

Stochastic Control Behavior of the Balancing Rider for Cycling Safety in Traffic Simulation

Christoph M. Konrad , Riender Happee , Jason K. Moore , Azita Dabiri 

Faculty of Mechanical Engineering (ME), Delft University of Technology

*Correspondence: c.m.konrad@tudelft.nl

This paper is a preprint that has not been peer reviewed. It has been submitted for peer review and is subject to change with revisions.

Abstract

Cyclist models with realistic behavior are essential for simulation-based safety testing of automated vehicles, assistance systems, traffic control, and infrastructure. We investigated the role of bicycle dynamics and human control for traffic processes in an experiment with ten cyclists following unanticipated, visual heading commands under varying time-pressure at multiple speeds, and propose a new Balancing Rider model for traffic simulation. It combines the Carvallo-Whipple bicycle model with full-state feedback to capture bicycle dynamics, human heading control, and balancing. We estimated control parameters for each individual maneuver, and described their distribution in a stochastic model. We tested trajectories generated from this distribution for predictive power and simulated obstacle avoidance to analyze safety prediction in comparison to a planar baseline model. The cycling experiment revealed a large variety of slow and quick control strategies for heading changes. The Balancing Rider model achieved excellent calibration and promising predictive performance covering the observed control strategies. A baseline planar point model couldn't capture details of cyclist maneuvers like counter-steering and heading overshoot, with path deviations reaching an order of meters. For obstacle avoidance, we project > 0.4 s Time-to-Collision difference between the models, showing that the added physical detail greatly impacts simulated safety assessment. This illustrates that bicycle dynamics and human control behaviors are important elements explaining causal cycling traffic processes and should be included in models for simulated safety testing. The large variety of control behaviors among the participants could potentially be an explanatory factor of traffic conflict outcomes, which future work should investigate further.

Keywords: Microscopic Traffic Simulation, Surrogate Safety Analysis, Cyclist Modeling, Vehicle Dynamics, Bicycle Control, Control Behavior

Highlights

- Lateral bicycle dynamics are required to explain cyclist maneuvers in detail
 - Human control behaviors during heading changes vary considerably between individuals
 - Considering bicycle dynamics and human control significantly impacts simulated safety
 - We propose a stochastic model of the balancing rider suitable for traffic simulation
-

1. Introduction

Technologies such as Automated Vehicles (AVs), Advanced Driver Assistance Systems (ADAS), or novel infrastructure require traffic simulation during the development phase or must be rigorously tested before public introduction. A particular point of interest is the safety of Vulnerable Road Users (VRUs), including cyclists. For cyclists, global fatality numbers are reducing less than for car occupants and pedestrians (ITF, 2023). Statistics from the EU show an over-proportional risk in traffic (European Commission, 2025), highlighting a particular need to ensure cycling safety in the development of traffic innovations. The interpretation of VRU actions is a significant challenge contributing to records of AV conflicts (Alozi and Hussein, 2023). To test potential remedies, simulations enable comprehensive scenario coverage, targeted testing of edge cases, and testing scenarios modeled after naturalistic conditions (Sun et al., 2022). For driver assistance systems and advanced driving functions, which are also expected on bicycles (Kapousizis et al., 2022; Berge et al., 2023), virtual randomized trials (Brunner et al., 2019; Denk et al., 2022), and counterfactual analysis (Bärgman et al., 2017) are solutions to systematically quantify benefits through simulation. Further use-cases include the microsimulation-based safety evaluation of infrastructure designs and intelligent traffic control (Konrad et al., 2025).

These purposes require traffic models to describe the traffic environment on a microscopic level, where tested technologies, infrastructure, road user behavior, their vehicles, and the environment are modeled explicitly and have causal dependencies. However, missing validations or an inappropriate model choice limit the predictive power of simulations, casting doubts on the ecological validity of the results. Bärgman et al. (2017) demonstrate conclusively how different choices of car driver behavior may affect the safety assessment based on counterfactual simulations due to their different effect on event causality. Specifically for cycling, we review existing studies using simulated cyclist traffic in our previous work (Konrad et al., 2025) and observe that the predominantly used microscopic traffic flow models inadequately describe cyclist behavior to predict safety. We conclude that predictive models must be rooted in a deep understanding of conflict processes and explicitly model the key aspects of conflict causality to ensure that simulated assessments foster cycling safety. Popular traffic flow models included in tools like VISSIM and SUMO describe road user states in terms of lateral and longitudinal position on predefined lanes and use different parameterizations of the same behavioral models for cyclists and drivers. A large body of work tackles improving tactical driver and rider behaviors like queuing, overtaking, negotiating priority, or path planning. To the best of our knowledge, no previous work has sufficiently considered accurate physical dynamics of the bicycle, including the coupling of roll (how the bicycle leans left or right) and steer angle, or the control behaviors for balancing and navigation of the rider. Related research bringing traffic simulation cyclist models closer to physical bicycle dynamics was conducted by Guo et al. (2021). They derive the desired speed of cyclists in turns of different radii from an estimate of the centrifugal force, but do not explicitly model roll dynamics. Huang et al. (2012) pair a two-wheeled car model with a social force model to simulate vehicle dynamics and driver steering control. This exemplifies how to introduce more complex vehicle dynamics to traffic models, but it does not consider cyclists. Adaptions of the social force framework for cyclists typically use purely planar formulations of the agents as accelerated, mass-less shapes, with examples being Yuan et al. (2019) and Li et al. (2021). Fu et al. (2023) and Wang et al. (2024) add the cyclist mass but otherwise keep the planar dynamics. Kathis (2023) and Hoogendoorn et al. (2021) separate longitudinal motion from turn rate to account for the two-wheeled motion constraints, but don't consider any physical relationships within the lateral motion. In our exploratory study, we propose a social force concept with simplified dynamics of the upright, balancing bicycle, but do not calibrate or validate our model (Schmidt et al., 2024).

The fundamental dynamics and human control of real bicycles are well understood and outlined in multiple works (Hess et al., 2012; Moore, 2012; Schwab and Meijaard, 2013). Due to the single-track two-wheel configuration with a mass balanced above the ground, similar to an inverted pendulum, bicycles are unstable at both low and high speeds, and need active control to be stabilized over a broad speed range. When disturbed, the rider ensures that the wheels move back under the center of gravity by steering in the direction of the fall. During a turn, the rider and vehicle lean into the turn. Initiating a turn requires briefly steering in the opposite direction to initiate the bicycle to fall and then quickly catching the fall by a compensating steer action in the desired direction. This is commonly referred to as countersteering.

These dynamic principles result in countersteer motions and persistent steering actions for stabilization, visible in slightly meandering cyclist trajectories. Cyclists must execute these strategies in all situations, ranging from straight riding over planned turns to emergency

evasive actions. Especially during the latter, the control requirements of cycling could affect the available action space and shape the causal chain of traffic conflicts. Exploratory research indicates that simulated interaction trajectories may differ significantly from generic models if accurate bicycle dynamics are accounted for (Schmidt et al., 2024). Hence, we expect bicycle dynamics to already impact the process leading up to a conflict. This could make bicycle dynamics relevant for modeling causal conflict processes in simulated safety assessment, which often focuses on road user conflicts as a surrogate measure for collisions and relies on accurate spatiotemporal relationships between road users. Additionally, Reijne (2025) suggests that physical bicycle models may be promising to explain individual cyclist falls.

Previous work in bicycle control has primarily focused on the control performance of bicycle robots, theoretical stability, and understanding the human control of the bicycle (Schwab and Meijaard, 2013). Moore (2012) develops and tests a control model with explicit neurological feedback loops and neuromuscular dynamics. His model reproduces measured lateral dynamics in perturbation experiments well, correctly predicts all neurological assumptions, and highlights the expected speed dependency of control parameters. Dialynas et al. (2023) add sensory delays and state prediction, and equally achieve a good lateral fit. Reijne (2025) proposes an optimal control model leveraging significant domain knowledge, reducing the required number of control parameters to a single speed-dependent gain for stabilization after lateral disturbance. He tests the predictive performance of the average fit to multiple riders across a speed range of $2 - 7 \text{ m s}^{-1}$ and achieves acceptable results. To the best of our knowledge, no previous work has explicitly investigated the distribution of different control behaviors over multiple riders. In contrast, simulated safety assessment is typically stochastic, creating slightly different scenario conditions by randomly generating traffic exposure and sampling from distributions of behavioral parameters (Brunner et al., 2019; Bärghman et al., 2017).

Balancing the roll of the bicycle is ignored in traffic models of interactions with infrastructure and other road users, but has been extensively studied for modeling rider behavior in reaction to mechanical disturbances. This study aims to bridge the gap between bicycle dynamics research and traffic modeling to bring a validated, physical model of the upright bicycle and the balancing rider to stochastic traffic simulation. We identify the control parameters of a rider-bicycle model for multiple riders in a cycling experiment and fit parametric distributions to the behavioral parameters. We then test trajectory predictions based on draws from this distribution against previously unseen samples to demonstrate the predictive power of our model. We compare our results to a simple planar point model without explicit bicycle dynamics to test whether the added detail is relevant for simulated safety assessment and to gain an understanding of its contribution to conflict causality.

The following section (Sec. 2) outlines our methodological pipeline consisting of our choice of cyclist model, the design of a riding experiment, data processing, control parameter identification, predictive modeling, and our testing procedures. Section 3 shows the experimental, identification, modeling, and testing results, while Section 5 discusses their implications, and Section 6 provides a conclusion.

2. Method

We aim to observe, model, and predict cyclist control behavior for simulated safety assessment. We pursue this goal in four steps outlined in Figure 1. This section first describes our cyclist models and then follows the steps indicated by the figure. For notations, we use bold lower-case symbols for vectors and bold capital letters for matrices, tildes ($\tilde{\cdot}$) to indicate measurements, and hats ($\hat{\cdot}$) to indicate commands.

2.1. Cyclist Modeling

For an accurate physical representation of bicycle motion, we need to model the time-varying pose of the upright bicycle and rider. It consists of the roll angle φ , the roll rate $\dot{\varphi}$, the steer angle δ , the steer rate $\dot{\delta}$, and the yaw angle ψ of the bicycle, together forming the lateral dynamics. A mechanical multi-body model of the bicycle describes the relationship between these quantities. A controller emulates balancing and navigation. The control input is the desired heading or commanded yaw angle $u(t) = \hat{\psi}(t)$, enabling us to expose reactive bicycle control below the tactical planning and navigation levels. The lateral cyclist state $x_{\text{lat}}(t) = [\varphi(t), \delta(t), \dot{\varphi}(t), \dot{\delta}(t), \psi(t)]^T$ is the model output. To focus on lateral dynamics, we do not model longitudinal speed modulation and instruct experiment participants

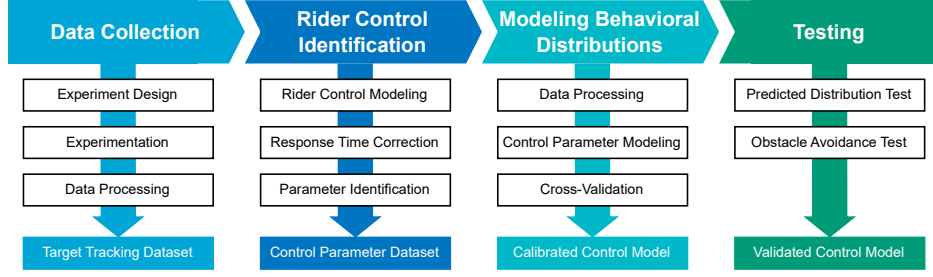


Figure 1: Calibration and testing pipeline for the bicycle-rider control models of this study.

to ride at a constant speed. To account for the strong speed dependency of the models we include small speed variations by using the measured instantaneous speed $v(t) = \dot{v}(t)$ as input. We argue that speed modulation is a mostly tactical element that cyclists use to achieve safe riding according to environmental conditions and otherwise prefer to maintain their momentum. While cyclists use speed modulation for evasive actions in severe conflicts (Liang et al., 2021) or in lane-based traffic among motor vehicles (Johnson et al., 2010), speeds remain nearly constant during moderate conflicts (Liang et al., 2021) and in unstructured environments that provide sufficient space (Marbus, 2025; Zhang et al., 2017), justifying the constant speed choice.

The desired heading input ensures that the model can replace the generic dynamics of a wide variety of traffic simulation frameworks. Examples are social force models (Kaths, 2023; Schmidt et al., 2024), game-theoretical models (Hoogendoorn et al., 2021), or even the lane-based traffic flow simulators, where the desired heading may be derived from the lane centerline. To analyze the impact of bicycle dynamics, we compare the dynamic model, hereafter referred to as the Balancing Rider model (BR), to an established planar model, hereafter referred to as the Planar Point model (PP).

To calculate a path and describe the forward motion, we integrate the heading and instantaneous speed:

$$\dot{p}_x = v(t) \cos(\psi(t)) \quad (1)$$

$$\dot{p}_y = v(t) \sin(\psi(t)). \quad (2)$$

Here, $[p_x, p_y]^T$ describes the position of the rear-wheel contact point in the xy -plane of a reference frame fixed to the road surface.

Balancing Rider Model

Our previous work (Schmidt et al., 2024) models the roll of the bicycle using a simple inverted pendulum bicycle model and shows an example where considering the vertical stability of the bicycle impacts simulated trajectories. This insight motivates us to choose a state-of-the-art bicycle dynamics model for our present study. The linearized Carvallo-Whipple model (Meijaard et al., 2007) is the simplest model that captures all significant characteristics, including self-stability and countersteering. It considers the bicycle and rider as rigid bodies in three-dimensional space and has become the conventional method for describing the lateral dynamics of a bicycle. We hypothesize that these characteristics are essential for a meaningful simulation of conflicts. Observations of cyclists in natural traffic reveal little to no upper-body lean (Kooijman et al., 2009), and marginal impacts of flexible upper bodies on the rider-bicycle open-loop dynamics (Schwab et al., 2008) justify fixing the rider rigidly to the bicycle. The model's state-space formulation (Meijaard et al., 2007; Moore, 2012) is

$$\dot{\mathbf{x}}_{\text{lat}}(t) = \begin{bmatrix} 0_{2 \times 2} & 1_{2 \times 2} & 0 \\ -\mathbf{M}^{-1}(g\mathbf{K}_0 + v(t)^2\mathbf{K}_2) & -\mathbf{M}^{-1}v(t)\mathbf{C}_1 & 0 \\ 0 & v(t)l^{-1}\cos\lambda & 0 \end{bmatrix} \mathbf{x}_{\text{lat}}(t) + \begin{bmatrix} 0_{2 \times 2} \\ \mathbf{M}^{-1} \\ 0 \end{bmatrix} \mathbf{u}_{\text{lat}}(t). \quad (3)$$

Refer to Meijaard et al. (2007) for definitions of the physical rider and bicycle parameters $\mathbf{M}, \mathbf{K}_0, \mathbf{K}_2, \mathbf{C}_1 \in \mathbb{R}^{2 \times 2}$, and λ, c, l . Realizing that steering is the predominant control action (Sharp, 2008; Schwab and Meijaard, 2013), we only allow steer torque input $\mathbf{u}_{\text{lat}}(t) = [T_\delta(t) \ 0]^T$ to limit the complexity of the controller.




Figure 2: The Balancing Rider model (BR) model consisting of the multi-body Carvallo-Whipple bicycle dynamics and a full-state feedback controller (a) and the Planar Point (PP) model (b) as a comparison baseline. All quantities are functions of time.

To mimic the control exercised by the human rider, we implement a full-state feedback controller (Figure 2a) to control steer torque with the gains $\mathbf{K}_x = [k_\phi, k_\delta, k_{\dot{\phi}}, k_{\dot{\delta}}, k_\psi]$. By choosing suitable full-state feedback gains, the system dynamics can be modified such that its closed loop poles move to any desired location for any controllable speed of interest. This gives us ample flexibility to emulate the observed rider control behavior.

For the scope of traffic simulation, we name the combination of Carvallo-Whipple dynamics and full-state feedback control the Balancing Rider model (BR) and describe it as *dynamic* to highlight that it considers forces, masses, and inertia. Simulations with this model require finding solutions to the differential equations, which can be done efficiently using numerical integration.

