



Finanziato
dall'Unione europea
NextGenerationEU



Ministero dell'Università e della Ricerca



Global Stability of Road Vehicle Motion - STAVE

Deliverable 3 Stability and Path Tracking

*Document in fulfilment of
MILESTONE 3*

“Stable path tracking for motorsport or road vehicles performed”



This project has received funding from Italian Ministero del Merito e della Ricerca, PRIN: Progetti di ricerca di rilevante interesse nazionale – Bando 2022 PNRR Prot. P2022KSN9Z

Table of Contents

Table of Contents.....	2
Introduction.....	4
1 Early Detection of the loss of control of motorsport vehicles.....	5
1.1 Degree of Stability (DoS) index: Driving simulator tests.....	5
1.1.1 Test setup: Single Lane changes	5
1.1.2 Estimation of the Saddle Limit Cycle	6
1.1.3 Application of the DoS index to Single-Lane change maneuvers	7
1.2 References.....	9
2 Global stability of vehicle-with-driver dynamics via Sum-of-Squares programming	10
2.1 Introduction.....	10
2.1.1 Literature Review.....	10
2.1.2 Contributions and organization.....	11
2.2 ROA estimation via Lyapunov functions.....	12
2.2.1 General formulation	12
2.2.2 SOS formulation	13
2.2.3 Iterative SOS formulation	15
2.2.4 2.4 Safe subset with state constraints.....	16
2.3 SOS-friendly vehicle-with-driver model	18
2.3.1 Vehicle-with-driver model.....	18
2.3.2 Polynomial approximation.....	19
2.4 Vehicle-with-driver safe invariant set.....	20
2.4.1 Tested scenarios and SOS procedure set-up	20
2.4.2 Results.....	23
2.4.3 Discussion.....	24
2.5 Conclusions.....	25
2.6 References.....	25
3 Disturbance-Aware Minimum-Time Trajectory Planning: Evidence from Tests on a Dynamic Driving Simulator.....	28
3.1 Introduction.....	28
3.2 Open-loop disturbance-aware planning	30
3.2.1 Vehicle and disturbance model	31
3.2.2 Mean and covariance propagation	31
3.2.3 Collocation discretization and H -step predictions	32
3.2.4 Open-loop disturbance-aware OCP.....	32

3.2.5	Constraint back-offs and planner variants.....	32
3.2.6	Parameters selection.....	33
3.3	Vehicle model identification procedure.....	33
3.3.1	DriSMi layout and vehicle dynamics model.....	33
3.3.2	Identification procedure.....	34
3.4	Experimental Setting and Test Protocol.....	37
3.4.1	Drivers, vehicle and track.....	37
3.4.2	Reference trajectories and baseline.....	37
3.4.3	Metrics.....	38
3.5	Results.....	39
3.5.1	Lap-time distributions.....	39
3.5.2	Steering effort.....	40
3.5.3	Time vs effort Pareto structure.....	41
3.5.4	RMS tracking errors and activity.....	42
3.5.5	Summary tables (medians and IQR).....	43
3.5.6	Discussion.....	44
3.6	Conclusion.....	45
3.7	References.....	45
4	Conclusions.....	48

Introduction

This deliverable presents three complementary research activities supporting Milestone 3, *Stable path tracking for motorsport or road vehicles*. The first contribution focuses on **stability monitoring in motorsport**, where early detection of loss-of-control conditions is essential for safety and performance. Building on the theoretical elements introduced in Deliverable 2, a stability-sensing methodology is further developed through the analysis of trajectories near saddle-type limit cycles and the associated Floquet-inspired Degree of Stability (DoS) indicator. Recent driving-simulator experiments at Politecnico di Milano demonstrate the capability of the DoS metric to distinguish recoverable from non-recoverable maneuvers, even when recovery appears successful but remains dynamically unsafe.

The second contribution addresses **mathematically certified stability for vehicle-with-driver systems**. Extending the preliminary results reported in Sec. 3.2 of Deliverable 2 (Milestone 2), Lyapunov-based Sum-of-Squares techniques are employed to estimate safe operating regions endowed with formal guarantees. This approach provides a systematic alternative to heuristic stability thresholds and offers a rigorous foundation for supervisory control architectures in both road-vehicle scenarios and, in perspective, high-performance driving.

The third contribution examines the **execution of robust, disturbance-aware reference trajectories** and their influence on actual driving performance. The approach has been already introduced in Sec. 3.3 of Deliverable 2 (Milestone 2). In collaborative tests involving Politecnico di Milano and Università di Pisa, three planned trajectories—nominal time-optimal (NOM), track-limit-robust (TLC), and friction-limit-robust (FLC)—are evaluated in a dynamic driving simulator with professional drivers. The results highlight a clear trade-off between lap time and steering effort, and show that robustness-embedded trajectories, particularly the FLC variant, can provide fast yet stable driving with reduced control demand compared with nominal or unguided driving.

Taken together, these three research lines form a coherent pathway: from **detecting instability**, to **certifying stability**, to **planning and experimentally validating robust trajectories**. This integrated perspective reinforces both the theoretical foundations and practical applicability of stable path tracking for motorsport and advanced road-vehicle contexts.

1 Early Detection of the loss of control of motorsport vehicles

In section 2.3 of Deliverable 2, we have shown that the loss of vehicle control is directly linked to the existence of a **saddle-type limit cycle**, and we have proposed a novel methodology for vehicle stability sensing based on the mathematical analysis of trajectories near the saddle cycle. By exploiting Floquet theory, we have thus defined a **Degree of Stability (DoS) index** to recognize whether the driver will ultimately be able to recover the steady-state condition. Even when maneuvers are technically "recovered", they may feature strong oscillations and thus be considered unsafe. This effect is captured by the DoS index, representing an important advantage over traditional instability detection techniques.

The effectiveness of the DoS index was initially demonstrated using a simple vehicle-and-driver model (Section 1.1 of Deliverable 2). In our latest work (Della Rossa et al. 2025), we have tested the index using the **driving simulator** facility available at Politecnico di Milano.

1.1 Degree of Stability (DoS) index: Driving simulator tests

1.1.1 Test setup: Single Lane changes

At the dynamic driving simulator, a series of **single lane changes** have been performed by an actual human driver. The vehicle model used at the driving simulator is the VI-CarRealTime model that has been described as Model 1 in Deliverable 1, in both the oversteering and understeering configuration. The setup for each maneuver is shown in Fig. 1.

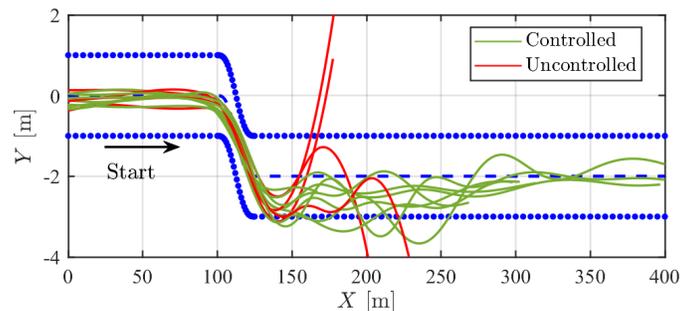


Figure 1: Setup for single lane change maneuvers at the dynamic driving simulator. Two rows of cones (blue) delimit the trajectory that the driver has to follow. The driver is instructed to drive the virtual vehicle between the cones, although the latter do not constitute a physical barrier. The forward speed is kept constant by using a cruise control. The center of gravity trajectories followed by vehicle-and-driver while performing the maneuver are plotted in green (controlled maneuver) and in red (uncontrolled maneuver). Oversteering vehicle, $u_{ref} = 140$ km/h.

The setup consists of cones which are initially placed along two parallel rows on the virtual road. After 100 m, a transient section (25 m in length in the longitudinal direction) leads to the target lane. The two rows of cones are laterally displaced by 2 m relative to their original positions, simulating a lane change. The road surface is smooth and the coefficient of friction is set to $\mu = 1$.

Throughout the maneuver, the driver is instructed to keep the vehicle between the cones, acting exclusively on the steering wheel. Cruise control is used to maintain a constant forward speed, which is set to 140 km/h for maneuvers with the oversteering vehicle and to 220 km/h for the understeering vehicle. The cones do not represent physical barriers, they are only placed to guide the driver during the maneuver. Indeed, the spacing between the two rows of cones (2 m) is intentionally set close to the width of the vehicle (about 1.6 m).

The tests have all been conducted by the same driver. Nonetheless, **there are occasions when the driver loses control of the vehicle and other cases when he completes the maneuver in a controlled manner.** The outcome of each maneuver mostly depends on how the driver initiates the lane change.

1.1.2 Estimation of the Saddle Limit Cycle

The trajectories in the state space (orbits) traced by vehicle-and-driver during all the performed maneuvers are shown in Fig. 2 in the (β, r) -plane. **The motion of the (virtual) complex vehicle driven by a human driver is clearly influenced by a saddle-type limit cycle**, both in the understeering (Fig. 2a) and oversteering (Fig. 2b) configuration. Whether the motion is controlled or not is determined by how the trajectory approaches the saddle cycle.

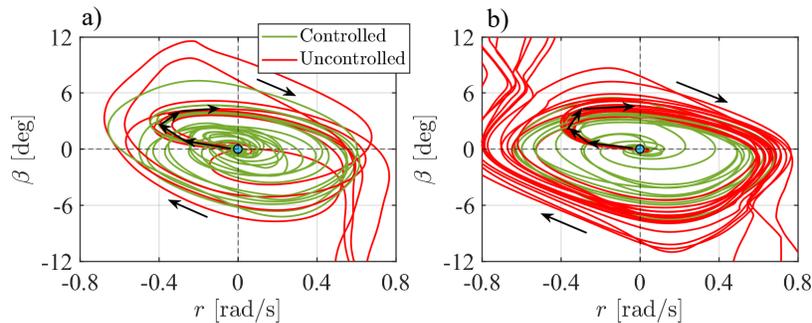


Figure 2: Single lane changes performed by a human driver at the dynamic driving simulator. The orbits of all maneuvers are represented in the (β, r) -plane (red: uncontrolled motion, green: controlled motion). a): tests executed with the oversteering vehicle ($u_{ref} = 140$ km/h); b): tests executed with the understeering vehicle ($u_{ref} = 220$ km/h). A saddle-type limit cycle influences the motion of the system in both cases.

The exact limit cycle is impossible to compute, since we have a real human driver in the loop. The behaviour of the real driver can, however, be reproduced using the simple mathematical model described in section 2.1 of Deliverable 1, after properly adjusting the steering control parameters (k_p, k_d, T_p, τ).

The VI-CarRealTime vehicle model considered in the analysis represents an additional limitation. Indeed, performing a bifurcation analysis of a full 14 DoF vehicle model is too computationally expensive. A faster computational model is highly appropriate for bifurcation analyses, therefore a **feedforward artificial neural network (ANN)** was developed to **emulate the dynamics of the complex vehicle model**. To train the ANN, several open loop maneuvers (step steer, ramp steer, pulse steer and swept steer) and closed loop maneuvers (double lane changes) have been performed with the complex vehicle model in VI-CarRealTime, considering a smooth road and a friction coefficient of $\mu = 1$. The maneuvers were selected to achieve a sufficiently complete exploration of the state space of the dynamical system, in order to characterize the behavior of the vehicle under a wide range of conditions.

A bifurcation analysis was performed on the combined vehicle-and-driver model (ANN vehicle model + simple driver model), to estimate the saddle-type limit cycle observed at the driving simulator. The limit cycles obtained for the oversteering and understeering vehicle configuration are plotted in Fig. 3. It can be clearly seen that the trajectories of both uncontrolled and controlled motion are well represented by the limit cycle.

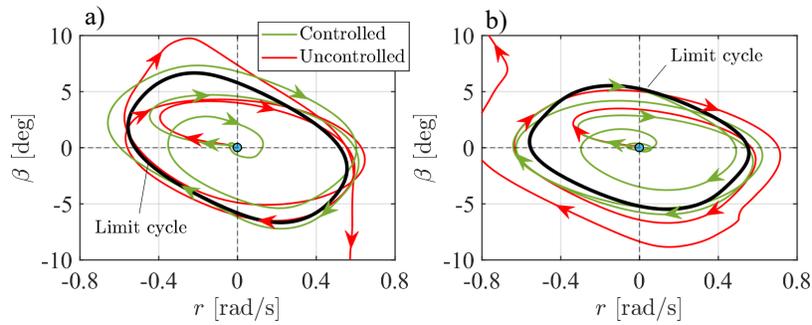
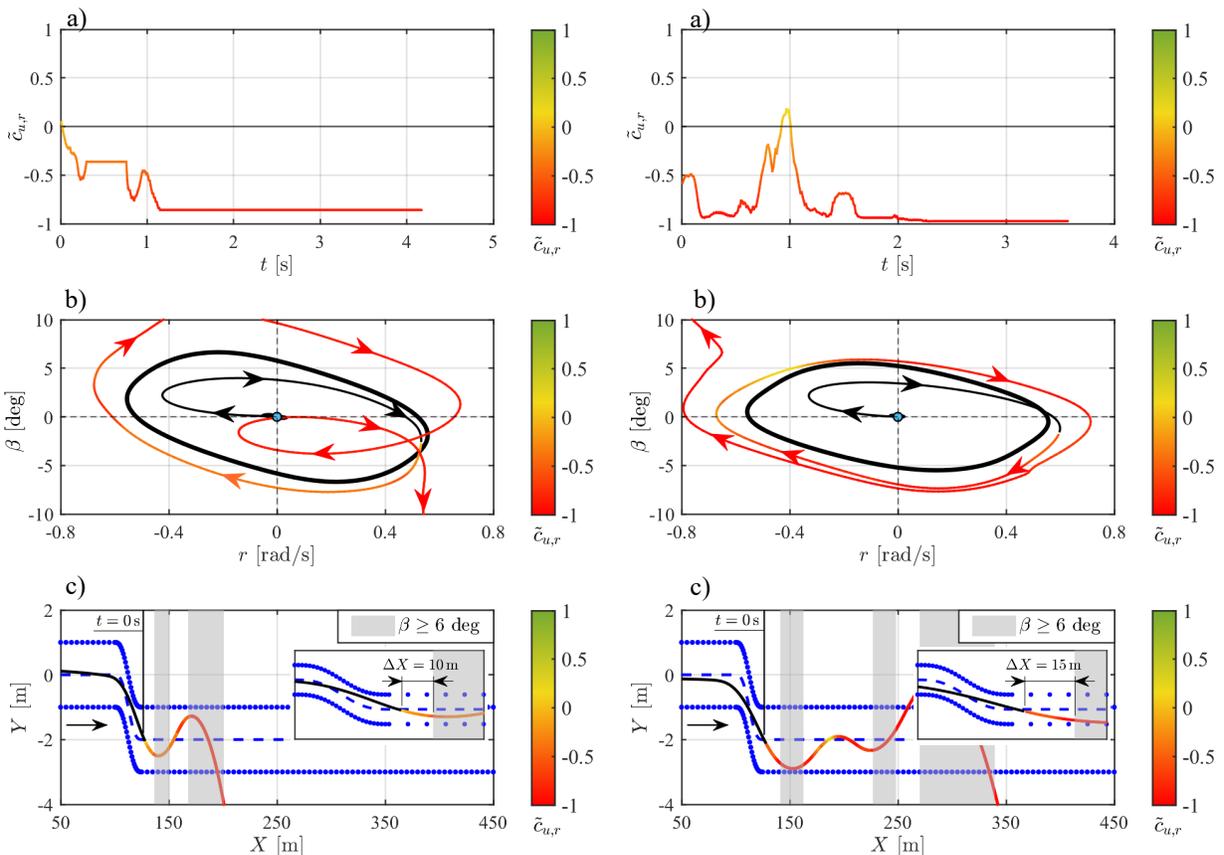


Figure 3: Single lane changes performed by a human driver at the driving simulator (green: controlled maneuver, red: uncontrolled maneuver). a): oversteering vehicle ($u_{ref} = 140$ km/h); b): understeering vehicle ($u_{ref} = 220$ km/h). The motion of vehicle-and-driver is influenced by a saddle-type limit cycle, whose approximation is plotted in black. The limit cycle was computed using the neural network vehicle model coupled with the simple driver model.

