Adaptive-Horizon Model Predictive Control for Modeling Anticipative Behavior in Cyclist Interaction

Yuke Huang

MSc Mechanical Engineering
BioMechanical Design Track
Delft University of Technology

Supervision
Daily Supervisor: Christoph M. Konrad
Dr. Jason K. Moore
Dr. Azita Dabiri

ABSTRACT

Anticipation is a vital component of cyclist behavior. Empirical evidence consistently suggests that cyclists do not merely react to immediate stimuli, but also actively predict future developments and adjust their motion accordingly. And such anticipative behavior varies across interaction scenarios. However, existing model-based studies provide only limited representations of how cyclists anticipate and how this anticipation varies, often relying on heuristic perception ranges, if-then decision rules, constant preview distance or fixed prediction horizons. Furthermore, anticipation-related parameters are rarely calibrated independently using real traffic data.

To address these gaps, this thesis proposes an adaptive-horizon Model Predictive Control (MPC) framework to explicitly represent cyclists' anticipative behavior. The proposed framework builds upon an existing MPC formulation for cyclist interaction [1] and extends it by introducing an adaptive prediction horizon that varies with the inter-cyclist distance. We formulate the prediction horizon as a bounded and distance-dependent variable that can continuously adapt.

We calibrate the model parameters using two-cyclist overtaking interactions extracted from the TUMDOT dataset [2]. The objective is to optimize MPC parameters to match the measured trajectory of the overtaking cyclist via a two-step Bayesian Optimization (BO) procedure. To disentangle the effects of the horizon-varying parameters from those of the remaining model parameters (background parameters), we implement the calibration sequentially. In step 1, we assume constant-horizon MPC and jointly calibrate the background parameters across all cyclist pairs within a given group. In step 2, the optimal background parameters are fixed, and the horizon-varying parameters are calibrated individually for each cyclist pair. We repeat both calibration steps under 6 BO configurations to reduce sensitivity to BO settings and to mitigate convergence to suboptimal local solutions.

The optimal constant-horizon model calibrated in step 1 achieves an average trajectory error of 0.3691 m relative to the measured trajectories, whereas the varying-horizon model calibrated in step 2 reduces this error to 0.16821 m, corresponding to an improvement of approximately 54.43%. In the correlation analysis of calibrated parameters, the practically used minimum prediction horizon is strongly correlated with its theoretical lower bound, indicating that this bound is often reached in practice. Also, the minimum horizon in all samples occurs at the overtaking point.

Furthermore, we observe moderate correlation between the calibrated horizon-related parameters and certain trajectory features. Higher speed is associated with smaller prediction horizon, indicating a shift toward more reactive behavior when cyclists move faster. And larger minimum inter-cyclist distance is associated with larger prediction horizon, suggesting that stronger anticipation may be associated with maintaining safer spatial margins during overtaking.

ACKNOWLEDGEMENTS

I would like to sincerely thank my supervisors for their invaluable support throughout this research. I am especially grateful to Christoph for his daily guidance, encouragement, and understanding. Whenever I got stuck, he consistently inspired me with constructive ideas that inspired me to find a way forward. When things did not go as planned, he was always kind and supportive, and when I felt discouraged by disappointing results, he reminded me that some progress had already been made. He also helped me stay focused on the bigger picture when I sometimes got lost in side branches. I would also like to express my sincere gratitude to Prof. Jason K. Moore for his prompt responses, thorough feedback, and constant availability. His every detailed comments on my early drafts, even when they were still far from polished, greatly improved the quality of my work, and his organized way of guiding both research and the lab gave me a clear sense of structure and direction. In addition, I am very grateful to Prof. Azita Dabiri for her invaluable advice on BO, MPC settings, and HPC tools. Her guidance on many occasions played a decisive role in pushing the research direction forward. As someone who started this work with very limited knowledge of some of these topics, I was able to grow quickly thanks to her support. Her advice was particularly valuable in helping me identify the limitations of my current model and guided me toward better research directions when I might otherwise have settled for suboptimal ones. I also appreciate Prof. Reinder Happee for joining my earlier meeting and for offering insightful suggestions, particularly regarding a simplifying assumption later used in the model.

I would like to express my deepest gratitude to my parents, whose love and support have always been unconditional. Throughout my life, they have tried to give me everything I needed with the simplest expectation for me to be healthy and happy. I am sincerely grateful for the happy childhood they gave me and for their continued love, encouragement, and support as I grew up. I would also like to thank my dearest cat, Baobao, for bringing so much comfort and gentleness into my life. When I used to stay up late at home, Baobao would often stay up with me, and the magical frequency of his purring always heals me. He has been a precious companion to us, caring for each family member in his own sweetest way. Finally, I am very grateful to my close friends for all the joy and support. The happy moments we shared, and the way we supported each other through stress and uncertainty mean a lot to me.

I would also like to thank myself for choosing courage over familiarity. I could have chosen a more comfortable path by staying closer to what I already knew in my undergraduate studies. Instead, I kept choosing courses and projects that were new, demanding but deeply interesting to me. That choice asked a great deal of me, but it also gave me more than I expected: new knowledge, new ways of thinking, and a deeper sense of what it means to grow. I am grateful not only for what I learned, but also for the moments when I became completely absorbed in the process, losing track of time and finding a quiet satisfaction in the work itself.

CONTENTS

1

INTRODUCTION

Urban cycling has become increasingly important in modern transportation systems [3]. As cyclists form a growing share of urban traffic, understanding their behavior is essential not only for improving cycling safety and infrastructure design, but also for developing microscopic traffic simulation models that can realistically reproduce cyclist interactions.

Among the many aspects of cyclist behavior, anticipation plays a particularly important role. Empirical studies suggest that cyclists often begin to adapt their speed or trajectory before an immediate conflict occurs, indicating that their behavior cannot be explained purely as a reaction to current stimuli. However, existing studies capture this anticipatory component only partially, for example through spatial perception ranges, rule-based decisions, or predictive control with a fixed planning horizon. This chapter therefore introduces the background of cyclist anticipation, reviews the main empirical and modeling studies related to this topic, and identifies the research gaps that motivate the adaptive-horizon MPC framework proposed in this thesis.

1.1 Background and Motivation of Studying How Cyclists Anticipate

1.1.1 Relevance of Studying Cyclist Behavior

The relevance of studying cyclist behavior is multifaceted. From a scientific perspective, the accuracy and reliability of microscopic traffic simulation models depend strongly on how well road users' behavior and interactions are understood and represented [4]. A better understanding of cyclist behavior can therefore help refine microscopic traffic models and improve their ability to reproduce realistic cyclist movement and interaction in urban traffic environments. From an engineering perspective, studying cyclists' anticipation is essential for designing safer and smarter transportation systems. A better understanding of cyclist behavior directly supports the development of safer infrastructure and the implementation of efficient safety measures [5].

Furthermore, studying cyclist behavior is important because it is closely related to cyclist safety [6], and therefore has significance in economic and environmental aspects. Economically, bicycle-related injuries in the United States resulted in a total societal cost of $237 billion between 1999 and 2013, including $209 billion from non-fatal injuries and $28 billion from fatal cases [7]. Also, bicycle-related traumatic brain injuries (TBI) impose a substantial burden on healthcare systems. In the Netherlands, the total medical cost of bicycle-related TBIs was estimated at €74.5 million annually [8], covering hospital care, emergency services, long-term rehabilitation, and productivity losses due to work absence [9]. Conversely, safe cycling can lead to positive economic outcomes: in Copenhagen, where 420,000 residents cycle annually, the city reports an estimated $253 million per year in healthcare savings due to improved public health [10]. Environmentally, cycling is a cornerstone of sustainable transport, yet its adoption depends heavily on perceived safety. As cities increasingly prioritize sustainability and transport efficiency, bicycles have become a significantly more prominent and relevant mode of urban mobility [11]. In a microsimulation study evaluating mode shifts in urban transport, Kilani and Bennaya reported that replacing private car trips with bicycles for short-distance, single-person travel can yield substantial environmental benefits. Specifically, they found that a modal shift of 27% from car use to cycling results in a notable decrease in pollutant gas emissions and fuel consumption, with reductions ranging from 15% to 19% [12]. Research also shows that concerns over safety (e.g., potential conflicts with motorists or general feelings of vulnerability) discourage some cyclists from choosing to cycle [13]. Therefore, improving cycling safety not only protects road users but also promotes greener transportation choices, contributing to urban sustainability goals.

1.1.2 Anticipation as Vital Component in Cyclist Behavior

According to Butz et al., anticipatory behavior refers to a process or behavior that is not only based on past and present information, but also on predictions, expectations, or beliefs about the future [14]. It goes beyond merely looking ahead and emphasizes how current actions are shaped by predictions or expectations of the future. Huesmann and Loffing further defined anticipation in the context of sports as a perceptual-cognitive process that informs action by enabling athletes to predict near-future events and align their responses accordingly, which helps to reduce motor costs under spatiotemporal pressure and ensures effective performance in dynamic environments [15].

Anticipation is an important component of cyclist behavior. In psychology and cognitive science, anticipation research suggests that actions are guided by internally generated expectations rather than triggered solely by external stimuli [16]. In cycling, anticipation is evident in many everyday interaction scenarios. For example, when avoiding an obstacle, cyclists typically do not wait until the last moment to react; instead, they begin adjusting their trajectory in advance based on the anticipated future conflict. Empirical studies also indicate that cyclist behavior exhibits clear anticipatory characteristics. In a controlled experiment, researchers introduced a physical bottleneck in the track layout and expected to observe cyclist queue formation patterns. However, such patterns did not occur, because cyclists began braking early in the preceding road segment after anticipating the upcoming narrowing, leading to larger inter-cyclist spacing, smaller densities and lower speeds than expected by the researchers [17]. These observations suggest that cyclist behavior cannot be understood as purely reactive, but must also account for anticipatory components.

1.1.3 Insufficient Understanding of How Cyclists Anticipate

Microscopic traffic simulation mainly focuses on the operational level (automatic, subconscious actions) and tactical level (conscious actions that allow road users to adjust behavior in response to traffic conditions) [18]. And researchers have pointed out that in the field of microscopic traffic simulation, the mechanisms behind cyclists' tactical decision-making are poorly understood in contrast to the well-studied operational level behavior [19]. Also, as noted in the broader anticipation literature [20], anticipation in sports is frequently acknowledged but rarely quantified. In other words, although anticipation is widely recognized as an important aspect of human behavior, our understanding of how cyclists anticipate in dynamic traffic situations remains incomplete. A more detailed examination of this limitation is therefore needed. The following section reviews existing studies related to cyclists' anticipation in order to clarify what current research has revealed and where important gaps remain.

1.2 Existing Research on Cyclists' Anticipation

1.2.1 Empirical Studies

Dummy Obstacle Avoidance: Oliver Lee et al. studied cyclists' comfortable zones during obstacle avoidance maneuvers through a series of controlled experiments. The participants were instructed to cycle toward a dummy obstacle placed 60 m ahead, at three speeds representing different stability levels (12 km/h for laterally instable, 17 km/h for self-stable and 22 km/h for mildly unstable) [21]. They were instructed to perform either comfortable braking (stopping before the dummy) or comfortable steering (overtaking the dummy) as they would in real traffic.

During comfortable braking, faster cyclists started braking further from the obstacle, while deceleration profiles remained similar across speeds. This suggests that cyclists may compensate for higher speed by braking earlier rather than braking harder, indicating a longer anticipation horizon at higher speeds. In comparison, during comfortable steering, cyclists started steering at a relatively consistent distance from the obstacle regardless of speed. A limitation of this experiment is that the dummy obstacle may not fully represent the dynamics of real traffic conditions.

Gaze Distribution Study: Eye-tracking studies provide further insight into the perceptual foundations of cyclists' anticipatory behavior. In a controlled experiment with nine conditions (three lane widths × three speeds), Vansteenkiste et al. analyzed cyclists' gaze allocation patterns and identified several consistent tendencies across individuals [22]. The dominant regions of fixation (fixation refers to a brief period during which the eyes remain relatively stable to acquire visual information) were path fixations (approximately 40%) and goal fixations (approximately 41%), indicating that cyclists allocate visual attention to both near and far regions of the environment. The results suggest that cyclists integrate near and distant cues during locomotion, which is consistent with the core concepts of reactive and perceptive range models [23] emphasizing the integration of reacting and predicting during cycling.

In addition, gaze distance was found to increase with cycling speed, implying that cyclists tend to look further ahead when riding faster. One explanation proposed by the authors is the concept of a visual buffer in human locomotion, referring to the temporal interval required for visual information to be processed and translated into motor action. This buffer typically ranges between 0.80–2.00 seconds [24]. These findings indicate that visual perception operates within a temporally structured anticipation window rather than a purely instantaneous feedback loop.

Moreover, lane width was shown to influence gaze allocation. On narrower lanes, gaze shifted closer to the near pathway, suggesting increased reliance on foveal (central) vision for precise lane-keeping. This observation is consistent with findings from vehicle-driving studies [25], where increased task demands lead to reduced dependence on peripheral vision and greater focus on central visual information.

Based on these observations, Vansteenkiste et al. proposed a gaze constraint model for cyclists, inspired by the two-level steering framework used in vehicle control. The model distinguishes between direct control processes for maintaining stability and maneuvering, and anticipatory processes for navigation and hazard identification. Attention demand increases with both the need for precise control and the need to anticipate upcoming environmental changes, as shown in Figure 1.1.

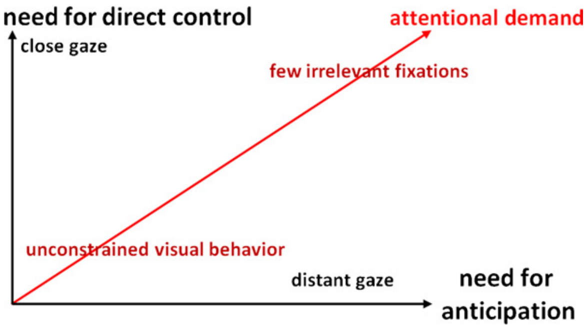

Figure 1.1: Gaze constraint model from [22] with axes for direct control and anticipation needs

These findings suggest that cyclists' visual perception is not a fixed sensory process, but one that is shaped by environmental context, task demands, speed, and physiological constraints of the visual system. Variations in gaze allocation and fixation distance therefore reflect adaptive adjustments to perceptual limitations, which in turn influence the effective scope and horizon length of anticipation during cycling.