Planar Point Model

A planar point model, a common choice in various traffic simulation frameworks such as Hoogendoorn et al. (2021) or Kathis (2023), serves as the baseline for our model comparison. It describes road users as massless points on the road surface and mimics bicycle dynamics by separating forward motion and turning. Typically, the yaw rate calculation involves the difference between the desired and the instantaneous yaw, modulated by a time constant. Such models are first-order heading trackers, closely related to the original social force formulation (Helbing and Molnár, 1995). For consistency with the Balancing Rider model, we substitute the time constant with the gain parameter k_ψ , leading to

$$\dot{\psi}(t) = k_\psi(\psi(t) - \hat{\psi}(t)), \quad (4)$$

with the lateral state $x_{\text{lat}}(t) = \psi(t)$ (Figure 2b). Refer to Appendix A for the derivation of this model from Helbing and Molnar's original and further justification of why this model is a suitable baseline.

2.2. Experiment Design

Both models are parameterized with gains representing the rider control behavior that must be calibrated to enable predictive simulation of traffic conflicts. Additionally, the BR model requires physical measurements of the bicycle and rider (e.g., dimensions and inertial properties). To collect data for identifying behavioral parameters and calibrating our models, we design a controlled experiment.

Participants

We recruited participants through advertisements in the university environment. Ten agreed to partake via a written consent form. The TU Delft Human Research Ethics Committee approved our experiment (No. 4689). The participants (4 women, 6 men) had an average age of 29.0 ($\sigma = 3.9$). Participants cycled an average of 4.0 h ($\sigma = 1.7$) per week for commuting (8), general mobility (7), and leisure (7). The group included six experienced cyclists (>10 years experience) and four cyclists with low to medium experience (2-10 years) regularly riding city bicycles (9), trekking bicycles (3), race bicycles (2), city e-bicycles (2), as well as cargo and mountain bicycles (both 1). The group is diverse in terms of cycling experience and habits, and is accustomed to various bicycle types. We expect this to manifest in varying cycling styles within our experiment, which our predictive model must subsequently reproduce.

Test Area and Run Plan

In our experiment, five lights at the end of a 25 m test track indicate a commanded destination and thereby a commanded heading $\hat{\psi}(t)$ that depends on the current cyclist position and the position of the active light (Figs. 3c and 3d). We instruct subjects to cycle towards the active light at fixed speeds and change their heading in accordance with light changes. This mimics step responses in the desired heading. To prevent anticipating commands, the active lights change to random positions, at random frequencies sampled uniformly from $[f_c - \Delta f, f_c + \Delta f]$ around the command frequency f_c with jitter Δf . Following our pilot trials, we chose two different command frequencies, corresponding to around 3 s between changes in the desired heading and 1.5 s. The lower frequency enables complete command execution without time pressure. The higher frequency is close to what we perceived as possible without requiring extraordinary riding skills or dangerous maneuvers. The step amplitude depends on the light positions, the cyclist's position, and their orientation and varies between 0 and 50°. For safety, participants were instructed to ride normally as they would in regular traffic, prioritize bicycle stabilization over executing commands, and abort a maneuver if they fear a loss of control. We assured participants their performance would not be scored or ranked to limit task-induced pressure. Participants had to be adults, fit to cycle, and have experience riding bicycles in public traffic. Three different speed commands (8, 11, 14 km h⁻¹) were communicated verbally before each run, and participants could use a handlebar speed display to check their speed. We instructed them to accelerate before entering the evaluation area and approximately maintain the speed in the area without checking the display to avoid missing heading commands.

Participants completed each combination of speed and command frequency three times, resulting in 18 runs in two blocks with a break in between (Tab. 1). After traversing the evaluation area of the track marked by traffic cones (Figure 3a), participants circled to begin a new run.

Table 1: Run plan for 18 runs with different command frequencies and speeds executed by each participant.

Block 1				Block 2			
Run #	f_c [Hz]	Δf [Hz]	\hat{v} [km/h]	Run #	f_c [Hz]	Δf [Hz]	\hat{v} [km/h]
1 - 3	0.33	0.2	8	10 - 12	0.66	0.2	8
4 - 6	0.33	0.2	11	12 - 15	0.66	0.2	11
7 - 9	0.33	0.2	14	16 - 18	0.66	0.2	14

We conducted the experiment during two dry days across September and October 2024 to collect data for rider control identification. The anonymized data is publicly available at DOI: [10.4121/f881dd80-b9f5-4322-9fd5-192034c9717f](https://doi.org/10.4121/f881dd80-b9f5-4322-9fd5-192034c9717f). We collected data for 188 runs across all participants. The exact number of runs deviated from the target because of counting errors.

Instrumented Bicycle, Command Lights and Rider Parameters

All participants rode a prototype Grenoble C8 HMB electric bicycle (Royal Dutch Gazelle, Dieren, The Netherlands) for which we measured the dimensions, mass, and inertia following the approach outlined by Moore (2012) and the corresponding BicycleParameters software package¹. The bicycle was developed in a previous project (Alizadehsaravi and Moore, 2023) and has pedal support as well as a balance assist system, both of which were deactivated and had no effect on our test. It is equipped with an inertial measurement unit (IMU), a steer encoder, and a speedometer to measure the speed, roll angle, and steer angle at a rate of 100 Hz. We additionally mounted a satellite positioning unit (SwiftNav Piksi Multi RTK GNSS²) on the rack (Figure 3b). A nearby base station of the Dutch permanent GNSS Array³ enabled us to correct measurements for atmospheric distortion, leading to centimeter-accuracy measurements at 10 Hz. An Arduino controlled the heading commands for which a second GNSS unit supplied high-accuracy global timestamps.

We did not collect the rider's mass, height, and body inertia to keep the experiment execution feasible. Instead, we use the publicly available properties of the Jason benchmark rider provided as part of the BicycleParameters package and rescale them linearly to represent an average human. We roughly estimate the average human height and weight to 166 cm and 69 kg (incl. 1 kg for clothes) from global data on the human height (Roser et al., 2021) and BMI (Ritchie and Roser, 2017).

¹<https://bicycleparameters.readthedocs.io>

²https://www.swiftnav.com/sites/default/files/piksi_multi_product_summary.pdf

³<https://gnss1.tudelft.nl/dpga/>

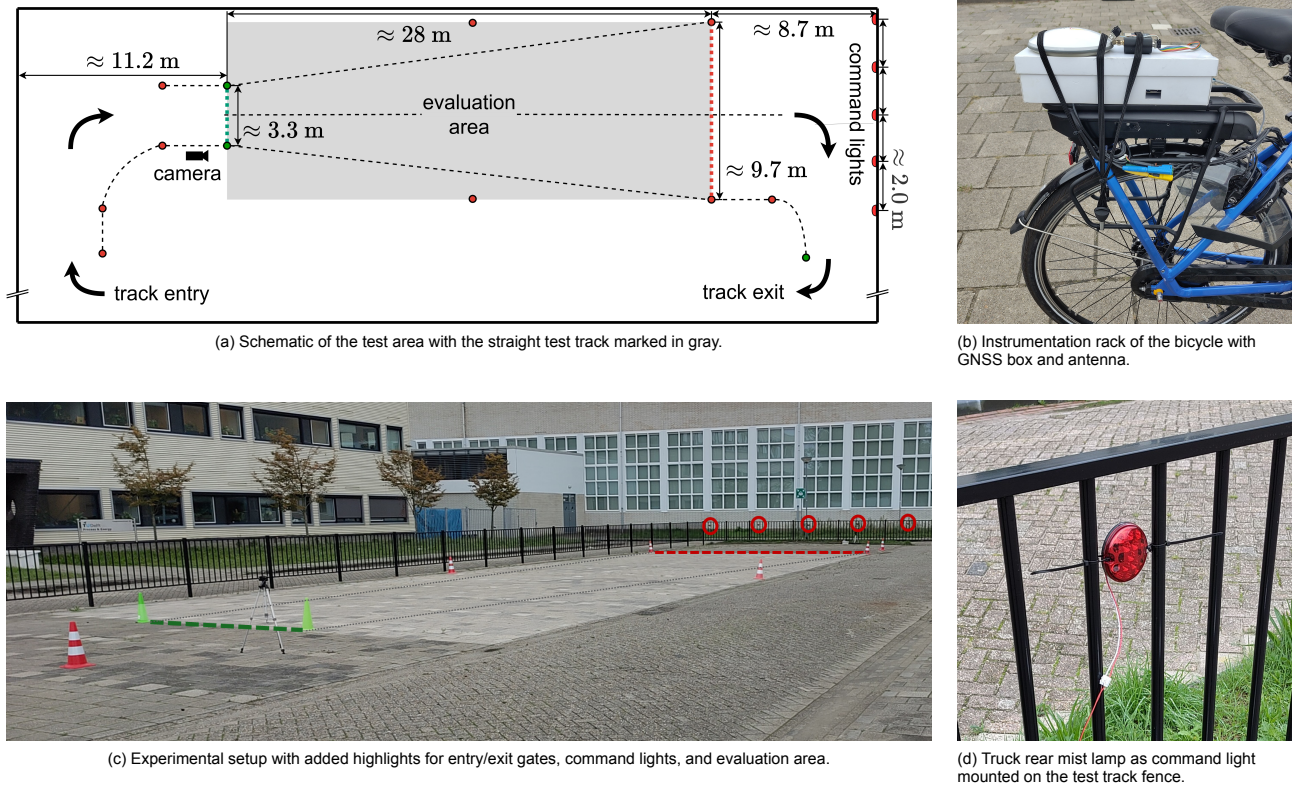


Figure 3: The experimental setup on a converted parking lot on Delft University of Technology Campus.

2.3. Data Processing

After acquisition, we process the raw data to prepare samples for rider-control parameter identification (Figure 4). Following GNSS correction, identical motion cues in the bicycle kinematics enable time synchronization between the bicycle IMU and the GNSS based on total acceleration. The light commands receive high-accuracy GNSS timestamps during recording and do not need further alignment. We then apply an Unscented Kalman Filter⁴ to fuse and filter GNSS and bicycle data. The filter uses the previously mentioned Carvallo-Whipple model at the instantaneous velocity, assuming zero steer torque and constant acceleration. The known coordinates of the evaluation area enable us to extract individual runs from experiment trajectories.

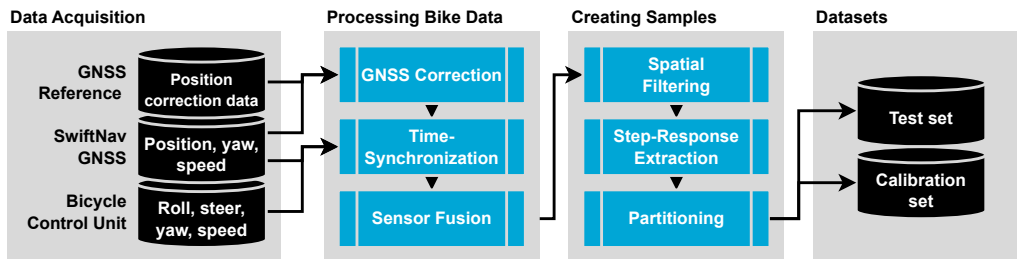



Figure 4: Data postprocessing steps (blue) with raw data from our measurement devices and resulting data sets in black.

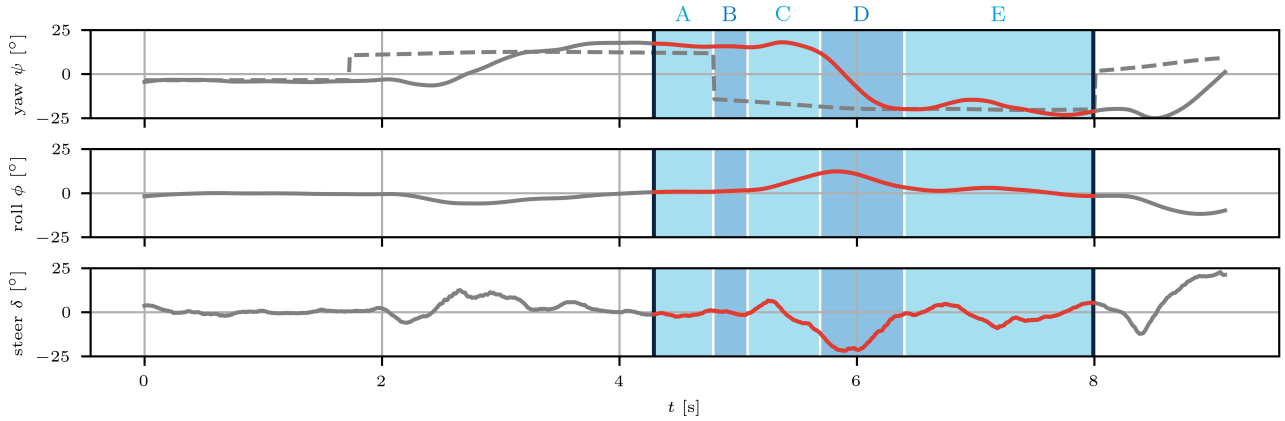
Responses to individual commands differ in response time and transient characteristics. Hence, we further extract the individual step responses from each run (Figure 5). A step sample starts 0.5 s before the command and ends at the following command. The period before the step increases the robustness of the parameter identification because it enables the optimizer to compensate for small measurement errors at the time of the command. For some samples of the high-frequency runs in block two, the time between commands is shorter than 2 s. For these, we include two commands and force a minimum sample duration of 2.5 s to incentivize stable results. To limit the impact of the second command on the identification, we only retain samples where the first change in commanded yaw satis-

⁴<https://filterpy.readthedocs.io>

fies $\Delta\psi_0 \geq 8^\circ$ and the second satisfies $\Delta\psi_1 < 8^\circ$. Further, we filter out samples with variations between maximum and minimum speed $\Delta v > 3.0 \text{ km h}^{-1}$ because we expect larger variations to distort the identification of single parameter sets per command as a result of the speed-dependent dynamics (See Eq. 2.1.1). We then randomly partition the step response samples into a training and a test dataset using an 80/20 split. In total, we extract 218 step responses from all runs, which leads to 173 samples (79.4 %) in the training dataset and 45 samples (20.6 %) in the test dataset.



(a) Trajectory on the road surface.



(b) Lateral dynamics.

Figure 5: Example of a run across our evaluation area (gray line) with a step-response extracted from the full run (marked red segment). The dashed line shows the heading command, derived from the current position and the light location. The step response consists of five phases representing the initial conditions before the command (A), the response delay between command and the first rider action (B), the countersteer motion (C), the heading correction (D), and the settling phase with oscillations for stabilization (E).

2.4. Rider Control Identification

Next, we identify the missing gain parameters of our rider-control model for all step responses in the test and the training datasets. The fixed-in-space command lights create destination commands that must be converted into heading commands. We close the position loop and derive the commanded heading $\hat{\psi}(t)$ from the command light location $\hat{\mathbf{p}}(t) = [\hat{p}_x(t), \hat{p}_y(t)]^T$ to embed our rider-bicycle control models with heading-input into the experimental conditions (Figure 6):

$$\hat{\psi}(t) = \arctan\left(\frac{p_y(t) - \hat{p}_y(t)}{p_x(t) - \hat{p}_x(t)}\right) \quad (5)$$

This augments the total system state to $\mathbf{x}(t) = [x_{\text{lat}}(t), p_x(t), p_y(t)]$ with the total system dynamics $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), \hat{\mathbf{p}}(t), v(t))$ becoming a nested function of the lateral dynamics, the forward motion, and the experimental setup. To simulate synthetic data or predictions, we apply midpoint integration using the timestep $\Delta t = 0.01 \text{ s}$.




Figure 6: The dynamic system used for control parameter identification with the rider control model embedded into a position feedback loop.

Response Time Estimation

Human responses to visual input in traffic exhibit considerable delay times (Jurecki and Stańczyk, 2014). To prevent distorting the fit of the model to the measurements, we compensate the response time of each step by shifting the destination commands $\hat{p}'_{x,y}(t) = \hat{p}_{x,y}(t - \tau)$. For the Balancing Rider model, we define the response time $\tau := \tau_{cs}$ as the onset of the countersteer motion in the measured response when first $\delta(t) < \epsilon_{\delta} = 22.9^{\circ} s^{-1}$. Since the Planar Point model does not capture countersteering and immediately corrects yaw in the right direction, using the same delay would deteriorate the performance of the Planar Point model (Figure 7b). Setting $\tau := \tau_{ec}$ to the moment when the current heading exceeds the previous command $\psi(t) > \hat{\psi}_{prev}$ leads to a fair comparison. To avoid detecting actions other than the response, we limit the search to $\tau_{cs} \in [0.1 \text{ s}, 0.6 \text{ s}]$ and $\tau_{ec} \in [0.4 \text{ s}, 1.5 \text{ s}]$ and assign the median of all other results to samples that do not pass the threshold within the search range.