1.1.3 Application of the DoS index to Single-Lane change maneuvers

The DoS index can be computed using the estimated saddle-cycle to detect motion instability during the lane-change maneuvers performed by the real driver at the driving simulator. Figs. 4 and 5 present the analysis of two maneuvers: one executed with the oversteering vehicle and one with the understeering vehicle. The objective of the analysis is to assess whether **the DoS index provides relevant information on stability**.



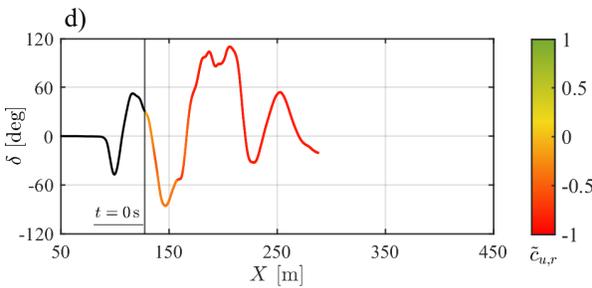


Figure 4: Single lane change performed by a human driver at the driving simulator. Oversteering vehicle, $u_{ref} = 140$ km/h. a): time history of the DoS index $\tilde{c}_{u,r}$ during the maneuver. b): orbit traced by vehicle-and-driver in the (β, r) -plane. The limit cycle is shown as a thick black line. c): center of gravity trajectory followed by vehicle-and-driver while performing the lane change. d): driver steering wheel input (positive counterclockwise). All data gathered during the maneuver is colored based on the value of $\tilde{c}_{u,r}$. Here, $t = 0$ s marks the end of the transient part of the maneuver, where the vehicle runs from the starting lane to the target lane. The loss of control is detected by the DoS index ($\tilde{c}_{u,r} < 0$) immediately after the disturbance has acted, and 10 m before the slip-angle threshold ($\beta = 6$ deg) is first exceeded.

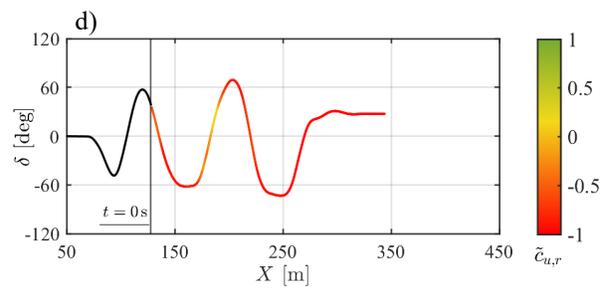


Figure 5: Single lane change performed by a human driver at the driving simulator. Understeering vehicle, $u_{ref} = 220$ km/h. See Fig. 4 for explanation of the general plot structure. The loss of control is detected by the DoS index ($\tilde{c}_{u,r} < 0$) immediately after the disturbance has acted, and 15 m before the slip-angle threshold ($\beta = 6$ deg) is first exceeded. The index $\tilde{c}_{u,r}$ briefly climbs towards positive values as the driver appears to recover control of the vehicle, at about $t = 1$ s ($X = 180$ m), before sharply decreasing again when control is fully lost.

When compared to the classical stability/controllability limit ($\beta = 6$ deg) (Mastinu and Ploechl 2014), **the DoS index demonstrates superior performance in identifying dangerous conditions**. The DoS index can detect unstable behavior well before the slip-angle threshold is exceeded (about 10 m in advance in Fig. 4c, and 15 m in advance in Fig. 5c). The detection occurs at the end of the disturbance. In practice, a few milliseconds are needed to compute the index, assuming that the limit cycle is already available.

Fluctuations of the DoS index observed throughout the maneuvers are partly due to the driver's real-time adjustments in response to instability cues, which cannot be predicted by the index. Indeed, the DoS index is based on an estimated limit cycle obtained with constant driver control parameters. Nevertheless, the continuous computation of the **DoS index enables capturing the driver's real-time performance** more accurately and could prove to be valuable in identifying errors or effective control actions executed by the driver. For instance, in Fig. 5a, the index briefly climbs towards positive values as the driver appears to recover control of the vehicle, before sharply decreasing again.

In the actual case of a road vehicle running on a generic road, a digital twin of the vehicle should be available to compute the limit cycle. Additionally, the driver's parameters should be somehow estimated. This makes the computation of the DoS index quite complex at the moment. In this context, **motorsport may represent a possible environment for using the DoS index, since both vehicles and drivers can be well characterized**.

1.2 References

(Della Rossa et al. 2025) Della Rossa, F., Fontana, M., Giacintucci, S., Gobbi, M. et al., "Early Detection of the Loss of Control of Motorsport Vehicles," SAE Int. J. Veh. Dyn., Stab., and NVH 9(3):473-499, 2025.

(Mastinu and Ploechl 2014) Mastinu, G., Ploechl, M., (Eds). Road-and off-road Vehicle System Dynamics Handbook, CRC, Boca Raton, USA, 2014.

2 Global stability of vehicle-with-driver dynamics via Sum-of-Squares programming

2.1 Introduction

Vehicle stability is a cornerstone of road safety because any departure from stable motion can quickly escalate into accidents and injuries. Engineers have therefore devoted substantial effort to early detection of the onset of instability and to designing supervisory layers—such as ESC, active yaw-rate or torque-vectoring controllers, and coordinated brake- or steer-by-wire actuation—that prevent dangerous operating conditions. Even advanced supervisory strategies still depend on empirically tuned thresholds and controller parameters, often lacking actual stability certificates grounded on mathematical guarantees. The role of the driver further complicates this picture: its influence on the coupled dynamics is frequently ignored, even though driver actions can trigger instability (Mastinu et al. 2023). Developing rigorous model-based stability certificates that delineate safe operating regions across driving maneuvers, driver behaviors, and hardware configurations would help close this gap while directly informing controller design. These needs motivate the adoption of advanced estimation techniques for safe operative regions endowed with provable guarantees.

2.1.1 Literature Review

In the scientific literature there are several approaches to vehicle stability. Many application-oriented studies adopt empirical limits on vehicle state variables, their thresholds being chosen experimentally or based on engineering practice. It is quite common to impose tunable bounds on the vehicle sideslip angle to stabilize the vehicle, for example in the design of ESC systems (Chen and Kuo 2014) or torque vectoring controllers (Lenzo et al. 2021). Alternatively, some methods avoid relying on the vehicle sideslip angle, whose direct measurement is challenging and typically requires estimation. For instance, in (Daher et al. 2017) the authors proposed a vehicle stability controller that sets bounds on the vehicle longitudinal and lateral accelerations, i.e., on the g-g diagram.

On the other hand, a large body of work takes a more theoretical approach by relying on simplified mathematical models of the vehicle dynamics—typically single-track or double-track models that describe the lateral and yaw motion of the vehicle—and analyzing their stability as dynamical systems. It is standard textbook practice (Guiggiani 2023; Mastinu and Ploechl 2014) to assume some constant inputs (e.g., fixed longitudinal speed and steering angle) and to linearize the vehicle dynamics around the corresponding steady-state motion, i.e., straight-ahead or steady-cornering conditions. The eigenvalues analysis of the linearized system allows to assess the stability of the corresponding equilibrium point.

Because tire characteristics introduce nonlinearities in the vehicle dynamics, multiple equilibrium points may coexist. Early contributions in this area by Pacejka (Pacejka 2012) already showed that a two-state single-track model with nonlinear axle characteristics may exhibit three distinct steady-state points. The study of how the number and stability of these equilibria change as a function of vehicle parameters is known as bifurcation analysis (Strogatz 2019). Comprehensive analyses using continuation methods have identified various bifurcation phenomena—including saddle-node, Hopf, homoclinic, heteroclinic, and Bogdanov-Takens bifurcations—by varying driver inputs and tire characteristics (Della Rossa et al. 2012). Complementary approaches combined phase-plane and handling diagram methods to investigate how vehicle parameters such as roll stiffness and center of mass position can also induce bifurcation (Farroni et al. 2013), and to compare the stability performance of different steering control architectures (Lai et al. 2021).

For a stable and desirable steady-state motion, a crucial question concerns the size of its Region of Attraction (ROA), i.e., the set of initial conditions from which trajectories converge asymptotically to the

equilibrium. Estimating the ROA allows to quantify how large deviations from the steady-state can be tolerated before the vehicle loses stability. Early efforts to compute the ROA analytically applied Lyapunov's direct method with quadratic Lyapunov functions to single-track models with cubic polynomial axle characteristics (Johnson and Huston 1984). More recent study have conducted ROA estimation with more complex tire models. Linear parameter-varying (LPV) formulations with LuGre tire dynamics have been used to derive Lyapunov-based stability conditions (Hashemi et al. 2016). Alternative methods based on local linearization at multiple operating points have also been proposed (Huang et al. 2020). Other works have focused on quantitative stability indicators, such as Lyapunov exponents to measure convergence rates (Sadri and Wu 2013) or energy dissipation as a measure of stability (Meng et al. 2022).

For systems with polynomial dynamics, Lyapunov-based ROA estimation can be performed efficiently using Sum-of-Squares (SOS) programming techniques (Jarvis-Wloszek et al. 2005; Topcu et al. 2008). This approach exploits the fact that a sufficient condition for polynomial Lyapunov candidates to be positive-definite can be imposed by requiring them to be written as SOS polynomials. This SOS requirement can be formulated as a convex constraint in semidefinite programs (SDPs), enabling systematic computation of polynomial Lyapunov functions. Several applications to vehicle lateral dynamics have already used SOS techniques to estimate stability regions (Imani Masouleh and Limebeer 2018; Tamba and Nazaruddin 2018; Ribeiro et al. 2020; Ribeiro et al. 2022; Zhu et al. 2022).

2.1.2 Contributions and organization

In the present work we extend the application of SOS programming to estimate the safety sets for a vehicle-with-driver system. While SOS techniques have already been successfully applied to vehicle lateral dynamics (Imani Masouleh and Limebeer 2018; Tamba and Nazaruddin 2018; Ribeiro et al. 2020; Ribeiro et al. 2022; Zhu et al. 2022), our contribution focuses on two important aspects. First, we explicitly account for the driver's action, performing a *global stability analysis* of the coupled vehicle-with-driver system. Recent works (Mastinu et al. 2023; Mastinu et al. 2024) have shown that the driver can be a source of instability for the overall system. Although human driver modeling is a challenging task, simplified models available in the literature (Plöchl and Edelmann 2007) allow for preliminary estimates. While certainly not exact, these models enable qualitative reasoning and provide results that are also quantitatively consistent with experimental observations (Mastinu et al. 2024). Second, we explore how to constrain the Lyapunov function search to identify *safe* invariant subsets of the ROA. Unlike traditional SOS approaches that focus on asymptotic convergence to equilibrium, we require that trajectories remain within prescribed state boundaries throughout their entire time evolution. In our setting, constraining the trajectories to these safe invariant subsets guarantees that the polynomial model used in the SOS formulation remains valid and thus provides an accurate local approximation of the original nonlinear dynamics.

The remainder of the paper is organized as follows. Sec. 2 introduces the SOS framework on a two-state benchmark, detailing each step of the iterative Lyapunov search and illustrating how the resulting level sets approximate a prescribed safe region. Sec. 3 presents the seven-state vehicle-with-driver system, specifies the polynomial approximations required to apply the SOS framework, and highlights the state constraints needed to preserve fidelity to the original non-polynomial model. Sec. 4 reports two case studies—understeering and oversteering scenarios—examining the intersections of the estimated safe sets with relevant slices of the state space and comparing the SOS-certified regions with simulation-based target envelopes. Finally, Sec. 5 summarizes the main findings and outlines directions for future research.

2.2 ROA estimation via Lyapunov functions

2.2.1 General formulation

We first illustrate the ROA estimation procedure on a two-state polynomial system, i.e., the *time-reversed* Van der Pol oscillator, whose dynamics $\dot{\mathbf{x}} = f(\mathbf{x})$ is given by

$$\begin{cases} \dot{x} = -y \\ \dot{y} = x + \mu(x^2 - 1)y \end{cases} \quad (1)$$

The system has an asymptotically stable equilibrium at the origin and an unstable (saddle-type) limit cycle that marks the boundary of the ROA for the origin. The distortion of this region increases with greater values of μ : here we pick $\mu = 1$, since it already provides a challenging ROA shape to approximate. The phase portrait of system (1) is illustrated in Figure 1 along with the unstable limit cycle (black curve).

A Lyapunov function $V(\mathbf{x})$ is a differentiable energy-like function whose value decreases along the system trajectories. As is well known (Slotine and Li 1991), *if one can find* a function $V(\mathbf{x})$ satisfying (i) $V(\mathbf{x}) > 0$ and (ii) $\dot{V}(\mathbf{x}) = \nabla V(\mathbf{x})f(\mathbf{x}) < 0$, then the origin is asymptotically stable. This is the case for the *time-reversed* Van der Pol oscillator. Here, $V(\mathbf{x}) > 0$ means that V is positive definite—i.e., $V(\mathbf{0}) = 0$ and $V(\mathbf{x}) > 0$ in a neighborhood \mathcal{B}_0 —while $\dot{V}(\mathbf{x}) < 0$ indicates negative definiteness.

Moreover, *whenever* such a Lyapunov function exists, any sublevel set $\Omega = \{\mathbf{x}: V(\mathbf{x}) \leq \rho\}$ in which $\dot{V}(\mathbf{x}) < 0$ for all $\mathbf{x} \neq \mathbf{0}$, by the Local Invariant Set theorem (Slotine and Li 1991, 69), constitutes an inner estimate of the ROA, since it guarantees that all the trajectories starting within Ω converge to the origin—as it is the largest invariant set in Ω for system (1). The brown curves in Figure 2 depict increasingly accurate approximations of the ROA boundary for system (1). At convergence, the red curve delineates the boundary of the final ROA estimate—the entire pink region—and it lies essentially on top of the true limit cycle.

The problem of estimating the ROA can be cast as the search for the *largest sublevel set* Ω , i.e. with highest volume. In particular, one can seek a function $V(\mathbf{x})$, a level ρ , and a parameterized set S and pose the problem as follows:

$$\begin{aligned} & \underset{V(\mathbf{x}), \rho, S}{\text{maximize}} && \text{vol}(S) \\ & \text{s.t.} && V(\mathbf{x}) > 0 \\ & && \dot{V}(\mathbf{x}) < 0 \\ & && \Omega \subseteq \{\mathbf{x}: \dot{V}(\mathbf{x}) \leq 0\} \\ & && S \subseteq \Omega. \end{aligned} \quad (2)$$

Here, S denotes the shape set (typically a non-degenerate ellipsoid $S = \{\mathbf{x}: \mathbf{x}^T \mathbf{P} \mathbf{x} \leq 1\}$) whose volume can be computed analytically. Introducing S in this way provides an explicit surrogate objective $\text{vol}(S)$ for $\text{vol}(\Omega)$, allowing to systematically seek the largest certified inner approximation of the ROA by the chain of inclusions in (2).