Overtaking Interaction: Sarosh I. Khan et al. studied the characteristics of overtaking interactions and bi-directional interactions using data collected on 3-meter-wide bicycle-exclusive paths [26]. During overtaking interactions, cyclists were observed to travel an average distance of 91.4 m, indicating a long-range anticipatory process rather than a reactive, last-moment decision. Within a single overtaking event, the researchers assumed constant speed for the overtaking and overtaken cyclists, and the average speed of the overtaking cyclists during the entire maneuver exceeded that of the overtaken cyclists by 9.38 km/h for the 29 observed overtaking events. As shown in Figure 1.2, across different overtaking events, a positive correlation was found between the speeds of the overtaking and overtaken cyclists, with the fitted relationship $U_{\text{overtaking}} = 16.67 + 0.50U_{\text{overtaken}}$. And a negative correlation was found between the speed of the overtaken cyclists and their relative speed. This can be explained by two possible mechanisms: faster overtaken cyclists may increase the perceived difficulty of the overtaking task, prompting earlier anticipation of overtaking cyclists and resulting in longer overtaking durations; alternatively, slower overtaken cyclists are more laterally unstable, which encourages overtaking cyclists to adopt a higher relative speed to maintain their own lateral stability. In addition, the study identified a comfortable speed threshold for overtaking: if the initial speed difference was less than 5 ft/s, overtaking cyclists tended to accelerate within the first 1–2 seconds to ensure a sufficient speed gap. However, if the speed difference exceeded 5 ft/s within the first 1.5 seconds, overtaking cyclists typically did not increase their speed further. Furthermore, during overtaking interactions, overtaken cyclists shifted closer to the right edge of the bicycle path, reducing their average lateral distance from 0.86 m to 0.58 m.

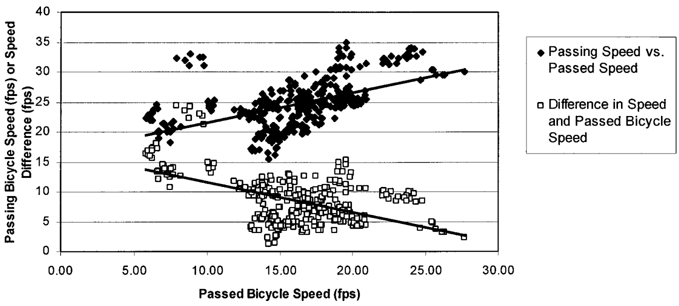

Figure 1.2: Passed cyclist speed vs. passing speed and speed difference measured in [26]

Bi-directional Interaction: Yufei Yuan et al. investigated one-on-one bi-directional cyclist interactions in controlled experiments [27], where they quantified how much cyclists deviated during bi-directional interactions using unhindered cycling as a baseline scenario. Both cyclists began evasive maneuvers when they were around 30 meters apart. This indicates that cyclists anticipate potential collisions well in advance rather than reacting at the last moment. Also, the fact that both cyclists initiated adjustments suggests a mutual anticipatory strategy, which is different from pedestrian interactions studied in [28], with typically only one pedestrian deviating.

During the encounters, the rate of change for longitudinal speed and speed magnitude were negligible, whereas the rate of change for lateral speed was noticeable. And the maximum lateral speeds in both directions were of similar magnitude. This suggests that cyclists prefer to preserve their pace and rely on side movements to avoid collision.

This study has two main limitations. First, the sample size is relatively small, involving only twelve participants. Second, a learning effect was observed during the experiment: as participants watched others perform, they may have adjusted their own behavior accordingly. This could have introduced bias into the trajectory data and should be considered when interpreting the results.

1.2.2 Model-based Studies

Perception Combined with Rule-based Decision-making: X. Liang et al. proposed a conceptual framework consisting of sight range, reactive range and perceptive range [29]. Sight range is defined by field of view $\theta_i(t)$ (FOV) and fixation point $l_i(t)$ (FP). As shown in Figure 1.3a and Figure 1.3b, these two parameters play a crucial role in shaping the geometry of the reactive range and the perceptive range. According to equation (1.1), when speed $v_i(t)$ increases, filed of view becomes narrower while fixation point increases (μ_i denotes the fixation change coefficient).

$$\begin{aligned} \theta_i(t) &= 179.4e^{-0.05022v_i(t)} \\ l_i(t) &= \mu_i \|\vec{v}_i(t)\| + a_i \end{aligned} \tag{1.1}$$

Reactive range refers to the spatial area in which cyclists take evasive actions to avoid collisions with intruding obstacles or other road users. This range is modeled as a asymmetric convex curve as shown in Figure 1.3a. According to equation (1.2), its semi-major axis length $L_i(t)$ expands with $v_i(t)$, reaction time t_i, reaction coefficient α_i and a_i, and its semi-minor axis length $b_i(t)$ increases with $v_i(t)$, safe reaction coefficient β_i and b_i.

$$\begin{aligned} L_i(t) &= \alpha_i t_i \|\vec{v}_i(t)\| + a_i \\ b_i(t) &= \beta_i \|\vec{v}_i(t)\| + b_i \end{aligned} \tag{1.2}$$

Perceptive range refers to the spatial area in which cyclists sense the states of obstacles and other road users. As shown in Figure 1.3b, the perceptive range is modeled as an elliptical region with the cyclist's pronasale as one focus. As shown in equation (1.3), where $S_i^{\text{safe}}(t)$ (the variable denoting the local half-width of the ellipse at the focus) increases with $v_i(t)$, b_i and personal sensory coefficient γ_i. And FP $l_i(t)$, as is already defined in sight range, equals to the sum of the ellipse semi-major axis length a_i' and center-to-focus distance c_i'. Therefore, when the cyclist rides faster, the perceptive range extends in both longitudinal and lateral direction. Here, the longitudinal direction refers to the direction of the cyclist's movement, while the lateral direction is perpendicular to it.

$$S_i^{\text{safe}}(t) = \gamma_i \|\vec{v}_i(t)\| + b_i \tag{1.3}$$

We can find concepts similar to the elliptical-like reactive range and perceptive range in other model-based papers. For example, Johannes S. Brunner et al. modeled cyclists' decision-making process applying the concept of safety region (the clearance area that the cyclists tend to maintain) and consideration range (the space considered for cyclists for their operational-level behavior) [30]. Both safety region and consideration range enlarge as the cyclist rides faster, as shown in Figure 1.4, where l_b, ω_b denote the longitudinal and lateral size of the bicycle-rider unit, v_i is velocity and α, β denote coefficients.

The integration of these spatial ranges into rule-based decision-making and strategy-selection mechanisms typically begins with a perceptual scan of the environment and unfolds through sequential decision-making stages that ultimately culminate in movement execution. This process is illustrated in a more detailed way as follows.

In the microscopic dynamic model, Liang et al. integrated the perception ranges into a layered decision-making structure [23]. As shown in Figure 1.5), cyclists scan within

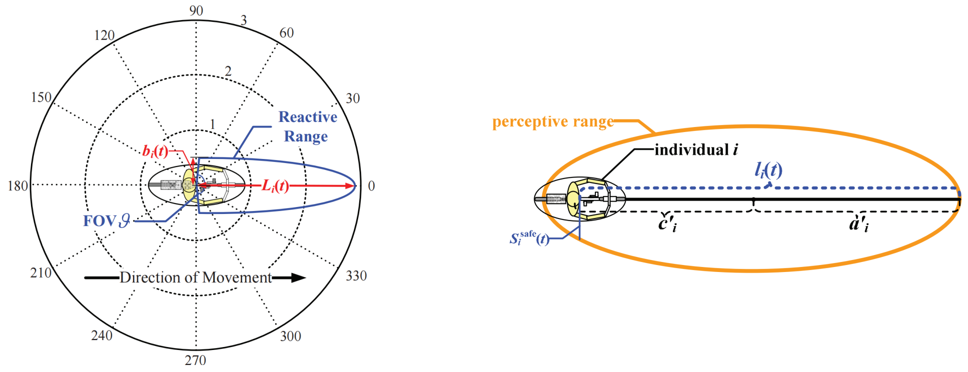

(a) Reactive range modeled as an asymmetric convex curve extending from the cyclist, shaped by FOV θ, with semi-major axis $L_i(t)$ and semi-minor axis $b_i(t)$.

(b) Perceptive range modeled as an ellipse with the cyclist's pronasale as one focus, with $S_i^{\text{safe}}(t)$ denoting the local half-width of the ellipse at the focus, semi-major axis a'_i, and center-to-focus distance c'_i.

Figure 1.3: Illustrations of cyclist interaction-related ranges adopted from [29]: (a) reactive range; (b) perceptive range.

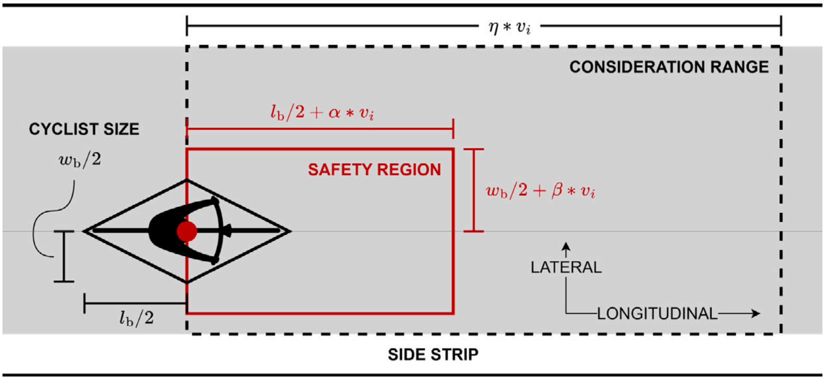

Figure 1.4: Safety region and consideration range from [30], both expanding with cyclist speed v_i

the perceptive range, evaluate spatial parameters such as flow density when an object or obstacle enters the reactive range, and make decisions such as moving directions and desired speeds. Once the decision is made, it is implemented through the psychological-physiological force (PPF) model governing social forces among cyclists. For example, the driving force is influenced by the desired speed or group speed, and the collision avoidance force is activated when obstacles enter the cyclist's reactive range and acts in the direction determined by the selected riding strategy. In summary, anticipatory decisions are the direct outcome of perceptual input and cognitive processing, which guide the actions that follow. This layered control structure, consisting of perception, strategy selection, force generation and motion, demonstrates that anticipation functions as a cognitive pre-processor. It enables the selection of optimal response parameters such as direction and speed, ultimately determining how the cyclist moves.

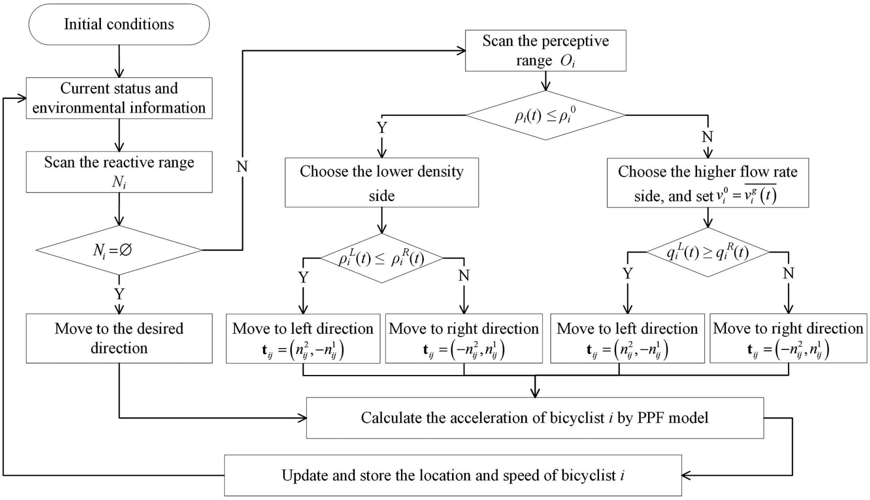

Figure 1.5: Flowchart of cyclists' riding strategy selection and motion update based on input in the reactive and perceptive ranges and the psychological-physiological force (PPF) model from [23]

In the safety region and consideration range model, Brunner et al. built a three-level decision and action framework to describe how cyclists use anticipation to guide both their motion planning and execution ([30]). As shown in Figure 1.6, early perceptual decisions at Level A (maneuver) shape movement strategies at Level B (movement planning), which in turn determine the parameters for actions at Level C (acceleration). At Level A, cyclists scan their consideration range and classify surrounding riders based on speed in order to decide whether overtaking is necessary. If so, they identify downstream lateral gaps that afford the most efficient overtaking. These inputs are passed to Level B, where cyclists compute feasible lateral speeds and select appropriate leaders to follow to reach the anticipated position, adjusting their plan based on constraints like achievable lateral speeds. Finally, the outputs from the previous stages feed into calculating the longitudinal acceleration or deceleration required in Level C. For example, using the Necessary Deceleration Model (NDM), cyclists adjust their speed to maintain safe distances. The parameters used are determined by earlier anticipatory assessments (e.g., acceleration in Level C is a function of the distance to the leader, and the leader is chosen in Level B based on the perception in Level A).

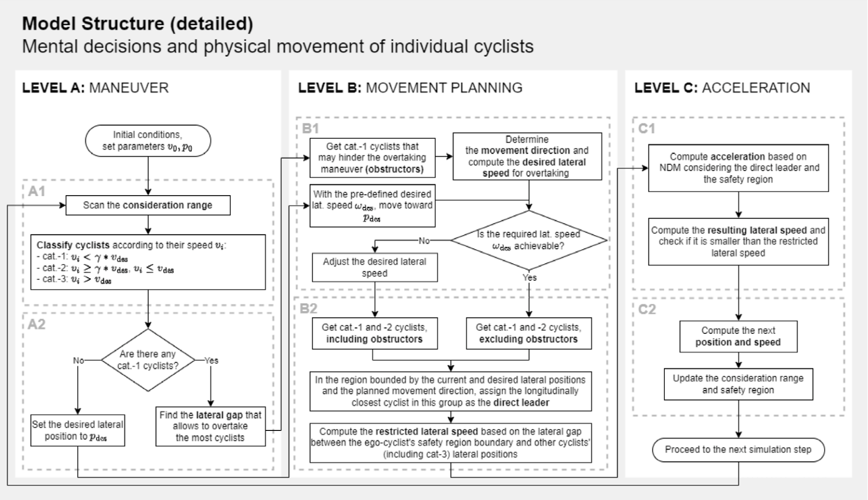

Figure 1.6: Three-level cyclist decision model, comprising of Level A (maneuver), Level B (planning) and Level C (acceleration) (source: [30])

Preview Control Framework: Although this thesis focuses on cyclist interactions, the modeling framework is rooted in control-theory concepts, which have been widely studied in the broader vehicle dynamics and control literature. Therefore, we want to draw methodological inspiration from these related domains, while adapting the framework to the specific anticipatory and interactive characteristics of cyclists.

The preview control approach models anticipation by minimizing preview error, which is defined as the deviation from the desired states after a preview time. This error typically includes components such as current deviation, rate of change of deviation, and correction terms based on road curvature, preview distance, and lateral/angular offsets. Control mechanisms use these errors to generate anticipatory steering commands via various architectures, including linear feedback loops (e.g., PID), fuzzy logic controllers, sliding-mode control, and hierarchical control models. The following five studies illustrate how different preview control strategies have been applied to model cyclists' or drivers' anticipation under various scenarios.

