLC and TLC columns in the same figure reveals remarkable effort relief: for VP, FLC reaches the lowest median SE, while MM exhibits the ordering $TLC < FLC < NOM$ with partial overlap between FLC and TLC quartiles.

The rightmost column (NO-REF for MM) should be read together with the LT information: its elevated median and wide spread in Figure 8, combined with the slow and dispersed LT box in Figure 7, illustrate that free driving simultaneously degrades time and effort metrics relative




Figure 9: Overall Lap Time–Steer Energy scatter plot (test types as colors, both drivers). FLC typically lies near a Pareto-efficient band; NOM is *generally*² fastest but costly; TLC is lowest effort but slower.

to the reference-guided alternatives, especially for NOM and FLC.

5.3 Time vs effort Pareto structure

Figures 9 and 10 condense the box-plot insights into two scatter plots that directly show the LT–SE trade-off. In Figure 9, markers colored by the reference-guided trajectories fall into three recognizable clusters: NOM points accumulate along the lower boundary in LT but sit high in SE, TLC marks gravitate toward the upper-left (slow but economical), and FLC forms an elongated ribbon close to an empirical lower-left Pareto front linking intermediate LTs with markedly lower SE. Separating the runs by driver in Figure 10 clarifies how this ribbon behaves: the boxed markers for MM lie near the fast end of the FLC band, whereas the circle markers for VP extend the same band toward slightly longer LTs while maintaining a considerable effort gap to NOM. As shown in Figures 9 and 10, the NO-REF trajectory confirms the picture anticipated in the previous section: VP exhibits a smaller dispersion in lap time at a moderate steering-energy level, with his runs clustering near the middle of the scatter plot, whereas MM populates the right-hand side of the plot at higher steering-energy levels with a much larger spread in lap times.

5.4 RMS tracking errors and activity

Figure 11 reports, for each reference condition, two four-bar clusters—left for VP and right for MM. Within each cluster, four quantities are shown: (i) $\text{RMS}(e_y)$, (ii) $\text{RMS}(e_v)$, (iii) $\text{RMS}(\beta_{\text{drv}})$ with the corresponding $\text{RMS}(\beta_{\text{ref}})$ overlaid in black, and (iv) a light-colored bar for $\text{RMS}(\delta)$ as a qualitative proxy for SE.

²This is because only the lowest lap-time sample corresponds to TLC.




Figure 10: Lap Time–Steer Energy scatter plot with driver markers (colors = test types; markers distinguish driver VP and driver MM). FLC retains the favorable band for both drivers; NO-REF for driver MM sits at high SE and high LT.

Across conditions, FLC typically pulls $\text{RMS}(e_y)$ and $\text{RMS}(e_v)$ below NOM, while simultaneously reducing $\text{RMS}(\dot{\delta})$; the effect is clearest for VP but remains visible for MM. Under TLC both drivers see reduced $\text{RMS}(\beta_{\text{drv}})$ and $\text{RMS}(\dot{\delta})$ —consistent with its more relaxed trajectory shaping—but at the price of increased *Lap Time* and tracking errors: VP experiences a noticeable rise in speed error $\text{RMS}(e_v)$, whereas MM shows larger lateral error $\text{RMS}(e_y)$ relative to NOM. According to the drivers’ qualitative feedback, this is because the TLC path feels somewhat unnatural to follow, tending to keep close to the lane center rather than supporting the corner-cutting strategies they would normally adopt. The larger RMS errors observed for MM under both FLC and TLC correlate with his faster-than-planner laptimes (visible in Figure 7): by deviating more from the planned trajectory than VP, MM exploits lower margins and achieves superior performance, indicating that the planner’s conservative calibration leaves room for skilled drivers to operate closer to the vehicle limits.

Beyond these aggregate trends, the NO-REF cluster also highlights differences in driving style and background between the two drivers: VP, who is less experienced and more conservative, exhibits lower $\text{RMS}(\beta_{\text{drv}})$ and $\text{RMS}(\dot{\delta})$ than in the NOM condition, where he attempts to push harder, whereas MM, the more experienced driver, operates closer to the vehicle limits and displays a rally-style behavior with larger $\text{RMS}(\beta_{\text{drv}})$ and $\text{RMS}(\dot{\delta})$ under NO-REF than when following reference trajectories. This further indicates that, when a reference is provided, both drivers deliberately move away from their natural style. Crucially, the *natural* behavior observed in NO-REF is not a *comfort zone*, especially for MM: it yields worse performance and higher energy expenditure, because the free-run trajectories they gravitate toward are, paradoxically, less stable and demand more corrective action—hence a larger steering effort SE.