The problem (2) is an optimization over a function space, and is intractable in its general form. To obtain a computationally feasible formulation, we rely on sum-of-squares (SOS) programming (Jarvis-Wloszek et al. 2005). The key idea is to restrict the search for $V(\mathbf{x})$ to SOS polynomials and reformulate (2) as SOS conditions.

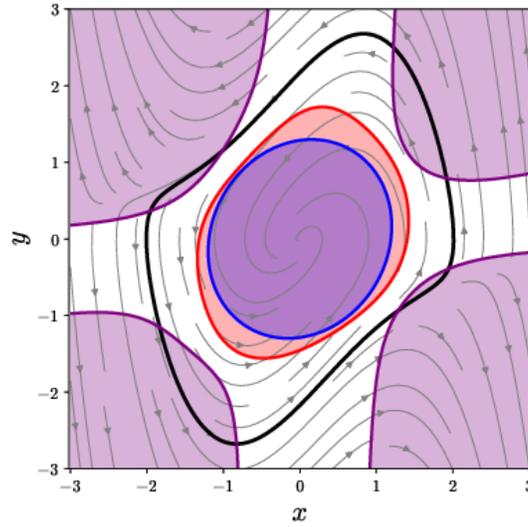


Figure 1. First iteration of the SOS-iterative procedure described in Sec. 2.3: $\Omega = \{\mathbf{x}: V(\mathbf{x}) \leq \rho\}$ (pink region, red boundary), elliptical shape function (blue region), $\{\mathbf{x}: \dot{V}(\mathbf{x}) \geq 0\}$ (purple region). Unstable limit cycle in solid black line.

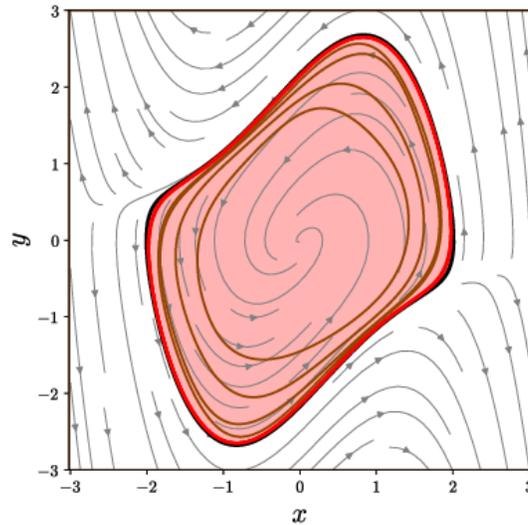


Figure 2. Final ROA estimate (pink region) with ROA boundary (red line) lying almost on top of the limit cycle (solid black line). Increasingly tight approximation of the ROA boundaries during iterations are represented as brown lines.

2.2.2 SOS formulation

A polynomial $p(\mathbf{x})$ is SOS if it can be written as $p(\mathbf{x}) = \sum_i q_i^2(\mathbf{x})$ for some polynomials $q_i(\mathbf{x})$. We denote by $\Sigma[\mathbf{x}]$ the space of SOS polynomials in \mathbf{x} . Note that $p(\mathbf{x}) \in \Sigma[\mathbf{x}]$ implies that $p(\mathbf{x})$ is a positive polynomial

(a.k.a. nonnegative), i.e., $p(\mathbf{x}) \geq 0$ for all \mathbf{x}^1 . Checking (or enforcing) that a polynomial $p(\mathbf{x})$ is SOS, i.e. $p(\mathbf{x}) \in \Sigma[\mathbf{x}]$, amounts to solving a semidefinite program (SDP) whose decision variables are the coefficients of $p(\mathbf{x})$, collected through its *Gram-matrix* representation.

Specifically, any polynomial $p(\mathbf{x})$ can be decomposed as $p(\mathbf{x}) = \mathbf{z}(\mathbf{x})^T \mathbf{Q} \mathbf{z}(\mathbf{x})$, where $\mathbf{z}(\mathbf{x})$ is the vector of monomials and $\mathbf{Q} = \mathbf{Q}^T$ is the Gram matrix. As an example, a fourth-degree polynomial in two variables x and y can be written as

$$p(x, y) = \sum_{i=1}^6 \sum_{j=1}^6 Q_{ij} z_i(x, y) z_j(x, y),$$

where $\mathbf{z}(x, y) = [1 \ x \ y \ x^2 \ xy \ y^2]^T$, $z_i(x, y)$ is the i -th element of $\mathbf{z}(x, y)$, and Q_{ij} the (i, j) -th element of $\mathbf{Q} \in \mathbb{R}^{6 \times 6}$.

The condition $p(\mathbf{x}) \in \Sigma[\mathbf{x}]$ is then equivalent to the linear matrix inequality $\mathbf{Q} \succcurlyeq 0$, which is a convex conic constraint for the SDP. Moreover, to certify that a polynomial $p(\mathbf{x})$ is *strictly positive definite* we can subtract a small positive term $\varepsilon \|\mathbf{x}\|^2$, with $\varepsilon > 0$, and require that the auxiliary polynomial $a(\mathbf{x}) = p(\mathbf{x}) - \varepsilon \|\mathbf{x}\|^2$ belongs to $\Sigma[\mathbf{x}]$. Denoting by $\Sigma^+[\mathbf{x}]$ the subset of strictly-positive SOS polynomials, the existence of such an ε then suffices to conclude that $p(\mathbf{x}) \in \Sigma^+[\mathbf{x}]$.

The problem (2) can be recast as

$$\begin{aligned} & \underset{V(\mathbf{x}), \rho, \mathbf{P}, \lambda(\mathbf{x}), \mu(\mathbf{x})}{\text{minimize}} && \text{tr}(\mathbf{P}) \\ & \text{s.t.} && V(\mathbf{x}) \in \Sigma^+[\mathbf{x}] \\ & && V(\mathbf{0}) = 0 \\ & && \|\mathbf{x}\|^{2d} (V(\mathbf{x}) - \rho) + \lambda(\mathbf{x}) \dot{V}(\mathbf{x}) \in \Sigma[\mathbf{x}] \\ & && \|\mathbf{x}\|^{2d_1} (\mathbf{x}^T \mathbf{P} \mathbf{x} - 1) + \|\mathbf{x}\|^{2d_2} \mu(\mathbf{x}) (V(\mathbf{x}) - \rho) \in \Sigma[\mathbf{x}] \end{aligned} \quad (3)$$

First, we notice that (2) involves an ellipsoidal shape function. The volume of an ellipsoid in the form is $\mathbf{x}^T \mathbf{P} \mathbf{x} = 1$ is proportional to $\frac{1}{\sqrt{\det(\mathbf{P})}}$. To inflate the ellipsoidal shape function, as cost function we minimize

$\text{tr}(\mathbf{P})$, which serves as a linear proxy of $\sqrt{\det(\mathbf{P})}$ and helps formulating the problem as convex. Second, the positive definite constraint on $V(\mathbf{x})$ is now given by $V(\mathbf{x}) \in \Sigma^+[\mathbf{x}]$ and $V(\mathbf{0}) = 0$. Finally, the set inclusion constraints are enforced via the S-procedure (Jarvis-Wloszek et al. 2005), which provides a sufficient condition through polynomial multipliers. Consider, for instance, the first constraint: if there exists a multiplier $\lambda(\mathbf{x})$ such that condition holds, then wherever $\dot{V}(\mathbf{x}) = 0$, we must have $V(\mathbf{x}) \geq \rho$. This implies that the boundary of the region $\{\mathbf{x}: \dot{V}(\mathbf{x}) \leq 0\}$ (in purple in Figure 1) lies outside $\Omega = \{\mathbf{x}: V(\mathbf{x}) \leq \rho\}$ (red region), thus proving the first inclusion. The factor $\|\mathbf{x}\|^{2d}$ with $d \in \mathbb{N}^+$ ensures the validity of the S-procedure also at the origin, since $\dot{V}(\mathbf{0})$ is null but clearly the origin does not lie outside the ROA.

Similarly, the second constraint is enforced via multiplier $\mu(\mathbf{x})$. The scaling factors $\|\mathbf{x}\|^{2d_1}$ and $\|\mathbf{x}\|^{2d_2}$ are not strictly required to guarantee the inclusion, but are introduced to balance the degree of the two addends. The coefficients $d_1, d_2 \in \mathbb{N}^+$ are chosen heuristically, typically between 0 and 4: without this scaling the terms $(\mathbf{x}^T \mathbf{P} \mathbf{x} - 1)$ would remain quadratic while $\mu(\mathbf{x})(V(\mathbf{x}) - \rho)$ inherits the higher degree of V , often leading to ill-conditioned constraints. In our experience, the proposed factors improve the feasibility of the SDP.

¹ The converse does not hold, as exemplified by the Motzkin polynomial (Motzkin 1967), which is globally nonnegative yet admits no representation as a sum of squares.

Unfortunately, the products $\lambda(\mathbf{x})\dot{V}(\mathbf{x})$ and $\mu(\mathbf{x})(V(\mathbf{x}) - \rho)$ in (2) make the problem bilinear in the coefficients of $V(\mathbf{x})$, $\lambda(\mathbf{x})$, and $\mu(\mathbf{x})$, rendering the optimization problem non-convex. To address this issue, we propose an alternating optimization procedure, described below.

2.2.3 Iterative SOS formulation

There are several examples in the scientific literature of multi-step formulations for the described problem (Jarvis-Wloszek et al. 2005; Topcu et al. 2008). Here we adopt a *three-step iterative procedure*. The initial guess of the procedure is obtained from a feasibility SDP, in which we search for a global polynomial Lyapunov for the linearized system around the origin. The constraints of this initialization problem are $V_0(\mathbf{x}) \in \Sigma^+[\mathbf{x}]$ and $-\dot{V}_{0,\text{lin}}(\mathbf{x}) = -\nabla V_0(\mathbf{x}) \frac{\partial f(\mathbf{x})}{\partial \mathbf{x}}|_{\mathbf{x}=\mathbf{0}} \in \Sigma[\mathbf{x}]$. It is also useful to fix the scaling of the Lyapunov polynomial by constraining its value at one or more points of the state space, e.g., $V_0(1,2) = 1$.

Algorithm 1 Iterative SOS procedure

```

1: Fix Lyapunov degree  $n_V$ 
2: Solve the initialization SDP to obtain  $V_0$ 
3: for  $k = 0$  to  $k_{\text{max}}$  do
    Step 1:
4:   repeat
5:     try
6:        $(\lambda_{k+1}, \rho_{k+1}) \leftarrow \lambda\text{-step}(V_k, n_\lambda)$             $\triangleright (\lambda, \rho)$  update, Solve Problem (5)
7:       if not successful then
8:          $n_\lambda \leftarrow n_\lambda + 1$                                 $\triangleright$  increase degree and retry
9:       end if
10:    until  $\lambda\text{-step}$  succeeds or  $n_\lambda > n_\lambda^{\text{max}}$ 
11:     $V_k \leftarrow V_k / \rho_{k+1}$                                     $\triangleright$  level set normalization
    End Step 1

    Step 2:
12:   repeat
13:     try
14:        $\mu_{k+1} \leftarrow \mu\text{-step}(V_k, n_\mu)$                     $\triangleright (\mu, P)$  update, Solve Problem (6)
15:     ...
16:   until  $\mu\text{-step}$  succeeds or  $n_\mu > n_\mu^{\text{max}}$ 
    End Step 2

    Step 3:
17:    $(V_{k+1}, P_{k+1}) \leftarrow V\text{-step}(\lambda_{k+1}, \mu_{k+1})$         $\triangleright (V, P)$  update, Solve Problem (4)
18:   break criterion on  $\text{tr } P_{k+1}$  convergence                  $\triangleright$  with fixed  $\lambda_{k+1}, \mu_{k+1}$ 
19:   return  $V_{k+1}$ 
    End Step 3
20: end for
    
```

Algorithm 1: Iterative SOS procedure

A pseudocode for three-step procedure is given in Algorithm 1. In the λ -step (**Step 1**) the Lyapunov candidate $V(\mathbf{x})$ is kept fixed. We search for the largest level set by maximizing the level value ρ while solving for the multiplier $\lambda(\mathbf{x})$:

$$\begin{aligned}
 & \underset{\rho, \lambda(\mathbf{x})}{\text{maximize}} && \rho \\
 & \text{s.t.} && \|\mathbf{x}\|^{2d}(V(\mathbf{x}) - \rho) + \lambda(\mathbf{x})\dot{V}(\mathbf{x}) \in \Sigma[\mathbf{x}]
 \end{aligned} \tag{4}$$

The multiplier degree n_λ is chosen a priori; starting from $n_\lambda = 0$, if the resulting SDP is infeasible we increase n_λ and retry, since higher degrees often improve feasibility in practice. Once ρ is optimized, the Lyapunov function is normalized so that the newly identified level set Ω corresponds to the unit value.

In the subsequent μ -step (**Step 2**), we keep the normalized $V(\mathbf{x})$ fixed and find the multiplier $\mu(\mathbf{x})$ in so as to enlarge the ellipsoidal set $S = \{\mathbf{x}: \mathbf{x}^T \mathbf{P} \mathbf{x} \leq 1\}$ within the previously obtained Ω :

$$\begin{aligned} & \underset{\mathbf{P}, \mu(\mathbf{x})}{\text{minimize}} && \text{tr}(\mathbf{P}) \\ & \text{s.t.} && \|\mathbf{x}\|^{2d_1} (\mathbf{x}^T \mathbf{P} \mathbf{x} - 1) + \|\mathbf{x}\|^{2d_2} \mu(\mathbf{x}) (V(\mathbf{x}) - 1) \in \Sigma[\mathbf{x}] \end{aligned} \quad (5)$$

The multiplier degree n_μ is chosen a priori, and the same try-and-update logic as in the λ -step is implemented.

Finally, in the V -step (**Step 3**) we solve a Lyapunov-update problem identical to (3) but with the multipliers $\lambda(\mathbf{x})$ and $\mu(\mathbf{x})$ fixed to the values computed in the first two steps (thereby eliminating the bilinear terms) to update $V(\mathbf{x})$ for the next iteration. The described procedure, applied to the benchmark system (1), allows to approximate the origin's ROA very accurately using a sixth-degree Lyapunov function. As shown in Figure 2, the boundary of the estimated ROA (pink region with red boundary) closely tracks the system's limit cycle (black line), i.e., the boundary of the exact ROA.

2.2.4 2.4 Safe subset with state constraints

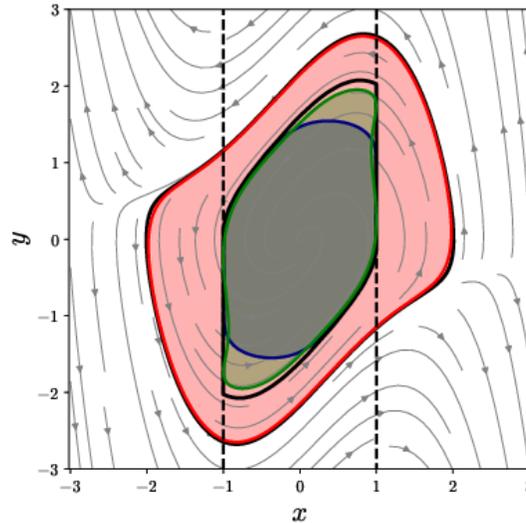


Figure 3. Classical ROA estimate with no constraints (pink region, red boundary). State constraints $x^2 \leq 1$ are introduced (dashed black lines). The boundary of the true safe ROA is illustrated in solid black lines. The safe ROA estimate, computed enforcing state constraints, is given by the blue region (blue boundary). The safe ROA estimate given by the green region (green boundary) is obtained introducing also anchor points as described in the hybrid approach in Sec 2.4.