Cho et al. integrated the preview control mechanism with a closed-loop human–vehicle control system to model how a driver anticipates and steers a vehicle along curved paths [31]. As shown in Figure 1.7a, the preview error $\eta(t+t_a)$ is defined as the lateral deviation between where the driver is and where the driver is supposed to be on the desired path after a preview distance L. The desired path is the curved path with radius R. As shown in equation 1.4 where t_a is preview time, the preview error $\eta(t+t_a)$ consists of $\eta(t)$ (current deviation), $\frac{L}{u_x}\dot{\eta}(t)$ (increment in deviation assuming cyclists are moving at constant speed u_x for the preview distance) and correction term $-R\left(1-\cos\frac{L}{R}\right)$. It should be noted that a specific formula for the preview distance L is not provided (i.e., it is assumed to be a fixed value in the simulations), but the authors suggested that L increases with the square of velocity and is constrained by visibility limits and road geometry.

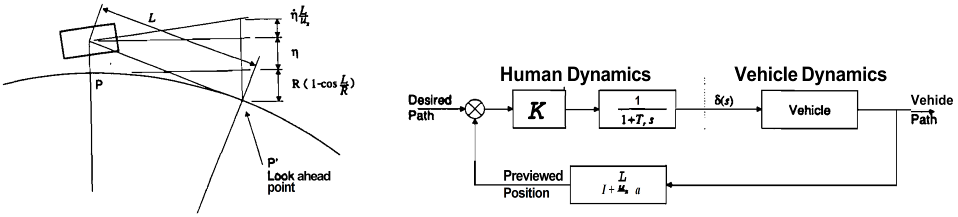

(a) Preview error on a curved path as current deviation plus anticipatory correction.

(b) Preview control model with feedback from previewed position to steering.

Figure 1.7: Illustration of the preview control framework from [31]: (a) preview error on a curved path; (b) block diagram of the preview control model.

$$\eta(t+t_a) = \eta(t) + \frac{L}{u_x}\dot{\eta}(t) - R\left(1 - cos\frac{L}{R}\right) \tag{1.4}$$

$$T_r\dot{\delta}(t) + \delta(t) = -K\eta(t+t_a) \tag{1.5}$$

This preview error is processed by the driver who is modeled as a first-order linear feedback controller with a proportional gain K and reaction lag T_r. And the steering input is given by equation 1.5. As shown in Figure 1.7b, in the feedback loop, the steering angle δ is dynamically adjusted based on the preview error. In short, the preview control mechanism provides the anticipatory component, while the closed-loop system ensures that the executed action (steering) is dynamically corrected in response to this forecast. Although the driver model is formally implemented as a feedback control loop, it effectively incorporates feedforward characteristics, since the steering input is based on a previewed lateral deviation at a future time. This anticipatory correction reflects a hybrid control strategy, where future error is estimated and acted upon before it occurs, rather than responding solely to current deviation.

T.-K. Da et al. implemented preview control within a two-level control architecture [32]: an inner loop ensuring roll-angle stability, and an outer loop managing path tracking using preview information. Specifically, the outer loop uses a fuzzy logic controller (FLC) to compute the desired roll angle based on preview errors. As shown in Figure 1.8, preview error consists of lateral error e_L (i.e., the distance from a preview point ahead along the current velocity vector to the target path) and angular error e_θ (i.e., the angle between the bicycle's current heading and the tangent of the path at the preview point). The preview distance L_{pre} is determined by: $L_{\text{pre}} = V_x T_{\text{pre}}$, where V_x is the bicycle speed and T_{pre} is a fixed preview time, ensuring that the system adapts to different speeds. The FLC processes these preview errors using 169 fuzzy rules to generate a reference roll angle θ_{ref}, which is fed into the inner control loop. Meanwhile, the inner loop employs a sliding mode controller (SMC) that uses θ_{ref} to generate steering torques that stabilize the roll angle. This inner controller guarantees convergence through a Lyapunov-based sliding surface, ensuring robustness to disturbances and speed variation. Overall, this two-level structure enables anticipatory path-following behavior: while the outer loop utilizes forward-looking spatial information to anticipate path deviations, the inner loop ensures that the corresponding balancing action is executed precisely.

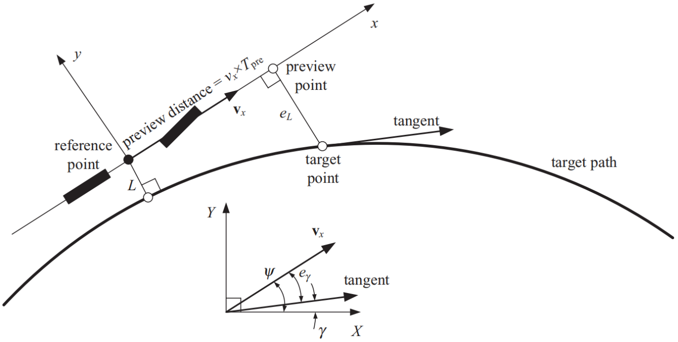

Figure 1.8: Preview errors defined by lateral and angular deviations between the reference and preview points on a curved path (source: [32])

In [33], preview control is incorporated into a model predictive control (MPC) framework to improve path tracking. When the driver is at the position (X_o^k, Y_o^k), instead of selecting the nearest point (X_r^k, Y_r^k) on the reference path to be the control reference state point, the controller selects a point (X_r^{k+i}, Y_r^{k+i}) located at a preview distance L_{pre} ahead of the vehicle. This approach allows the controller to anticipate future deviations rather than react to current errors, effectively mimicking human-like look-ahead behavior in the prediction process.

The preview distance is defined as $L_{\text{pre}} = V_{OX} \cdot t_{\text{pre}}$, where V_{OX} is the longitudinal velocity of the vehicle and t_{pre} is the preview time. This preview time is not fixed. Rather, it adapts according to the current lateral deviation and road curvature. Specifically, the maximum and minimum bounds of the preview time are set as $t_{\text{pre}\,\max} = 0.02V_{OX}$ and $t_{\text{pre}\,\min} = 0.016V_{OX}$, respectively. The actual preview time t_{pre} is calculated using the Equation 1.6, where e_1 denotes the lateral tracking error, $e_{\max}$ is the maximum allowable lateral error, ρ is the instantaneous road curvature, and $\rho_{\max}$ is the maximum expected curvature. The coefficients K_1 and K_2 are weighting factors that define the relative influence of lateral deviation and curvature on the preview time, and they satisfy the condition $K_1 + K_2 = 1$. Therefore, when the lateral error is large or the curvature is sharp, the preview time becomes shorter, allowing faster control responses.

$$
\begin{aligned}
t_{\text{pre}\,\max} &= 0.02V_{OX} \\
t_{\text{pre}\,\min} &= 0.016V_{OX} \\
t_{\text{pre}} &= \begin{cases} t_{\text{pre}\,\max} - K_1\, t_{\text{pre}\,\max} \cdot \left(\left|\dfrac{e_1}{e_{\max}}\right|\right) - K_2\, t_{\text{pre}\,\max} \cdot \left(\left|\dfrac{\rho}{\rho_{\max}}\right|\right), & \text{if } t_{\text{pre}} > t_{\text{pre}\,\min} \\ t_{\text{pre}\,\min}, & \text{if } t_{\text{pre}} \le t_{\text{pre}\,\min} \end{cases} \qquad (1.6) \\
L_{\text{pre}} &= V_{\text{OX}}\, t_{pre}
\end{aligned}
$$

To evaluate the effectiveness of the proposed preview-based MPC strategy, the authors conducted co-simulations using Simulink/CARSIM. The preview-enhanced MPC was compared against traditional MPC under identical simulation conditions. Results demonstrated that the proposed method significantly reduced lateral tracking error and achieved smoother steering inputs, particularly under curved path scenarios. However, the reference trajectory used in the simulations was predefined and the paper does not describe how such a reference trajectory should be generated.

Edelmann et al. proposed a hierarchical rider control model [34] to represent cyclists' path tracking tasks as shown in Figure 1.9, which integrates an anticipatory feed-forward layer, a compensatory feedback layer, and an inner stabilizing control loop, each addressing different aspects of rider behavior from preview-based planning to real-time dynamic stabilization. More specifically, anticipatory feed-forward control uses the road curvature κ as input to predict the required steady-state roll angle φ_{ff}, allowing the rider to prepare for upcoming turns. Simultaneously, the compensatory feedback control level compares the predicted lateral deviation y_B obtained from an internal model of the bicycle to the desired lateral deviation y_{Bd}, which is assumed to be zero, to calculate a deviation error e_y. This error is translated into a corrective roll angle φ_{fb}, which compensates for external disturbances or model inaccuracies. Both anticipatory and feedback roll angles are summed to form the total desired roll angle φ_d. To realize this desired roll, the stabilizing feedback control level calculates a steering torque M_δ, which is applied to the bicycle model using full-state feedback of the actual system variables denoted as $\underline{x}$. The bicycle's output includes the actual lateral position y, which is used in subsequent deviation predictions, closing the feedback loop.

This control strategy captures key aspects of human motor control (i.e., anticipation, correction and stabilization), making it a suitable framework for simulating realistic rider behavior in dynamic environments. The authors evaluated the performance of the proposed controller in a cornering scenario involving a gradual entry into a constant-radius turn and a gradual exit, modeled using clothoid segments. Throughout the maneuver at a velocity of 3 m/s, the maximum lateral deviation did not exceed approximately 0.3 m, demonstrating that the controller can effectively anticipate and respond to smooth and continuous changes in road curvature.

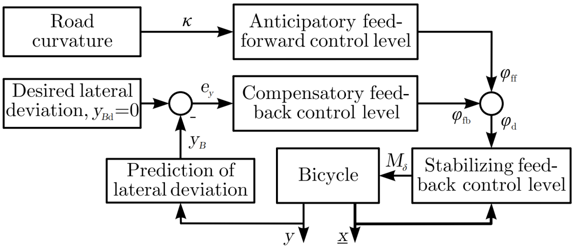

Figure 1.9: Rider control model with anticipatory, compensatory, and stabilizing loops (source: [34])

Utility/Optimization-based Representation Another class of model-based approaches represents anticipation through optimization principles. In these models, cyclists are assumed to evaluate future outcomes by maximizing a utility function or minimizing a cost function subject to constraints. Anticipation is therefore embedded in the objective formulation itself: future states, potential conflicts, and motion efficiency are quantified and weighed before an action is selected.

Jianqiang Li et al. developed a two-layer framework to predict cyclists' trajectories [35]. In the trajectory generation layer, candidate trajectories are provided by first applying kinematic laws to determine feasible end points, followed by deep learning to generate complete trajectories. In the trajectory selection layer, the cyclists are modeled as agents making decisions by maximizing utility. As shown in equation 1.7, U_{risk} denotes path risk value, which involves risks arising from static obstacles, dynamic obstacles and the perceived risk of intrusion. U_{effi} denotes path efficiency value defined as the velocity component in the intended direction of movement. φ is the weight of risk-efficiency preference (i.e., to what extent cyclists favor obstacle-avoidance over the pursuit of efficiency). And the trajectory S_i that maximizes the overall utility U is chosen.

$$\arg\max U(S_i) = \varphi U_{\text{risk}} + (1-\varphi)U_{\text{effi}} \tag{1.7}$$

Building on the previous study [35], Jianqiang Li et al. further developed a self-learning framework in their follow-up work [36], where 2 hidden psychological variables are incorporated as the self-learning parameters: aggressiveness level (AL) and rationality level (RL). To be more specific, AL corresponds to φ in equation 1.7. Rather than being the constant 0.5 in [35], in this work it is identified by comparing the observed trajectory with candidate trajectories, selecting the closest match, and inferring which value makes this trajectory utility-optimal. AL is then identified as the value of φ that maximizes the utility of this specific trajectory. In addition, RL measures the extent to which the cyclist fails to select the utility-maximizing option. And ut is identified by comparing the chosen trajectory's utility ranking among all candidates' utility (higher-ranked choices indicate higher rationality). Their results suggest that cyclists do not always behave as utility-maximizing agents, but are instead constrained by bounded rationality.

Hoogendoorn et al. used Model Predictive Control (MPC) to study cyclist interaction [1], where cyclist optimize control inputs by minimizing a cost function that encodes behavioral priorities such as effort, efficiency, and collision avoidance, subject to constraints. Anticipation is thus operationalized through the forward evaluation of predicted future states within the planning horizon. Their results indicate that increasing the prediction horizon leads to earlier and stronger evasive maneuvers, resulting in safer interactions. However, within each individual simulation, the prediction horizon is predefined and remains fixed throughout the interaction. Although different horizon lengths are compared across scenarios, the horizon does not adapt dynamically during the simulation. In this thesis, the MPC framework is extended in three aspects: the transition dynamics are represented using bicycle dynamics, the prediction horizon is adapted according to interaction scenarios, and the model parameters are calibrated based on measurement to learn about the horizon-varying parameters.

1.3 Gaps in Existing Studies and Positioning of This Thesis

Despite growing research related to cyclist anticipation, several important gaps remain in the existing literature. First, empirical studies have provided valuable insights into the anticipatory nature of cyclist behavior, such as early braking before bottlenecks and gaze distribution patterns. These studies demonstrate that cyclists often adjust their actions in advance based on anticipated future conditions and such adjustment is scenario-specific. However, the internal anticipation mechanism cannot be directly measured and most of the empirical observations alone do not provide a clear computational representation of how anticipation is generated and translated into motion.

Second, many existing model-based approaches represent anticipation through spatial perception concepts, such as reactive ranges, perceptive ranges, safety regions, or consideration zones. While these frameworks capture how cyclists sense and evaluate surrounding traffic conditions, they often rely on predefined decision rules or heuristic strategies to determine actions. As a result, anticipation is represented indirectly through rule-based logic rather than through a continuous predictive planning mechanism.

Third, preview control approaches provide an explicit representation of anticipation by introducing a look-ahead mechanism. In these frameworks, control actions are generated based on predicted future deviations at a preview point ahead of the current position. While preview control successfully captures the idea that human operators respond to anticipated future states rather than only current errors, multiple existing formulations still assume a predefined preview distance or preview time. As a result, the anticipation horizon is often treated as a fixed parameter, rather than as a dynamic quantity that adapts to changing interaction conditions.

Finally, optimization-based approaches such as Model Predictive Control (MPC) incorporate anticipation by evaluating predicted future states within a planning horizon. However, in existing studies the prediction horizon is predefined and remains fixed during the interaction. This assumption neglects the possibility that cyclists may dynamically adjust how far ahead they plan depending on the context of the interaction.

Finally, although some existing modeling frameworks have been calibrated using empirical data, anticipation-related factors are rarely calibrated explicitly as distinct model components. As a result, it remains unclear how the anticipation mechanisms embedded in these models correspond to actual cyclist behavior observed in real interactions.