Figure 11: RMS metrics by test and driver (driver order: VP then MM). FLC eases tracking vs. NOM and reduces steering activity; TLC is most “relaxed” but slowest; NO-REF for MM combines poor tracking with elevated SE and slow LTs.

5.5 Summary tables (medians and IQR)

To complement the description above with concise descriptors, in Tables 3 and 4 we report *medians* and *interquartile ranges* (IQR) that are extracted from Figures 7 and 8. The median captures the central lap-time or steer-energy tendency even when distributions are skewed, while the IQR (the span between the 25th and 75th percentiles) reflects how tightly each controller concentrates its performance. Because our sample sizes per condition are modest, the IQR is especially informative: a narrow band signals repeatable behavior, whereas a wide one warns that the observed median may be sensitive to session-to-session variability.

Taken together, Tables 3–5 sharpen the picture sketched by the box plots. For VP, FLC trades a modest 2% increase in LT for an 81% reduction in SE relative to NOM, while TLC pushes LT up by a further 3% with only a slight SE gain over FLC. For MM, FLC essentially preserves LT (a 0.3% improvement vs. NOM) while still cutting SE by about 60%, whereas TLC delivers an additional 34% SE reduction at the cost of a 2% LT increase. The NO-REF rows confirm that free driving is systematically dominated: both drivers incur longer LTs than under NOM/FLC and substantially higher SE, with MM in particular combining the slowest and most variable times with the largest effort.

Table 3: Lap Time (s): medians and IQR by driver and test.

	NOM	TLC	FLC	NO-REF
VP Median [IQR]	85.47 [1.58]	89.92 [1.15]	87.32 [1.49]	87.60 [0.76]
MM Median [IQR]	84.87 [0.55]	86.63 [0.53]	84.65 [0.69]	86.11 [1.21]

Table 4: Steer Energy (rad^2/s): medians and IQR by driver and test (log-scale plots).

	NOM	TLC	FLC	NO-REF
VP Median [IQR]	552 [232]	142 [93]	103 [77]	234 [212]
MM Median [IQR]	526 [180]	138 [88]	210 [161]	605 [236]

Table 5: Pairwise median differences (FLC vs NOM; TLC vs FLC). Negative Δ in SE denotes effort reduction.

	VP		MM	
	ΔLT [s]	ΔSE [rad^2/s]	ΔLT [s]	ΔSE [rad^2/s]
FLC – NOM	+1.85 (+2%)	−449 (−81%)	−0.22 (−0.3%)	−316 (−60%)
TLC – FLC	+2.60 (+3%)	+39 (+38%)	+1.98 (+2%)	−72 (−34%)

5.6 Discussion

Finding 1 (drivability trade-off, preliminary). Taken together, Figures 7, 8, and 11 position FLC on a “reduced-effort / comparable-time” band: for driver MM the LT gap to NOM is within the quartiles of the benchmark box plot, whereas for driver VP the penalty reaches 1–2s yet coincides with the lowest SE box in Figure 8 and a marked drop in RMS indicators. Given the modest number of laps per condition, this constellation should be read as an encouraging trend rather than a definitive performance guarantee.

Finding 2 (conservatism lever). The combination of Figures 7 and 8 indicates that TLC achieves low SE medians (the lowest among all conditions for driver MM), but at the cost of a marked increase in LT. In motorsport, where lap time is the primary objective, this effort-time trade-off is generally unacceptable. Moreover, TLC’s margin-keeping geometry discourages decisive corner cutting, lengthens paths, and depresses apex speeds—yielding trajectories that feel less natural to fast drivers despite the reduced effort.

Finding 3 (driver modulation). Figures 10 and 7 jointly indicate that the LT penalty under FLC depends on the driver: MM keeps lap times close to NOM (clustered at the fast end of the FLC ribbon), whereas driver VP trades about 1–2s for the same controller. More strikingly, MM exceeds the planner’s reference laptime for both FLC and TLC (as indicated by the grey lines in Figure 7), demonstrating that skilled drivers can operate with lower margins than those assumed during the planner’s calibration phase. This performance advantage comes at the cost of trajectory fidelity: MM’s higher RMS tracking errors (Figure 11) reflect his deliberate deviations from the planned path to exploit tighter margins, revealing a trade-off between strict trajectory following and laptime performance that the planner’s conservative calibration does not capture. Free driving of MM (NO-REF) reinforces this driver effect: in both figures the corresponding points occupy the high-SE, medium-high-LT side, suggesting that the reference simplifies the exploitation of the vehicle envelope.