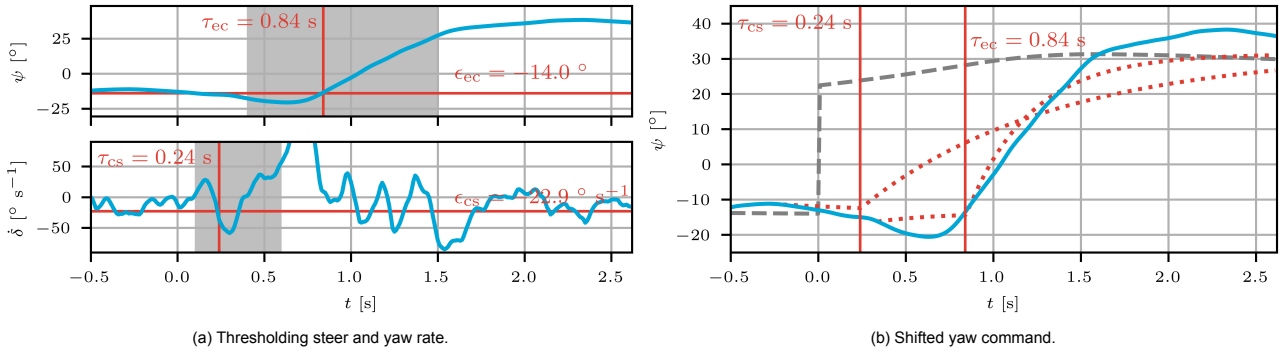


Figure 7: Example of response delay estimation by thresholding the steer rate to detect the countersteer onset τ_{cs} and by thresholding the yaw angle to detect the moment τ_{ec} when the heading is corrected. Shaded areas mark the search range for response delays. The right figure illustrates the effect of different command shifts for the Planar Point model (dotted red) fitted to the measured yaw (blue). A delay to the countersteer onset underestimates the rise time, while a delay to the heading correction leads to a better fit.

Parameter Identification

To determine the behavioral parameters and calibrate our model, we identify a set of gains for each individual step response. Separate identification for each step response enables compensation for the step-specific response time and accounts for varying behaviors between steps. The desired gains K_x are those that best reproduce the measured trajectories when exciting the model with the response-time-corrected input. To identify those gains, we formulate a non-linear optimization problem constrained by the system dynamics:

$$\begin{aligned}
 & \min_{\mathbf{K}_x} J(\mathbf{y}(t), \tilde{\mathbf{y}}(t)) \\
 \text{subject to } & \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t), v(t), \mathbf{K}_x, \hat{\mathbf{p}}(t)) && \text{(dynamics)} \\
 & \mathbf{y}(t) = \mathbf{C}_{\text{trk}} \mathbf{x}(t) && \text{(tracked states)} \\
 & k^{\min} \leq k \leq k^{\max} \quad \forall k \in \{k_\varphi, k_\delta, k_{\dot{\varphi}}, k_{\dot{\delta}}, k_\psi\} && \text{(gain bounds)} \\
 & p_x(0) = \tilde{p}_x(0), p_y(0) = \tilde{p}_y(0), \psi(0) = \tilde{\psi}(0) && \text{(initial conditions)} \\
 & \hat{p}_x(t) = \hat{\tilde{p}}_x(t), \hat{p}_y(t) = \hat{\tilde{p}}_y(t), v(t) = \tilde{v}(t) && \text{(command and speed measurements)}
 \end{aligned}$$

Here, T is the step response duration. $k^{(\cdot)}$ denotes the upper / lower bound of a gain k . The matrix \mathbf{C}_{trk} is chosen such that it selects the tracked states $\mathbf{y}(t)$ from the state vector $\mathbf{x}(t)$. Refer to Table 2 for detailed definitions for the individual models. The objective function is

$$J(\mathbf{y}(t), \tilde{\mathbf{y}}(t)) = \int_0^T \|\mathbf{y}(t) - \tilde{\mathbf{y}}(t)\|^2 dt \quad (6)$$

and we denote J^* for the objective value of a solution. We solve the problem using `opty`⁵, a Python package for trajectory optimization based on direct collocation (Moore and Bogert, 2018), that applies the midpoint method to discretize the problem and interfaces the `Ipopt` solver (Wächter and Biegler, 2006). We provide feasible and dynamically stable initial guesses to improve convergence. Multi-start optimization with n_{guesses} different initial guesses is deployed to minimize the risk of obtaining a local optimum.

The complete full-state feedback variant of the Balancing Rider model (BR_0) has five gain parameters. While this enables arbitrary locations of the dynamic poles, it could exceed human capabilities. Reducing the number of parameters leads to smaller complexity and less susceptibility to overfitting. Hess et al. (2012) introduce a bio-mechanical rider control model considering visual, vestibular, and proprioceptive feedback loops represented by roll angle, roll rate, steer angle, and yaw angle feedback. Analogously, we test a model variant (BR_1) where we drop k_δ . Additionally, we limit the gains to a fixed sign based on the signs of the full-state feedback optimal controller designed by Schwab et al. (2008).

Table 2: Control parameter identification configuration for the different dynamic models

	BR_0	BR_1	PP
lateral dynamics	Carvallo-Whipple (Eq. 3)	Carvallo-Whipple (Eq. 3)	Planar Point (Eq. 4)
state vector $\mathbf{x}(t)$	$[\varphi, \delta, \dot{\varphi}, \dot{\delta}, \psi, p_x, p_y]^T$	$[\varphi, \delta, \dot{\varphi}, \dot{\delta}, \psi, p_x, p_y]^T$	$[\psi, p_x, p_y]^T$
number of states	7	7	3
tracked states $\mathbf{y}(t)$	$[\varphi, \delta, \psi]^T$	$[\varphi, \delta, \psi]^T$	ψ
identified parameters	$k_\varphi, k_\delta, k_{\dot{\varphi}}, k_{\dot{\delta}}, k_\psi$	$k_\varphi, k_\delta, k_{\dot{\varphi}}, k_{\dot{\delta}}$	k_ψ
fixed parameters	-	$k_\delta = 0$	-
gain bounds	$k_{\varphi, \delta, \dot{\varphi}, \dot{\delta}, \psi} \in [-250, 250]$	$k_\varphi \in [-250, 0]$ $k_{\delta, \psi} \in [0, 250]$ $k_{\dot{\varphi}} \in [-250, 250]$	$k_\psi \in [0, 20]$
n_{guesses}	30	30	10
response time	countersteer onset τ_{CS}	countersteer onset τ_{CS}	heading error correction τ_{EC}

2.5. Modeling Behavioral Distributions

Rider control identification on all step responses returns 218 gain value combinations. From these, we design a stochastic model that enables drawing synthetic gain combinations to predict cyclist behavior in new scenes. However, the effect of specific gain values on

⁵<https://github.com/csu-hmc/opty>

the system dynamics is not straightforward, and small changes may lead to large trajectory differences. A more interpretable alternative for modeling control behaviors are the closed-loop poles of the system (i.e., the zeros of the characteristic polynomial). Since our rider-bicycle models are linear full-state feedback models, there is a unique bi-directional mapping between gains and poles. The pole placement technique finds the gains corresponding to a set of poles. For the models with Carvallo-Whipple dynamics, we use the mean speed of a step response for this conversion.

Data Preparation

We take several steps to prepare the identification results for predictive modeling (Figure 8). Firstly, we extract the poles from the characteristic polynomial of the rider-bicycle control models for all 218 gain sets. This results in five poles with two pairs of complex poles and one real-valued pole for each set of identified gains of the Balancing Rider models. A small number of identification results with different pole configurations are excluded as outliers. We sort the poles $s_i = \sigma_i \pm j\omega_i$ according to their complex frequency so that $|\Im(s_i)| < |\Im(s_{i+1})|$. We obtain $n_s = 5$ pole features $s_{BR} = [\sigma_0, \sigma_1, \omega_1, \sigma_2, \omega_2]^T$ for pole modeling. The Planar Point model has $n_s = 1$ real-valued pole which leads to $s_{PP} = [\sigma_0]$




Figure 8: Data preparation for predictive pole modeling from the identification results.

Secondly, identification returns results for each step response, but not all are robust representations of the observed behavior. For example, if the response is slow and the sample duration is short, the available data may not include the decaying nature of the transient response. As a result, the optimizer may favor unstable or slowly decaying oscillations over robust solutions if they fit the initial response well. Moreover, if the participant performed steering actions that did not correspond to the command, the measurements do not fit the input and the identified parameters are distorted. This gets amplified if the threshold-based response time identification catches on the wrong steering input. Finally, in exceptional cases, the optimizer can choose vanishing magnitudes of the heading gain k_ψ and reproduce the transient behavior exclusively based on roll and steer. For predictive simulation, the existence of k_ψ is necessary to ensure that the cyclist reacts to an input. To treat the outliers in both the calibration and test set, we filter the results to only include those satisfying conditions on the objective value (BR: $J^* < 2 \times 10^{-4}$, PP: $J^* < 5 \times 10^{-3}$), stability ($\sigma_{0,1,2} < 0$), and gains not hitting the identification bounds $|k - k_{\min, \max}| < 0.01$ (only considering non-zero bounds except for k_ψ). Additionally, we simulate the step responses of the system using the identified parameters of all models to a large $\hat{\psi} = 45^\circ$ command step and remove samples for which the yaw angle and lateral deviation exceed three standard deviations of the overall distribution. Applying all filters results in a split of 123 calibration and 31 test samples.

Thirdly, we apply a non-linear transformation to the real parts to focus on the important poles close to the imaginary axis. These decay slowly and strongly impact the transient behavior. However, their real part σ_i varies within a range that is small compared to poles with the largest negative real parts. The log-shift transformation stretches the distribution of small negative σ_i out and condenses the distribution of the less important large negative σ_i .

$$\sigma_i^* = \log(a_i - \sigma_i) \quad (7)$$

We choose a_i to be 0.9 times the largest observed value for each σ_i . We then apply the Yeo-Johnson power transformation⁶ to improve normality and symmetry of the feature distributions (Yeo and Johnson, 2000). After sampling predictions, the inverse transformations must be applied again. Then, the shift parameter results in the upper limit of the sampled poles $\sigma_i < a$, guaranteeing asymptotic stability of the rider-bicycle control model.

⁶<https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.PowerTransformer.html>

Multivariate Gaussian Mixture Model

The distribution of poles, which will be discussed in detail in the following section (Fig. 17), shows multiple modes, and individual features are heavily correlated. Multivariate Gaussian Mixture Models (GMMs) are well suited to capture these multiple modes and correlations (Bishop, 2006, ch. 2.3.9). Hence, the distribution of the random vector \mathcal{S} representing the n_s pole features can be described as a GMM:

$$p(\mathcal{S}) = \sum_{q=0}^Q \pi_q \mathcal{N}(\mathcal{S} | \boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q). \quad (8)$$

$\boldsymbol{\mu}_q \in \mathbb{R}^{n_s}$, $\boldsymbol{\Sigma}_q \in \mathbb{R}^{n_s \times n_s}$, and π_q are the mean, covariance and mixing coefficient of the q -th component in the mixture model such that $\sum \pi_q = 1$. A sufficiently high number of components Q enables describing arbitrary distributions. We estimate this distribution for the set of identified pole features \mathcal{S}_{BR0} , \mathcal{S}_{BR1} , and \mathcal{S}_{PP} of each of our three models.

Previous works have shown a speed dependency of the parameters of rider-bicycle control models (Moore, 2012; Reijne, 2025). To account for this dependency, we estimate three additional models with augmented feature vectors including speed. For this, we add the average measured speed of each step response to the previously introduced pole features of each step response, resulting in $\mathbf{s}^* = [\bar{v}, \mathbf{s}^T]^T$ and the corresponding random vector \mathcal{S}^* . To predict new pole feature samples for given speeds, one must find the conditional distribution $p(\mathcal{S} | v)$ from the distribution of the augmented Gaussian mixture model $p(\mathcal{S}^*)$:

$$p(\mathcal{S} | v) = \frac{p(\mathcal{S}, v)}{p(v)} = \frac{p(\mathcal{S}^*)}{\int p(\mathcal{S}^*) d\mathbf{s}}. \quad (9)$$

Chehade et al. (2022) show that this distribution is again a Gaussian mixture model whose moments can be directly inferred from $p(\mathcal{S}^*)$ given samples of \bar{v} and provide the necessary calculations. This results in three additional pole distributions based on the sets of augmented features \mathcal{S}_{BR0}^* , \mathcal{S}_{BR1}^* , and \mathcal{S}_{PM}^* of each of our three models.

Calibration and Cross-Validation

We fit the stochastic model using Scikit-Learn's implementation of the expectation-maximization algorithm⁷ and perform grid-search to find the optimal number of components Q . The performance metric is the mean negative log-likelihood (NLL) of the training datasets $\mathcal{S}_{BR0, BR1, PP}^{\text{train}}$ under the fitted distribution, where the GMM-NLL is defined by

$$\text{GMM-NLL}(\mathcal{S}) = -\frac{1}{M} \sum_{m=0}^M \log(p(\mathcal{S} = s_m | \boldsymbol{\mu}_0, \dots, \boldsymbol{\mu}_Q, \boldsymbol{\Sigma}_0, \dots, \boldsymbol{\Sigma}_Q, \boldsymbol{\pi}_0, \dots, \boldsymbol{\mu}_Q,)) \quad (10)$$

with $M = |\mathcal{S}|$ and s_m is the m -th sample in the set \mathcal{S} . To prevent overfitting to a limited number of command step samples, we employ a 10-fold cross-validation scheme and select the parameters with the lowest NLL for testing. The final predictive model is calibrated on the full calibration dataset.

2.6. Testing

We perform two tests outlined below to explore the predictive power of our calibrated rider-control models and investigate the impact of bicycle dynamics on traffic processes.

Predicted Variance Test

Our model can predict a distribution of trajectories based on the distribution of behavioral parameters for a given scene, defined by some initial conditions and input trajectories. To demonstrate the predictive power of our model, we generate 1000 trajectory predictions per scene in the test dataset and separately test whether our model accurately predicts the variance of this distribution and whether it accurately predicts individual measured test trajectories.

⁷<https://scikit-learn.org/stable/modules/generated/sklearn.mixture.GaussianMixture.html>

Each test trajectory represents only one sample of the underlying true response distribution of the trajectory's scene. Hence, the target variance for our model per scene is not known. In a single scene, the measured sample may correctly appear at the tail of the predicted distribution because the measurement could be an unlikely sample of the true distribution. However, over the whole test data, these edge cases should be rare. The negative log-likelihood is an established metric for trajectory prediction (Rudenko et al., 2020) which considers this trade-off by measuring how likely the observed data is, assuming it was generated by the predicted distribution for each scene. To evaluate the NLL, we take the lateral deviation $d[n]$, $\forall n$, $n < N_m$, $nT \in \{1.0 \text{ s}, 1.25 \text{ s}, \dots, 2.0 \text{ s}\}$ of each predicted trajectory in the m -th scene at five points. Figure 9 explains the calculation of the lateral deviation metric. Denoting $\mathbf{d} = [d[n_0], d[n_1], \dots, d[n_H]]$ and the corresponding random vector \mathbf{D} , we can estimate the distribution of predicted lateral deviations $p_m(\mathbf{D})$ for each scene using SciPy's Kernel Density Estimation⁸. The LD-NLL of scene m is the likelihood of observing the measured lateral deviation $\tilde{\mathbf{d}}_m$ of the test trajectory under the predicted distribution, given by

$$\text{LD-NLL}_m = -\log(p_m(\mathbf{D} = \tilde{\mathbf{d}}_m)). \quad (11)$$

The final metric is the median over all scenes. Note that our calculation differs from the popular NLL introduced by Ivanovic and Pavone (2019), who estimate individual distributions for a larger number of time steps rather than one joint distribution per scene. Our approach additionally considers the correlation between time steps. For reference, we also report their approach (referred to as T-NLL) using the median within the prediction interval $t \in [\tau_{cs} + 0.5 \text{ s}, \tau_{cs} + 2.0 \text{ s}]$. We disregard the first 0.5 seconds after the response, where the predicted distributions are very narrow. This period would result in extremely large NLL values, even if measurement and prediction show negligible errors. The final T-NLL result is the median over all test scenes.