To address these gaps, this thesis proposes an adaptive-horizon Model Predictive Control (MPC) framework to explicitly represent cyclists' anticipative behavior. Instead of assuming a fixed prediction horizon, the horizon is formulated as a bounded and distance-dependent variable that can continuously adapt. The model parameters are calibrated using overtaking interaction in real traffic, allowing the anticipation mechanism to be quantitatively inferred from observed cyclist behavior. In this way, the proposed framework aims to bridge empirical observations of anticipative behavior and optimization-based modeling approaches, providing a more interpretable and realistic representation of cyclists' anticipation.

The remainder of this thesis is organized as follows. Chapter 2 presents the adaptive-horizon MPC formulation, data collection and calibration framework. Chapter 3 illustrates calibration results and analyzes the behavioral implications of the identified parameters. Chapter 4 concludes the thesis and discusses limitations and directions for future research.

2

METHODOLOGY

We investigate cyclists' anticipation behavior within a model predictive control (MPC) framework. We use MPC because it explicitly captures anticipation through the prediction horizon, and its cost function provides a natural way to represent the multiple weighted objectives underlying cyclist behavior. By assigning different weights to components such as obstacle avoidance, goal reaching, and control effort minimization, the model can reflect different behavioral tendencies in how cyclists balance safety, progress, and maneuver smoothness.

The MPC framework is based on the work of Hoogendoorn et al. [1] with several modifications introduced. In particular, the original unicycle transition dynamics are replaced by bicycle-rider dynamics, the fixed prediction horizon is extended to an adaptive prediction horizon, and the model is calibrated using real-world trajectory data.

Anticipation is represented by an adaptive prediction horizon, which serves as a behavioral control knob spanning cyclists' responses from purely reactive behavior to strongly anticipatory behavior. This reflects the limited and context-dependent foresight of cyclists in real traffic scenarios. The parameters governing the variation of prediction horizon are identified through calibration against measured overtaking cyclist trajectories extracted from interaction pairs in the TUMDOT dataset [2] using a two-step Bayesian Optimization (BO) procedure. In the first step, we assume constant-horizon MPC and calibrate parameters unrelated to horizon adaptation (background parameters), which are taken to be shared across all interaction pairs in a group. In the second step, the optimal background parameters are fixed, and we calibrate horizon-varying parameters individually for the overtaking cyclist in each interaction cyclist pair.

The proposed MPC and calibration framework is based on the following assumptions:

1. We assume the parameters unrelated to the prediction horizon to be the same across all cyclists. Accordingly, individual differences in behavior are captured solely through horizon-related parameters.

2. We allow the prediction horizon to vary over time, while we assume all other MPC parameters to remain constant during an interaction. This assumption is motivated by prior work [36], which relates parameters weighting safety (e.g., obstacle avoidance) and efficiency (e.g., goal reaching) to people's psychological traits (e.g., conservative versus aggressive behavior). We assume that such traits remain constant over the short duration of the cyclist interactions considered in this study (approximately 10 seconds).

The following sections describe the main components of the proposed methodology. First, the MPC framework is introduced, including its general formulation, the rider–bicycle dynamics and state-transition model, the cost function design, and the imposed constraints. Next, the adaptive prediction horizon model is presented as the key mechanism for representing cyclists' anticipative behavior. This is followed by a description of the MPC workflow, including the solution procedures and pseudo code representations. The methodology chapter then introduces the data collected for model development and evaluation. Finally, the calibration procedures based on Bayesian Optimization (BO) is described, including the BO setup, workflow, and the two-step BO procedures.

2.1 MPC Framework

2.1.1 MPC General Formulation

In general, MPC solves a finite-horizon optimal control problem at each time step by optimizing a sequence of control inputs $\vec{U}$ to minimize a cumulative cost function J, subject to system dynamics and state and input constraints. The standard MPC formulation can be written as in Equation (2.1), where $\vec{X}$ denotes the system state and $\vec{U}$ the control input. The system is subjected to equality constraints $F(\vec{X},\vec{U}) = 0$ and inequality constraints $G(\vec{X},\vec{U}) \leq 0$. The cost function J is accumulated over the prediction horizon N. At each time step, only the first element of the optimal control sequence is applied, following the receding horizon principle [37].

$$\begin{aligned} &\underset{\vec{U}}{\arg\min}\; J(\vec{X},\vec{U}) \\ &\text{s.t.} \begin{cases} F(\vec{X},\vec{U}) = 0, \\ G(\vec{X},\vec{U}) \leq 0 \end{cases} \end{aligned} \tag{2.1}$$

In the following, the general MPC formulation in Equation (2.1) is instantiated for the cyclist behavior modeling problem by specifying the dynamics, cost design, constraints, adaptive prediction horizon, and workflow.

2.1.2 Rider–Bicycle Dynamics and State Transition Model

The system state and control input of the rider-bicycle system are defined as in Equation (2.2). The state vector $\vec{X}$ consists of the cyclist's roll angle ϕ, roll rate $\dot{\phi}$, steer angle δ, steer rate $\dot{\delta}$, yaw angle ψ, position (x,y) and velocity v. The control input vector $\vec{U}$ consists of the cyclist's acceleration a and reference yaw angle ψ_c. The derivation of the state dynamics $\dot{\vec{X}} = f(\vec{X},\vec{U})$ is explained as follows.

$$\begin{aligned}\vec{X} &= \begin{bmatrix}\phi & \delta & \dot{\phi} & \dot{\delta} & \psi & x & y & v\end{bmatrix}^{\mathrm{T}},\\ \vec{U} &= \begin{bmatrix}\psi_c & a\end{bmatrix}^{\mathrm{T}}.\end{aligned} \tag{2.2}$$

The position and speed (last 3 elements of $\vec{X}$) evolve according to standard kinematic relations given by Equation (2.3). And the lateral dynamics (first 5 elements of $\vec{X}$) evolve based on the linearized Carvallo–Whipple dynamics which is widely regarded as the canonical benchmark model for analyzing bicycle stability and control [38], combined with a full-state feedback controller [39].

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix}x\\ y\\ v\end{bmatrix} = \begin{bmatrix}v\cos\psi\\ v\sin\psi\\ a\end{bmatrix} \tag{2.3}$$

More specifically, according to Whipple–Carvallo benchmark model linearized about the upright, straight-running equilibrium [38], the dynamics governing yaw angle ϕ and steer angle δ is shown in Equation (2.4). The generalized coordinate vector is $[\phi,\ \delta]^{\top}$ and the generalized forcing term is $[T_\phi,\ T_\delta]^{\top}$ consisting of the roll and steer torques. The matrix M is the symmetric mass matrix and captures the inertial coupling between roll and steer. The damping-like matrix vC_1 is linear in speed v and reflects skew-symmetric gyroscopic couplings arising from roll rate and steer rate. The stiffness matrix comprises a gravity-driven part gK_0 and a speed-induced part v^2K_2 that captures gyroscopic and centrifugal effects. Given that steering dominates bicycle control [40], we restrict the lateral control input to steer torque only ($T_\phi = 0$). The kinematic relation linking yaw rate $\dot{\psi}$ to steer angle δ and steer rate $\dot{\delta}$ under the no-slip rolling constraints is shown in Equation (2.5) [38], where w denotes the wheelbase (distance between the rear and front contact points), c is the trail (the offset between the steer-axis ground intersection and the front contact point), and λ is the steer-axis tilt angle.

$$M\begin{bmatrix}\ddot{\phi}\\ \ddot{\delta}\end{bmatrix} + vC_1\begin{bmatrix}\dot{\phi}\\ \dot{\delta}\end{bmatrix} + \left(gK_0 + v^2K_2\right)\begin{bmatrix}\phi\\ \delta\end{bmatrix} = \begin{bmatrix}T_\phi\\ T_\delta\end{bmatrix}. \tag{2.4}$$

$$\dot{\psi} = \left(\frac{\cos\lambda}{w}\,v\right)\delta + \left(\frac{\cos\lambda}{w}\,c\right)\dot{\delta}. \tag{2.5}$$

Let $\vec{X}_{\mathrm{lat}} = [\phi,\ \delta,\ \dot{\phi},\ \dot{\delta},\ \psi]^{\top}$ denote the lateral state vector. Combining Equations (2.4)–(2.5) yields Equations (2.6)– (2.7) governing the lateral state dynamics. Here, M^{-1}_{12} and M^{-1}_{22} denote the entries in the first-row second-column and second-row second-column positions of M^{-1} respectively.

$$\dot{\vec{X}}_{\mathrm{lat}}(t) = A(v)\vec{X}_{\mathrm{lat}}(t) + B_\delta(v)T_\delta(t), \tag{2.6}$$

with

$$A(v) = \begin{bmatrix} 0 & 0 & 1 & 0 & 0\\ 0 & 0 & 0 & 1 & 0\\ \multicolumn{2}{c}{\left[-M^{-1}(gK_0+v^2K_2)\right]_{2\times2}} & \multicolumn{2}{c}{\left[-M^{-1}vC_1\right]_{2\times2}} & \begin{matrix}0\\ 0\end{matrix}\\ 0 & \dfrac{\cos\lambda}{w}v & 0 & \dfrac{\cos\lambda}{w}c & 0 \end{bmatrix}, \qquad B_\delta(v) = \begin{bmatrix}0\\ 0\\ -M^{-1}_{12}\\ -M^{-1}_{22}\\ 0\end{bmatrix} \tag{2.7}$$

To enable heading tracking, we employ a full-state feedback law [39] with a reference scaling gain as shown in Equation (2.8). Therefore, substituting (2.8) into the lateral state dynamics (2.7) gives the closed-loop dynamics from reference yaw ψ_c to lateral state $\vec{X}_{\text{lat}}$ as shown in Equation (2.9), where A_{cl} is the closed-loop state matrix describing the lateral dynamics of the controlled bicycle system, and B_{cl} is the input matrix mapping the reference yaw input ψ_c to the system dynamics.

$$T_\delta(t) = -K_x\,\vec{X}_{\text{lat}} + K_u\,\psi_c(t), \tag{2.8}$$

$$\dot{\vec{X}}_{\text{lat}} = \underbrace{\left(A(v) - B_\delta(v)K_x\right)}_{A_{\text{cl}}}\,\vec{X}_{\text{lat}} + \underbrace{\left(B_\delta(v)K_u\right)}_{B_{\text{cl}}}\,\psi_c(t). \tag{2.9}$$

The feedback gain $K_x(v)$ is computed using pole placement such that the eigenvalues of $A_{\text{cl}}(v)$ match a prescribed set of desired closed-loop poles. We adopt the heuristic pole set used in the open-source implementation *cyclistsocialforce* [41], accompanying the work of Schmidt et al. [42]: $\lambda \in \{-3.3+9.5j, -3.3-9.5j, -1.3+2.5j, -1.3-2.5j, -4.0\}$. More recent versions of *cyclistsocialforce* use a pole distribution model identified from experimental data as described in [39]. However, at the time of developing the present code base, that pole model was not yet available; therefore, the constant heuristic pole set is adopted here. Given the general validity of the underlying Carvallo–Whipple dynamics and the comparable pole configuration, this approximation is expected to yield reasonable closed-loop bicycle behavior. In addition, to compute K_u, we first set $K_u = 1$ and simulate the closed-loop response to a unit-step heading command $\psi_c(t) = 1$. Let $\psi_\infty^{(1)} \triangleq \lim_{t\to\infty}\psi(t)$, denoting the resulting steady-state heading (with $K_u = 1$). We then set $K_u = \frac{1}{\psi_\infty^{(1)}}$ so that a unit command yields $\psi_\infty \approx 1$ in steady state. For the stable SISO case, this procedure is equivalent to the DC-gain expression $K_u = \left(-C(A - B_\delta K_x)^{-1}B_\delta\right)^{-1}$.

Therefore, the full state dynamics are written as in Equation (2.10). Note that the linearized Carvallo–Whipple model is derived under constant speed assumption, which means that the speed v appears as a parameter in the lateral state-space matrices $A(v)$ and $B_\delta(v)$ rather than as a dynamic state. In this work, we accommodate time-varying speed via a quasi-steady (gain-scheduled) approximation: at each integration step, the lateral dynamics are evaluated using the matrices corresponding to the instantaneous speed. Practically, we precompute the closed-loop matrices $A_{\text{cl}}(v)$ and $B_{\text{cl}}(v)$ on a discretized speed grid and obtain intermediate values through linear interpolation (a lookup table). During the MPC simulation, the instantaneous speed v at each time step is used to query this lookup table, yielding the scheduled matrices $A_{\text{cl}}(v)$ and $B_{\text{cl}}(v)$ that are then applied to propagate the lateral states. This strategy preserves the validated constant-speed lateral dynamics locally while enabling moderate speed variations over the MPC horizon.

$$\dot{\vec{X}} = f(\vec{X}, \vec{U}) = \begin{bmatrix} \left[A_{\text{cl}}(v)\,\vec{X}_{\text{lat}} + B_{\text{cl}}(v)\,\psi_c\right]_{5\times 1} \\ v\cos\theta \\ v\sin\theta \\ a \end{bmatrix} \tag{2.10}$$

The bicycle–rider parameters used are summarized in Table 2.1. We use the `BicycleParameters` parameter set `meijaard2007_browser_jason` (Browser bicycle with Jason rider), which provides the reduced canonical matrices M, C_1, K_0, and K_2. The geometric constants w, c, and λ are taken from the Browser benchmark dataset distributed with `BicycleParameters`. All values are adopted as provided by the corresponding dataset/documentation [43], [44], [45].

Table 2.1: Bicycle–rider parameters used in this thesis (Browser bicycle with Jason rider).

Symbol	Value	Unit
M	$\begin{bmatrix} 102.78013216 & 1.53582801 \\ 1.53582801 & 0.24890226 \end{bmatrix}$	$\mathrm{N{\cdot}m{\cdot}s^2}$
C_1	$\begin{bmatrix} 0 & 26.3947333 \\ -0.4503006 & 1.037066 \end{bmatrix}$	kg
K_0	$\begin{bmatrix} -89.32195981 & -1.74159477 \\ -1.74159477 & -0.67769624 \end{bmatrix}$	kg·m
K_2	$\begin{bmatrix} 0 & 74.12543 \\ 0 & 1.57021553 \end{bmatrix}$	kg/m
w	1.121	m
c	0.0685808540382	m
λ	0.399680398707	rad
g	9.81	$\mathrm{m/s^2}$

2.1.3 MPC Cost Design

Following [1], we define the objective function J as the time-discounted accumulation of the stage cost L over the prediction horizon $[t_k, t_K + T]$, as shown in Equation (2.11). Here, the prediction horizon length is given by $T = N\,dt$, where $dt = 0.1\,\mathrm{s}$ is the discrete time step and N is the number of prediction steps. For convenience, N is hereafter referred to as the prediction horizon. The exponential discount factor $e^{-\eta t}$ reflects decreasing confidence in predictions further into the future, such that near-term behavior is weighted more strongly than long-term predictions over the horizon. The stage cost L is composed of multiple terms capturing different behavioral objectives illustrated in Equation (2.12) following [1].

$$J = \int_{t_k}^{t_k+T} e^{-\eta t}\, L\,\mathrm{d}t \tag{2.11}$$

$$L = L_{\mathrm{goal}} + L_{v,\mathrm{keep}} + L_{\dot{\psi}} + L_{\mathrm{ctrl}} + L_{\mathrm{avoid}} \tag{2.12}$$

As shown in Equation (2.13), the goal-reaching cost L_{goal} penalizes deviations from a reference position using the weight w_{goal} following [1]. The longitudinal and lateral deviations, denoted by e_s and e_d, respectively, are defined as the differences between the cyclist's current position (x, y) and the reference position (x_d, y_d) on the reference path. The reference position is selected as the closest point on the reference path to the cyclist, and the local orientation of the path at this point is characterized by the slope angle θ, as illustrated in Figure 2.1. The error terms e_s and e_d are computed according to Equation (2.14). How to define the reference path is illustrated in subsection 2.2.1.