Implications and open questions. The qualitative structure in Figures 7–11 hints that FLC biases closed-loop operation away from abrupt friction-utilization changes, whereas NOM remains sensitive to modeling mismatch and TLC slows the car by maintaining distance from geometric limits.

Future experiments should probe whether these patterns persist with more participants, alternative tracks, longer exposure, and richer effort proxies (including cognitive-load measures), and should explicitly quantify sequencing effects that are not controlled in the present dataset.

6 Conclusion

The experiments demonstrate that disturbance-aware planned trajectories yield distinct operating regimes. We recall that NOM denotes the nominal time-optimal trajectory, TLC a track-limit-robust time-optimal trajectory obtained by tightening margins to the track edges, and FLC a friction-limit-robust time-optimal trajectory obtained by tightening against axle/tire saturation. Apart from a single sample, NOM retains the lap-time benchmark overall but imposes the highest steer-energy expenditure among the reference-following conditions; TLC markedly lowers effort yet elongates trajectories, eroding its usefulness in lap-time-driven contexts; and FLC remains close to the empirical LT-SE Pareto band, delivering substantial effort savings with only marginal penalties, especially for driver MM. Notably, driver MM achieved lap-times faster than the planner’s reference for both FLC and TLC, indicating that the planner’s calibration margins may be conservative for skilled drivers who can operate closer to vehicle limits. This performance advantage correlates with MM’s higher RMS tracking errors, reflecting his deliberate deviations from the planned trajectory to exploit tighter margins—a finding that suggests the calibration phase can be refined to better match driver capabilities. Free driving (NO-REF) offers no significant advantage: medians for MM reach 86.11 s and 605 rad²/s, confirming that unguided laps are simultaneously slower and costlier.

These observations position FLC as the most promising compromise for applications that must balance pace and driver workload, including endurance stints, acclimation to unfamiliar circuits, or vehicle configurations with pronounced uncertainty.

The study remains exploratory—two drivers and limited laps per condition—so the findings should be validated with broader participant pools, additional tracks, and richer workload indicators (cognitive or physiological). Such extensions will clarify how persistent the favorable FLC band and the dominated NO-REF regime are when exposure time, learning effects, and driver diversity increase.

Funding

This work is supported by project PRIN 2022 PNRR “Global Stability of road vehicle motion - STAVE” CUP I53D23005670001.

References

- [1] B.-C. Chen and C.-C. Kuo, “Electronic stability control for electric vehicle with four in-wheel motors,” *Int.J Automot. Technol.*, vol. 15, no. 4, pp. 573–580, 2014.
- [2] B. Lenzo, M. Zanchetta, A. Sorniotti, P. Gruber, and W. De Nijs, “Yaw Rate and Sideslip Angle Control Through Single Input Single Output Direct Yaw Moment Control,” *IEEE Trans. Contr. Syst. Technol.*, vol. 29, no. 1, pp. 124–139, 2021.
- [3] F. Della Rossa, G. Mastinu, and C. Piccardi, “Bifurcation analysis of an automobile model negotiating a curve,” *Vehicle System Dynamics*, vol. 50, no. 10, pp. 1539–1562, 2012.
- [4] E. Hashemi, M. Pirani, A. Khajepour, and A. Kasaiezadeh, “A comprehensive study on the stability analysis of vehicle dynamics with pure/combined-slip tyre models,” *Vehicle System Dynamics*, vol. 54, no. 12, pp. 1736–1761, 2016.
- [5] Y. Huang, W. Liang, and Y. Chen, “Stability Regions of Vehicle Lateral Dynamics: Estimation and Analysis,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 143, no. 051002, 2020.