Figure 9: The lateral deviation metric (LD) describes the distance from the line connecting the cyclist position at the time of the counter-steer onset τ_{cs} and the command light position. For all turn directions, deviations to the outside are assigned negative values. Red ticks mark the evaluation points for the features of the LD-NLL metric.

To test the prediction accuracy with respect to individual measured trajectories, we record the single best Average Displacement Error of all predicted trajectories (1-minADE) per scene, calculated as the average Euclidean distance between predicted and measured position trajectories (Rudenko et al., 2020), and report the median over all scenes.

Obstacle Avoidance Test

We design a simulated obstacle avoidance test to demonstrate the impact of the Balancing Rider model compared to the Planar Point baseline in a setting close to traffic safety simulation scenarios. A simulated cyclist follows a destination sequence reflecting an obstacle avoidance maneuver. We measure the Time-to-Collision (TTC) between the cyclist and the obstacle using a diamond-shaped footprint for the cyclist and linear extrapolation of the current position with fixed heading and speed.

3. Results

After data collection, we apply the processing pipeline outlined in Section 2 to analyze the observed behavior, calibrate our control behavior model, and test its predictive power. Detailed lists of parameter values required to use our models for further research are available in Appendix B.2.

⁸https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.gaussian_kde.html

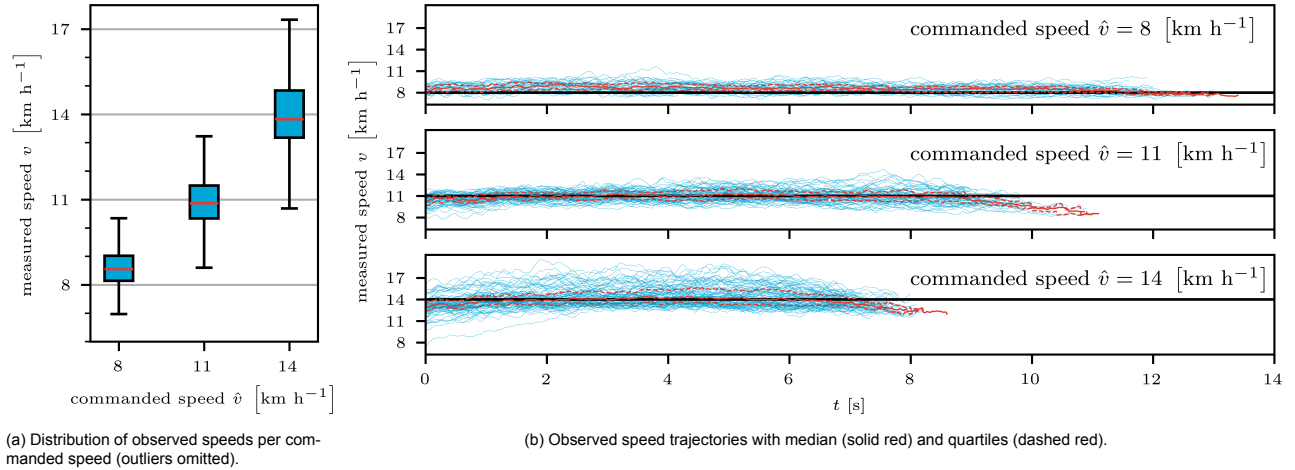


Figure 10: Measured speed trajectories and distributions for runs over the entire evaluation area.

3.1. Measured Behaviors

The experiment enables various observations on cyclist control behavior, as described below.

Measured Cycling Speeds

Figure 10 shows that most participants executed the instructed speeds accurately and did not significantly accelerate or decelerate during runs. Small drops in speed during each run may be due to the approaching turn directly after the end of the evaluation area. Two participants consistently and significantly exceeded the commands and accelerated during the run. We attribute most of their runs to the fastest commanded speed category, regardless of the actual command, to prevent distortion of the speed categories. Since their trajectories still cover relevant speed ranges, we include them in the analysis.

Measured Lateral Dynamics

The riding task of our experiment exposes different control styles among the riders. Figure 11 shows examples of two runs of different riders for both command frequencies. The blue rider performs sharp and quick heading corrections, while the purple rider employs a smooth strategy that creates wide bends. Blue arrives at a stationary heading for all commands in the low frequency case with less time pressure (left) and for most cases with shorter intervals between commands (right). The heading of purple, however, strongly overshoots the commands and doesn't settle to the commanded direction for either of the command frequencies. The next command arrives while purple is still cornering, which results in seemingly nonlogical responses where the rider first increases the deviation of their heading from the command after a coincidentally good fit at the time of the command (see step 5, right). This results from the requirement to correct the roll of the steady turn to an upright position. Until the bicycle is upright again, the heading deviation increases, requiring another correction and rolling in the opposite direction. For purple, this process is slower than the time between commands, requiring the same process for the next command and creating oscillations.

The cumulative distributions of the measured lateral dynamics further highlight large individual differences in control behavior (Figure 12a). While the blue and dark blue participants with sharp course corrections mostly exhibit small yaw, and roll angles, the smooth trajectories of the purple and green participants result in more time spent at higher roll angles, higher yaw angles, and higher difference to the commanded direction. Regardless of these trends, control strategies appear to be individual. While green and purple are very similar in terms of roll angle and steer rate, their roll rate and steer angle differ. Some strategies stand out through very small steer rates and angles (dark blue) and large steer angles (light green). The measured steer angle and roll angles change slightly with higher temporal pressure due to less time spent in straight riding (Figure 12b). In contrast, temporal pressure appears to have only a minimal effect on the corresponding rates as well as the yaw error. With faster speed, steer angles get lower and roll angles get higher (Figure 12c). Com-







Figure 11: Examples of full runs across the experimental area for two riders (blue squares and purple circles) with a large difference in control behavior. Numbered markers indicate new heading commands. The dashed lines point in the direction of the heading command. Participants try to follow the indicated command from the marked moment till the next marker. Due to the random command sequences, the displayed runs are responses to different inputs for the two riders. Left shows runs with low time pressure (command frequency $f_c = 0.3$ Hz) and right high time pressure ($f_c = 0.6$ Hz).




(a) Lateral dynamics grouped by participant. Colors represent different riders and match the colors from Figure 11.



(b) Lateral dynamics grouped by command frequency.



(c) Lateral dynamics grouped by commanded speed.



(d) Lateral dynamics grouped by command step height.

Figure 12: Cumulative distributions of the lateral dynamic measurements for all time steps in all step responses grouped for participants, and command frequencies, speeds, and command step heights. This includes steps where the corresponding identification results are classified as outliers (see Section 2.5.1).

pared to the differences between participants, both effects are minor. The effect of command step height on steer and roll, particularly in their rates, is similarly minor compared to the differences between participants (Figure 12d). Large effects of the command step height on the yaw error and angle are a logical consequence of grouping by commanded yaw.




Figure 13: Response time histograms based on countersteer onset for the Balancing Rider models (left) and rising yaw flank for the Planar Point models (right). Red lines show exponentially modified Gaussian distributions fitted to the response time samples. The box shows the distribution parameters.

Command Response Time

Response time estimation based on countersteer onset for the Balancing Rider models and rising yaw flank for the Planar Point model leads to two sets of response times (see Section 2.4.1). Figure 13 shows the results excluding unrealistic detections outside the search range. The countersteer onset lies in a range comparable to response times of car drivers to suddenly appearing obstacles (Jurecki and Stańczyk, 2014). The mean moment of beginning heading correction is 0.69 s later, exposing the delay between the initial reaction of the rider and its effect on the path. Varying response intensities widen the temporal spread of the heading correction compared to the countersteer onset.

3.2. Identified Rider Control Parameters

Identifying control parameters for all samples and all models requires solving the optimization problem 17,440 times. We ran the optimization in parallel on our server (2x Intel Xeon E5-2680 v2, 20C/40T, 128 GB RAM), completing the identification within 8 h 20 m.

The Balancing Rider models achieve an excellent fit exceeding 90% VAF for roll, yaw, and y-position. Model BR₀ outperforms the reduced BR₁ in all metrics (Table 3). The fit across participants varies slightly, with the fits of three of seven participants notably worse than the rest. The variation within models across command frequencies is small, indicating similar good fit for the two different time-pressure scenarios. The Planar Point model scores significantly lower in all metrics across all sample groups, indicating issues reproducing the measured dynamics. As an example, Figure 14 shows a fit with good results for all models. With VAF_φ and VAF_ψ scores over 90 %, the two Balancing Rider models capture the overall dynamics excellently. The lateral dynamics translate to a well-fitting spatial trajectory, although the position is not part of the objective. The Planar Point model cuts the countersteer motion, leading to a small offset in the spatial trajectory while decently capturing the rising flank. Figure 15 shows a challenging example. After beginning with a large residual error from the previous step, the measured step response shows very slow dynamics and extreme overshoot. The full BR₀ model captures these dynamics very well. The reduced BR₁ model performs worse and accumulates errors in steer and roll, indicating that the maneuver may not be feasible with the bounded gains derived from optimal control. However, results still translate to a well-fitting spatial trajectory. The Planar Point model does not fit the overshoot, leading to a path offset of approximately 1 m at the end of the trajectory.

The identified gains of the full model BR₀ are scattered widely, and $k_φ$ and $k_δ$ exist with both signs (Figure 16). Most results still match the theoretical expectation for optimal full-state feedback control shown by Schwab et al. (2008) in terms of sign and align with the steer-into-the-fall principle (negative roll rate gain). The gain on the steer rate $k_δ$ vanishes for the majority of samples, justifying fixing it to zero for model BR₁. With tighter bounds, BR₁ exhibits much denser clusters exposing tendencies of higher gains for lower speeds that can be expected based on previous work (Schwab et al., 2008). The results in Figure 15 represent a case where the identified gain ($k_δ ≈ -94.9 << 0$) of BR₀ does not align with the theoretical results ($k_δ > 0$) of Schwab et al. (2008) and the reduced model fits notably worse. This suggests that not all human control strategies adhere to theoretical optimality.

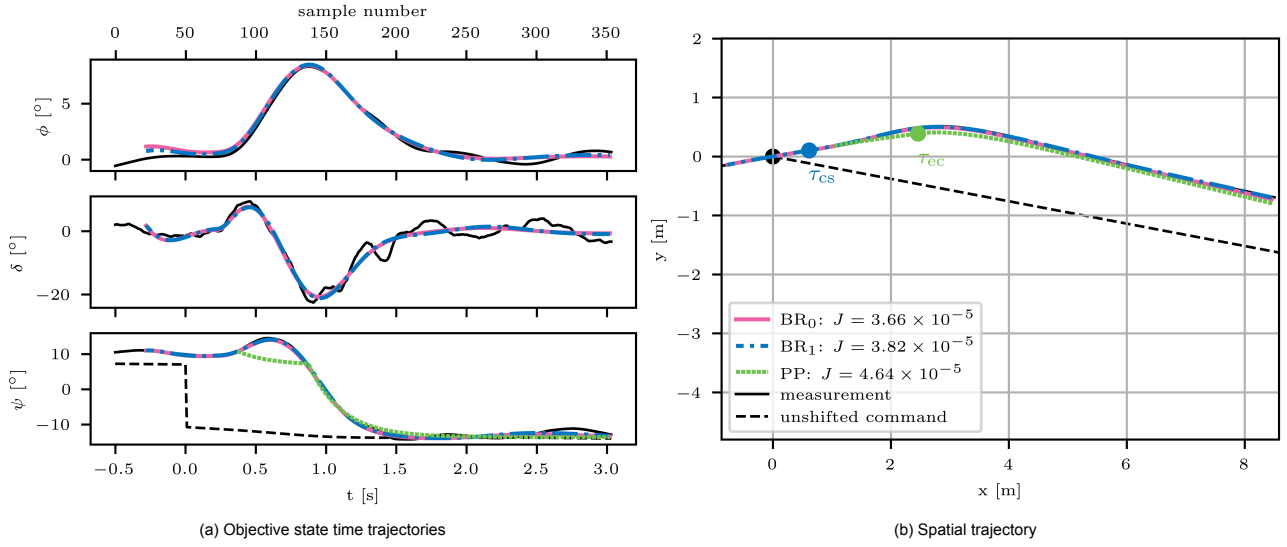


Figure 14: Example of a good identification fit for all models. Filled dots mark the moment of the unshifted command (black) and the shifted commands at τ_{cs} for the Balancing Rider models (dash-dotted blue, solid pink) and τ_{rf} for the Planar Point model (dotted green). J is the objective value und T the step response duration.

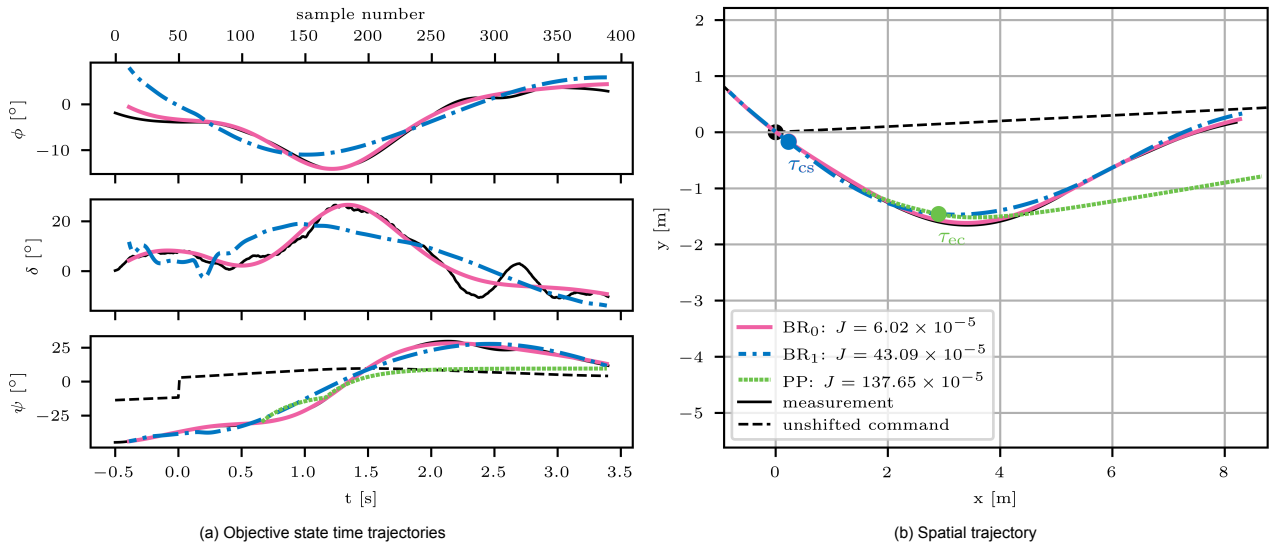


Figure 15: Example of identification fits for a challenging step response with slow dynamics and extreme overshoot. Dots mark the moment of the unshifted command (dashed black) and the shifted commands at τ_{cs} for the Balancing Rider models (dash-dotted blue, solid pink) and τ_{rf} for the Planar Point model (dotted green). J is the objective value und T the step response duration.

3.3. Modeled Behavioral Distributions

After identification, we sort the pole samples generated by the three models by their imaginary part and remove outliers (see Section 2.5.1).

Figure 17 illustrates for the BR₁ model that the imaginary part separates the pole samples into two groups. For some samples, the real-valued poles are located far in the negative half-plane, indicating little influence on the dynamics. The pole distribution of BR₀ is qualitatively similar but stretches far further into the negative half-plane.