In many control applications the operating envelope is constrained not only by stability criteria, but also by safety or comfort requirements that bound the admissible trajectories of the system. We therefore target, for the Van der Pol oscillator, an invariant subset that satisfies a prescribed collection of state inequalities, defined as safe ROA. The SOS framework accommodates for state constraints with minor adjustments to the alternating scheme described above.

To illustrate the idea, we enforce the constraint $x^2 \leq 1$ to our benchmark system (1), i.e., we require that trajectories originating inside the estimated safe set must remain within the strip delimited by the dashed black lines in Figure 3. Adopting again the S-procedure, the state constraint is given by

$$(V(\mathbf{x}) - \rho) + \eta(\mathbf{x})(x^2 - 1) \in \Sigma[x],$$

and is added to the λ - and V -steps (Step 1 and Step 3 respectively) of the iterative procedure. During the λ -step (Step 1), we introduce and optimize the new polynomial multiplier $\eta(\mathbf{x})$ tied to the set inclusion $\Omega \subseteq \{\mathbf{x}: x^2 \leq 1\}$. In the subsequent V -step (Step 3), $\eta(\mathbf{x})$ is frozen to the value found in the λ -step and the same SOS constraint is enforced while updating $V(\mathbf{x})$.

The blue region in Figure 3 is the estimated safe ROA resulting from the constrained SOS computation: it is an invariant subset of the unconstrained ROA (red region) that also fulfils the additional state bounds $x^2 \leq 1$.

For comparison, the solid black line tangent to the estimated safe ROA represent the exact safe ROA boundary, obtained via exhaustive simulations starting from sampled initial conditions all over the state space. Although more conservative with respect to the ROA estimation, the estimated safe ROA still tracks the target region closely and preserves its main geometric features.

Furthermore, we explored a hybrid method to tighten the estimation of the safe ROA. While an exhaustive simulation sweep over the state space is often impractical, a small set of sampled boundary points of the safe ROA can efficiently guide the SOS solution to higher levels of accuracy. The hybrid method requires a preliminary step, during which a few points on the searched safe set are identified. For example, by sampling initial conditions on the set $y = 2x$ and computing the resulting trajectories, we identify two target points $\mathbf{x}_1 = (1, 2)$ and $\mathbf{x}_2 = (-1, -2)$, laying approximately near the top right and bottom left vertices of the target safe ROA.

As an example, we use \mathbf{x}_1 and \mathbf{x}_2 as *anchor points* for the estimated safe ROA by reformulating the V -step as

$$\begin{aligned} & \underset{V(\mathbf{x}), \mathbf{P}, \gamma}{\text{minimize}} && w_1 \text{tr} \mathbf{P} + w_2 \gamma \\ & \text{s.t.} && V(\mathbf{x}) \in \Sigma^+[\mathbf{x}] \\ & && V(\mathbf{0}) = 0 \\ & && V(\mathbf{x}_1) = \gamma \\ & && V(\mathbf{x}_2) = \gamma \\ & && \|\mathbf{x}\|^{2d} (V(\mathbf{x}) - 1) + \lambda(\mathbf{x}) \dot{V}(\mathbf{x}) \in \Sigma[\mathbf{x}] \\ & && \|\mathbf{x}\|^{2d_1} (\mathbf{x}^T \mathbf{P} \mathbf{x} - 1) + \|\mathbf{x}\|^{2d_2} \mu(\mathbf{x}) (V(\mathbf{x}) - 1) \in \Sigma[\mathbf{x}] \\ & && (V(\mathbf{x}) - \rho) + \eta(\mathbf{x})(x^2 - 1) \in \Sigma[x] \end{aligned} \tag{6}$$

where γ is a slack variable representing the symbolic value of the unknown Lyapunov function at the anchor points, and w_1 and w_2 are weights for the cost function. In our example these are set to $w_1 = 0.9$ and $w_2 = 0.1$, respectively. Intuitively, the boundary of the Lyapunov level set $V(\mathbf{x}) = 1$ is “pulled” toward the sampled anchor points lying on the exact target region. As γ decreases toward 1, the level set $V(\mathbf{x}) = 1$ progressively tightens around these anchor points. The state constraints in (6) prevents the boundary of the safe set from extending beyond them. The outcome of the hybrid method is illustrated by the green region in Figure 3, which provides a more accurate approximation than the basic version shown in blue.

2.3 SOS-friendly vehicle-with-driver model

2.3.1 Vehicle-with-driver model

The SOS procedure is applied to the vehicle-with-driver model introduced in (Mastinu et al. 2023), which is briefly summarized here for completeness. Despite its simplicity, this model already offers quantitative insight into the stability of the coupled driver–vehicle dynamics, as validated in (Mastinu et al. 2024).

The system relies on a single-track model with fixed longitudinal speed u and uses a delayed input of the front steering angle δ that drives the vehicle to follow a reference straight path (see Figure 4), with a time delay of $\tau = 0.2$ s. The delayed steering logic mimics a human driver by correcting the lateral error e , with respect to the reference path, of a lookahead point P . The lookahead distance L_P of P from the vehicle centre of mass G is computed as $L_P = \frac{u}{T_P}$, with the preview time T_P fixed at 0.5 s.

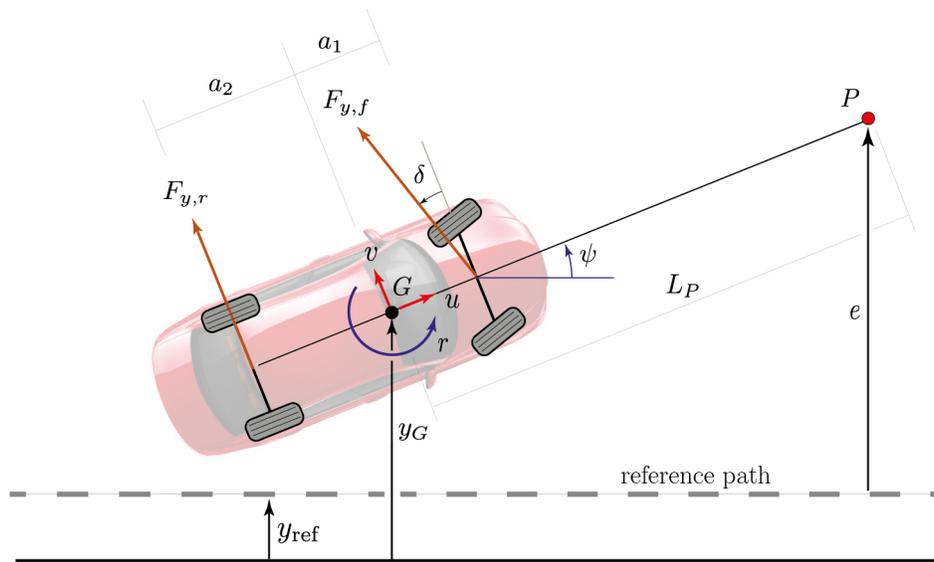


Figure 4. Vehicle-with-driver model. The delayed steering logic mimics a human driver by correcting the lateral error e with respect to the reference path (grey dashed line) of the lookahead point P (red).

The dynamics of the system is described by the following set of equations:

$$\begin{aligned}
 m \dot{v} &= -m u r + F_{y,f} \cos \delta + F_{y,r} \\
 J \dot{r} &= a_1 F_{y,f} \cos \delta - a_2 F_{y,r} \\
 \dot{y}_G &= u \sin \psi + v \cos \psi \\
 \dot{\psi} &= r \\
 \dot{\delta} &= \delta_1 \\
 \dot{\delta}_1 &= \delta_2 \\
 \dot{\delta}_2 &= \frac{6}{\tau^3} \left(-\delta - \tau \dot{\delta}_1 - \frac{\tau^2}{2} \dot{\delta}_2 - k e - k_d \dot{e} \right)
 \end{aligned}$$

The descriptions of the variables and parameters appearing above are summarized in Table [1](#).

Symbol	Description	Symbol	Description
m	vehicle mass	δ	front axle steering angle
J	yaw moment of inertia	δ_1	front axle steering rate
a_1	distance of G from the front axle	δ_2	front axle steering angular acceleration
a_2	distance of G from the rear axle	$F_{y,f}$	lateral force acting on the front axle
u	longitudinal speed	$F_{y,r}$	lateral force acting on the rear axle
v	lateral speed	e	lateral error of preview point P
r	yaw rate	τ	steering delay time
y_G	lateral displacement of G	k	proportional steering gain
ψ	yaw angle	k_d	derivative steering gain

Table 1: Variables and parameters of the vehicle-with-driver model considered.

In the vehicle-with-driver equations a delayed PD steering law is implemented on the lateral error $e = y_G - y_{\text{ref}} + L_p \sin\psi$, with proportional gain k and derivative gain k_d :

$$\delta(t + \tau) = -k e - k_d \dot{e}.$$

The left hand term of steering angle jerk is approximated through the third-order Taylor expansion

$$\delta(t + \tau) \approx \delta(t) + \dot{\delta}(t) \tau + \frac{\ddot{\delta}(t)}{2} \tau^2 + \frac{\dddot{\delta}(t)}{6} \tau^3.$$

Deeper insights on the accuracy of the adopted approximation can be found in (Mastinu et al. 2023).

The described vehicle-with-driver system has a fixed point at the origin, i.e., for $\mathbf{x} = [v \ r \ y_G \ \psi \ \delta \ \delta_1 \ \delta_2]^T = \mathbf{0}$. The stability of this equilibrium depends on the system's parameters; for realistic settings one can usually identify a critical longitudinal speed below which the origin is asymptotically stable (Mastinu et al. 2023). In what follows, we focus on such operating conditions and estimate a safe invariant subset of the ROA of the stable equilibrium at the origin.

2.3.2 Polynomial approximation

The SOS procedure for ROA estimation introduced in Sec. 2 presumes polynomial system dynamics, a condition that the above vehicle-with-driver model clearly violates. Two sources break the polynomial structure: the trigonometric terms in ψ and δ (noting that $e = y_G - y_{\text{ref}} + L_p \sin\psi$ carries one as well), and the tire constitutive law, which we model through a standard Magic Formula (MF)

$$\Phi_i(\alpha_i) = D_i \sin\{C_i \arctan[B_i \alpha_i - E_i (B_i \alpha_i - \arctan(B_i \alpha_i))]\}, \quad i = \{f, r\}$$

mapping the axle slip angles α_f and α_r to the lateral front and rear forces $F_{y,f}$ and $F_{y,r}$, respectively. The axle slip angles depend linearly on the state variables (u is constant),

$$\alpha_f = \delta - \frac{v + a_1 r}{u}, \quad \alpha_r = -\frac{v - a_2 r}{u},$$

yet the Magic-Formula relationships $F_{y,f} = \Phi_f(\alpha_f)$ and $F_{y,r} = \Phi_r(\alpha_r)$ remain non-polynomial, which prevents a direct SOS treatment. Including longitudinal dynamics would introduce an additional non-polynomial term in the congruence equations, since the denominator u would become a state rather than a parameter.

A polynomial approximation of that model is thus required, noting that higher polynomial degrees entail increased computational cost in the SOS procedure. We proceed in two steps. First, trigonometric terms in δ and ψ are replaced by Taylor expansions. Given that $\delta \in [-5^\circ, 5^\circ]$, we set $\cos\delta \approx 1$. Larger yaw

excursions motivate third-order approximations, $\sin\psi \approx \psi - \psi^3/6$ and $\cos\psi \approx 1 - \psi^2/2$, which remain accurate for $\psi \in [-50^\circ, 50^\circ]$. Second, the axle characteristics are approximated using cubic fits of $\Phi_f(\alpha_f)$ and $\Phi_r(\alpha_r)$.

Because polynomials cannot reproduce the saturation inherent in the Magic Formula, the fitting domain is limited to 95% of the peak lateral force, yielding slip-angle thresholds $\bar{\alpha}_f$ and $\bar{\alpha}_r$. Cubic fits are then computed over $\alpha_f \in [-\bar{\alpha}_f, \bar{\alpha}_f]$ and $\alpha_r \in [-\bar{\alpha}_r, \bar{\alpha}_r]$, as illustrated in Figure 5.

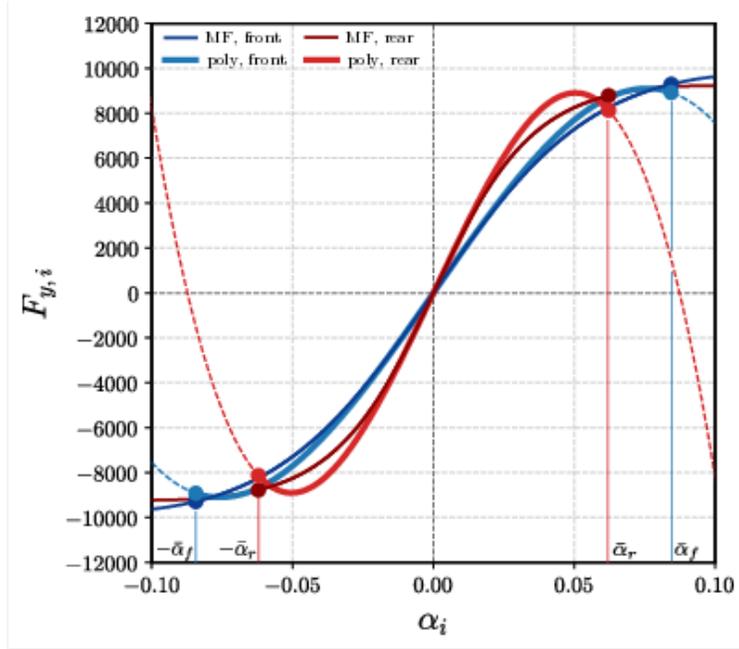


Figure 5. Polynomial fitting of axle characteristics in the understeering setup. Dark blue/red thin curves show the Magic Formula (MF) for the front/rear axles (α_i in rad, $F_{y,i}$ in N), while the light counterparts depict the cubic fits (poly), solid within the $[-\bar{\alpha}_i, \bar{\alpha}_i]$ range used for fitting and dashed outside it. Vertical blue/red lines mark the axle slip angles limits corresponding to 95% of the respective Magic Formula peak forces.

The resulting polynomials provide an accurate representation only within these bands. Consequently, when computing constrained invariant subsets of the ROA we enforce as validity constraints

$$|\alpha_f(\mathbf{x}(t))| \leq \bar{\alpha}_f, \quad |\alpha_r(\mathbf{x}(t))| \leq \bar{\alpha}_r, \quad \forall t \geq 0.$$

This ensures that the certified ROA remains entirely within the region where the polynomial approximation of the axle characteristics is reliable.

No additional bounds are imposed on the steering angle δ or the yaw angle ψ , as the invariant subsets of the ROA remained within their validity limits in all examined cases. If needed under different operating conditions, constraints analogous to those on the axle slip angles could be introduced. With the above substitutions, the polynomial approximation of the vehicle-with-driver model is of third degree.

2.4 Vehicle-with-driver safe invariant set

2.4.1 Tested scenarios and SOS procedure set-up

Results of the proposed procedure on the vehicle-with-driver model described in Sec. 3 are presented for two different scenarios: an oversteering (OV) and an understeering (UN) vehicle. In both cases the

longitudinal speed is fixed at $u = 90 \text{ km/h}$, while the axle characteristics $\Phi_f(\alpha_f)$ and $\Phi_r(\alpha_r)$ are described by the corresponding Magic Formula parameters listed in Table 2. The controller gains k and k_d also differ between scenarios, with values listed in Table 2. These gains were obtained by fitting experimental data (Mastinu et al. 2023) and depend on both driver behavior and vehicle parameters.