- [6] A. M. Ribeiro, A. R. Fioravanti, A. Moutinho, and E. C. De Paiva, “Sum of squares approach for ground vehicle lateral control under tire saturation forces,” *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 14387–14393, 2020.
- [7] Z. Zhu, Y. Zhang, Y. Huang, and Y. Qin, “Sum-Of-Squares Based Vehicle Dynamic Stability Method and Its Applications in ADAS,” in *2022 IEEE Intelligent Vehicles Symposium (IV)*, (Aachen, Germany), pp. 247–254, IEEE, 2022.
- [8] S. Sadri and C. Wu, “Stability analysis of a nonlinear vehicle model in plane motion using the concept of Lyapunov exponents,” *Vehicle System Dynamics*, vol. 51, no. 6, pp. 906–924, 2013.
- [9] S. Zhang, C. Zhao, Z. Zhang, and Y. Lv, “Driving simulator validation studies: A systematic review,” *Simulation Modelling Practice and Theory*, vol. 138, p. 103020, Jan. 2025.
- [10] Y. Maruyama and F. Yamazaki, “Driving simulator experiment on the moving stability of an automobile under strong crosswind,” *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 94, pp. 191–205, Apr. 2006.
- [11] D. Wanner, M. Kreußlein, B. Augusto, L. Drugge, and A. Stensson Trigell, “Single wheel hub motor failures and their impact on vehicle and driver behaviour,” *Vehicle System Dynamics*, vol. 54, pp. 1345–1361, Oct. 2016.
- [12] H. Cho, Y. Kim, S. Oh, and I. Yun, “Development of a Stability Index for Evaluating Drivers’ Psychological Stability During Truck Platooning,” *Applied Sciences*, vol. 15, p. 5429, May 2025.
- [13] G. Mastinu, D. Biggio, F. Della Rossa, and M. Fainello, “Straight running stability of automobiles: Experiments with a driving simulator,” *Nonlinear Dynamics*, vol. 99, pp. 2801–2818, Mar. 2020.
- [14] G. Mastinu, F. Della Rossa, G. Previati, M. Gobbi, and M. Fainello, “Global stability of road vehicle motion with driver control,” *Nonlinear Dyn*, vol. 111, no. 19, pp. 18043–18059, 2023.
- [15] G. Mastinu, G. Previati, F. Della Rossa, M. Gobbi, and M. Fainello, “How Drivers Lose Control of the Car,” *SAE Int. J. Veh. Dyn., Stab., and NVH*, vol. 8, no. 1, pp. 10–08–01–0007, 2024.
- [16] G. Previati, G. Mastinu, and M. Gobbi, “Driver–steering wheel interaction during cornering,” *SAE Int. J. Veh. Dyn., Stab., and NVH*, vol. 8, no. 3, pp. 427–442, 2024.
- [17] S. Kharrazi, B. Augusto, and N. Fröjd, “Vehicle dynamics testing in motion based driving simulators,” *Vehicle System Dynamics*, vol. 58, pp. 92–107, Jan. 2020.
- [18] E. Musso, G. Previati, and G. Mastinu, “Analytical and experimental handling performance of ultra-efficient lightweight vehicles,” *SAE Int. J. Adv. & Curr. Prac. in Mobility*, vol. 6, no. 4, pp. 2049–2058, 2024.
- [19] M. Asperti, M. Vignati, and E. Sabbioni, “Design of electric power steering control for compensating the torque steer effect due to torque vectoring control,” *SAE Int. J. Veh. Dyn., Stab., and NVH*, vol. 9, no. 2, pp. 177–193, 2025.
- [20] F. Alfatti, M. Montani, T. Favilli, C. Annicchiarico, L. Berzi, M. Pierini, L. Pugi, and R. Capitani, “Implementation and Performances Evaluation of Advanced Automotive Lateral Stability Controls on a Real-Time Hardware in the Loop Driving Simulator,” *Applied Sciences*, vol. 13, p. 6592, May 2023.

- [21] T. Novi, A. Liniger, R. Capitani, M. Fainello, *et al.*, “The influence of autonomous driving on passive vehicle dynamics,” *SAE Int. J. Veh. Dyn., Stab., and NVH*, vol. 2, no. 4, pp. 285–295, 2018.
- [22] F. Della Rossa, M. Fontana, S. Giacintucci, M. Gobbi, *et al.*, “Early detection of the loss of control of motorsport vehicles,” *SAE Int. J. Veh. Dyn., Stab., and NVH*, vol. 9, no. 3, pp. 473–499, 2025.
- [23] M. Gulisano, M. Masoni, M. Gabiccini, and M. Guiggiani, “Disturbance-aware minimum-time planning strategies for motorsport vehicles with probabilistic safety certificates,” *Vehicle System Dynamics*, pp. 1–25, Oct. 2025.
- [24] PoliMi, “Drismi (official website).” <https://www.drismi.polimi.it>, 2024. Accessed Nov. 2025.
- [25] M. Guiggiani, *The Science of Vehicle Dynamics: Handling, Braking, and Ride of Road and Race Cars*. Springer International Publishing, 3rd ed., 2023.
- [26] VI-grade GmbH, “Vi-grade (official website).” <https://www.vigrade.com>, 2024. Accessed Nov. 2024.
- [27] G. Mastinu and M. Ploechl, eds., *Road and Off-Road Vehicle System Dynamics Handbook*. CRC Press, 1st ed., 2014.
- [28] H. B. Pacejka, *Tire and Vehicle Dynamics*. Elsevier, 3rd ed., 2012.