We then fit Gaussian mixture models to the pole samples and perform cross-validation to find the optimal number of GMM-components (Table 4). The distributions are best modeled by one to two Gaussian components. Including speed improves the fit for BR₁ and PP. As the speed increases, so do the damping and the complex frequency, resulting in faster responses. The fitted marginals (Figure 18) further reveal that the two Gaussian components mostly separate samples with large and small real-valued poles. The step responses corresponding to the mean poles of each GMM component (Figure 19) reveal different countersteering intensities, rise times, overshoot,

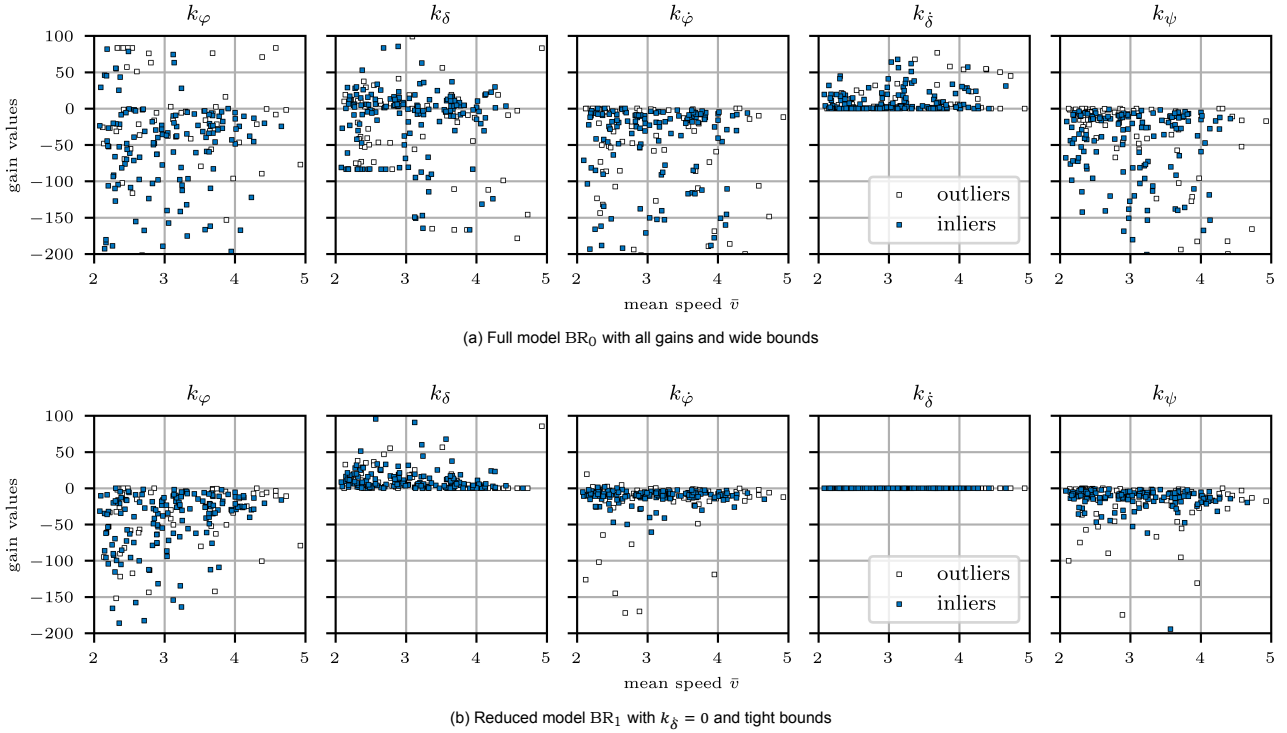


Figure 16: Identified gains versus the mean speed per command response sample for the two different Balancing Rider models. Note that a few samples lie outside the visualized limits of the vertical axis. Empty markers represent results classified as outliers for pole modeling later.

Table 3: Rider-bicycle control model fit after control parameter identification in terms of Variance Accounted For (VAF) and Mean Absolute Error (MAE). We show the results for all samples, give statistics across participants (part.), and group results by command frequency. IQR denotes the Inter-Quartile Range across participants. Best results are highlighted for all samples.

model	samples	objective J	VAF $_{\delta}$ [%]	VAF $_{\phi}$ [%]	VAF $_{\psi}$ [%]	VAF $_{y}$ [%]	MAE $_{\delta}$ [rad]	MAE $_{\phi}$ [rad]	MAE $_{\psi}$ [rad]	MAE $_{y}$ [m]
BR $_0$	all	5.42×10^{-5}	90.8	95.4	99.3	99.5	0.028	0.008	0.011	0.025
BR $_1$	all	7.21×10^{-5}	88.6	93.7	99.1	98.8	0.031	0.010	0.015	0.038
PP	all	54.71×10^{-5}	-	-	82.1	83.2	-	-	0.078	0.125
BR $_0$	part. median	3.69×10^{-5}	92.9	98.0	99.6	99.8	0.025	0.006	0.009	0.018
BR $_0$	part. IQR	2.57×10^{-5}	2.8	0.6	0.2	0.1	0.008	0.004	0.003	0.011
BR $_0$	part. min	1.11×10^{-5}	84.3	95.3	99.2	99.7	0.014	0.004	0.005	0.006
BR $_0$	part. max	8.08×10^{-5}	96.9	99.0	99.8	100.0	0.036	0.014	0.017	0.031
BR $_1$	part. median	4.73×10^{-5}	91.4	97.0	99.6	99.7	0.028	0.008	0.012	0.026
BR $_1$	part. IQR	4.85×10^{-5}	3.0	1.4	0.4	0.3	0.008	0.006	0.007	0.012
BR $_1$	part. min	1.52×10^{-5}	78.1	93.6	98.5	99.0	0.016	0.005	0.005	0.007
BR $_1$	part. max	12.30×10^{-5}	96.6	98.8	99.9	100.0	0.047	0.016	0.022	0.053
PP	part. median	23.29×10^{-5}	-	-	82.0	90.4	-	-	0.067	0.105
PP	part. IQR	44.25×10^{-5}	-	-	13.0	13.5	-	-	0.065	0.138
PP	part. min	4.78×10^{-5}	-	-	70.0	61.3	-	-	0.030	0.043
PP	part. max	176.27×10^{-5}	-	-	93.2	98.0	-	-	0.167	0.247
BR $_0$	$f_c = 0.3$ Hz	2.91×10^{-5}	92.7	98.1	99.7	99.9	0.023	0.005	0.009	0.016
BR $_0$	$f_c = 0.6$ Hz	5.77×10^{-5}	94.3	97.9	99.6	99.8	0.031	0.008	0.010	0.019
BR $_1$	$f_c = 0.3$ Hz	3.89×10^{-5}	90.8	97.2	99.6	99.8	0.028	0.007	0.011	0.019
BR $_1$	$f_c = 0.6$ Hz	6.63×10^{-5}	92.8	97.8	99.5	99.7	0.032	0.009	0.013	0.024
PP	$f_c = 0.3$ Hz	12.55×10^{-5}	-	-	85.6	92.9	-	-	0.051	0.088
PP	$f_c = 0.6$ Hz	23.48×10^{-5}	-	-	82.6	87.2	-	-	0.063	0.090

and higher oscillation frequencies between the two GMM components of model BR $_1$. Appendix B.2 lists the full parameters of the behavioral model.




Figure 17: Identified and sorted poles s_0 to s_4 after outlier removal with the BR_1 model. Each result consists of one real pole (circles) and two pairs of complex poles, one with larger imaginary part (triangles) and one with smaller imaginary part (rectangles).

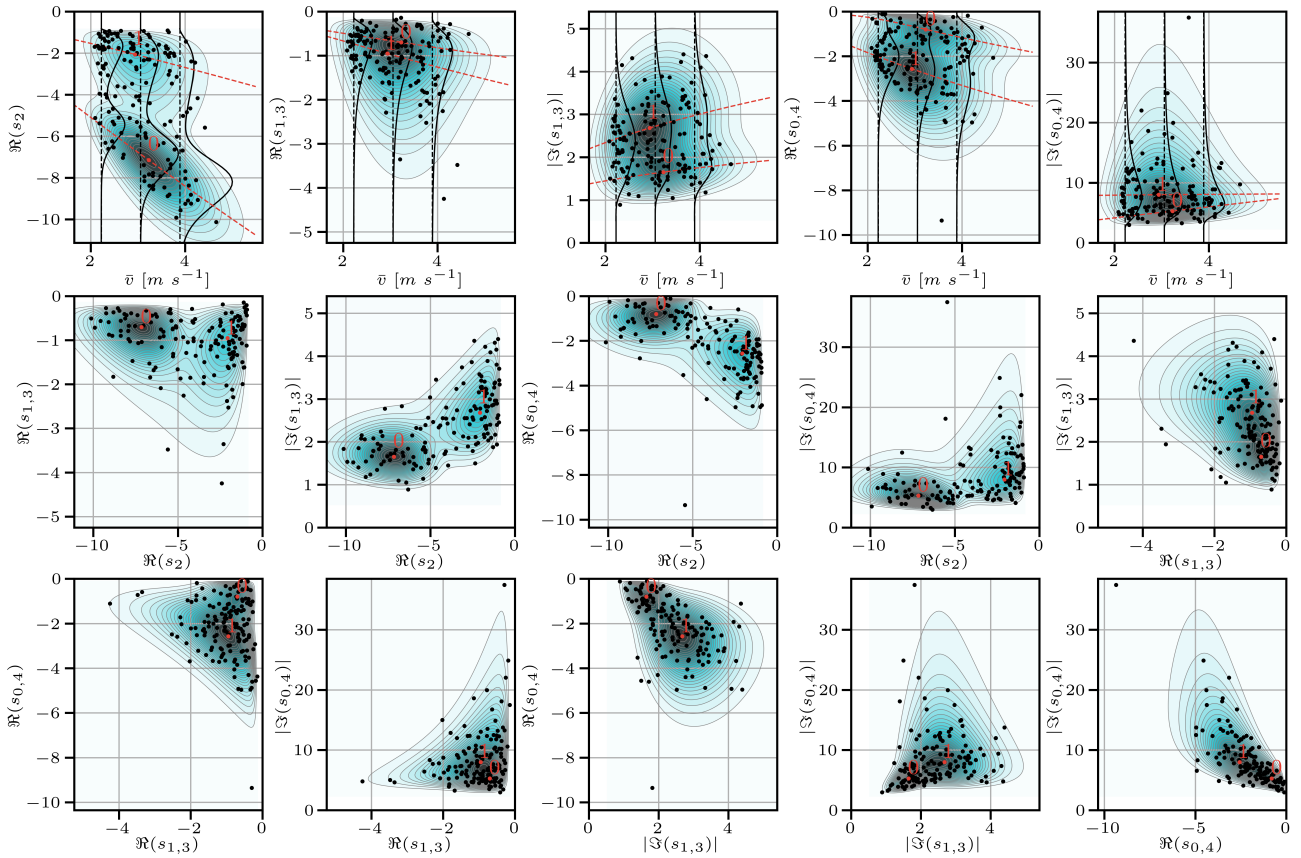


Figure 18: Marginal distributions of the augmented BR_1 pole features including speed. Turquoise contours show the fitted Gaussian mixture pole model. Black lines show the model's marginal distributions of pole locations for speeds fixed at 8, 11, and 14 km h^{-1} . Black dots represent the identification results used for training. Red dots mark the mean of individual Gaussian components, and red dashed lines mark the component means of the marginal distribution as a function of speed. Samples are distributed across multiple skewed clusters and with strong correlations. The Gaussian distributions fit the skewed clusters because of the power and log-shift transformations.

3.4. Test results

For testing the predicted trajectory distributions of the calibrated models and their different impact on safety assessment, we retain the three conditional models due to their superior performance in predicting the pole locations (Tab. 3). The following section shows the results of both tests.

Table 4: Gaussian mixture pole model training and test results showing the optimal number of components and the resulting number of parameters after 10-fold cross-validation, as well as the negative-log-likelihood evaluated on the test and the full training dataset. Results between the Balancing Rider and Planar Point models vary in magnitude due to the different number of features. All pole models accurately predict the locations of the test poles.

model	condition	#(GMM comp.) K	# parameters	NLL_{xval}	NLL_{test}
BR ₀	-	1	20	6.40	5.78
BR ₀	speed	1	27	6.41	5.74
BR ₁	-	2	41	6.12	5.31
BR ₁	speed	2	55	5.51	4.35
PP	-	1	2	1.42	1.57
PP	speed	1	5	1.40	1.50

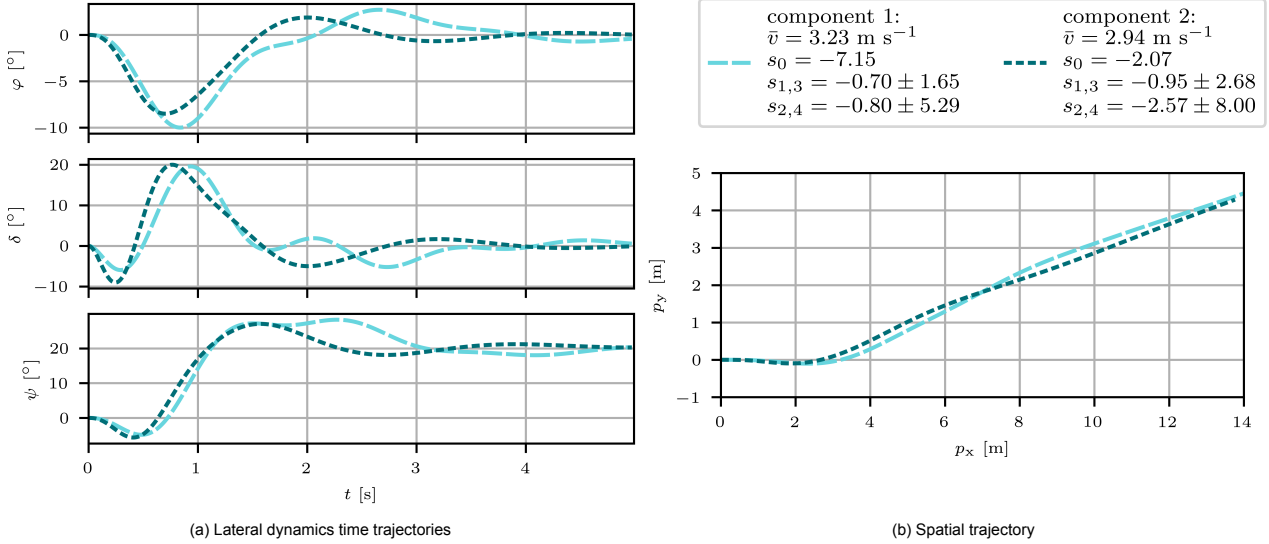


Figure 19: The responses to a $\hat{\psi} = 20^\circ$ command using the mean poles of the two GMM components of the speed-conditional BR₁ model representing the different control behaviors encoded by the components of the behavioral model.

Predicted Trajectory Distribution

Model BR₁ outperforms the more complex BR₀ model in the T-NLL of the predicted trajectory distribution and in terms of the average displacement error of the best predicted trajectory (1-minADE) (Table 5). However, BR₀ achieves a better LD-NLL. This reveals that the trajectories generated by BR₀ show better temporal consistency with the measurements, while the lower T-NLL shows that the spatial distribution per timestep is better captured by BR₁. The Planar Point model performs considerably worse in all metrics. The 1-minADE is four times that of BR₁ and shows considerably more variation between samples. This implies even larger errors for some individual scenes and trajectory segments. The NLL results are several orders of magnitude off, indicating a severe mismatch of the predicted distribution with the measured samples.

A look at examples of predicted response distributions to large command steps explains these results (Figure 20). BR₀ predicts the widest trajectory distribution, including very oscillatory and very slow reactions. The large variation enables a good fit for samples with rare dynamics (e.g., bottom row). The distributions of BR₁ are narrower with less extreme samples. At each time step, this captures the underlying distribution more accurately but shows less agreement with rare dynamics, such as the slow turn in the bottom row. Model PP predicts very narrow distributions that don't align well with the measurement, resulting in large penalties.

The examples further illustrate the capabilities of the Balancing Rider models to predict a diverse range of behaviors. In the third row, the real rider performed a smooth and spacious maneuver with a large heading overshoot. In contrast, in the second row, the measured turn is much sharper and reaches the commanded heading quickly. In the first case, the measurement is at the outer edge of the predicted distribution, indicating that most control strategies employed by our participants would have resulted in faster turns. The quicker turn in the second row is more towards the inner edge of the predicted distribution.