Scenario	B_f	C_f	D_f	E_f	B_r	C_r	D_r	E_r	k	k_d
OV	14.50	1.89	9778	0.29	13.50	1.45	9234	0.31	0.025	0.004
UN	9.86	1.87	9778	0.28	18.75	1.53	9234	0.30	0.010	0.008

Table 2: Parameters of tested scenarios

For each scenario, a safe invariant subset of the ROA (SOS-S-ROA) of the stable equilibrium at the origin is computed using the three-step SOS procedure described in Sec. 2.3. In all cases, this safe set is represented as the region enclosed by the unit level set $V(\mathbf{x}) = 1$ of a degree-4 polynomial Lyapunov function—this is the final outcome of the procedure. For a system with $n = 7$ states, higher Lyapunov degrees would lead to a prohibitive number of coefficients, since a polynomial in n variables of degree r has $\frac{(n+r)!}{n!r!}$ coefficients. The symmetry of the vehicle-with-driver model is further exploited by restricting $V(\mathbf{x})$ and the multipliers $\lambda(\mathbf{x})$ and $\mu(\mathbf{x})$ to even polynomials, thereby reducing the number of decision variables in each of the three steps in Sec. 2.3. The exponents d , d_1 , and d_2 introduced in (6) are set to 2, 1, and 0, respectively.

The state constraints on the axle slip angles are enforced via the S-procedure, i.e.,

$$(V(\mathbf{x}) - \rho) + \eta_i(\mathbf{x})(\alpha_i(\mathbf{x}) - \bar{\alpha}_i) \in \Sigma[\mathbf{x}], \quad i \in \{f, r\}.$$

where $\eta_i(\mathbf{x})$ are polynomial multipliers of degree $\deg(V) - 1$. The symmetry of $V(\mathbf{x})$ guarantees that also the lower boundary $\alpha_i(\mathbf{x}) \geq -\bar{\alpha}_i$ is fulfilled. Such constraints are enforced in both the λ -step (**Step 1**) and the V -step (**Step 3**). In **Step 1**, the multipliers $\eta_i(\mathbf{x})$ are optimized jointly with $\lambda(\mathbf{x})$. In **Step 3**, the same SOS constraints are imposed while updating $V(\mathbf{x})$, but the multipliers $\eta_i(\mathbf{x})$ are kept fixed at the values obtained in **Step 1**, mirroring the treatment of $\lambda(\mathbf{x})$.

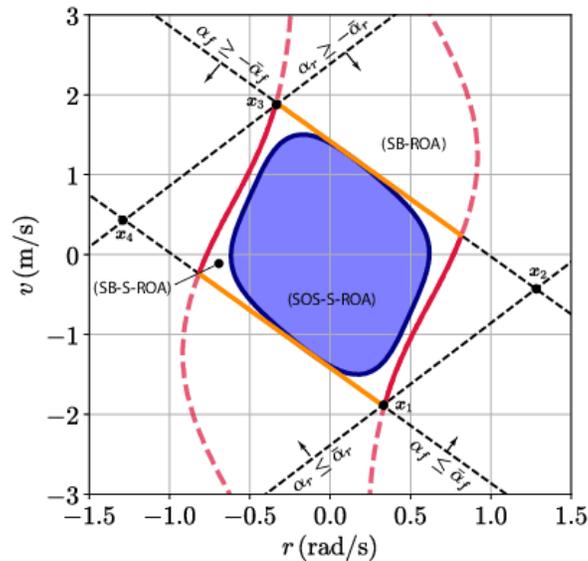


Figure 6. Oversteering vehicle (OV). Intersections with v - r plane

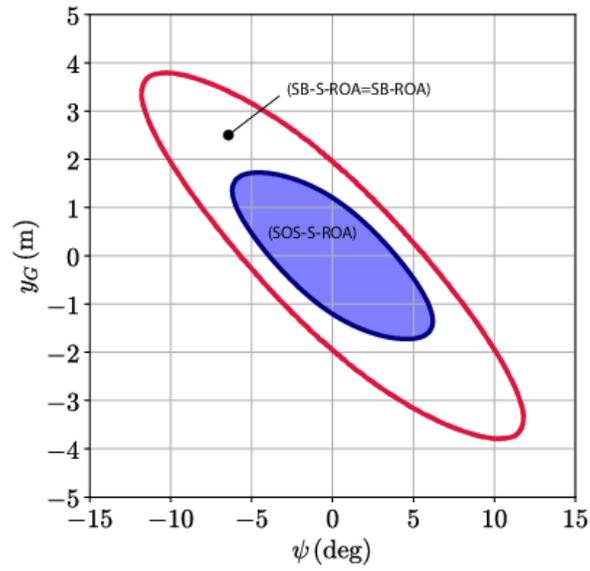


Figure 7. Oversteering vehicle (OV). Intersections with y_G - ψ plane

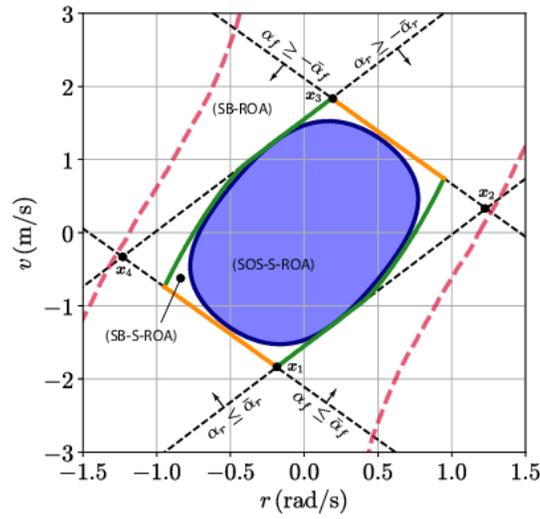
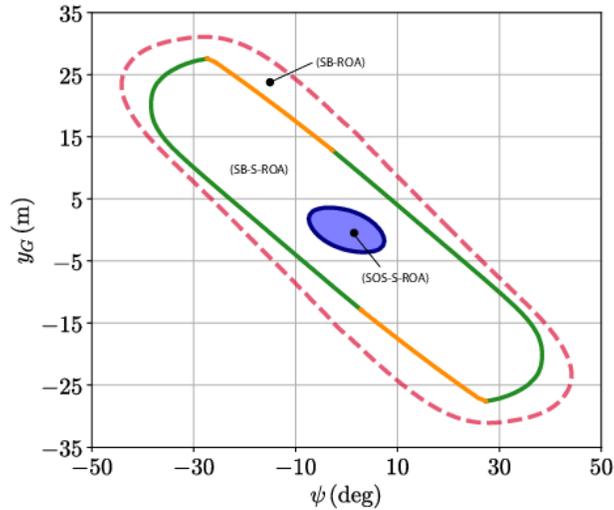


Figure 8. Understeering vehicle (UN). Intersections with v - r plane


 Figure 9. Understeering vehicle (UN). Intersections with y_G - ψ plane

We show the SOS-S-ROAs obtained for both OV and UN vehicles through their *intersections* with two specific planes of the state space: the v - r plane (Figures 6 and 8) and the y_G - ψ plane (Figures 7 and 9). On each panel, the blue region represents the intersection of the estimated SOS-S-ROA $\Omega = \{\mathbf{x}: V(\mathbf{x}) \leq 1\}$ with the corresponding plane.

On the v - r plane, the constraints on axle slip angles are shown to form a parallelogram (dashed black lines), where the bottom left side corresponds to $\alpha_f \leq \bar{\alpha}_f$, while the bottom right side corresponds to $\alpha_r \leq \bar{\alpha}_r$. The upper sides correspond to $\alpha_f \geq -\bar{\alpha}_f$ and $\alpha_r \geq -\bar{\alpha}_r$, respectively, and are parallel to their counterparts. In both OV and UN cases we use the vertices of the parallelogram \mathbf{x}_j (Figures 6 and Figures 8) in the cost functions of the V -step as anchor points for the safe set boundary, following the hybrid method described in (6) Sec. 2.2.4. Although these vertices do not necessarily lie on exact safe ROA, they are profitably used as readily available by pairwise intersections of the four slip-angle constraints. It is worth noting their role as anchor points: these vertices act as *pegs* to which the Lyapunov level set is attached and *stretched*.

2.4.2 Results

To interpret the SOS-S-ROA intersections, consider an initial condition $\mathbf{x}_0 = (v_0, r_0, 0, 0, 0, 0, 0)$ on the v - r plane, with (v_0, r_0) lying in the blue region of Figure 6. Such conditions may arise when the vehicle, initially at $\mathbf{x} = \mathbf{0}$, experiences lateral and yaw disturbances that primarily excite the v - r dynamics. Membership of \mathbf{x}_0 in the SOS-S-ROA can be immediately verified by checking $V(\mathbf{x}_0) \leq 1$, where $V(\mathbf{x})$ is the Lyapunov function obtained from the SOS procedure in Algorithm 1. This *certifies* i) convergence of the trajectory $\mathbf{x}(t)$ originated in \mathbf{x}_0 to the origin and ii) satisfaction of the slip-angle constraints for all $t \geq 0$, ensuring the validity of the polynomial approximation. It is worth remarking that, although the initial condition lies in the v - r plane, the trajectory $\mathbf{x}(t)$ evolves in the full seven-dimensional state space; nevertheless, its evolution remains confined within the certified SOS-S-ROA.

A similar interpretation applies to the y_G - ψ plane (Figure 7), where $\mathbf{x}_0 = (0, 0, y_{G0}, \psi_0, 0, 0, 0)$ represents a lateral displacement y_{G0} and an initial heading ψ_0 relative to the reference path. Such conditions may arise, for instance, from an abrupt lane-change manoeuvre in which the vehicle is laterally offset and exhibits a nonzero yaw angle at the onset.

For comparison, for the full nonlinear vehicle dynamics (with no polynomial approximations), the ROA was computed via a simulation-based approach (SB-ROA). Initial conditions were sampled on each plane

and the system was integrated over a 10 s horizon. A sampled point was classified as belonging to the simulation-based safe ROA, SB-S-ROA, if the corresponding trajectory converged to the equilibrium and satisfied all state constraints—specifically, the front and rear axle slip-angle bounds—for the entire time interval.

The resulting simulation-based boundaries of both the SB-ROA and the SB-S-ROAs are reported in Figures from 6 to 9. They also indicate which constraint is violated immediately outside the SB-S-ROA: loss of asymptotic stability (solid red), violation of the front slip-angle limit (yellow), or violation of the rear slip-angle limit (green). Along the solid red curves (see Figure 6), the boundaries of the SOS-S-ROA and of the SB-S-ROA coincide; the portions of the simulation-based ROA where slip-angle constraints are not enforced (simply SB-ROA) are shown as dashed red lines. The same observations apply to the understeering cases in Figures 8 and 9.

2.4.3 Discussion

Assuming the SB-S-ROAs as the exact estimate, the SOS-S-ROAs prove to be conservative in both scenarios. In the OV case, on the v - r plane (Figure 6) the SOS-S-ROs share boundaries with the SB-S-ROA. It is worth noting that stability boundaries (red) are more restrictive than the rear axle slip limits (dotted black lines). On the y_G - ψ plane (Figure 7), the intersection of the SB-S-ROA coincides exactly with the SB-ROA (or more precisely, with the simply connected portion of the SB-ROA containing the origin²). Here the SOS-S-ROA is visibly less tight than the exact estimate SB-S-ROA. The actual safe sets live in \mathbb{R}^7 while we display only planar sections, so the *sharpness* of the SB-S-ROA boundary may not be fully captured in these intersections. In addition, a fourth-degree Lyapunov polynomial lacks the flexibility to reproduce such sharp features. Therefore, in our view, this loss of accuracy stems from the polynomial-degree constraint required to keep the SOS problem tractable in the present formulation.

For the UN case, analogous considerations apply. On the v - r plane (Figure 8) the influence of the axle slip angle bounds—which are needed to preserve the validity of the polynomial approximation—is even more evident. Large portions of the SB-S-ROA boundary coincide with the constraint lines for the front axle; however, the rear axle slip angle threshold smooths the SB-S-ROA for higher values of $|r|$. For the y_G - ψ plane (Figure 9) we first point out the different axis scales with respect to Figure 7: the SOS-S-ROA has an extension comparable to that of the OV case, whereas the SB-S-ROA and the SB-ROA stretch considerably farther on this plane. As in the OV discussion, this discrepancy is mainly attributable to the limited flexibility of the fourth-degree Lyapunov polynomial and to the sharp features of the true boundary outside the displayed slice. Despite these limitations, the certified region still spans a practically relevant operating envelope for the understeering scenario.

The SOS procedure was implemented in Python and executed through the DRAKE (Drake 2019) optimization suite, employing MOSEK (Mosek 2025) as the underlying SDP solver. It was executed on a laptop equipped with an Intel i7 processor and 26 GB of RAM. The OV scenario required 301 s and six iterations, while the UN scenario required 287 s, also over six iterations. Each step of Algorithm (1) corresponds to a distinct problem with its own set of decision variables: the initialization step has 1296 variables, the λ -step (Step 1) has 55 894 variables, the μ -step (Step 2) has 723 variables, and the V -step (Step 3) has 57 007 variables.

The resulting Lyapunov polynomial $V(\mathbf{x})$ provides a global certification of the safe ROA in the full

² A more detailed analysis (omitted here) shows that, on this plane, the section of the SB-ROA forms a non-simply-connected set with islands. Only its simply-connected component containing the origin coincides with the SB-S-ROA.

seven-dimensional state space. Once computed, the safety assessment of the vehicle-with-driver system reduces to a single inequality check, namely evaluating whether $V(\mathbf{x}) \leq 1$ at the current state.

If the human driver is well represented by the controller structure adopted in the model, the availability of such a certified safe ROA (SOS-S-ROA) enables a fast and reliable stability assessment: evaluating $V(\mathbf{x})$ immediately determines whether the closed-loop driver-vehicle system will return to the equilibrium or is instead headed toward instability.

This capability could be used as a supervisory trigger for active safety controllers that override or modulate the driver's inputs when the system state approaches the boundary of the safe ROA. The same methodology could also be applied to characterize the safe ROA of a vehicle with ESP system by adjusting the time-delay model and including the control logic accordingly.

2.5 Conclusions

This work presents a preliminary application of SOS programming to estimate safe invariant sets for a vehicle-with-driver system. The proposed approach extends Lyapunov-based ROA estimation to a seven-dimensional coupled driver-vehicle model, explicitly incorporating the driver's action into the stability assessment.

The results show that SOS techniques can provide certified safety regions for the vehicle-with-driver system through the Lyapunov polynomial $V(\mathbf{x})$, allowing efficient online safety evaluation via a single function check. Comparison with simulation-based boundaries indicates that, despite their inherent conservatism, the SOS-based estimates capture the main features of the safe ROA and offer rigorous guarantees on both stability and state-constraint satisfaction.

Several limitations warrant further investigation. The polynomial approximation required for SOS tractability imposes restrictive slip-angle bounds, which significantly affect the estimated safe region, particularly in the understeering case. In addition, the simplified single-track model and basic driver representation reduce applicability to more complex operating conditions. These approximations are necessary to control computational complexity, but they limit the domain in which the polynomial dynamics remain accurate.

Despite these constraints, the results are promising and demonstrate the feasibility of the approach. The methodology provides a systematic framework for certifying safety that can be extended to richer vehicle models and more realistic driver behaviors. Future work should investigate higher-degree polynomial approximations, improved representations of tire forces, and applications to more complex maneuvers beyond straight-line operation. A particularly promising direction is the integration of such certified safe regions as supervisory layers in active control systems, potentially contributing to enhance active safety in passenger cars.