Symbols with overbar ($\bar{e_s}$, $\bar{e_d}$) in Equation (2.13) are normalization constants. We introduce them to make the different cost components dimensionless and comparable in magnitude. All cost terms are normalized so weights have equivalent effect. We define them as the interquartile ranges (i.e., the 75th percentile minus the 25th percentile) of e_s and e_d during the whole simulation time in the measurement. The same is for $\bar{v}$ in Equation (2.15) and $\bar{\dot{\psi}}$ in Equation (2.16). Their numerical values are reported in Table 2.3. In Equation (2.17), $\bar{a} = 5.0\,\mathrm{m/s^2}$ and $\bar{\psi}_c = \pi/2\,[rad]$ are used as fixed reference values to normalize the control inputs. Unlike the other normalization constants, these values are not derived from the measurement but are selected to represent reasonable scales for acceleration and reference yaw angle.

$$L_{\text{goal}} = w_{\text{goal}} \left[\left(\frac{e_s}{\bar{e}_s} \right)^2 + \left(\frac{e_d}{\bar{e}_d} \right)^2 \right] \tag{2.13}$$

$$\begin{cases} e_s = (x_d - x)\cos\theta + (y_d - y)\sin\theta \\ e_d = (y_d - y)\cos\theta - (x_d - x)\sin\theta \end{cases} \tag{2.14}$$

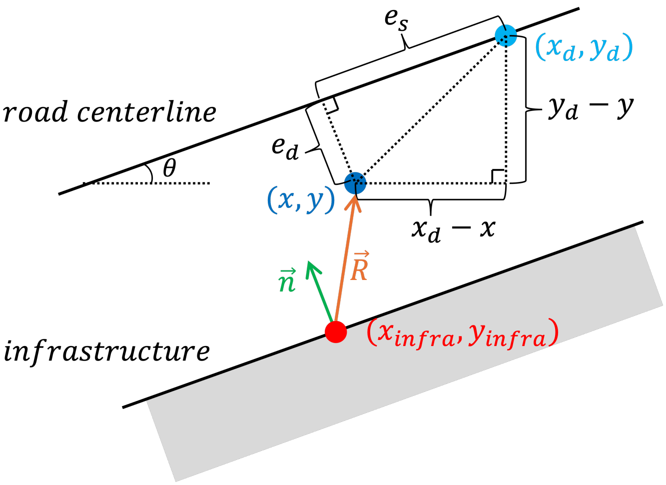

Figure 2.1: center and infra

As shown in Equations (2.15) and (2.16), the velocity-keeping cost $L_{v,\text{keep}}$ penalizes deviations of the current velocity v from the desired velocity v_d with weight $w_{v,\text{keep}}$, while the cost $L_{\dot{\psi}}$ penalizes yaw rate $\dot{\psi}$ with weight w_{ψ} following [1]. In practice, the desired velocity v_d is obtained from the measured trajectory data after noise filtering.

$$L_{v,\text{keep}} = w_{v,\text{keep}} \left(\frac{v - v_d}{\bar{v}} \right)^2 \tag{2.15}$$

$$L_{\dot{\psi}} = w_{\dot{\psi}} \left(\frac{\dot{\psi}}{\bar{\dot{\psi}}} \right)^2 \tag{2.16}$$

As shown in Equation (2.17), the control effort cost L_{ctrl} penalizes the magnitude of control inputs following [1]. Specifically, it penalizes deviations of the longitudinal acceleration a from zero acceleration and deviations of the reference yaw angle ψ_c from the desired yaw angle $\psi_{c,d}$. This term acts as a regularization that favors smooth and energy-efficient control actions in the absence of interaction-induced maneuvers. The desired yaw angle $\psi_{c,d}$ is computed as the local orientation of the reference path at the point closest to the cyclist.

$$L_{\text{ctrl}} = w_a \left(\frac{a}{\bar{a}}\right)^2 + w_{\psi_c} \left(\frac{\psi_c - \psi_{c,d}}{\bar{\psi_c}}\right)^2 \tag{2.17}$$

As shown in Equation (2.18), the proximity cost L_{avoid} penalizes the ego cyclist's closeness to other road users following [1]. As illustrated in Figure 2.2a based on [1], the center of mass (COM) of the ego cyclist (overtaking cyclist) is (x, y) and the COM of the opponent cyclist (overtaken cyclist) is at (x_O, y_O), with the Euclidean distance between them denoted by d. The exponential term e^{-d/R_0} ensures that the proximity cost decreases monotonically with increasing distance, such that nearby cyclists induce a stronger avoidance cost while distant cyclists have a weaker influence. The parameter $R_0 > 0$ controls the rate at which this influence decays with distance.

According to [1], directional sensitivity is introduced through an anisotropic weighting term. Let γ denote the orientation of the line connecting the ego cyclist to the opponent cyclist, and let ψ be the yaw angle of the ego cyclist. The relative angle $\beta = \psi - \gamma$ describes the position of the opponent cyclist in the reference frame fixed on the ego cyclist. The anisotropy factor $k \in [0, 1]$ determines the directional dependence of the proximity cost, with larger values emphasizing obstacles located in front of the cyclist. Figure 2.2b and 2.2c illustrate the spatial distribution of L_{avoid} in the isotropic case ($k = 0$) and the fully anisotropic case ($k = 1$). When $k = 0$, the repulsive field has the same magnitude in all directions, whereas for $k = 1$ the proximity cost is maximal in the forward direction and vanishes in the lateral directions. In this study, the fully anisotropic formulation is adopted to reflect cyclists' stronger sensitivity to obstacles located ahead. In overtaking scenarios, we assume that the proximity cost L_{avoid} is activated only when the opposing cyclist is ahead of the ego cyclist. This is enforced by evaluating the projection of the relative position of the opponent cyclist onto the ego cyclist's forward heading direction $x_{\text{OinE}} = (x_O - x)\cos\psi + (y_O - y)\sin\psi$, such that L_{avoid} is applied only when $x_{\text{OinE}} > 0$.

$$L_{\text{avoid}} = \begin{cases} w_{\text{avoid}}\, e^{-\frac{d}{R_0}} \left[k + (1-k)\dfrac{(1+\cos\beta)}{2} \right], & \text{if } x_{OinE} > 0 \\ 0, & \text{if } x_{OinE} \leq 0 \end{cases} \tag{2.18}$$

Note that to calculate the proximity cost of the ego cyclist, the positions of the opponent cyclist over the prediction horizon are required. We assume that the ego cyclist can perfectly predict the future positions of the opponent cyclist. So the opponent cyclist's future positions are obtained by directly extracting from measurement data, rather than being estimated or predicted through an additional model. We made this assumption to isolate the modeling error associated with the ego cyclist's optimization process. In other words, in the future calibration process, discrepancies between MPC result and measurement arise solely from the ego cyclist's decision model, rather than from uncertainties in predicting the opponent's future motion.

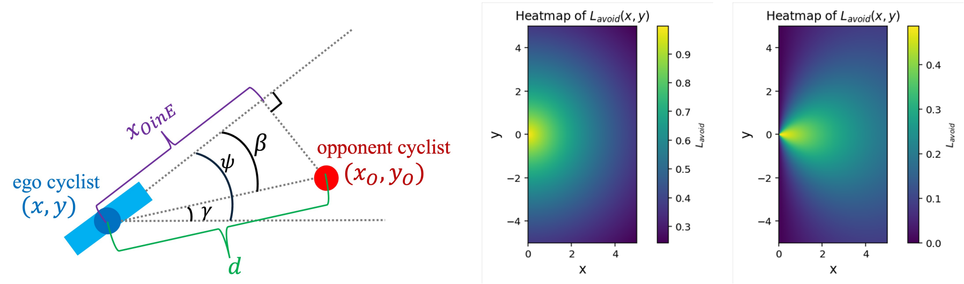

(a) Illustration of the parameters used in the proximity cost (b) Isotropic repulsive field (c) Fully anisotropic repulsive field

Figure 2.2: Illustration of the proximity cost formulation and its resulting obstacle repulsive fields, adapted from the model in [1]. The left picture shows the physical meaning of the geometric parameters used in the cost definition, while the middle and right pictures show the isotropic and fully anisotropic repulsive fields, respectively.

2.1.4 MPC Constraints

Equality Constraints: According to the state transition dynamics in Equation (2.10), the equality constraint in discretized time can be written in Equation (2.19), where the first equation specifies the initial state of the prediction horizon ($\vec{X}_{\text{init}}$), and the second equation describes the discrete-time state propagation. k denotes the discrete time-step index and dt is the sampling time. To ensure consistency, the Forward Euler method is used both in the Nonlinear Programming (NLP) dynamics equality constraints and in the state propagation. We select this method mainly for its shorter computational time. This is particularly important in this work, because a large number of MPC problems must be solved repeatedly in the later parameter calibration process. Nevertheless, this choice comes at the expense of numerical accuracy: compared with the midpoint method and Runge–Kutta methods, Forward Euler has lower-order accuracy and may therefore yield larger discretization errors [46]. Hence, the use of Forward Euler should be regarded as a practical but imperfect compromise between computational speed and integration accuracy.

$$\begin{cases} \vec{X}_0 = \vec{X}_{\text{init}}, \\ \vec{X}_{k+1} = \vec{X}_k + \vec{f}(\vec{X}_k, \vec{U}_k) \cdot dt, \qquad k = 0, 1, \ldots, N-1 \end{cases} \tag{2.19}$$

Inequality Constraints: In the MPC formulation, inequality constraints are introduced to ensure physically feasible, smooth and behaviorally consistent trajectories, which can be grouped into four categories: state bounds, control-input bounds, collision-avoidance constraints, and overtaking side-related constraints.

As shown in Equation (2.20), at time step k, the yaw angle ψ_k is constrained relative to the local reference path orientation $\theta_{\text{ref},k}$ to avoid unrealistically large heading deviations and excessively winding trajectories. The velocity v_k is bounded to remain within a realistic range of riding speed. The reference yaw angle $\psi_{c,k}$ is also constrained relative to the local reference path orientation $\theta_{\mathrm{ref},k}$, and the acceleration a_k is restricted to a feasible range. Also, to prevent the two cyclists from colliding, we introduce a minimum-distance requirement, where d_k denotes the distance between the two cyclists.

$$\begin{cases} \theta_{\mathrm{ref},k} - \dfrac{\pi}{2} \le \psi_k \le \theta_{\mathrm{ref},k} + \dfrac{\pi}{2}, \\ 0.5\ \mathrm{m/s} \le v_k \le 12.0\ \mathrm{m/s}, \\ \theta_{\mathrm{ref},k} - \dfrac{\pi}{2} \le \psi_{c,k} \le \theta_{\mathrm{ref},k} + \dfrac{\pi}{2}, & k = 0, \ldots, N-1 \\ -8.0\ \mathrm{m/s^2} \le a_k \le 8.0\ \mathrm{m/s^2}, \\ d_k \ge 0.8\ \mathrm{m} \end{cases} \tag{2.20}$$

Furthermore, in the subsequent calibration procedure, we use the measurement data from the overtaking cyclist's trajectory. We introduce another constraint to ensure that the simulated cyclist overtakes from the same side as observed in the measurements. This is because during the calibration process, different parameter combinations may lead the simulated cyclist to overtake either from the left or from the right. As the calibration aims to minimize trajectory error, parameter sets that produce a realistic overtaking trajectory but on the wrong side would still be penalized as poor solutions. This would lead to inefficient calibration trials. Therefore, we introduce an additional overtaking side-related constraint to ensure that the simulated cyclist overtakes from the same side as observed in the measurements.

The data was collected in Germany, where the right-hand rule applied (road users are expected to overtake from the left), but we observe that cyclists still overtake from the right-hand side in multiple overtaking pairs. This behavior may be influenced by local infrastructure and surrounding environmental conditions. For example, inspection of Google Maps suggests that there is a step-down boundary on the left-hand side of a bicycle lane, which may discourage cyclists from overtaking on the left-hand side because of the increased perceived risk. Rather than explicitly modeling all possible environmental causes behind this overtaking side preference, we adopt a simplified representation, i.e., overtaking is constrained to occur only on the side consistent with the measurement. To enforce this, a virtual infrastructure boundary is introduced on the prohibited side of the reference path. This boundary is defined as a line parallel to the reference path with an offset of 0.2 m. The offset accounts for small lateral oscillations and bicycle wobbling, such that small deviations do not render the optimization infeasible.

The implementation of this infrastructure constraint is illustrated in Figure 2.1. The point on the virtual infrastructure boundary closest to the cyclist is denoted by $(x_{k,\mathrm{infra}}, y_{k,\mathrm{infra}})$. The vector $\vec{n}_k$ represents the outward normal vector of the boundary at that point, and $\vec{R}_k$ is the vector from $(x_{k,\mathrm{infra}}, y_{k,\mathrm{infra}})$ to the cyclist position (x_k, y_k). The prohibited side of the boundary is shown as the shaded gray region in Figure 2.1. To ensure that the cyclist remains on the admissible side, the constraint is formulated as in Equation (2.21). Alternative formulations based on biased soft constraints (i.e., adding biased terms in the cost function) were explored but found insufficient to consistently reproduce the observed overtaking side preference. The hard inequality constraint adopted here in IPOPT solver acts as a barrier function, providing a robust and numerically stable mechanism for enforcing overtaking behavior from a certain side within the MPC framework.

$$\vec{n}_k \cdot \vec{R}_k \geq 0. \tag{2.21}$$

2.1.5 Adaptive Prediction Horizon

Classical MPC formulations often include a terminal cost and a terminal constraint set to guarantee stability and recursive feasibility under infinite-horizon assumptions. In this context, the terminal cost is commonly chosen to approximate the infinite-horizon cost-to-go beyond the finite prediction horizon, while the terminal constraint set ensures that the system can be steered to a region where a stabilizing control law exists. However, in the context of cyclist behavior modeling, we think that such assumptions are not appropriate. Cyclists do not plan infinitely ahead, nor do they target a well-defined terminal state. Instead, their anticipation horizon is finite and subject to perceptual limitations. For example, in the model of perceptive range and reactive range [29], how far ahead cyclists can perceive is shaped by cyclists' fixation point limited by riding speed. Therefore, terminal costs and terminal constraint sets are intentionally omitted in this work.