4. Obstacle Avoidance

The obstacle avoidance test shows how the models' differences in predicting cyclist trajectories translate into conflict outcome predictions (Figure 21). The Balancing Rider models predict how real cyclists must perform a countersteer maneuver to roll the bicycle into the corner, effectively causing a delay of the yaw angle correction. The Planar Point model steers directly in the right direction and thus reaches a higher yaw angle sooner most of the time, even with the additional response delay. The conflict point (point of minimum TTC) lies ≈ 1.27 m further away from the obstacle, ≈ 4 cm to the right, and exhibits an $\approx 1^\circ$ smaller yaw angle compared to the three-dimensional models. The median TTC differs by about ≈ 0.42 s. A Kruskal-Wallis test (Kruskal and Wallis, 1952) on the TTC distributions of all three models confirms that the predicted distributions are significantly different ($p \ll 0.01$). Highly oscillatory turns and very slow turns occur more often for BR_0 than for BR_1 . The latter create a small number of particularly small TTC values, with one of 1000 BR_0 predictions even ending in a collision with the obstacle.

Table 5: Results of the predicted trajectory distribution test in terms of between-sample median and Inter-Quartile Range (IQR). Best results highlighted in bold.

model	condition	1-minADE [m]		LD-NLL		T-NLL	
		median	IQR	median	IQR	median	IQR
BR_0	speed	0.018	0.018	-8.256	5.218	-4.067	2.200
BR_1	speed	0.015	0.019	-7.920	8.362	-4.268	2.887
PP	speed	0.060	0.087	5.90×10^5	2.18×10^6	1.73×10^5	5.06×10^5






Figure 20: Three test scenes (rows) with of trajectory distributions (black) generated by the conditional model variants BR_0 (left), BR_1 (middle), PP (right) compared to the measured trajectory (red). Dots mark the time of the command, the moment of countersteer onset τ_{cs} , and the moment of yaw error correction τ_{ec} . Filled dots indicate which delay is passed to the respective model. Ticks mark the LD-NLL evaluation moments. Green segments mark where the T-NLL is evaluated. The dashed line represents the commanded heading, with the LD being the perpendicular distance to this line.

5. Discussion

The following sections discuss the observed behaviors, the proposed models, and the resulting implications for traffic processes before elaborating on the limitations of our study.



(a) Simulated spatial trajectories. Gray crosses mark the commanded sequence of destinations to avoid the rectangular obstacle. A simulated cyclist reached a destination once they entered the dashed circle. Red dots mark the point of highest conflict severity (i.e., the conflict location), characterized by the minimum TTC of each interaction.



(b) The TTC versus the cyclist position and orientation at the conflict points (red dots in (a)), showing how the predicted conflict severity and conflict location vary between models.

(c) Minimum TTC distributions.

Figure 21: Obstacle avoidance test results with XY-trajectories (top) and the distributions of the minimum Time-to-Collision (TTC) (right) and the conflict point locations (left) for 1000 samples of behavioral parameters drawn from the predictive model at a fixed speed of 11 km s^{-1} . Dashed lines mark medians.

5.1. Observed Control Behaviors

We performed a cycling experiment in which participants followed quickly changing visual heading commands while cycling at a constant speed on a test track. We measure the kinematics of the bicycle and analyze the rider control behavior. Even within our small set of participants, we observe large differences in control strategies. While some participants perform smooth and spacious turns, others employ sharp and quick course corrections. The former leads to more time spent with larger deviations from the command, large heading overshoot, and the realized trajectories are oscillatory. While the sharp strategies enable riders to keep up with most commanded heading changes for long and short command intervals in our experiment, the smooth strategies often do not enable them to finish maneuvers in time.

Comparing the identified gains of the Balancing Rider models with the theoretical results of Schwab et al. (2008) shows that the general trend of control strategies follows the principles of theoretical optimality. However, individual samples may stray widely and even contradict theoretical optimality in terms of gain signs. Particularly, examples where the complete rider control model (BR_0) outperforms the reduced model (BR_1) and shows gains with different signs indicate that some riders do not always employ optimal strategies. Since cycling is an acquired skill, employing suboptimal strategies may result from limited experience, a lack of familiarity with the bicycle type, and individual motor skills. Previous work has determined that experienced cyclists, as well as younger cyclists, exhibit smaller steer angles and roll rates than unexperienced cyclists or older cyclists during straight riding (Kováčsová et al., 2016; Fonda et al., 2017). It is therefore reasonable to expect the same factors to also explain some of the variability in control strategies during turning observed in our experiment. Inconsistencies between the observed strategies highlight that the employed strategies are highly individual.

Another explanation may be different levels of task engagement among our participants, which we had no way of measuring. It was made clear that the task should not be seen as a challenge. Still, we expect that the consecutive commands created a sense of time pressure, and the participants were instructed to follow the commands as well as possible while maintaining safe riding. In accordance, the small differences in roll rate, steer rate, and yaw error between responses with different time pressure indicate that control strategies were generally close to a dynamic limit. We therefore suspect that control capacity and individual preferences play a larger role in the strategy choice than task engagement.

The large space required for slower and more oscillatory maneuvers raises the question of how the observations correspond to behaviors in real traffic. In our test, the random command timing and direction prevent anticipation and trigger purely reactive maneuvers. This is rare in natural settings, where the awareness and anticipation of the dynamic environment enable more maneuver planning. As all participants reported experience riding in public traffic, we expect that planning ahead enables maneuvers typically required in natural traffic within the dynamic control limits observed in our data. For example, [Lee et al. \(2020\)](#) measured cyclists' trajectories around static obstacles and reported that evasive maneuvers with 1 m lateral displacement start approximately 10 m before the obstacle. Since these radii greatly exceed our experiment conditions, we assume that such mild maneuvers should be well within the observed dynamic limits and control capacities. Measurements of cyclist conflicts in naturalistic conditions reported by [Zhang et al. \(2017\)](#) suggest that the average turn rates of full trajectories across an intersection lie below $12.5^\circ \text{ s}^{-1}$ for 85 % of cases, showing that more severe actions may rarely be needed during normal riding. The turn rates in our experiment exceed this by more than a factor of two, underlining that the observed control capacities should enable typical traffic interactions and that our setting is more akin to emergency evasive maneuvers in unexpected situations. Here, quickly correcting the heading without large spatial requirements is beneficial to avoid collisions and stay on the track.

Future work should explore the extent to which cyclists can execute planned and unexpected maneuvers within the dynamic limits on steer, roll, and yaw observed in our data to establish a connection between control capacity, traffic behavior, and the corresponding traffic safety outcome. Furthermore, since our test prescribed a constant speed, follow-up studies should analyze how speed modulation may enable keeping the dynamics within the observed limits for a specific maneuver.

5.2. The Stochastic Balancing Rider Model

Our Balancing Rider model uses the well-established linearized Carvallo-Whipple dynamics, which is the simplest model that captures all relevant dynamic phenomena of cycling and has been validated for upright riding with recovery from lateral disturbances ([Kooijman et al., 2008](#)). This lends great confidence in the accuracy of the fundamental results achieved through simulations with this model. However, using this model to describe bicycle dynamics in arbitrary scenarios may put it outside its validated domain. For example, narrow turns require a non-linear formulation for best accuracy ([Franke et al., 1990](#)), and significant acceleration needs to be included explicitly in the dynamic equations ([Meijaard and Schwab, 2006](#)). While the exact impact of these simplifications on simulated interactions of varying severity should be studied in future work, the improvements of the modeled dynamics with the Balancing Rider compared to the popular Planar Point model are already considerable. Especially for surrogate safety assessment, minor errors during extreme maneuvers may be acceptable, assuming that the situation passes criticality thresholds anyway if the model captures spatiotemporal proximity up to the maneuver well enough.

The model requires the physical parameters of the bicycle and rider, which we measure for our test bicycle and generalize to an average human. Representing all participants with the same average rider parameters introduces a bias related to the rider's physique, which may be one contributing factor to slightly worse model fits for some participants. Future studies should measure rider height and mass to scale the rider parameters accordingly and incorporate statistical knowledge of human body parameters for more accurate prediction. Nevertheless, the difference in fit quality between the Balancing Rider and the Planar Point models greatly exceeds the difference between participants, indicating that our model improves simulated trajectories even for badly-represented rider physiques.

For the first time, we explicitly model distributions of behavioral parameters for a control model of the balancing rider, which enables sampling behaviors for Monte Carlo-based traffic simulations. The excellent fit of our model to the measured data shows that the full-

state feedback approach provides enough degrees of freedom to compensate our simplifications and conveniently eliminates the need for traffic safety professionals to consider additional details like muscular dynamics. Filtering the measurement data with a similar model facilitates this fit and provides a clear view of the behavioral parameters. Modeling the distributions of the closed-loop poles rather than gains enables us to guarantee lateral stability.

We compensate the response time of our participants to the visual commands. This makes our models independent of the rider response delay and enables undisturbed insights into the lateral dynamics. It enables users to independently choose a response-time model suitable for their application scenario and delay the commanded heading accordingly.

Our models predict a large variance in observed behaviors, which aligns with the large variance in observed behaviors. Between the two model variants, the more complex BR_0 is able to capture the large variety of observed control behaviors slightly better than the reduced BR_1 . Consequently, its predicted trajectory distributions also better explain the test samples when considering each trajectory's temporal coherence (lower LD-NLL). However, it also creates a small number of extremely oscillatory or extremely slow responses that, due to their limited occurrence, don't significantly impact the LD-NLL score. BR_1 creates significantly fewer extremes and explains the spatial distribution of cyclist trajectories better when not considering temporal coherence. With the limited dataset size, this study can't definitively determine whether these extremes are realistic, safety-relevant extrapolations of the observed behaviors, or unrealistic samples generated from the multivariate behavioral distribution that stretches into space not supported by observations. Future studies should compare the trajectory predictions of both models to a larger set of real conflict trajectories to determine which model is more accurate. The current observations suggest that the more conservative BR_1 model may be more appropriate for applications where different degrees of normal riding are required, while BR_0 may enable uncovering more worst-case scenarios for safety assessment.

Our models do not distinguish variation between and within subjects. From examining the gain distributions qualitatively, both seem to be non-negligible, indicating different rider types and non-constant individual rider behavior. With this in mind, users of our models should sample new behavioral parameters per scene. For cyclists traversing a larger network, parameters must be resampled if the speed changes substantially. If stochastic control behavior is not required for a target application, the means of our parameter distributions describe representative behaviors (see Appendix B.2 for values).

Finally, it may be questionable if specific behaviors are likely to be chosen by any cyclist during any scenario. The true behavior may depend on more factors than the cyclist's speed, which we included in our model because the speed dependency of control parameters is thoroughly established in literature. We omitted command frequency because the effect on the measured steer and roll rate was small and there is no obvious equivalent for prediction in general traffic simulation scenarios. Unfortunately, the size of our dataset did not allow us to explore other factors. Future research should examine the impact of rider characteristics, such as experience and age, environmental factors (e.g., road conditions and perceived urgency of the maneuver), and other kinematic factors from the current bicycle state on control behaviors to differentiate control strategies for specific situations.

5.3. The Role of Bicycle Dynamics for Safety-Aware Traffic Modeling and Conflict Causality

We train two different rider-bicycle control model types based on the measured data. The Balancing Rider models describe the full dynamics of the upright bicycle, including the resulting instability, the need to lean into a corner, and the need to countersteer to initiate roll. The Planar Point model represents a strong abstraction of the bicycle as a first-order yaw tracker, commonly found in social-force traffic modeling.

Significantly worse calibration results (a factor of 7 difference between the mean absolute errors of the yaw angle, path errors reaching the order of meters) show that the Planar Point model cannot capture the trajectories of real bicycles. Its single degree of freedom only allows it to fit the rise time of the response. It neither captures countersteering nor the oscillations and overshoot that some participants create while stabilizing the bicycle after a course correction, and behaves more like a car. As a result, the fitted position trajectories deviate significantly from real observations. The distribution of complex pole pairs of our fitted Balancing Rider models shows that the oscillatory nature of the response is a key differentiator between control behaviors. Trajectory predictions with the Planar Point model fail to capture the features that distinguish these behaviors.

The simulated obstacle avoidance test highlights the significance of the position trajectory differences between the Balancing Rider and Planar Point models for safety assessment. The models predict different trajectory shapes, resulting in different spatiotemporal closeness of the rider to the obstacle. The larger minimum time to collision with the obstacle (TTC) for the Planar Point model indicates less severe simulated conflicts than for the Balancing Rider model. Additionally, the location of the most critical point shifts away from the obstacle. Small differences of the predicted lateral position (≈ 4 cm) and orientation ($\approx 1^\circ$) at the conflict points coincide with larger differences in TTC (≈ 0.42 s) and longitudinal position (≈ 1.27 m). This underlines the sensitivity of TTC and conflict location to the predicted heading resulting from the linear extrapolation of the position based on current heading and velocity. The TTC differences between the Planar Point and Balancing Rider models underscore that inaccurate models in simulated safety assessments can lead to incorrect conclusions about predicted safety. Typically, surrogate safety assessment considers only events below a threshold as potentially unsafe, and the difference between the models amounts to almost a third of the typical threshold value $\text{TTC} < 1.5$ s (Johnsson et al., 2018). The different dynamic models used for prediction can thus easily make the difference between predicting safe or unsafe situations. The range of TCC values predicted by the BR models fits empirical values from a similar real-world obstacle avoidance experiment performed by Lee et al. (2020) well (median TTC cycling experiment ≈ 2.3 s, median TTC our simulation 2.6 s) although the validity of this comparison is limited due to small differences in the setup.

Moreover, the observed and modeled behaviors shed light on the role of these dynamics in traffic processes. To turn, cyclists must first initiate a suitable roll angle by steering in the opposite direction. This effectively causes a time delay between the initial reaction of the cyclist and the beginning of correcting the heading in the desired direction, which amounts to more than 0.6 s in our experiment. Furthermore, only specific combinations of roll, steer, and yaw rate are feasible in certain situations, and these are described by the Balancing Rider model. For example, when a cyclist is committed to a turn, a certain roll angle is required, which must be adjusted if the desired direction changes. Thus, direction changes cannot be instantaneous but depend on the current state of the bicycle as well as the speed of the rider's control behavior. In consequence, the bicycle's physics profoundly limit cyclists' action space and form the foundation of traffic process causality. Comparing the mean absolute error of the y-position after fitting our Balancing Rider model ($\text{MAE}_y(\text{BR}_0) = 0.025$ m) and after fitting the more car-like Planar Point model ($\text{MAE}_y(\text{PM}) = 0.125$ m) shows how large this impact is. Future work should investigate where the dynamics of the bicycle contribute to safe traffic processes and where they create risks. This may be particularly interesting for interactions with other road users, where visible dynamic cues could facilitate short-term anticipation and response delays inherent to the dynamics of the inert rider-bicycle system could impede effective evasive actions.

5.4. Limitations

We identify the following limitations of our work. First, we ran the experiment on only ten participants. Our behavioral distributions thus cannot be used to generalize to the population level. However, they still show that bicycle control strategies are personal characteristics that vary between individuals. Second, the conditions of our controlled experiment may impact participant behavior, limiting the transferability to real-world traffic situations. Concrete issues may arise from the perceived level of urgency of the maneuver, the track layout with fences on some sides, the unfamiliar bicycle type, or the requirement to keep a constant speed. While we cannot quantify any of these factors, we are confident that the generality of the underlying physical bicycle model enables us to transfer our qualitative observations robustly to naturalistic scenarios. Third, although well-reasoned, several methodological choices may cause unquantified errors in our results. The response time estimation is not robust to behaviors other than direct responses to the command. The sample duration is generally short, and samples with very slow dynamics, particularly slow asymptotic decay, may not be identified correctly. Simplifying assumptions in the physical model, like neglecting roll torque input, fixing the rider's body, and generalizing rider physical parameters, introduce errors. Nevertheless, strong fitting results indicate that these errors may be small, and the large differences between the three-dimensional and particle dynamics give us confidence that these errors do not qualitatively affect the conclusions. Fourth, our models assume constant speed. We did not explore how behavioral parameters change during substantial speed changes within a short time horizon. As a result, a set of sampled parameters is only valid over time-horizons with reasonably constant speed. When the speed changes, parameters need to be resampled, leading to different cyclist behavior. Lastly, we present an elaborate model that combines the equations of motion of a multi-body system with a multivariate stochastic model of behavioral parameters. This results in a large

number of parameters (see Appendix B.2), which may be a hurdle for adopting the model on a larger scale within traffic simulation. However, our results indicate that including such elaborate dynamic models provides crucial details for understanding and predicting safety-relevant traffic processes.

6. Conclusion

Traffic simulations with the power to predict cycling safety are of great value for testing and development of transportation technologies. However, existing cyclist models employ planar abstractions of the bicycle that neither account for the unique motion constraints of bicycles nor the cyclists' related control behaviors. Whether these details help to explain and predict traffic processes has not been previously investigated.