2.6 References

- (Mosek 2025) MOSEK. 2025. *The MOSEK Python Fusion API Manual. Version 11.0.*
- (Chen and Kuo 2014) Chen, B.-C., and C.-C. Kuo. 2014. "Electronic Stability Control for Electric Vehicle with Four in-Wheel Motors." *Int.J Automot. Technol.* 15 (4): 573–80.
- (Daher et al. 2017) Daher, Alaa M., Carine R. Bardawil, and Naseem A. Daher. 2017. "Vehicle Stability Based on g-g Diagram Through Braking and Driveline." *2017 American Control Conference (ACC)* (Seattle, WA, USA), 309–14.

- (Della Rossa et al. 2012) Della Rossa, Fabio, Giampiero Mastinu, and Carlo Piccardi. 2012. “Bifurcation Analysis of an Automobile Model Negotiating a Curve.” *Vehicle System Dynamics* 50 (10): 1539–62.
- (Farroni et al. 2013) Farroni, Flavio, Russo, Russo, Terzo, and Francesco and Timpone. 2013. “A Combined Use of Phase Plane and Handling Diagram Method to Study the Influence of Tyre and Vehicle Characteristics on Stability.” *Vehicle System Dynamics* 51 (8): 1265–85.
- (Guiggiani 2023) Guiggiani, Massimo. 2023. *The Science of Vehicle Dynamics: Handling, Braking, and Ride of Road and Race Cars*. Springer International Publishing.
- (Hashemi et al. 2016) Hashemi, Ehsan, Mohammad Pirani, Amir Khajepour, and Alireza Kasaiezadeh. 2016. “A Comprehensive Study on the Stability Analysis of Vehicle Dynamics with Pure/Combined-Slip Tyre Models.” *Vehicle System Dynamics* 54 (12): 1736–61.
- (Huang et al. 2020) Huang, Yiwen, Wei Liang, and Yan Chen. 2020. “Stability Regions of Vehicle Lateral Dynamics: Estimation and Analysis.” *Journal of Dynamic Systems, Measurement, and Control* 143 (051002).
- (Imani Masouleh and Limebeer 2018) Imani Masouleh, Mehdi, and David J. N. Limebeer. 2018. “Region of Attraction Analysis for Nonlinear Vehicle Lateral Dynamics Using Sum-of-Squares Programming.” *Vehicle System Dynamics* 56 (7): 1118–38.
- (Johnson and Huston 1984) Johnson, D. B., and J. C. Huston. 1984. “Nonlinear Lateral Stability Analysis of Road Vehicles Using Liapunov’s Second Method.” *1984 SAE International Off-Highway and Powerplant Congress and Exposition* (Milwaukee, Wisconsin, United States), 841057.
- (Jarvis-Wloszek et al. 2005) Jarvis-Wloszek, Zachary, Ryan Feeley, Weehong Tan, Kunpeng Sun, and Andrew Packard. 2005. “Control Applications of Sum of Squares Programming.” In *Positive Polynomials in Control*, edited by Didier Henrion and Andrea Garulli. Springer.
- (Lai et al. 2021) Lai, Fei, Chaoqun Huang, and Chengyue Jiang. 2021. “Comparative Study on Bifurcation and Stability Control of Vehicle Lateral Dynamics.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 6 (1): 35–52.
- (Lenzo et al. 2021) Lenzo, Basilio, Mattia Zanchetta, Aldo Sorniotti, Patrick Gruber, and Wouter De Nijs. 2021. “Yaw Rate and Sideslip Angle Control Through Single Input Single Output Direct Yaw Moment Control.” *IEEE Trans. Contr. Syst. Technol.* 29 (1): 124–39.
- (Mastinu et al. 2024) Mastinu, Giampiero R. M., Giorgio Previati, Fabio Della Rossa, Massimiliano Gobbi, and Marco Fainello. 2024. “How Drivers Lose Control of the Car.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 8 (1): 10-08-01-0007.
- (Mastinu and Ploechl 2014) Mastinu, Gianpiero, and Manfred Ploechl, eds. 2014. *Road and Off-Road Vehicle System Dynamics Handbook*. CRC Press.
- (Mastinu et al. 2023) Mastinu, Gianpiero, Fabio Della Rossa, Giorgio Previati, Massimiliano Gobbi, and Marco Fainello. 2023. “Global Stability of Road Vehicle Motion with Driver Control.” *Nonlinear Dyn* 111 (19): 18043–59.
- (Meng et al. 2022) Meng, Fanyu, Shuming Shi, Boshi Zhang, Yunxia Li, and Nan Lin. 2022. “Global Characteristics Analysis for Vehicle System with Driving Torque Based on Dissipation of Energy.” *Int.J Automot. Technol.* 23 (6): 1609–19.
- (Motzkin 1967) Motzkin, Theodore S. 1967. “Inequalities.” In *Studies in Pure Mathematics*, edited by Oved Shisha. Academic Press.
- (Pacejka 2012) Pacejka, Hans B. 2012. *Tire and Vehicle Dynamics*. 3rd ed. Elsevier.
- (Ribeiro and Fioravanti 2020) Ribeiro, Alexandre M., André R. Fioravanti, Alexandra Moutinho, and Ely C. De Paiva. 2020. “Sum of Squares Approach for Ground Vehicle Lateral Control Under Tire Saturation Forces.” *IFAC-PapersOnLine* 53 (2): 14387–93.
- (Plochl and Edelman 2007) Plöchl, Manfred, and Johannes Edelmann. 2007. “Driver Models in Automobile Dynamics Application.” *Vehicle System Dynamics* 45 (7-8): 699–741.
- (Ribeiro et al. 2022) Ribeiro, A. M., Fioravanti, Moutinho, and E. C. and de Paiva. 2022. “Nonlinear State-Feedback Design for Vehicle Lateral Control Using Sum-of-Squares Programming.” *Vehicle System Dynamics* 60 (3): 743–69.

(Sadri and Wu 2013) Sadri, Sobhan, and Christine Wu. 2013. “Stability Analysis of a Nonlinear Vehicle Model in Plane Motion Using the Concept of Lyapunov Exponents.” *Vehicle System Dynamics* 51 (6): 906–24.

(Slotine and Li 1991) Slotine, Jean-Jacques E., and Weiping Li. 1991. *Applied Nonlinear Control*. 1st ed. Prentice Hall.

(Strogatz 2019) Strogatz, Steven. 2019. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Second edition, first issued in hardback. A Chapman & Hall Book. CRC Press.

(Tamba and Nazaruddin 2018) Tamba, Tua A., and Yul Y. Nazaruddin. 2018. “Nonlinear Stability Analysis of Vehicle Side-Slip Dynamics Using SOS Programming.” *2018 5th International Conference on Electric Vehicular Technology (ICEVT)*, 1–6.

(Topcu et al. 2008) Topcu, Ufuk, Andrew Packard, and Peter Seiler. 2008. “Local Stability Analysis Using Simulations and Sum-of-Squares Programming.” *Automatica* 44 (10): 2669–75.

(Zhu et al. 2022) Zhu, Zhewei, Yu Zhang, Yiwei Huang, and Yechen Qin. 2022. “Sum-Of-Squares Based Vehicle Dynamic Stability Method and Its Applications in ADAS.” *2022 IEEE Intelligent Vehicles Symposium (IV)* (Aachen, Germany), 247–54.

(Drake 2019) Tedrake, Russ, and the Drake Development Team. 2019. *Drake: Model-Based Design and Verification for Robotics*.

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

3.1 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 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 (Chen and Kuo 2014; Lenzo et al. 2021). 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 (Della Rossa et al. 2012).

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 (Hashemi et al. 2016) or multi-point linearization techniques (Huang et al. 2020). 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 (Ribeiro et al. 2020; Zhu et al. 2022).

Alternative quantitative indicators include Lyapunov exponents to measure convergence rates to capture how disturbances decay during aggressive maneuvers (Sadri and Wu 2013). 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 (Zhang et al. 2025) 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 (Maruyama and Yamazaki 2006), 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 (Wanner et al. 2016), 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 (Cho et al. 2025) 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 (Mastinu et al. 2020), 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 (Mastinu et al. 2023), 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 (Mastinu et al. 2024), 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 (Previati et al. 2024), 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 (Kharrazi et al. 2020), 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 (Musso et al. 2024). 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 (Asperti et al. 2025), 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 (Alfatti et al. 2023), 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 (Novi et al. 2018), 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 (Della Rossa et al. 2025) 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 (Della Rossa et al. 2025) touches on motorsport applications, most of the cited simulator-based studies are framed around road-vehicle applications, such as crosswind 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 (Gulisano et al. 2025) and brings it to the DriSMi driving simulator (PoliMi 2024), 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 (Gulisano et al. 2025) 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 this chapter is organized as follows. Section 3.2 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 3.3 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 3.4 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 3.5 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 3.6 summarizes the main findings and discusses limitations and future research directions.

3.2 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 (Gulisano et al. 2025).

3.2.1 Vehicle and disturbance model

The planner relies on the nonlinear Single-Track model (Guiggiani 2023) with nonlinear tire characteristics and open differential schematically depicted in Figure 1.

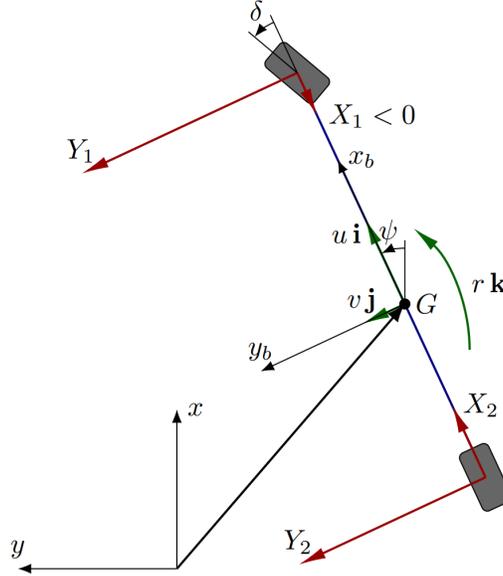


Figure 1: Single-track model used as a basis for our stochastic vehicle dynamics model.

The state vector $\mathbf{x} = (u, v, r, x_G, y_G, \psi)^T$ collects longitudinal and lateral velocities, yaw rate, and the pose of the vehicle with respect to the center of mass. The control input $\mathbf{u} = (X_{2,a}, X_{2,b}, \delta)^T$ 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 $\mathbf{w}(t)$,

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

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

3.2.2 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

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

while the covariance evolves through the Lyapunov differential equation

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

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. 3.2.5. A more detailed treatment is provided in (Gulisano et al. 2025).

3.2.3 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 $\bar{\mathbf{P}}_0 \succcurlyeq 0$, 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 $\boldsymbol{\Omega}_k(\cdot)$.

3.2.4 Open-loop disturbance-aware OCP

The disturbance-aware reference generator solves the nonlinear program

$$\begin{aligned}
 & \underset{\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{P}_k, \mathbf{z}_k}{\text{minimize}} && J_k(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k) \\
 & \text{s.t.} && \mathbf{0} = \boldsymbol{\Psi}_k^\mu(\boldsymbol{\mu}_{k-1}, \boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{z}_k), && k = 1, \dots, N \\
 & && \boldsymbol{\mu}_0 = \bar{\boldsymbol{\mu}}_0 \\
 & && \mathbf{0} = \boldsymbol{\Psi}_k^P(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{P}_{k-1}^{j-1}, \mathbf{P}_k^j, \boldsymbol{\Sigma}_k^j, \mathbf{z}_k), && \begin{matrix} k=1, \dots, N; \\ j=1, \dots, H \end{matrix} \\
 & && \mathbf{P}_k^0 = \bar{\mathbf{P}}_0 \succeq 0 && k = 1, \dots, N \\
 & && \mathbf{0} = \boldsymbol{\Omega}_k(\boldsymbol{\mu}_k, \boldsymbol{\xi}_k, \mathbf{u}_k, \mathbf{z}_k), && k = 0, \dots, N \\
 & && 0 \geq h_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{z}_k) + \beta_i(\boldsymbol{\mu}_k, \mathbf{u}_k, \mathbf{P}_k^H, \mathbf{z}_k), && k = 1, \dots, N; i \in \mathcal{I}
 \end{aligned} \tag{1}$$

where J_k is the standard minimum-lap-time cost (i.e. our planning objective). Equations above collect the collocation residuals, initial conditions, and algebraic constraints. Inequality constraints in the last row 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. 3.2.5.

3.2.5 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

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

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

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

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. 3.5. A thorough theoretical treatment can be found in our previous work (Gulisano et al. 2025).

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

3.2.6 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 (1). 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.

3.3 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.

3.3.1 DriSMi layout and vehicle dynamics model

The dynamic simulator employed is located at Politecnico di Milano (DriSMi lab (PoliMi 2024)). 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 (VI-grade GmbH 2024) (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 (VI-grade GmbH 2024; Mastinu and Ploechl 2014; Mastinu et al. 2020).

3.3.2 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 (Mastinu et al. 2020, 2023; Mastinu et al. 2024). Tire forces are modeled via a Pacejka Magic Formula 6.1 (VI-grade GmbH 2024). 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.

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

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

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 (Pacejka 2012). In particular, for each axle we assume

$$F_y(\alpha, F_z; \mathbf{p}) = D_y \sin(C_y \arctan(B_y \alpha) - E_y (B_y \alpha - \arctan(B_y \alpha))), \quad (2)$$

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

$$\bar{\mathbf{p}} = \operatorname{argmin}_{\mathbf{p}} \sum_{k=1}^{N_{\text{train}}} (\bar{F}_{y,k} - F_y(\bar{\alpha}_k, \bar{F}_{z,k}; \mathbf{p}))^2, \quad (3)$$

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 (2), for the same axle slip angle $\bar{\alpha}_k$ and vertical load $\bar{F}_{z,k}$. The optimization (3) 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.

	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

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

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.

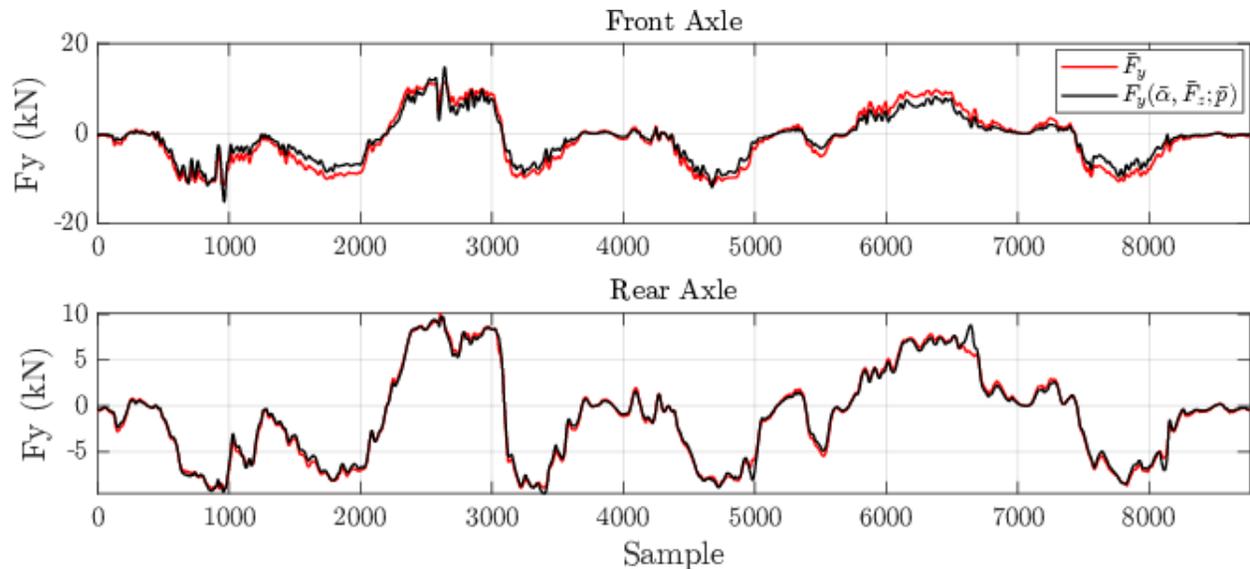


Figure 2: Comparison between axle lateral forces obtained from the simulator telemetry (red) and from the calibrated Magic Formula tire model (black) for the front and rear axles.