As discussed in the previous chapter, cyclists' anticipation depends on multiple contextual factors. Focusing on pairwise cyclist interactions, we model anticipation through an adaptive prediction horizon N that depends on the inter-cyclist distance d. Specifically, we assume that N is a smooth and bounded function of d, given by Equation (2.22). The proposed formulation is a stretched and shifted hyperbolic tangent function, ensuring that the prediction horizon is bounded between a minimum value $N_{\min}$ and a maximum value $N_{\max}$. And d_0 is a distance scale related to cyclists' behavioral transition from reactive to predictive. When $d \ll d_0$, $N(d)$ reduces with the reduction in d slowly and saturates at the lower bound $N_{\min}$; Around $d \approx d_0$, $N(d)$ is most sensitive to changes in d, corresponding to the mid-slope region of the adaptation curve; When $d \gg d_0$, $N(d)$ increases with the increment in d slowly and saturates at the upper bound $N_{\max}$.

$$N(d) = \frac{N_{\max} - N_{\min}}{2} \tanh(d - d_0) + \frac{N_{\max} + N_{\min}}{2} \tag{2.22}$$

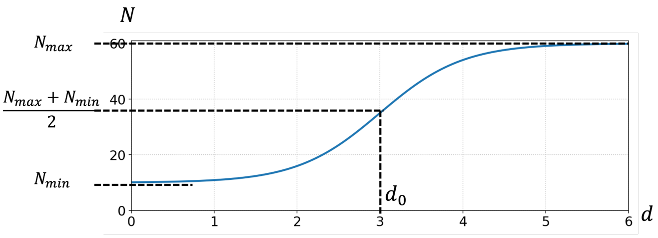

Figure 2.3: Prediction horizon N as a function of inter-cyclist distance d

Figure 2.3 illustrates the case where $N_{\min} = 10, N_{\max} = 60, d_0 = 3.0$. The prediction horizon N increases monotonically with distance d. When cyclists are closer to each other, the interaction is more urgent and collision risk is higher, leading to more reactive behavior and a shorter prediction horizon. Conversely, at larger distances, cyclists exhibit more anticipatory behavior.

When the prediction horizon increases between successive time steps, the warm-start initialization is handled by extending the previous optimal control sequence with repetitions of its last element; when it decreases, the sequence is truncated to match the shorter horizon.

Since we mainly want to investigate how cyclists adapt their anticipation horizon across different interaction scenarios, the parameters governing the horizon adaptation, namely $N_{\min}$, $N_{\max}$, and d_0, are identified from data using the calibration.

2.1.6 MPC Workflow

Up to this point, the MPC formulation has been defined as given in Equation (2.23). This section describes how the optimization problem is solved. To begin with, Algorithm 1 outlines the construction of the nonlinear programming (NLP) solver used in the MPC formulation with CasADi/IPOPT. Over a prediction horizon of length N (with sampling interval dt), the decision variables are the stacked control-input sequences $\vec{U} = \{\vec{U}_k\}_{k=0}^{N-1}$, where each stage input is $\vec{U}_k = [\psi_{c,k}, a_k]^\top$. The state is initialized with $\vec{X} \leftarrow \vec{X}_0$, and the objective and constraint containers are initialized as $J \leftarrow 0$ and $\vec{g} \leftarrow \emptyset$. For each prediction step $k = 0, \ldots, N-1$, the state is propagated forward using the discrete-time dynamics, $\vec{X}^+ = \vec{X} + dt \cdot f(\vec{X}, \vec{U})$, and the stage cost is accumulated via $J \leftarrow J + L(\vec{X}, \vec{U})$. In parallel, inequality constraints are appended to the stacked constraint vector $\vec{g}$, including state bounds, overtaking side-related constraints, and the minimum inter-cyclist distance constraint. After constructing the stage-wise decision variables and constraints along the prediction horizon, they are stacked into single vectors as $\vec{U} \leftarrow \mathrm{vec}(\vec{U})$ and $\vec{g} \leftarrow \mathrm{vec}(\vec{g})$, allowing the problem to be formulated as a standard nonlinear program (NLP) for the solver. All time-varying and scenario-dependent quantities (e.g., initial state, projected point on the reference trajectory, and predicted opponent positions) are collected in a parameter vector $\vec{p}$. Finally, the NLP is assembled in CasADi as the dictionary $\mathscr{D} = \{\texttt{x} : \vec{U}, \texttt{f} : J, \texttt{p} : \vec{p}, \texttt{g} : \vec{g}\}$, and an IPOPT-based solver is created through `casadi.nlpsol`, which is iteratively used during the MPC simulation.

$$
\begin{aligned}
&\min_{\{\vec{U}_k\}_{k=0}^{N-1}} \sum_{k=0}^{N-1} L(\vec{X}_k, \vec{U}_k) \\
&\text{s.t.} \begin{cases}
\vec{X}_0 = \vec{X}_{\text{init}}, & \\
\vec{X}_{k+1} = \vec{X}_k + f(\vec{X}_k, \vec{U}_k) \cdot dt, & \\
\psi_{c,\min} \le \psi_{c,k} \le \psi_{c,\max}, & \\
a_{\min} \le a_k \le a_{\max}, & \\
\psi_{\min} \le \psi_k \le \psi_{\max}, & k = 0, \ldots, N-1 \\
v_{\min} \le v_k \le v_{\max}, & \\
\vec{n}_k \cdot \vec{R}_k \ge 0, & \\
d_{\min} - d_k \le 0 &
\end{cases}
\end{aligned}
\tag{2.23}
$$

Algorithm 2 describes the implemented MPC loop for the whole trajectory. When solving the problem, IPOPT returns not only the optimal primal solution (the optimal control sequence $\vec{U}_A^*$), but also the optimal multipliers $(\lambda^\star_{x,A}, \lambda^\star_{g,A})$ as part of its Karush-Kuhn-Tucker (KKT) conditions [47]. At the beginning, both primal and dual warm-start variables are initialized: the previous control sequence $\vec{U}_{A,\text{last}}$ is used as an initial guess for the optimizer. And $(\lambda_{x0,A}, \lambda_{g0,A})$ denote the Lagrange multipliers associated with the control input bounds and the stacked inequality constraints, respectively. After applying only the first input $\vec{U}_A^\star(0)$, the warm-start sequence for the next MPC step is obtained by shifting the previous optimum as shown in Equation (2.24), i.e., discarding the applied control, moving the remaining inputs one step forward, and padding the last element to keep the horizon length constant. The dual variables are warm-started by reusing the previous multipliers, $(\lambda_{x0,A}, \lambda_{g0,A}) \leftarrow (\lambda^\star_{x,A}, \lambda^\star_{g,A})$. The primal–dual warm start leverages the similarity between consecutive MPC NLPs and, in practice, often reduces the number of interior-point iterations compared to cold-starting [48].

$$\vec{U}_{A,\text{last}} \leftarrow \texttt{Shift}(\vec{U}_A^\star) = \left[\vec{U}_A^\star(1), \ldots, \vec{U}_A^\star(N-1), \vec{U}_A^\star(N-1)\right], \tag{2.24}$$

Warm start can sometimes keep the solver in the same local basin. We therefore performed a small multi-start test around the warm-start initial guess. In some cases it did reach a lower cost, but the computation time increased by several times. Since we run many MPC calls inside Bayesian Optimization under a fixed time budget, we finally use a single-start (warm-started) solve to keep the evaluations fast and allow more parameter trials. Moreover, a better local optimum corresponding to a lower MPC cost does not necessarily imply a more realistic cyclist behavior, since cyclists may not act as perfectly utility-maximizing agents. Prior work suggests that the realized action selection depends on the rider's degree of rationality, and the chosen trajectory may deviate from the strict utility optimum [49].

For each simulation step $i \in \{0, \ldots, n_{\text{sim}} - 1\}$, the measured trajectory of the opponent cyclist (i.e. overtaking cyclist B) over the next N steps is extracted and aligned with the current prediction horizon window of the overtaking cyclsit A, yielding $xy_B^{0:N-1} = \{(x_B(k), y_B(k))\}_{k=i}^{i+N-1}$ (since overtaking cyclist A's proximity cost is accumulated over the prediction horizon, the predicted positions of the overtaken cyclist B over the same horizon are required). The parameter vector $\vec{p}_A$ is then assembled by packing the current state $\vec{X}_A(i)$ together with all other parameters, including the reference position (x_d, y_d) and the infrastructure geometry $(x_{\text{infra}}, y_{\text{infra}})$. Given $\vec{p}_A$, the NLP is solved using the CasADi/IPOPT solver with bounds $(lbx_A, ubx_A, lbg_A, ubg_A)$. Here, lbx_A and ubx_A denote lower and upper bounds on the control input sequence (ψ_c, a), whereas lbg_A and ubg_A denote lower and upper bounds on the stacked constraint vector $\vec{g}(\cdot)$, which includes state bounds, overtaking side constraint, and the minimum inter-cyclist distance requirement. Note that we use $\vec{U}_{A,\text{last}}$ as the primal initial point and $(\lambda_{x0,A}, \lambda_{g0,A})$ as the dual initial point to accelerate convergence.

After solving, we obtain the optimal control sequence $\vec{U}_A(i : i+N-1)$ from $\text{sol}_A.x$, and the dual variables are updated by setting $\lambda_{x0,A} \leftarrow \text{sol}_A.\lambda_x$ and $\lambda_{g0,A} \leftarrow \text{sol}_A.\lambda_g$ for the next iteration. Following the receding-horizon principle, we only apply the first control input, i.e., $\vec{U}_{A,\text{applied}}(i) = \texttt{First}(\vec{U}_A(i : i+N-1))$, and we propagate the system state by 1 step using the discrete-time dynamics $\vec{X}_A(i+1) = \texttt{Propagate\_One\_Step}(\vec{X}_A(i), \vec{U}_{A,\text{applied}}(i))$,

Algorithm 1 NLP Solver Construction (CasADi/IPOPT)

```
Require: : $(N, x_{\text{center}}, y_{\text{center}})$
Ensure: : solver
Define decision variables $\vec{U} = \{u_k\}_{k=0}^{N-1}$, with $u_k = [\psi_{c,k}, a_k]^\top$
$\vec{X} \leftarrow \vec{X}_0$;   $J \leftarrow 0$;   $\vec{g} \leftarrow \emptyset$
for $k = 0, 1, \ldots, n_{\text{sim}} - 1$ do
    $\vec{X}^+ \leftarrow \vec{X} + dt \cdot f(\vec{X}, \vec{U})$
    $J \leftarrow J + L(\vec{X}, \vec{U})$
    Append_Inequality_Constraint_1($\vec{g}$)                ▷ state bounds
    Append_Inequality_Constraint_2($\vec{g}$)                ▷ overtaking side requirement
    Append_Inequality_Constraint_3($\vec{g}$)                ▷ min distance between cyclists
    $\vec{X} \leftarrow \vec{X}^+$
end for
Stack variables: $\vec{U} \leftarrow \text{vec}(\vec{U})$;   stack constraints: $\vec{g} \leftarrow \text{vec}(\vec{g})$;   Pack parameters: $\vec{p}$
Assemble NLP dict: $\mathcal{D} \leftarrow$ {"x" : $\vec{U}$, "f" : $J$, "p" : $\vec{p}$, "g" : $\vec{g}$}
Create solver: Solver ← casadi.nlpsol(solver_name, "ipopt", $\mathcal{D}$, opts); return Solver
```

dt). Finally, the warm-start control sequence is updated by shifting the previously optimal sequence, $\vec{U}_{A,\text{last}} \leftarrow \text{Shift}(\vec{U}_A(i : i + N - 1))$.

2.2 Data Collection for Calibration

This section describes how overtaking cyclist pairs were extracted from the public TUMDOT dataset [2]. The dataset provides trajectory data for different types of road users observed in real traffic scenes, stored in `.csv` files. Each trajectory is associated with a unique track identifier (ID) and a class label (cyclist, pedestrian, vehicle, motorcycle, or trailer). For each tracked road user, time-resolved information is available, including the timestamp and the corresponding global (x, y) coordinates. The dataset also includes an `.xodr` map file describing the road layout and lane geometries.

Since we calibrate the MPC framework on two-agent overtaking interactions, only clean (not disturbed by a third road user) cyclist–cyclist overtaking pairs are considered. The calibration focuses on minimizing the prediction error of the overtaking cyclist, i.e., how the MPC trajectory of the overtaking cyclist deviates from the measurement. The overtaken cyclist is not used as a calibration target, because the overtaken cyclist typically remains close to the lane centerline and on many occasions has little lateral deviation.

2.2.1 Extraction of Bicycle-lane Centerlines and Reference Path

The bicycle lane centerline was extracted from the OpenDRIVE `.xodr` map file through the following steps. First, we parse the `.xodr` file and read the `planView` and `lanes` information for each road. The `planView` defines the road reference geometry, while the `lanes` element provides the lane offset, lane section, and lane width definitions. Next, we reconstruct the road reference line from the geometric segments in `planView`, including `line`, `arc`, `spiral`, and `paramPoly3`, such that the global position and head-

Algorithm 2 MPC Implementation Loop

```
Require: initial state $\vec{X}_A(0)$; prediction horizon $T = N \cdot dt$; time interval $dt$; number
    of simulation steps $n_{\text{sim}}$; measured opponent trajectory over the prediction horizon
    $xy_B^{0:N-1}$; solver SolverA; bounds $(lbx_A, ubx_A, lbg_A, ubg_A)$
Ensure: executed states $\{\vec{X}_A(i)\}_{i=0}^{n_{\text{sim}}}$ and applied control inputs $\{\vec{U}_A(i)\}_{i=0}^{n_{\text{sim}}-1}$
Initialize primal warm-start $\vec{U}_{A,\text{last}}$ and dual warm-start $(\lambda_{x0,A}, \lambda_{g0,A})$
for $i = 0, 1, \ldots, n_{\text{sim}} - 1$ do
    $xy_B^{0:N-1} = \{(x_B(k), y_B(k))\}_{k=i}^{i+N-1} \leftarrow$ ExtractMeasuredXY$(i, N)$
    $\vec{p}_A \leftarrow$ PackParams$\big(\vec{X}_A(i),\ \vec{\mathbf{p}}_B^{0:N-1},\ x_{\text{d}}, y_{\text{d}}, x_{\text{infra}}, y_{\text{infra}}\big)$
    sol$_A \leftarrow$ Solver(x0 = $\vec{U}_{A,\text{last}}$, p = $\vec{p}_A$, lbx = $lbx_A$, ubx = $ubx_A$, lbg = $lbg_A$, ubg = $ubg_A$, lam_x0 = $\lambda_{x0,A}$, lam_g0 = $\lambda_{g0,A}$)
    $\vec{U}_A(i : i+N-1) \leftarrow$ (sol$_A$.x)
    $\lambda_{x0,A} \leftarrow$ sol$_A$.lam_x                                ▷ warm start
    $\lambda_{g0,A} \leftarrow$ sol$_A$.lam_g                                ▷ warm start
    $\vec{U}_{A,\text{last}} \leftarrow$ Shift$(\vec{U}_A(i : i+N-1))$          ▷ warm start
    $\vec{U}_{A,\text{applied}}(i) \leftarrow$ First$\big(\vec{U}_A(i : i+N-1)\big)$   ▷ receding horizon principle
    $\vec{X}_A(i+1) \leftarrow$ Propagate_One_Step$\big(\vec{X}_A(i), \vec{U}_{A,\text{applied}}(i), dt\big)$
end for
```

ing could be evaluated for any longitudinal position s. Within each `laneSection`, only lanes with `type="biking"` are kept. Then, we sample the road uniformly along s at a fixed interval of 0.1 m. For each sampled longitudinal position, we obtain the bicycle lane center by laterally shifting the road reference line to the center of the target `biking` lane. This lateral shift consists of the lane offset, the widths of the lanes between the reference line and the target lane, and half of the target lane width. Finally, we obtain the global coordinates of the bicycle-lane center by shifting each reference-line point laterally along the normal direction by the computed offset. Repeating this process for all sampled positions yields a sequence of points representing the bicycle lane centerline $(x_{\text{center}}, y_{\text{center}})$.