For the first time, we present a calibrated and thoroughly tested model with physically accurate dynamics of the upright bicycle and stochastic control behavior of the balancing rider. The model uses the established linearized Carvallo-Whipple dynamics to describe the bicycle and full-state feedback control to describe control behavior. It enables drawing random new samples of control behavior parameters for stochastic traffic simulation. We calibrate our model using data from a controlled riding experiment where ten participants executed heading step commands on a test track, achieving excellent fitting results. Testing predicted trajectory distributions against a withheld set of measurements demonstrates the predictive power of our model.

Comparisons between the Balancing Rider model and the common abstract representation of cyclists as a point in planar space indicate that the dynamics of the bicycle play a fundamental role in the formation of cyclist trajectories. Real trajectories exhibit countersteering, heading overshoot, oscillations for stabilization, and dependency on current roll and steer that the Planar Point model can't capture. The projected median TTC difference of 0.42 s between the calibrated Planar Point model and our Balancing Rider model is approximately one third of standard conflict detection thresholds, which strongly suggests that these dynamic effects are relevant for surrogate safety assessment. Consequently, cyclist models for the predictive simulation of safety-relevant traffic scenarios should consider accurate physical dynamics of the bicycle.

The rider control behaviors in our experiment vary significantly between individuals. Observed strategies range from quick, abrupt heading corrections to spacious, smooth turns. The determining factors of different control behaviors and the implications of control capacity on maneuvers in more complex traffic scenarios may further explain causal traffic processes, but require further research.

Our small number of participants means that the proposed model is not representative of rider behavior on a population level. However, the large improvements compared to the commonly used particle models warrant expanding the study to a more representative sample. For this, traffic modeling researchers and engineers may use the methods presented in this study to incorporate accurate physical bicycle dynamics and stochastic rider control behavior into their traffic simulation frameworks. Considering the aforementioned limitations, developers of automated driving functions, advanced assistance systems, and intelligent traffic systems may apply our calibrated model directly to improve cyclist behavior in limited simulated testing during early development.

Acknowledgments

The authors would like to thank Hidde Boekema for his advice on test methods for predicted trajectory variance.

Data Availability

The data from the cycling experiment conducted as part of this study, along with Python software for data processing and analysis, as well as the analysis results, are available at 4TU.ResearchData (DOI: [10.4121/f881dd80-b9f5-4322-9fd5-192034c9717f](https://doi.org/10.4121/f881dd80-b9f5-4322-9fd5-192034c9717f)).

To facilitate convenient reuse of our models, we additionally list all parameters in the Appendix B.2 and publish software on Github ([chris-konrad/stochastic-balancing-rider](https://github.com/chris-konrad/stochastic-balancing-rider)).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of Generative AI in Scientific Writing

During the preparation of this work, the authors used Grammarly to improve spelling, grammar, and readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

CRedit Authorship Contribution Statement

Christoph M. Konrad: Conceptualization, methodology, formal analysis, investigation, visualization, writing - original draft, writing - review & editing. **Riender Happee:** Conceptualization, validation, supervision, writing - review & editing. **Azita Dabiri:** Conceptualization, validation, supervision, writing - review & editing. **Jason K. Moore:** Conceptualization, methodology, supervision, writing - review & editing.

Funding

This research was funded internally by the Faculty of Mechanical Engineering, Delft University of Technology.

References

- Alizadehsaravi, L., Moore, J.K., 2023. Bicycle balance assist system reduces roll and steering motion for young and older bicyclists during real-life safety challenges. *PeerJ* 11, e16206. doi:[10.7717/peerj.16206](https://doi.org/10.7717/peerj.16206).
- Alozi, A.R., Hussein, M., 2023. Active Road User Interactions With Autonomous Vehicles: Proactive Safety Assessment. *Transportation Research Record* 2677, 74–89. doi:[10.1177/03611981221151032](https://doi.org/10.1177/03611981221151032).
- Berge, S.H., de Winter, J., Hagenzieker, M., 2023. Support systems for cyclists in automated traffic: A review and future outlook. *Applied Ergonomics* 111, 104043. doi:[10.1016/j.apergo.2023.104043](https://doi.org/10.1016/j.apergo.2023.104043).
- Bishop, C.M., 2006. *Pattern recognition and machine learning*. Information science and statistics, Springer, New York. URL: <https://link.springer.com/book/9780387310732>.
- Brunner, P., Denk, F., Huber, W., Kates, R., 2019. Virtual safety performance assessment for automated driving in complex urban traffic scenarios, in: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), IEEE, Auckland, New Zealand. pp. 679–685. doi:[10.1109/ITSC.2019.8917517](https://doi.org/10.1109/ITSC.2019.8917517).
- Bärgman, J., Boda, C.N., Dozza, M., 2017. Counterfactual simulations applied to SHRP2 crashes: The effect of driver behavior models on safety benefit estimations of intelligent safety systems. *Accident Analysis & Prevention* 102, 165–180. doi:[10.1016/j.aap.2017.03.003](https://doi.org/10.1016/j.aap.2017.03.003).
- Chehade, A., Savargaonkar, M., Krivtsov, V., 2022. Conditional Gaussian mixture model for warranty claims forecasting. *Reliability Engineering & System Safety* 218, 108180. doi:[10.1016/j.res.2021.108180](https://doi.org/10.1016/j.res.2021.108180).
- Denk, F., Brunner, P., Huber, W., Margreiter, M., Bogenberger, K., Kates, R., 2022. Assessment of traffic safety interventions using virtual randomized controlled trials: potential of connected and automated driving including V2X for collision reduction at urban intersections, in: 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), Macau, China. pp. 1183–1190. doi:[10.1109/ITSC55140.2022.9921764](https://doi.org/10.1109/ITSC55140.2022.9921764).
- Dialynas, G., Christoforidis, C., Happee, R., Schwab, A., 2023. Rider control identification in cycling taking into account steering torque feedback and sensory delays. *Vehicle System Dynamics* 61, 200–224. doi:[10.1080/00423114.2022.2048865](https://doi.org/10.1080/00423114.2022.2048865).
- European Commission, 2025. *Annual statistical report on road safety in the EU 2025*. Governmental Report. European Road Safety Observatory, European Commission, Directorate General for Transport. Brussels. URL: https://road-safety.transport.ec.europa.eu/european-road-safety-observatory/data-and-analysis/annual-statistical-report_en.
- Fonda, B., Sarabon, N., Li, F.X., 2017. Bicycle rider control skills: expertise and assessment. *Journal of Sports Sciences* 35, 1383–1391. doi:[10.1080/02640414.2015.1039049](https://doi.org/10.1080/02640414.2015.1039049).
- Franke, G., Suhr, W., Riess, F., 1990. An advanced model of bicycle dynamics. *European Journal of Physics* 11, 116. doi:[10.1088/0143-0807/11/2/010](https://doi.org/10.1088/0143-0807/11/2/010).
- Fu, L., Zhang, Y., Qin, H., Shi, Q., Chen, Q., Chen, Y., Shi, Y., 2023. A modified social force model for studying nonlinear dynamics of pedestrian-e-bike mixed flow at a signalized crosswalk. *Chaos, Solitons & Fractals* 174, 113813. doi:[10.1016/j.chaos.2023.113813](https://doi.org/10.1016/j.chaos.2023.113813).

- Guo, N., Jiang, R., Wong, S., Hao, Q.Y., Xue, S.Q., Hu, M.B., 2021. Bicycle flow dynamics on wide roads: Experiments and simulation. *Transportation Research Part C: Emerging Technologies* 125, 103012. doi:[10.1016/j.trc.2021.103012](https://doi.org/10.1016/j.trc.2021.103012).
- Helbing, D., Molnár, P., 1995. Social force model for pedestrian dynamics. *Physical Review E* 51, 4282–4286. doi:[10.1103/PhysRevE.51.4282](https://doi.org/10.1103/PhysRevE.51.4282).
- Hess, R., Moore, J.K., Hubbard, M., 2012. Modeling the Manually Controlled Bicycle. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 42, 545–557. doi:[10.1109/TSMCA.2011.2164244](https://doi.org/10.1109/TSMCA.2011.2164244).
- Hoogendoorn, S., Gavriilidou, A., Daamen, W., Duives, D., 2021. Game theoretical framework for bicycle operations: A multi-strategy framework. *Transportation Research Part C: Emerging Technologies* 128, 103175. doi:[10.1016/j.trc.2021.103175](https://doi.org/10.1016/j.trc.2021.103175).
- Huang, W., Fellendorf, M., Schönauer, R., 2012. Social Force based Vehicle Model for 2-dimensional Spaces, in: *Transportation Research Board 91st Annual Meeting*, Transportation Research Board, Washington D.C.
- ITF, 2023. Road Safety Annual Report 2023. OECD Publishing, Paris, France. URL: <http://dx.doi.org/10.1787/8654c572-en>, doi:[10.1787/8654c572-en](https://doi.org/10.1787/8654c572-en).
- Ivanovic, B., Pavone, M., 2019. The Trajectron: Probabilistic Multi-Agent Trajectory Modeling With Dynamic Spatiotemporal Graphs, in: *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, IEEE, Seoul, South Korea. pp. 2375–2384. doi:[10.1109/ICCV.2019.00246](https://doi.org/10.1109/ICCV.2019.00246).
- Johnson, M., Charlton, J., Oxley, J., Newstead, S., 2010. Naturalistic Cycling Study: Identifying Risk Factors for On-Road Commuter Cyclists. *Annals of Advances in Automotive Medicine* 54, 275–283. URL: <https://pubmed.ncbi.nlm.nih.gov/21050610/>.
- Johnsson, C., Laureshyn, A., De Ceunynck, T., 2018. In search of surrogate safety indicators for vulnerable road users: a review of surrogate safety indicators. *Transport Reviews* 38, 765–785. doi:[10.1080/01441647.2018.1442888](https://doi.org/10.1080/01441647.2018.1442888).
- Jurecki, R.S., Stańczyk, T.L., 2014. Driver reaction time to lateral entering pedestrian in a simulated crash traffic situation. *Transportation Research Part F: Traffic Psychology and Behaviour* 27, 22–36. doi:[10.1016/j.trf.2014.08.006](https://doi.org/10.1016/j.trf.2014.08.006).
- Kapousizis, G., Ulak, M.B., Geurs, K., Havinga, P.J.M., 2022. A review of state-of-the-art bicycle technologies affecting cycling safety: level of smartness and technology readiness. *Transport Reviews* 43, 1–23. doi:[10.1080/01441647.2022.2122625](https://doi.org/10.1080/01441647.2022.2122625).
- Kaths, H., 2023. A movement and interaction model for cyclists and other non-lane-based road users. *Frontiers in Future Transportation* 4. doi:[10.3389/ffutr.2023.1183270](https://doi.org/10.3389/ffutr.2023.1183270).
- Konrad, C.M., Dabiri, A., Schulte, F., Moore, J.K., Happee, R., 2025. Cycling safety assessment in microscopic traffic simulation: A review and methodological framework. *Transportation Research Interdisciplinary Perspectives* 34, 101734. doi:[10.1016/j.trip.2025.101734](https://doi.org/10.1016/j.trip.2025.101734).
- Kooijman, J.D.G., Schwab, A.L., Meijaard, J.P., 2008. Experimental validation of a model of an uncontrolled bicycle. *Multibody System Dynamics* 19, 115–132. doi:[10.1007/s11044-007-9050-x](https://doi.org/10.1007/s11044-007-9050-x).
- Kooijman, J.D.G., Schwab, A.L., Moore, J.K., 2009. Some Observations on Human Control of a Bicycle, in: *Proceedings of the ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Volume 4: 7th International Conference on Multibody Systems, Nonlinear Dynamics, and Control, Parts A, B and C., ASME, San Diego, USA. pp. 2021–2028. doi:[10.1115/DETC2009-86959](https://doi.org/10.1115/DETC2009-86959).
- Kováčsová, N., de Winter, J.C.F., Schwab, A.L., Christoph, M., Twisk, D.A.M., Hagenzieker, M.P., 2016. Riding performance on a conventional bicycle and a pedelec in low speed exercises: Objective and subjective evaluation of middle-aged and older persons. *Transportation Research Part F: Traffic Psychology and Behaviour* 42, 28–43. doi:[10.1016/j.trf.2016.06.018](https://doi.org/10.1016/j.trf.2016.06.018).
- Kruskal, W.H., Wallis, W.A., 1952. Use of Ranks in One-Criterion Variance Analysis. *Journal of the American Statistical Association* 47, 583–621. URL: <https://www.jstor.org/stable/2280779>, doi:[10.2307/2280779](https://doi.org/10.2307/2280779).
- Lee, O., Rasch, A., Schwab, A.L., Dozza, M., 2020. Modelling cyclists' comfort zones from obstacle avoidance manoeuvres. *Accident Analysis & Prevention* 144, 105609. doi:[10.1016/j.aap.2020.105609](https://doi.org/10.1016/j.aap.2020.105609).
- Li, Y., Ni, Y., Sun, J., 2021. A modified social force model for high-density through bicycle flow at mixed-traffic intersections. *Simulation Modelling Practice and Theory* 108, 102265. doi:[10.1016/j.simpat.2020.102265](https://doi.org/10.1016/j.simpat.2020.102265).
- Liang, X., Meng, X., Zheng, L., 2021. Investigating conflict behaviours and characteristics in shared space for pedestrians, conventional bicycles and e-bikes. *Accident Analysis & Prevention* 158, 106167. doi:[10.1016/j.aap.2021.106167](https://doi.org/10.1016/j.aap.2021.106167).
- Marbus, A., 2025. Cyclist conflict behavior in shared spaces. Master's thesis. Delft University of Technology, Delft, Netherlands. URL: <http://resolver.tudelft.nl/uuid:092f3b70-2d97-436e-b193-139a593e09c7>.
- Meijaard, J., Papadopoulos, J.M., Ruina, A., Schwab, A., 2007. Linearized dynamics equations for the balance and steer of a bicycle: A benchmark and review. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 463, 1955–1982. doi:[10.1098/rspa.2007.1857](https://doi.org/10.1098/rspa.2007.1857).
- Meijaard, J.P., Schwab, A.L., 2006. Linearized equations for an extended bicycle model, in: *Motasoares, C.A., Martins, J.A.C., Rodrigues, H.C., Ambrósio, J.A.C., Pina, C.A.B., Motasoares, C.M., Pereira, E.B.R., Folgado, J. (Eds.), III European Conference on Computational Mechanics*, Springer, Dordrecht, Netherlands. pp. 772–772. doi:[10.1007/1-4020-5370-3_772](https://doi.org/10.1007/1-4020-5370-3_772).
- Moore, J., 2012. Human Control of a Bicycle. Doctoral dissertation. University of California, Davis, Davis, USA. URL: <https://moorepants.github.io/dissertation/index.html>.