(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.

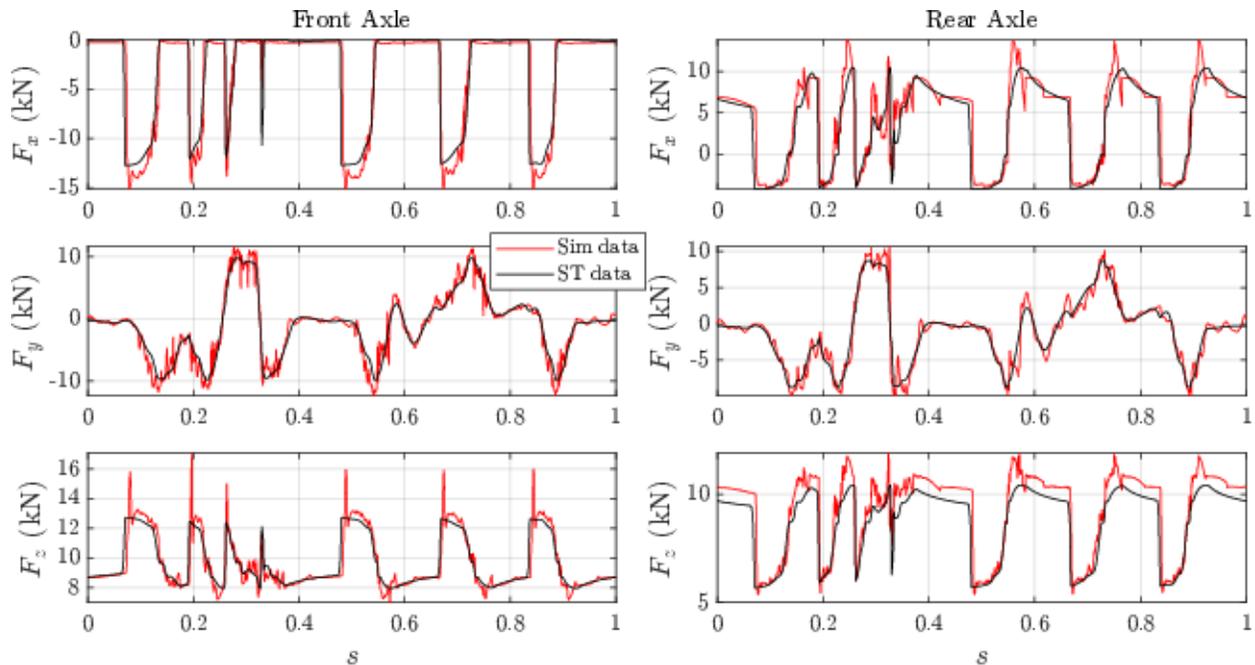


Figure 3: Comparison between longitudinal, lateral, and vertical axle forces obtained from the simulator (Sim data, red) and from the single-track minimum-lap-time optimization (ST data, black) while attempting to follow the optimized trajectory, plotted against the normalized distances along the track.

Figure 4 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.

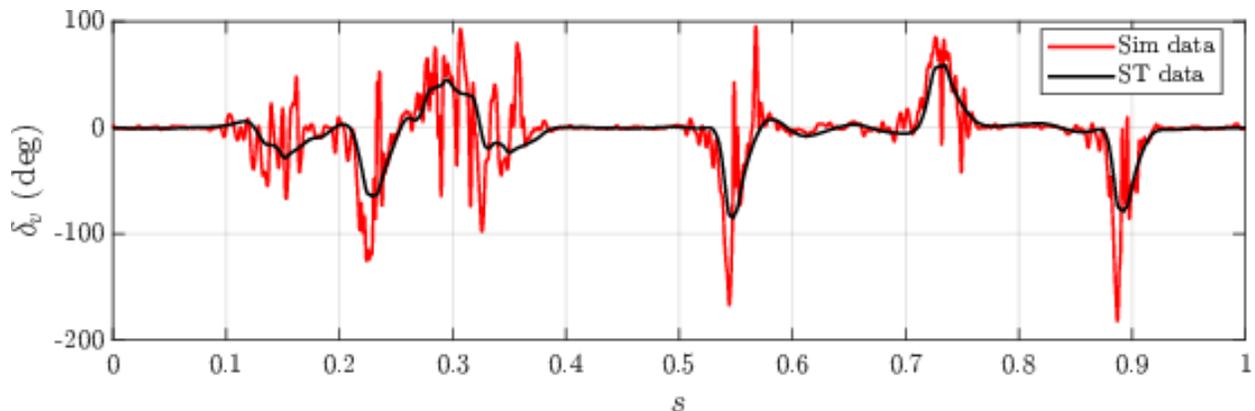


Figure 4: Comparison between steering-wheel angle recorded at the driving simulator (Sim data, red) and steering input from the single-track minimum-lap-time optimization (ST data, black), plotted against the normalized distance along the track while attempting to follow the optimized trajectory.

3.4 Experimental Setting and Test Protocol

3.4.1 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.



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

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.

3.4.2 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:

- 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 speed.

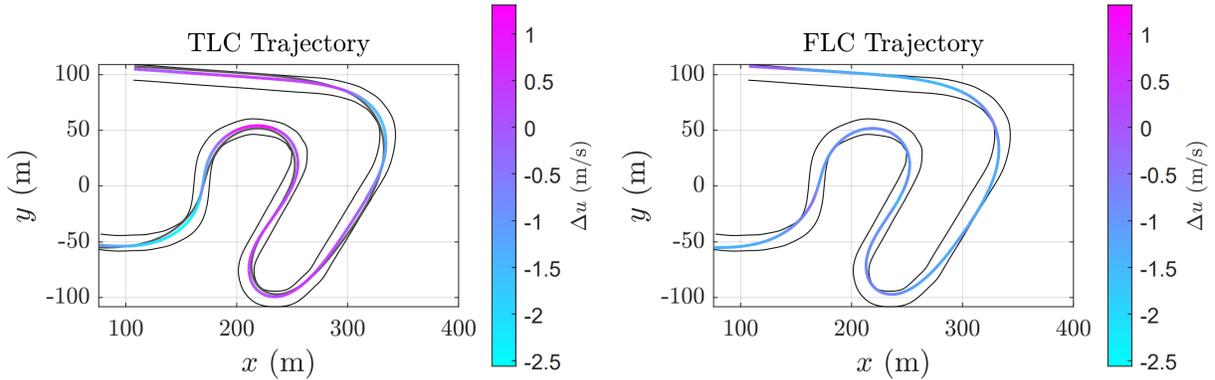


Figure 6: Comparison of robust planning trajectories (TLC and FLC) overlaid on the nominal trajectory (NOM) for a representative track sector. Trajectory color indicates longitudinal speed difference relative to the nominal. The TLC trajectory (left) deviates more from NOM and achieves higher speeds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

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 5.2 m in the driver's corrective actions. The grid spacing and the reset covariance assumptions were calibrated from pilot tests so that the induced tightening 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.

3.4.3 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

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

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

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

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:

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

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

Style indicators. RMS side-slip usage:

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

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

3.5 Results

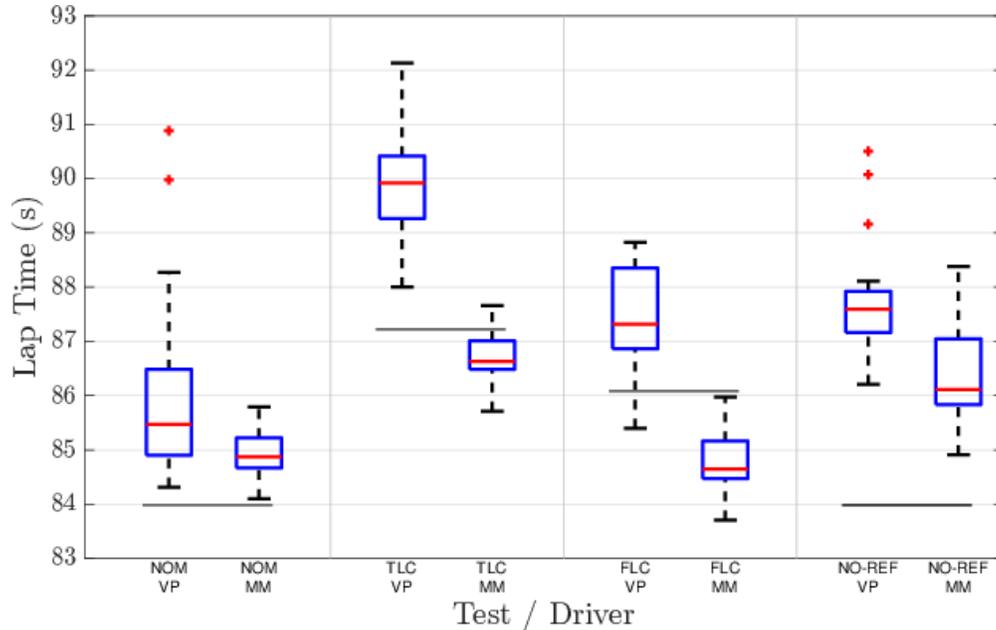


Figure 7: Lap times by test and driver (box plots). Grey lines indicate the reference lap time for each trajectory type. For driver VP, NOM is fastest and FLC is moderately slower; for driver MM, NOM and FLC yield comparable medians; TLC is slower for both; NO-REF for driver MM is slower and more variable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

3.5.1 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 driver 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 practically aligned: 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 lower lap times than the planner's reference (indicated by the horizontal 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 3.5.4: MM's lower trajectory-following fidelity (higher RMS errors, as visible in Figure 11) enables these faster lap times by deviating from the planned path to exploit tighter margins. The FLC column for VP shifts vertically by roughly 1–2 s relative to NOM yet preserves a compact Interquartile Range (IQR), signaling 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.

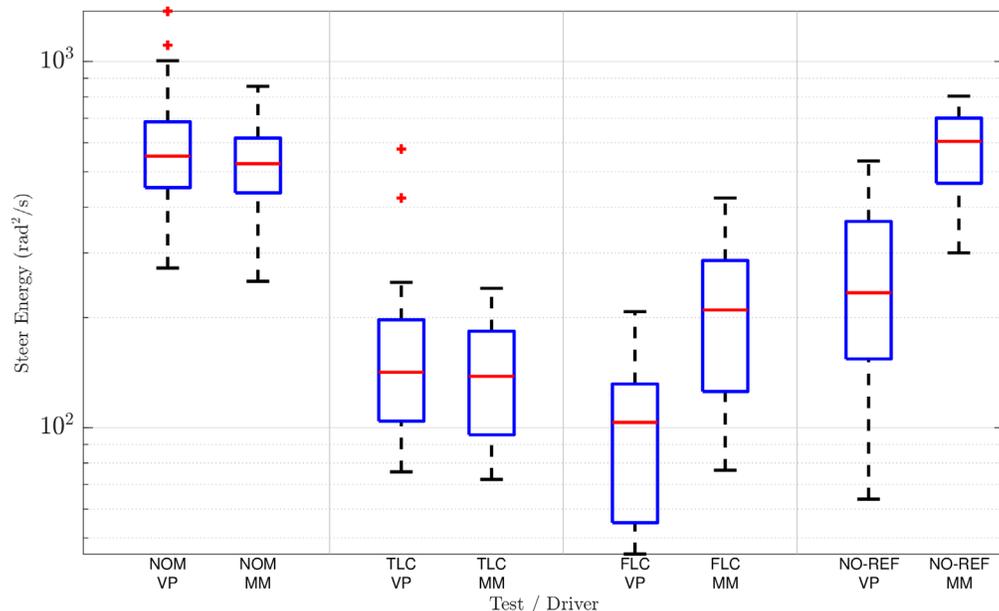


Figure 8: Steer Energy by test and driver (log scale). NOM requires the highest effort among reference-following tests; FLC and TLC markedly reduce effort. For driver MM, NO-REF shows *high* SE coupled with slow LTs.

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

3.5.3 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.

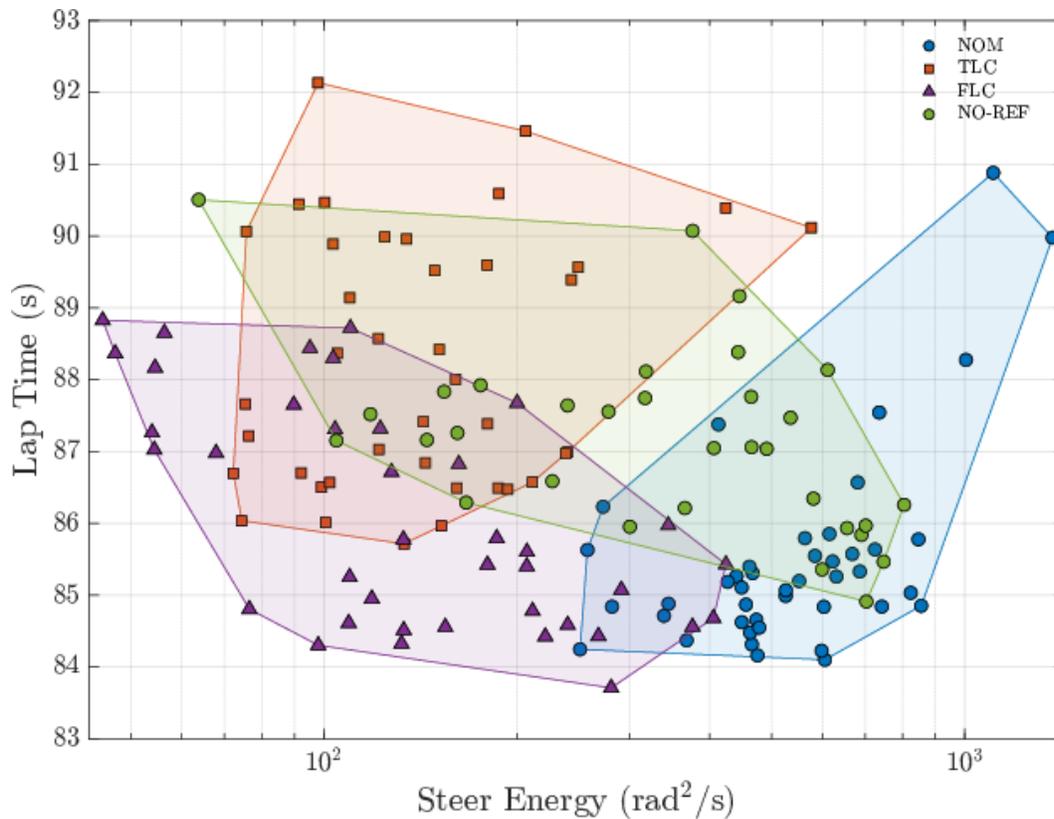


Figure 9: Overall Lap Time–Steer Energy scatter plot (test types as colors, both drivers). FLC typically lies near a Pareto-efficient band; NOM is *generally* fastest but costly; TLC is lowest effort but slower.