Note that the extracted bicycle-lane centerline $(x_{\text{center}}, y_{\text{center}})$ was not always used directly as the MPC reference path $(x_{\text{refer}}, y_{\text{refer}})$. Instead, the reference path was determined from the observed trajectory pattern in each overtaking scenario. In most cases, we observe that the overtaking cyclist follows a nominal riding line before the overtaking maneuver, temporarily deviates from it laterally during the maneuver, and then returns to the same line afterwards. This nominal riding line typically coincides with the trajectory of the overtaken cyclist, who has very little lateral deviations throughout the interaction. In such cases, we laterally shift the extracted bicycle-lane centerline, when necessary, to align with this common riding line, and we use the shifted centerline as the reference path for the overtaking cyclist. In some cases, no such shift is required because the extracted lane centerline already matched the observed riding line sufficiently well. However, a small number of cases have a different pattern: the overtaking cyclist follows one approximately straight riding line before the maneuver but continues on a different parallel line afterwards, rather than returning to the original one. In these cases, we cannot identify the reference path of the overtaking cyclist, since it is unclear when and how the intended nominal path changes. These cases are excluded from the data collection.

2.2.2 Extraction of Clean Two-Cyclist Overtaking Pairs

After determining the reference path, the cyclist trajectories are further processed to extract clean two-cyclist overtaking pairs. The extraction procedure consists of three filtering steps based on the defining characteristics of overtaking events ([50]).

First, a temporal overlap filter is applied. For all possible cyclist pairs on a certain bicycle lane, we identify the first and last timestamps of their trajectories. And we only keep a pair if the two cyclists' trajectories overlap for at least 2 s, ensuring that both cyclists are present in the scene long enough for a potential interaction.

Second, an orientation filter is applied to remove cyclists moving in clearly different directions. For each candidate pair, we compute the mean heading direction during the overlap time interval. And we only keep pairs whose heading difference is smaller than 30°.

Finally, the filter of potential overtaking point existence is applied. We project the two cyclists' trajectories onto bicycle-lane centerline to obtain the longitudinal coordinate s. For each pair, the quantity $\Delta s(t) = s_A(t) - s_B(t)$ is evaluated over the overlapping time interval. We only keep pairs for which the sign of Δs at least changes once between two consecutive timestamps, indicating that there is a potential overtaking point.

Thirdly, for each detected overtaking pair, a short time window around the estimated overtaking moment is extracted and visualized together with the bicycle-lane centerline. These plots are used for quick manual inspection and for selecting representative two-cyclist overtaking scenarios for subsequent MPC calibration.

After applying these filters, we plot the resulting candidate pairs to further examine if a genuine overtaking interaction occurs. In some cases, although the pair has passed the three filters, the two cyclists travel in parallel without exhibiting noticeable evasive maneuver. Such cases are excluded, as no meaningful interaction affecting their motion is observed.

Finally, we examine the environment around each candidate pair to ensure that the overtaking scenario is not influenced by additional road users, such as cars, pedestrians, trailers, motorcycles, or other cyclists. And we only keep pairs that are visually confirmed to represent isolated two-cyclist overtaking interactions.

Following the extraction procedures, 30 clean 2-cyclist overtaking pairs are selected for calibration. Detailed information for each extracted pair, including cyclist ID and overtaking point timestamp, is provided in Table 2.2.

2.3 Calibration (Bayesian Optimization)

In this section, how cyclists adapt their anticipation horizon (what are the values governing the change in prediction horizon N, which are N_{min}, N_{max}, d_0) is explored by calibration. More specifically, we want to find out the N_{min}, N_{max}, d_0 which leads the simulated trajectory calculated by the MPC to be as close to the measured trajectory as possible. In other words, we want to find out optimal parameter combination $\vec{P}^*$ which minimizes ERR (the error compared to measurement) as a shown in Equation (2.25).

Table 2.2: Clean overtaking pairs extracted from TUMDOT dataset [2] for calibration

pair	A (Overtaking Cyclist) ID	B (Overtaken Cyclist) ID	$t_{\text{overtaking}}$ [s]
1	5577	5622	3032.96
2	14097	14069	7605.44
3	19861	20023	10713.52
4	7888	7998	4242.32
5	1782	1821	949.45
6	21097	21061	11282.56
7	13813	13647	7457.52
8	1745	1608	885.52
9	20728	20718	11026.72
10	20181	20067	10779.12
11	52	80	-163.12
12	6896	6713	3719.12
13	3721	3611	2020.40
14	10708	10714	5780.96
15	15720	15721	8503.44
16	11898	11861	6419.68
17	21960	21853	11593.04
18	13466	13417	7269.12
19	18734	18711	10073.36
20	16184	16234	8806.56
21	18207	18315	9890.88
22	4384	4324	2339.52
23	819	798	364.24
24	814	812	299.76
25	17486	17347	9466.40
26	21214	21093	11275.20
27	12008	11971	6458.80
28	19196	19201	10278.96
29	2274	2276	1164.24
30	15971	15968	8645.20

$$\vec{P}^* = \arg\min ERR(\vec{P}) \tag{2.25}$$

The error metric ERR is defined as in Equation (2.26), where the coordinate s is the longitudinal (arc-length) coordinate along the reference path, and s_i are n uniformly sampled points on a common s-grid. The terms $\ell_{\text{MPC}}(s_i)$ and $\ell_{\text{meas}}(s_i)$ are the signed lateral offsets of the MPC-simulated and measured trajectories at the same longitudinal position s_i, obtained by projecting each trajectory onto the reference path. The lateral offset $\ell(s)$ is defined as the signed perpendicular distance to the centerline, with the sign indicating on which side of the reference path the trajectory lies. Hence, ERR measures the root-mean-square lateral deviation between the simulated and measured trajectories evaluated at identical longitudinal locations along the reference path.

$$ERR = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(\ell_{\text{MPC}}(s_i) - \ell_{\text{meas}}(s_i))^2} \tag{2.26}$$

The definition of parameter combination $\vec{P}$ is shown in Equation (2.27), consisting of w_{avoid} (the weight of obstacle-avoidance cost), R_0 (the parameter characterizing how fast obstacle's repulsive field decays with distance), $w_{\dot{\psi}}$ (the weight penalizing yaw rate) and $N_{\min}, N_{\max}, d_0$ (parameters characterizing the change of prediction horizon). Of course, there are other parameters that influence trajectory shape, but to reduce the parameter space dimension, we keep them constant: $w_{\text{goal}} = 1.0$, $w_{v,\text{keep}} = 30$, $w_a = 0.1$, $w_{\psi_c} = 0.1$, $\eta = 0.05$. The following illustrates how we solve Equation (2.25) by Bayesian Optimization (BO).

$$\vec{P} = \left[w_{\text{avoid}},\, R_0,\, w_{\dot{\psi}},\, N_{\min},\, N_{\max},\, d_0\right]^{\mathsf{T}} \tag{2.27}$$

2.3.1 BO Setup: Surrogate Model, Initialization and Acquisition

Bayesian Optimization (BO) is a widely-used method to effectively find good solutions for black-box optimization problems where the underlying form is unknown [51] [52]. BO proceeds iteratively. First, an initial set of sample points is generated and evaluated to form the initial dataset. A surrogate model is then fitted to these observations, providing a probabilistic approximation of the objective function. Based on the surrogate posterior, an acquisition function is constructed to guide the search for promising evaluation points by balancing exploration and exploitation. Given the data $\mathscr{D}_n = \{(\vec{P}_i, ERR_i)\}_{i=1}^{n}$ collected up to iteration n, the next query is selected by maximizing the acquisition function, i.e., $\vec{P}_{n+1} = \arg\max_{\vec{P}} \alpha(\vec{P}; \mathscr{D}_n)$. We then evaluate ERR_{n+1} by running the MPC simulation with parameters $\vec{P}_{n+1}$ and computing the resulting trajectory error with respect to the measurements. The new observation $(\vec{P}_{n+1}, ERR_{n+1})$ is appended to the dataset, and the surrogate model is updated accordingly. This procedure is repeated until a predefined stopping criterion is met [53].

BO consists of a surrogate model, an acquisition function, and an initialization strategy. In this work, we use a Gaussian Process surrogate with a Matérn kernel, Sobol sequences for initialization, and the qLogNoisyExpectedImprovement acquisition function. More details of initialization, surrogate model setting and acquisition function are as follows.

The quality of the surrogate model is crucial for good optimization performance. And initialization plays a critical role in shaping the surrogate's predictive quality and guiding subsequent optimization [54]. Therefore, the choice of the initial sampling set can noticeably influence the final outcome of BO. We generate the initial samples using Sobol sequences (i.e., a low-discrepancy quasi-random design that covers the parameter space more uniformly than purely random sampling) with different random seeds to mitigate potential bias introduced by a single initialization.

Gaussian Processes (GPs) are widely used as surrogate models in BO because they are data-efficient and naturally capture predictive uncertainty [55]. A GP prior can be seen as a mapping that assigns a normal distribution to the function ERR at input parameters $\vec{P}$, i.e., $ERR(\vec{P}) \sim \mathcal{GP}(m(\vec{P}), k(\vec{P}, \vec{P}'))$, where $m(\vec{P})$ is the mean function and $k(\vec{P}, \vec{P}')$ is the covariance function [55][56]. We employ the Matérn covariance function, whose smoothness is controlled by the positive parameter ν. Smaller values of ν correspond to rougher functions, while larger values produce smoother functions. For half–integer values of ν, the Matérn kernel admits simple closed-form expressions. In practice, the cases $\nu = 3/2$ and $\nu = 5/2$ are commonly used in machine learning. As noted by Rasmussen and Williams [56], $\nu = 1/2$ leads to very rough sample functions, whereas for $\nu \geq 7/2$ it becomes difficult to distinguish different smoothness levels from finite noisy data. Therefore, $\nu = 3/2$ and $\nu = 5/2$ provide a practical balance between model flexibility and smoothness.

For acquisition function, we adopt qLogNoisyExpectedImprovement (qLogNEI) from the LogEI family mainly for numerical reasons. Prior work shows that canonical EI and its variants frequently suffer from numerically vanishing acquisition values and gradients over large regions of the search space, which can cause gradient-based acquisition optimization to degenerate into ineffective random search as the number of observations grows. In comparison, LogEI-style reformulations compute improvements in log-space and use smooth approximations, resulting in substantially more stable values and gradients and thus more reliable acquisition maximization [55]. The acquisition function is shown in Equation (2.28), where ERR^* represents the best (smallest) ERR value observed so far. The term $ERR^{(i)}(\vec{P})$ denotes the i^{th} Monte Carlo sample drawn from the Gaussian Process posterior at $\vec{P}$, and Q is the number of Monte Carlo samples used to approximate the expectation. The function $\text{softplus}(z) = \log(1 + e^z) \approx \max(z, 0)$ is a smooth approximation of the positive-part operator, which ensures numerical stability. In other words, the acquisition function estimates the expected improvement in the objective value ERR relative to the current smallest one ERR^*. By drawing Q Monte Carlo samples from the Gaussian Process posterior, the expression approximates how much the candidate parameter vector $\vec{P}$ is expected to reduce ERR. In each iteration of Bayesian optimization, the next parameter vector is selected by maximizing $\alpha(\vec{P})$.

$$\alpha(\vec{P}) = \log\left(\frac{1}{Q}\sum_{i=1}^{Q} \text{softplus}\Big(ERR^* - ERR^{(i)}(\vec{P})\Big)\right) \tag{2.28}$$

We do not employ an explicit stopping criterion based on convergence. Instead, the BO loop is run under a fixed computational budget dictated by the DelftBlue HPC scheduling constraints: each run is executed as a 24-hour job. During the run, we log the result of every trial, including the evaluated objective value and the corresponding acquisition value. After the job finishes, we assess convergence by analyzing the objective (best-so-far ERR) and the evolution of the acquisition values ACQ over the elapsed trials within the 24-hour window.

In summary, our BO pipeline uses a Gaussian Process surrogate with a Matérn covariance function and the qLogNoisyExpectedImprovement (qLogNEI) acquisition function. To assess robustness with respect to modeling and initialization choices, we implement BO under different configurations: the Matérn smoothness parameter is set to $\nu \in \{3/2, 5/2\}$, and the initial sampling is generated using Sobol sequences with multiple random seeds. We report results across these configurations to reduce sensitivity to a single initialization and to account for the impact of the GP smoothness assumption on BO performance.

2.3.2 BO Workflow

Figure 2.4 illustrates the workflow of the calibration procedures, where the subscripts denote the trial index. First, an initial dataset is generated by evaluating 80 parameter vectors $\vec{P}_1, \ldots, \vec{P}_{80}$ sampled using a Sobol sequence. For each parameter vector $\vec{P}_i$, a MPC simulation is executed and the objective value ERR_i is computed as the discrepancy between the simulated trajectory and the measured data. Based on these initial observations, a Gaussian Process surrogate model is fitted to approximate the mapping from the parameter vector $\vec{P}$ to the error ERR. The acquisition function is then maximized to propose the next candidate parameter vector, denoted as $\vec{P}_{81}$. The MPC simulation is evaluated at this new parameter vector to obtain ERR_{81}, and the new observation is added to the dataset. The surrogate model is subsequently updated and the procedure iterates, sequentially generating new candidate parameters $\vec{P}_{82}, \ldots, \vec{P}_M$ until the computational budget is exhausted.