- Moore, J.K., Bogert, A.D.v., 2018. *optry*: Software for trajectory optimization and parameter identification using direct collocation. *Journal of Open Source Software* 3, 300. doi:[10.21105/joss.00300](https://doi.org/10.21105/joss.00300).
- Reijne, M., 2025. *Modeling and Evaluation of Cyclist Fall Prevention Interventions: A Proactive Cycling Safety Approach*. Doctoral Dissertation. Delft University of Technology. Delft, Netherlands. doi:[10.4233/uuid:ada397a7-0c55-4284-9133-3c5ea515ea67](https://doi.org/10.4233/uuid:ada397a7-0c55-4284-9133-3c5ea515ea67).
- Ritchie, H., Roser, M., 2017. *Obesity*. *Our World in Data*. URL: <https://ourworldindata.org/obesity>. online, accessed 26.11.2025.
- Roser, M., Appel, C., Ritchie, H., 2021. *Human Height*. *Our World in Data*. URL: <https://ourworldindata.org/human-height>. online, accessed 26.11.2025.
- Rudenko, A., Palmieri, L., Herman, M., Kitani, K.M., Gavrilu, D.M., Arras, K.O., 2020. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research* 39, 895–935. doi:[10.1177/0278364920917446](https://doi.org/10.1177/0278364920917446).
- Schmidt, C.M., Dabiri, A., Schulte, F., Happee, R., Moore, J., 2024. *Essential Bicycle Dynamics for Microscopic Traffic Simulation: An Example Using the Social Force Model*, in: *The Evolving Scholar - BMD 2023, 5th Edition*, TU Delft Open, Delft, Netherlands. doi:[10.59490/65a5124da90ad4aecf0ab147](https://doi.org/10.59490/65a5124da90ad4aecf0ab147).
- Schwab, A.L., Kooijman, J.D., Meijaard, J.P., 2008. Some recent developments in bicycle dynamics and control, in: *4th European Conference on Structural Control, ECSC 2008, Saint-Petersburg, Russia*. URL: <https://repository.tudelft.nl/record/uuid:68567a34-5e5d-4de4-ba78-0faeed584c30>.
- Schwab, A.L., Meijaard, J.P., 2013. A review on bicycle dynamics and rider control. *Vehicle System Dynamics* 51, 1059–1090. doi:[10.1080/00423114.2013.793365](https://doi.org/10.1080/00423114.2013.793365).
- Sharp, R.S., 2008. On the Stability and Control of the Bicycle. *Applied Mechanics Reviews* 61. doi:[10.1115/1.2983014](https://doi.org/10.1115/1.2983014).
- Sun, J., Zhang, H., Zhou, H., Yu, R., Tian, Y., 2022. Scenario-Based Test Automation for Highly Automated Vehicles: A Review and Paving the Way for Systematic Safety Assurance. *IEEE Transactions on Intelligent Transportation Systems* 23, 14088–14103. doi:[10.1109/TITS.2021.3136353](https://doi.org/10.1109/TITS.2021.3136353).
- Wang, W., Zhou, H., Lo, J.T.Y., Lo, S.M., Wang, Y., 2024. A modified social force model for pedestrian-bicycle mixed flows and its application on evaluating the conflict risk in shared roads. *Physica A: Statistical Mechanics and its Applications* 643, 129788. doi:[10.1016/j.physa.2024.129788](https://doi.org/10.1016/j.physa.2024.129788).
- Wächter, A., Biegler, L.T., 2006. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106, 25–57. doi:[10.1007/s10107-004-0559-y](https://doi.org/10.1007/s10107-004-0559-y).
- Yeo, I., Johnson, R.A., 2000. A new family of power transformations to improve normality or symmetry. *Biometrika* 87, 954–959. doi:[10.1093/biomet/87.4.954](https://doi.org/10.1093/biomet/87.4.954).
- Yuan, Y., Goñi-Ros, B., Oijen, T.P.v., Daamen, W., Hoogendoorn, S.P., 2019. *Social Force Model Describing Pedestrian and Cyclist Behaviour in Shared Spaces*, in: Hamdar, S.H. (Ed.), *Traffic and Granular Flow '17*, Springer, Cham, Washington, DC, USA. pp. 477–486. doi:[10.1007/978-3-030-11440-4_52](https://doi.org/10.1007/978-3-030-11440-4_52).
- Zhang, R., Wu, J., Huang, L., You, F., 2017. Study of Bicycle Movements in Conflicts at Mixed Traffic Unsignalized Intersections. *IEEE Access* 5, 10108–10117. doi:[10.1109/ACCESS.2017.2703816](https://doi.org/10.1109/ACCESS.2017.2703816).

Appendix

A. Derivation of the Planar Point Model

To find a suitable baseline for our model comparisons, we resort to the common representation of road users as moving points and enforce the constant speed assumption of our study. We base the planar point model (PP) on the original social force model for pedestrians by [Helbing and Molnár \(1995\)](#):

$$\dot{\mathbf{v}}(t) = \frac{1}{T} (\hat{\mathbf{v}}(t) - \mathbf{v}(t)) + \text{repulsive forces}, \quad (\text{A.12})$$

where we use the notation introduced in this paper rather than the original notation of [Helbing and Molnár \(1995\)](#) and T is the relaxation time. $\mathbf{v}(t) = [v_x(t), v_y(t)]^T$ is the velocity of the bicycle pointing in the direction of the current heading. $\hat{\mathbf{v}}(t)$ is the desired velocity pointing in the direction of the desired heading. For the scope of our study, we fix the forward speed to a constant value $\|\mathbf{v}(t)\| = v$ and assume that the cyclist perfectly follows the desired speed $\|\hat{\mathbf{v}}(t)\| = \|\mathbf{v}(t)\|$. Additionally, we neglect repulsive forces to purely focus on the lateral control behavior. This allows us to simplify the dynamics till we arrive at the planar point model introduced in Equation 4. For brevity, we omit the notation of time dependencies in the following.

We begin by writing out Equation A.12 in terms of the unit vectors \mathbf{e}_x and \mathbf{e}_y in the two spatial directions of the road surface:

$$\dot{v}_x \mathbf{e}_x + \dot{v}_y \mathbf{e}_y = \frac{1}{T} [(\hat{v}_x - v_x) \mathbf{e}_x + (\hat{v}_y - v_y) \mathbf{e}_y]. \quad (\text{A.13})$$

Then, we convert the dynamics to polar coordinates using $v_x = v \cos \psi$ and $v_y = v \sin \psi$ as well as the constant speed ($\dot{v} = 0$) and perfect speed following ($v = \hat{v}$) assumptions. Thus, the velocity components become

$$\dot{v}_x = \dot{v} \cos \psi - v \dot{\psi} \sin \psi = -v \dot{\psi} \sin \psi \quad (\text{A.14})$$

$$\dot{v}_y = \dot{v} \sin \psi + v \dot{\psi} \cos \psi = v \dot{\psi} \cos \psi, \quad (\text{A.15})$$

which plugged into the left-hand side of Equation A.13 leads to

$$v \dot{\psi} (-\sin \psi \mathbf{e}_x + \cos \psi \mathbf{e}_y) = v \dot{\psi} \mathbf{e}_\psi. \quad (\text{A.16})$$

Here, \mathbf{e}_ψ is the unit vector perpendicular to the current heading. Equivalently, we convert the right-hand side of Equation A.13:

$$\frac{1}{T} [(\hat{v}_x - v_x) \mathbf{e}_x + (\hat{v}_y - v_y) \mathbf{e}_y] = \frac{v}{T} [\mathbf{e}_{\hat{v}} - \mathbf{e}_v], \quad (\text{A.17})$$

where \mathbf{e}_v is the unit vector in the direction of the heading and $\mathbf{e}_{\hat{v}}$ is the unit vector in the commanded direction. The two substitutions above lead to the constant-speed formulation of the social force model in polar coordinates

$$\dot{\psi} \mathbf{e}_\psi = \frac{1}{T} (\mathbf{e}_{\hat{v}} - \mathbf{e}_v). \quad (\text{A.18})$$

To expose the dynamics of the yaw angle, we take the magnitude of the vectorial differential equation. While the left-hand side is trivial, the right-hand side requires some rearrangements:

$$\frac{1}{T} \|\mathbf{e}_{\hat{v}} - \mathbf{e}_v\| = \frac{1}{T} \sqrt{(\cos \hat{\psi} - \cos \psi)^2 + (\sin \hat{\psi} - \sin \psi)^2} \quad (\text{A.19})$$

$$= \frac{1}{T} \sqrt{2 - 2(\cos \hat{\psi} \cos \psi + \sin \hat{\psi} \sin \psi)} \quad (\text{A.20})$$

$$= \frac{1}{T} \sqrt{2 - 2 \cos(\hat{\psi} - \psi)} \quad (\text{A.21})$$

$$= \frac{1}{T} \sqrt{2 - 2 \cos(\hat{\psi} - \psi)} \quad (\text{A.22})$$

$$= \frac{2}{T} \sin \frac{\hat{\psi} - \psi}{2} \quad (\text{A.23})$$

$$\approx \frac{1}{T} (\hat{\psi} - \psi) \quad \text{for small command steps.} \quad (\text{A.24})$$

In combination with the left-hand side, this leads to the familiar dynamics of a moving point.

$$\dot{\psi} = k_\psi (\hat{\psi} - \psi), \quad (\text{A.25})$$

where we substitute the relaxation time by the heading gain $T^{-1} = k_\psi$ to get a parameter consistent with the full state feedback control of our Balancing Rider models.

Note, that [Hoogendoorn et al. \(2021\)](#) and [Kath \(2023\)](#) use exactly this formulation to separate lateral and longitudinal dynamics and bicycle dynamics within a social force framework. Other than our study, they allow speed modulation.

This shows that, under the given speed assumptions, the dynamics of the PP model are equivalent to those underlying various established social force models. If speed modulations are disregarded, these models are first-order heading trackers that implicitly combine the motion constraints, inertia, and control behaviors of the bicycle and the cyclist in a single parameter. The mechanical models explicitly model these characteristics, greatly increasing the detail of the models. Due to the relationship between the PP and established social force models, and our experiment's ability to sufficiently ensure the aforementioned speed assumptions, the PP serves as a suitable baseline for comparison with the mechanical models.

B. Parameters

To make our model reusable, we provide all necessary parameter values in this section.

B.1. Physical Model Parameter

We give the full physical parameters of the bicycle and the rider in the GitHub code repository ([chris-konrad/stochastic-balancing-rider](https://github.com/chris-konrad/stochastic-balancing-rider)).

B.2. Behavioral Model Parameters

This section provides the parameters of the stochastic control behavior model (Section 2.5). To use the model in Python, the supplementary material contains the Python package `rcid` (Rider Control IDentification), which provides the module `rcid.pole_modeling` to load the pole models and draw samples.

To use the full stochastic model, the full parameters of the GMM are necessary (Section B.2.2). If stochastic behavior is not required, the means of the behavioral distributions provide representative samples of the behavior observed in our experiments. Section (Section B.2.1) lists the location of the component means as functions of speed.

In both cases, the user should be aware that the models are only valid within the speed range of the experimental samples (approximately $2-4 \text{ m s}^{-1}$).

Component Means

Tables B.6 and B.7 show the locations of the mean poles per Gaussian component of the mixture models as a function of the forward speed $s_i(v) = \sigma_i(v) + j\omega_i(v)$. Our calibration resulted in one Gaussian component for BR_0 and two components for BR_1 .

Table B.6: Mean pole locations of the fitted BR_0 GMM pole model as functions of the forward speed.

Component	Symbol	Expression
0	$\sigma_2(v)$	$7.4774 + \frac{-7.5896}{\text{m s}^{-1}} v$
	$\sigma_{1/3}(v)$	$-0.6067 + \frac{-0.1089}{\text{m s}^{-1}} v$
	$\omega_{1/3}(v)$	$1.7882 + \frac{0.0411}{\text{m s}^{-1}} v$
	$\sigma_{0/4}(v)$	$-1.3282 + \frac{-0.0267}{\text{m s}^{-1}} v$
	$\omega_{0/4}(v)$	$5.3271 + \frac{0.0891}{\text{m s}^{-1}} v$

Table B.7: Mean pole locations of the fitted BR_1 GMM pole model as functions of the forward speed.

Component	Symbol	Expression
0	$\sigma_2(v)$	$-1.8536 + \frac{-1.6315}{\text{m s}^{-1}} v$
	$\sigma_{1/3}(v)$	$-0.1796 + \frac{-0.1597}{\text{m s}^{-1}} v$
	$\omega_{1/3}(v)$	$1.1730 + \frac{0.1425}{\text{m s}^{-1}} v$
	$\sigma_{0/4}(v)$	$0.6329 + \frac{-0.4482}{\text{m s}^{-1}} v$
	$\omega_{0/4}(v)$	$2.3198 + \frac{0.9166}{\text{m s}^{-1}} v$
1	$\sigma_2(v)$	$-0.3523 + \frac{-0.5884}{\text{m s}^{-1}} v$
	$\sigma_{1/3}(v)$	$-0.1027 + \frac{-0.2875}{\text{m s}^{-1}} v$
	$\omega_{1/3}(v)$	$1.7382 + \frac{0.3111}{\text{m s}^{-1}} v$
	$\sigma_{0/4}(v)$	$-0.4710 + \frac{-0.7008}{\text{m s}^{-1}} v$
	$\omega_{0/4}(v)$	$7.8118 + \frac{0.0609}{\text{m s}^{-1}} v$

Full GMM and Preprocessing Parameters

Tables B.8 and B.9 list all parameters of the speed-conditional Gaussian Mixture Models. The GMM parameters are presented in the transformed domain, after the preprocessing pipeline has been applied. Preprocessing consists of log-shift transformation, power transformation, and standardization applied to the pole data in that order. To convert samples drawn from the GMM, the inverse transformations of the pipeline must be applied in reverse order.

The supplementary material contains `.yaml` files of the parameters of all models and Python classes to conveniently import the model¹ in to a Python workspace and sample from the distribution².

C. Simulation

Simulations of traffic scenarios with our model may be created using the Python packages `rcid`³ and `cyclistsocialforces`⁴.

¹`rcid.pole_modeling.PoleModel.import_from_yaml()`

²`rcid.pole_modeling.PoleModel.sample_poles()`

³<https://github.com/chris-konrad/stochastic-balancing-rider>

⁴<https://github.com/chris-konrad/cyclistsocialforce>

Table B.8: Parameters of the fitted speed conditional BR₀ GMM pole model describing the distribution of rider control behaviors.

Parameter	Symbol	Value
GMM Parameters		
Covariances	Σ_0	[1.0301 0.1491 0.1186 0.0438 0.0134 0.0292
		0.1491 1.0386 0.1205 -0.6750 -0.5796 -0.6957
		0.1186 0.1205 1.1073 0.0884 -0.0404 -0.3375
		0.0438 -0.6750 0.0884 1.0015 0.5804 0.6086
		0.0134 -0.5796 -0.0404 0.5804 1.0262 0.6032
		0.0292 -0.6957 -0.3375 0.6086 0.6032 0.9851]
Means	μ_0	(0.0086, 0.0379, 0.0183, -0.0415, -0.0295, -0.1009)
Weights	π_0	1.0000
Log-Shift Transform Parameters		
shift	a	[0.5728, 0.1355, 0.1346]
sign	ξ	[-1.0000, -1.0000, -1.0000]
Power Transform Parameters		
Yeo-Johnson parameters	λ	[-0.0662, 0.9623, 1.4264, -0.5742, 1.2846, -0.7307]
Standard Scaler Parameters		
Means	μ	[1.3314, 2.4933, -0.2121, 0.8050, 0.2876, 1.0323]
Variances	σ^2	[0.1416, 2.0151, 0.7116, 0.1196, 1.1444, 0.0800]

Table B.9: Parameters of the fitted speed conditional BR₁ GMM pole model describing the distribution of rider control behaviors.

Parameter	Symbol	Value
GMM Parameters		
Covariances	Σ_0	[0.9950 0.1989 0.2217 0.1604 0.3565 0.3194
		0.1989 0.0708 -0.0354 0.0190 0.0250 0.0674
		0.2217 -0.0354 0.5955 0.0988 -0.0301 -0.0839
		0.1604 0.0190 0.0988 0.3643 0.1587 0.2643
		0.3565 0.0250 -0.0301 0.1587 0.4205 0.3456
		0.3194 0.0674 -0.0839 0.2643 0.3456 0.6029]
	Σ_1	[0.9734 0.2305 0.2903 0.2344 0.3032 0.0116
		0.2305 0.5325 0.2740 -0.2508 -0.1449 -0.1174
		0.2903 0.2740 1.2504 0.2732 -0.4973 -0.4916
		0.2344 -0.2508 0.2732 0.6731 0.1181 -0.0699
		0.3032 -0.1449 -0.4973 0.1181 0.6152 0.4596
		0.0116 -0.1174 -0.4916 -0.0699 0.4596 0.8296]
Means	μ_0	(0.3129, 1.1012, -0.3150, -0.9855, -1.0365, -0.7959)
	μ_1	(-0.1532, -0.5498, 0.1450, 0.4617, 0.5038, 0.2730)
Weights	π_0	0.3471
	π_1	0.6529
Log-Shift Transform Parameters		
shift	a	[0.8215, 0.1285, 0.0923]
sign	ξ	[-1.0000, -1.0000, -1.0000]
Power Transform Parameters		
Yeo-Johnson parameters	λ	[-0.0662, 1.3979, 1.4648, -0.0704, 1.8492, -0.6904]
Standard Scaler Parameters		
Means	μ	[1.3314, 0.9410, -0.2841, 1.1485, 0.7365, 1.1078]
Variances	σ^2	[0.1416, 1.2984, 0.6867, 0.2091, 1.0041, 0.0836]