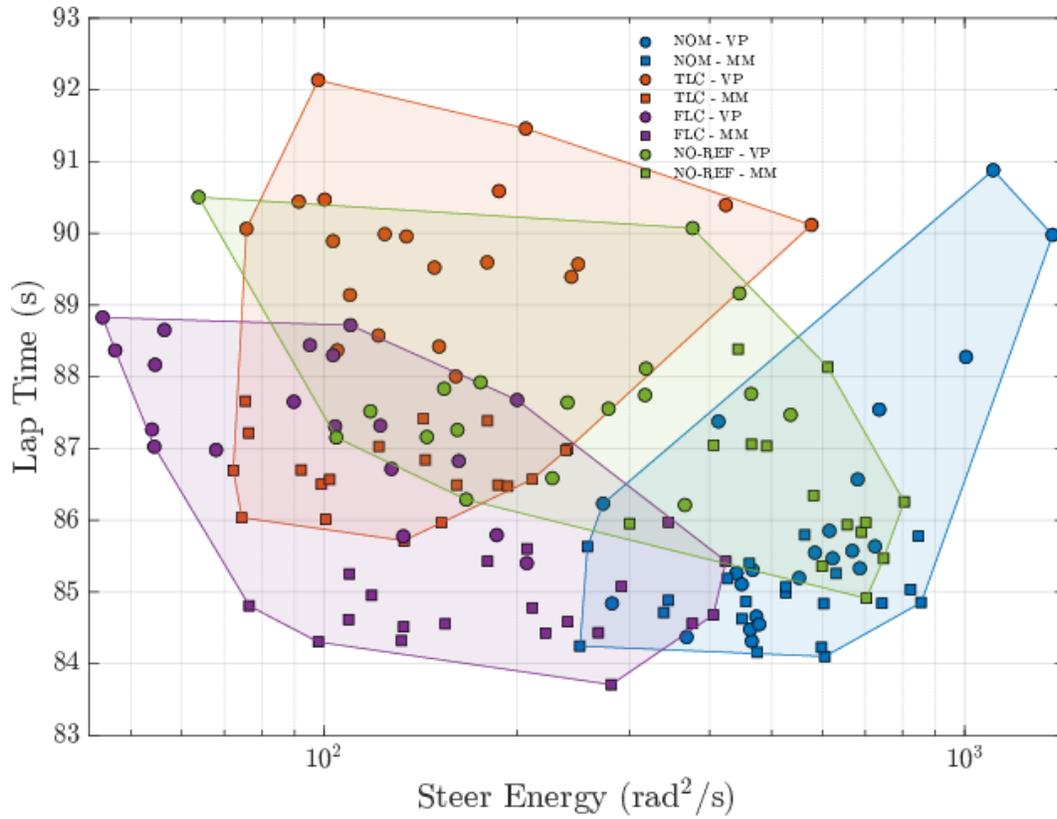


Figure 10: Lap Time–Steer Energy scatter plot with driver markers (colors = test types; markers distinguish driver VP and driver MM). FLC retains the favorable band for both drivers; NO-REF for driver MM sits at high SE and high LT.

3.5.4 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) RMS (e_y), (ii) RMS (e_v), (iii) RMS (β_{drv}) with the corresponding RMS (β_{ref}) overlaid in black, and (iv) a light-colored bar for RMS ($\dot{\delta}$) as a qualitative proxy for SE. Across conditions, FLC typically pulls RMS (e_y) and RMS (e_v) below NOM, while simultaneously reducing RMS ($\dot{\delta}$); the effect is clearest for VP but remains visible for MM. Under TLC both drivers see reduced RMS (β_{drv}) and 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 RMS (e_v), whereas MM shows larger lateral error 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 lap times (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 RMS (β_{drv}) and 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 RMS (β_{drv}) and 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.

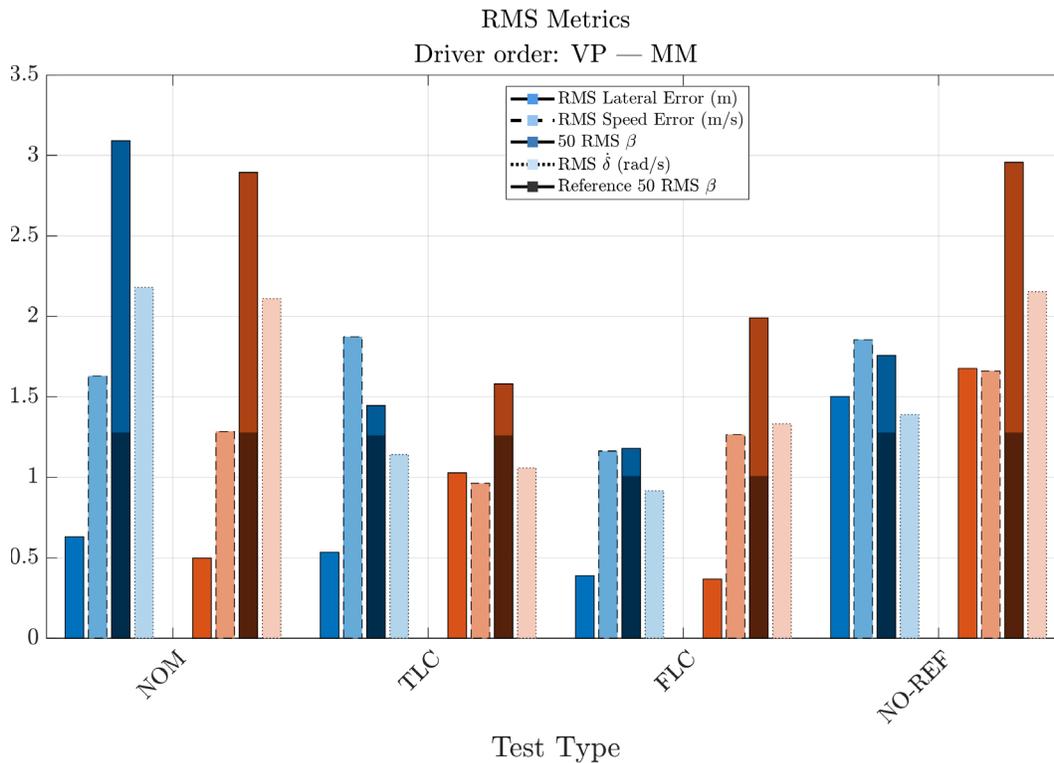


Figure 11: RMS metrics by test and driver (driver order: VP then MM). FLC eases tracking vs. NOM and reduces steering activity; TLC is most “relaxed” but slowest; NO-REF for MM combines poor tracking with elevated SE and slow LTs.

3.5.5 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.

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

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

	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%)

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

3.5.6 Discussion

1. 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 1–2 s 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.

2. 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.

3. 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 1–2 s for the same controller. More strikingly, MM exceeds the planner’s reference lap time 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.

4. *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.

3.6 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.

3.7 References

(Alfatti et al. 2023) Alfatti, Federico, Margherita Montani, Tommaso Favilli, et al. 2023. “Implementation and Performances Evaluation of Advanced Automotive Lateral Stability Controls on a Real-Time Hardware in the Loop Driving Simulator.” *Applied Sciences* 13 (11): 6592.

(Asperti et al. 2025) Asperti, Marco, Michele Vignati, and Edoardo Sabbioni. 2025. “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* 9 (2): 177–93.

(Chen and Kuo 2014) Chen, B.-C., and C.-C. Kuo. 2014. “Electronic Stability Control for Electric Vehicle with Four in-Wheel Motors.” *Int. J. Automot. Technol.* 15 (4): 573–80.

- (Cho et al. 2025) Cho, Hyonbae, Yejin Kim, SeokJin Oh, and Ilsoo Yun. 2025. “Development of a Stability Index for Evaluating Drivers’ Psychological Stability During Truck Platooning.” *Applied Sciences* 15 (10): 5429.
- (Della Rossa et al. 2025) Della Rossa, Fabio, Marco Fontana, Simone Giacintucci, Massimiliano Gobbi, et al. 2025. “Early Detection of the Loss of Control of Motorsport Vehicles.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 9 (3): 473–99.
- (Della Rossa et al. 2012) Della Rossa, Fabio, Giampiero Mastinu, and Carlo Piccardi. 2012. “Bifurcation Analysis of an Automobile Model Negotiating a Curve.” *Vehicle System Dynamics* 50 (10): 1539–62.
- (Guiggiani et al. 2023) Guiggiani, Massimo. 2023. *The Science of Vehicle Dynamics: Handling, Braking, and Ride of Road and Race Cars*. 3rd ed. Springer International Publishing.
- (Gulisano et al. 2025) Gulisano, Martino, Matteo Masoni, Marco Gabiccini, and Massimo Guiggiani. 2025. “Disturbance-Aware Minimum-Time Planning Strategies for Motorsport Vehicles with Probabilistic Safety Certificates.” *Vehicle System Dynamics*, October, 1–25.
- (Hashemi et al. 2016) Hashemi, Ehsan, Mohammad Pirani, Amir Khajepour, and Alireza Kasaiezadeh. 2016. “A Comprehensive Study on the Stability Analysis of Vehicle Dynamics with Pure/Combined-Slip Tyre Models.” *Vehicle System Dynamics* 54 (12): 1736–61.
- (Huang et al. 2020) Huang, Yiwen, Wei Liang, and Yan Chen. 2020. “Stability Regions of Vehicle Lateral Dynamics: Estimation and Analysis.” *Journal of Dynamic Systems, Measurement, and Control* 143 (051002).
- (Kharrazi et al. 2020) Kharrazi, Sogol, Bruno Augusto, and Niklas Fröjd. 2020. “Vehicle Dynamics Testing in Motion Based Driving Simulators.” *Vehicle System Dynamics* 58 (1): 92–107.
- (Lenzo et al. 2021) Lenzo, Bilio, Mattia Zanchetta, Aldo Sorniotti, Patrick Gruber, and Wouter De Nijs. 2021. “Yaw Rate and Sideslip Angle Control Through Single Input Single Output Direct Yaw Moment Control.” *IEEE Trans. Contr. Syst. Technol.* 29 (1): 124–39.
- (Maruyama et al. 2006) Maruyama, Yoshihisa, and Fumio Yamazaki. 2006. “Driving Simulator Experiment on the Moving Stability of an Automobile Under Strong Crosswind.” *Journal of Wind Engineering and Industrial Aerodynamics* 94 (4): 191–205.
- (Mastinu et al. 2024) Mastinu, Giampiero, Giorgio Previati, Fabio Della Rossa, Massimiliano Gobbi, and Marco Fainello. 2024. “How Drivers Lose Control of the Car.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 8 (1): 10-08-01-0007.
- (Mastinu et al. 2020) Mastinu, Gianpiero, Danilo Biggio, Fabio Della Rossa, and Marco Fainello. 2020. “Straight Running Stability of Automobiles: Experiments with a Driving Simulator.” *Nonlinear Dynamics* 99 (4): 2801–18.
- (Mastinu et al. 2023) Mastinu, Gianpiero, Fabio Della Rossa, Giorgio Previati, Massimiliano Gobbi, and Marco Fainello. 2023. “Global Stability of Road Vehicle Motion with Driver Control.” *Nonlinear Dyn* 111 (19): 18043–59.
- (Mastinu et al. 2014) Mastinu, Gianpiero, and Manfred Ploechl, eds. 2014. *Road and Off-Road Vehicle System Dynamics Handbook*. 1st ed. CRC Press.
- (Musso et al. 2024) Musso, Enrico, Giorgio Previati, and Gianpiero Mastinu. 2024. “Analytical and Experimental Handling Performance of Ultra-Efficient Lightweight Vehicles.” *SAE Int. J. Adv. & Curr. Prac. In Mobility* 6 (4): 2049–58.
- (Novi et al. 2018) Novi, Tommaso, Alexander Liniger, Renzo Capitani, Marco Fainello, et al. 2018. “The Influence of Autonomous Driving on Passive Vehicle Dynamics.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 2 (4): 285–95.
- (Pacejka et al. 2012) Pacejka, Hans B. 2012. *Tire and Vehicle Dynamics*. 3rd ed. Elsevier.
- (PoliMi 2024) PoliMi. 2024. *DriSMi (Official Website)*. <https://www.drismi.polimi.it>.
- (Previati et al. 2024) Previati, Giorgio, Gianpiero Mastinu, and Massimiliano Gobbi. 2024. “Driver–Steering Wheel Interaction During Cornering.” *SAE Int. J. Veh. Dyn., Stab., and NVH* 8 (3): 427–42.

(Ribeiro et al. 2020) Ribeiro, Alexandre M., André R. Fioravanti, Alexandra Moutinho, and Ely C. De Paiva. 2020. “Sum of Squares Approach for Ground Vehicle Lateral Control Under Tire Saturation Forces.” *IFAC-PapersOnLine* 53 (2): 14387–93.

(Sadri et al. 2013) Sadri, Sobhan, and Christine Wu. 2013. “Stability Analysis of a Nonlinear Vehicle Model in Plane Motion Using the Concept of Lyapunov Exponents.” *Vehicle System Dynamics* 51 (6): 906–24.

(VI-grade 2024) VI-grade GmbH. 2024. VI-Grade (Official Website). <https://www.vigrade.com>.

(Wanner et al. 2016) Wanner, Daniel, Maria Kreußlein, Bruno Augusto, Lars Drugge, and Annika Stensson Trigell. 2016. “Single Wheel Hub Motor Failures and Their Impact on Vehicle and Driver Behaviour.” *Vehicle System Dynamics* 54 (10): 1345–61.

(Zhang et al. 2025) Zhang, Siyang, Chi Zhao, Zherui Zhang, and Yecheng Lv. 2025. “Driving Simulator Validation Studies: A Systematic Review.” *Simulation Modelling Practice and Theory* 138 (January): 103020.

(Zhu et al. 2022) Zhu, Zhewei, Yu Zhang, Yiwei Huang, and Yechen Qin. 2022. “Sum-Of-Squares Based Vehicle Dynamic Stability Method and Its Applications in ADAS.” 2022 IEEE Intelligent Vehicles Symposium (IV) (Aachen, Germany), 247–54.

4 Conclusions

The deliverable focused on three planned areas:

- **Early detection of instability in motorsport vehicles**
- **Stability certification for vehicle–driver dynamics**
- **Robust path planning and tracking**

All objectives have been met, providing new methods to (i) compute thresholds above which disturbances make a vehicle–driver system unstable and (ii) support stability and path tracking in motorsport and road vehicles, with human or non-human drivers.

Early detection of instability.

A Floquet-based Degree of Stability (DoS) index has been validated on a dynamic simulator, offering an operational threshold for distinguishing recoverable from non-recoverable maneuvers and enabling early recognition of loss-of-control conditions.

Stability certification.

Lyapunov-based Sum-of-Squares tools have been employed to estimate safe operating regions with mathematical guarantees, extending preliminary results and providing formal thresholds for disturbance-induced instability in vehicle–driver systems.

Robust path planning and tracking.

Disturbance-aware reference trajectories have been tested with semi-professional drivers, showing that robustness-embedded planning enables fast yet stable laps and reduced steering effort compared with nominal or unguided driving.

Overall, the deliverable delivers a coherent set of advances - **detecting instability, certifying stability, and planning robust trajectories** - fully supporting Milestone 3, *Stable path tracking for motorsport or road vehicles performed*.