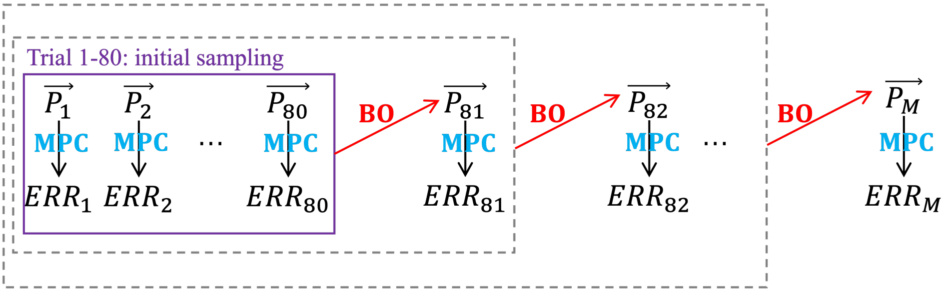

Figure 2.4: Bayesian Optimization Workflow

We implement the BO loop using the Ax platform [57] with the (modular) BoTorch backend [58]. Ax provides the higher-level experiment management and the generation strategy, while the BO surrogate modeling, acquisition computation, and candidate generation rely on BoTorch.

Algorithm 3 describes the BO loop pseudo code. We first create an `AxClient` and configure it with: the search space of $\vec{P}$, the objective ERR to be minimized, a predefined generation strategy (Center → Sobol → BoTorch), a fixed penalty value $ERR_{\text{bad}} = 100$ and a small observation noise $\sigma_{\text{sem}} = 0.01$ that is passed back to Ax when completing trials.

Algorithm 3 Bayesian Optimization Loop

```
Require: : Total number of trials $M$, Ax client $\mathscr{A}$, penalty value $ERR_{\text{bad}} = 100$, observation noise $\sigma_{\text{sem}} = 0.01$
Ensure: : A .csv log file with per-trial records and the best-so-far solution $(\vec{P}^*, ERR^*)$
$ERR^\star \leftarrow +\infty$
for $i \leftarrow 1$ to $M$ do
    $(\vec{P}$,trial_id$) \leftarrow \mathscr{A}$.get_next_trial()
    $\alpha^{\text{pre}} \leftarrow$ acq($\mathscr{A}, \vec{P}$)                    ▷ acquisition value before model update
    model_type $\leftarrow$ get_model_bridge_type($\mathscr{A}$)
    success $\leftarrow$ True
    $ERR \leftarrow$ evaluate_params($\vec{P}$)
    if evaluate_params failed then
        success $\leftarrow$ False
        $ERR \leftarrow ERR_{\text{bad}}$
    end if
    status $\leftarrow$ if success then ok else fail
    $y \leftarrow$ if (status=ok and $ERR$ finite) then $ERR$ else $ERR_{\text{bad}}$
    $\mathscr{A}$.complete_trial(trial_id, {"$ERR$" : ($y, \sigma_{\text{sem}}$)})
    $ERR^\star \leftarrow \min(ERR^\star, ERR)$
    log_csv($i, \vec{P}, ERR, \alpha^{\text{pre}}$,model_type,status)
end for
```

Before the loop starts, we initialize the best-so-far error as $ERR^\star \leftarrow +\infty$. For each iteration $i = 1, \ldots, M$, Ax proposes a new candidate by calling the method `get_next_trial()`, which returns the parameter vector $\vec{P}$ together with its `trial_id`. Since the generation strategy has already been specified when creating the `AxClient`, this method automatically follows the predefined sequence of generators. In our setting, the trial 0 corresponds to a center-point initialization, trials 1–80 are generated using Sobol sampling to obtain a space-filling initial design, and subsequent trials are produced by the BoTorch model using the GP surrogate and the qLogNEI acquisition function. Before evaluating the expensive MPC simulation, we query the current surrogate to record the acquisition value associated with this candidate (denoted as α^{pre}), and we also log the current model type via `get_model_bridge_type()` to keep track of which model or generation stage is active. We then evaluate the candidate parameters by running the MPC simulation and computing the resulting trajectory error ERR. During this evaluation step, any failure (e.g., simulator crash, infeasible run, or invalid/NaN output) is treated explicitly: we mark the evaluation as unsuccessful and replace the returned objective by a fixed penalty value ERR_{bad}. This guarantees that every proposed trial yields a valid scalar objective and can be safely reported back to Ax. After obtaining the objective value, the result is fed back to the BO model using method `complete_trial(trial_id,` {"ERR" : (y, σ_{sem})}), which updates the surrogate model with the new observation (including the specified observation noise). Finally, we update the best-so-far record as $ERR^\star \leftarrow \min(ERR^\star, ERR)$ and write all relevant information (iteration index, $\vec{P}$, ERR, α^{pre}, model type, and status) to `.csv` result file for offline analysis.

2.3.3 2-step BO

In practice, we don't calibrate all the parameters in $\vec{P}$ (i.e., w_{avoid}, R_0, $w_{\dot{\psi}}$, $N_{\min}$, $N_{\max}$, d_0) altogether simultaneously, as doing so may lead to parameter compensation and identifiability issues. More specifically, the effects of prediction-horizon-related factors ($N_{\min}$, $N_{\max}$, d_0) can become entangled with those of other factors (w_{avoid}, R_0, $w_{\dot{\psi}}$), such that different parameter combinations yield similar trajectories despite representing distinct behavioural interpretations. For example, an underestimated prediction horizon $N_{\min}$ and $N_{\max}$ may be offset by an overestimated avoidance-avoidance cost weight w_{avoid}, such that the combined effect produces trajectories that are close to the measurement despite misrepresenting the underlying anticipation behavior. In such cases, good trajectory fit alone does not guarantee that anticipation-related parameters are correctly identified. Therefore, the follwoing two-step BO procedure is adopted.

Step 1: Background Parameters Calibration: In Step 1, we assume constant prediction horizon (N_{const}) MPC and calibrate the background parameters, namely ($N_{\text{const}}, w_{\text{avoid}}$, R_0, $w_{\dot{\psi}}$). The goal in this step is to identify a shared set of background parameters that captures the general interaction behaviour across multiple overtaking scenarios. Therefore, we perform BO over the parameter space (N_{const}, w_{avoid}, $R_0, w_{\dot{\psi}}$), where each BO trial proposes a candidate parameter vector $\vec{P}_k$.

As illustrated in Figure 2.5, the 30 overtaking pairs are divided into two groups (pairs 1–18 and pairs 19–30). This grouping is motivated by systematic differences in both road environment and trajectory characteristics. Pairs 1–18 are located on an approximately horizontal main road segment, whereas pairs 19–30 occur on a nearly vertical main road segment, implying potentially different environment conditions (e.g., road width). Also, trajectory features are summarized using the interquartile range (IQR) (i.e., the difference between the 75th and 25th percentiles, which are shown in Table 2.3. In particular, the IQR of the lateral deviation from the centerline, $\bar{e_d}$, differs noticeably between the two groups (approximately 27.5%), indicating that the scenarios exhibit distinct lateral movement characteristics. Moreover, preliminary experiments in which all 30 pairs were calibrated jointly revealed an imbalanced fitting performance: the resulting parameters produced consistently good trajectory matches for pairs 19–30 but noticeably larger errors for pairs 1–18. This suggests that a single shared parameter set tends to favor one subset of scenarios and fails to represent the other sufficiently. To mitigate this bias and better capture group-specific interaction characteristics, the dataset is therefore divided into two groups and calibrated separately.

Table 2.3: Normalization constants for the two groups of cyclists

Trajectory group	$\bar{e}_s\,[m]$	$\bar{e}_d\,[m]$	$\bar{v}\,[m/s]$	$\bar{\dot{\psi}}\,[rad/s]$
pairs 1–18	29.822	1.013	0.804	0.0707
pairs 19–30	26.398	0.734	0.796	0.0836
Difference (%)	11.5	27.5	1.0	18.2

As shown in Figure 2.5, for a given BO trial k, the same candidate vector $\vec{P}_k^{1-18}$ is applied to all scenarios in Group 1. Each scenario is simulated independently using MPC, producing trajectory errors $ERR_k^1, \dots, ERR_k^{18}$. These errors are aggregated into a single objective value defined as the mean trajectory error $ERR_k^{1\text{–}18,\text{mean}} = \frac{1}{18}\sum_{j=1}^{18} ERR_k^j$. An analogous procedure is applied to Group 2 using $\vec{P}_k^{19\text{–}30}$, yielding $ERR_k^{19\text{–}30,\text{mean}} = \frac{1}{12}\sum_{j=19}^{30} ERR_k^j$. In Step 1, the BO objective is therefore to minimize the group-wise mean trajectory error, ensuring that all scenarios within each group share the same appropriate background parameters.

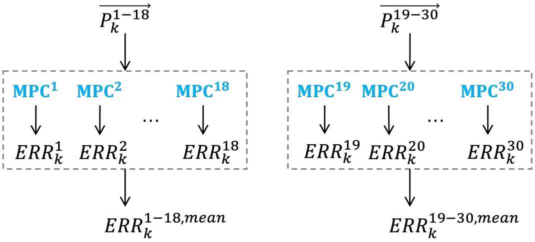

Figure 2.5: Mapping between $\vec{P}$ and ERR in step 1 BO (background paramters calibration)

Step 2: Horizon-varying Parameter Calibration: In Step 2, we calibrate the horizon-varying parameters of the MPC model, namely (N_{min}, N_{max}, d_0), while keeping the background parameters fixed to the optimal values obtained in Step 1. The purpose of this step is to capture individual anticipation behaviour that cannot be explained solely by the shared background parameters.

Unlike Step 1, where a single parameter vector is shared across scenarios within each group, the horizon-varying parameters are calibrated individually for each overtaking pair. As illustrated in Figure 2.6, for a given BO trial k, each pair j is assigned its own candidate parameter vector $\vec{P}_k^j$, which contains the horizon-related parameters (N_{min}, N_{max}, d_0) together with the optimal background parameters (w_{avoid}^*, R_0^*, w_ψ^*) obtained from Step 1. The MPC simulation is then performed for that specific pair, producing the corresponding trajectory error ERR_k^j ($j = 1, \dots, 30$). BO is conducted independently for the overtaking cyclist in each pair by minimizing its own trajectory error, allowing the optimal horizon-varying parameters to adapt to the characteristics of individual overtaking scenarios.

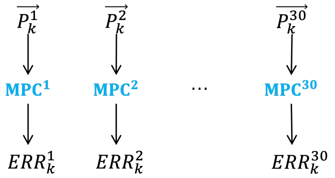

Figure 2.6: Mapping between $\vec{P}$ and ERR in step 2 BO (horizon-varying parameters calibration)

As discussed in Section 2.3.1, BO results can be different across difference initial sampling and choices of Matérn kernel smoothness parameter ν. Therefore, we repeat Step 1 under the 6 BO configuration settings in Table 2.4, And we repeat step 2 under 6 BO configuration settings in Table 2.5, defined by the combination of the Matérn kernel smoothness parameter $\nu \in \{3/2, 5/2\}$ and different Sobol initialization random-seed sets. The parameter space for the step 1 and step 2 are shown in Table 2.6.

Setting ID	Matérn ν	Initial Sampling Random Seed for All Pairs
S1	1.5	3
S2	1.5	6
S3	1.5	9
S4	2.5	3
S5	2.5	6
S6	2.5	9

Table 2.4: BO configurations in step 1 (background parameters calibration)

Setting ID	Matérn ν	Initial Sampling Random Seed for (pair 1, pair2, ..., pair 30)
S1	1.5	(1, 2, ..., 30)
S2	1.5	(101, 102, ..., 130)
S3	1.5	(201, 202, ..., 230)
S4	2.5	(1, 2, ..., 30)
S5	2.5	(101, 102, ..., 130)
S6	2.5	(201, 202, ..., 230)

Table 2.5: BO configurations in step 2 (horizon-varying parameters calibration)

Parameter	Domain
N_{const}	$\{10, 11, \ldots, 60\}$
w_{avoid}	$\{5, 10, \ldots, 100\}$
R_0	$\{0.5, 1.0, \ldots, 5.0\}$
w_{ψ}	$\{0.025, 0.050, \ldots, 0.500\}$
N_{min}	$\{10, 11, \ldots, 60\}$
N_{max}	$\{20, 21, \ldots, 100\}$
d_0	$\{0.0, 0.5, \ldots, 10.0\}$

Table 2.6: Parameter domains used in BO.

3

RESULTS AND DISCUSSION

In this chapter, we show and discuss the results of calibration step 1 (background parameters calibration) and calibration step 2 (horizon-varying parameters calibration). In all, varying-horizon MPC outperforms constant-horizon MPC by reducing 54.43% error compared to the measurement. Based on the calibrated results, we further analyze how the calibrated optimal anticipation-related factors are correlated with interaction features.

3.1 Results of Calibration Step 1 (Background Parameters Calibration)

In calibration step 1, a constant-horizon MPC framework is assumed, and the background parameters (N_{const}, w_{avoid}, R_0, $w_{\dot{\psi}}$) are jointly calibrated for two cyclist groups (i.e., the overtaking cyclists in pairs 1–18 and 19–30), aiming to minimize the mean ERR of the overtaking cyclist trajectory for the whole group. To mitigate the sensitivity of Bayesian Optimization (BO) to a single parameter setting, six configurations, denoted as S1–S6 listed in Table 2.4, are considered.

Figures 3.1 and 3.2 illustrate the evolution of the best-so-far ERR_mean (smallest mean of all pairs' error compared to measurement) and the acquisition value of the currently selected candidate over the optimization trials for the different Bayesian Optimization configurations (S1–S6). In each Bayesian Optimization configuration, the upper plots show the progression of the best ERR_mean achieved so far, revealing how the objective improves over time. And the lower plots display the acquisition value of the currently selected candidate parameter at each trial, reflecting the expected potential for further improvement.

Following Jones et al. [59], we can consider the convergence of BO to reach when the acquisition value becomes smaller than 1% of the current best objective value. In this study, however, we consider such a threshold too strict in view of the available computational budget and the high cost of running additional trials. For this reason, we adopt a more relaxed criterion of 10% to assess practical convergence. Under this loosened relative criterion (10%), we consider several configurations to have reached practical convergence based on $ERR_{\text{min}}/ACQ_{\text{min}} < 0.1$ as shown in the figures. For pair 1–18, we regard all configurations as converged and for pair 19–30, we consider configurations S1, S2, S5 and S6 as converged. In contrast, for pair 19–30, configurations S3 and S4 show larger minimum acquisition values ACQ_{min} than convergence requirement, suggesting that the search process may still have potential for further improvement.

Regardless of the convergence behavior of different BO configurations, we choose the final background parameter set as the one yielding the lowest mean ERR among all 6 BO configurations. Accordingly, for pair 1–18 and pair 19–30, the adopted parameter combinations are $(N_{\text{const}} = 51,\ w_{\text{avoid}} = 70,\ R_0 = 1.0,\ w_{\dot{\psi}} = 0.35)$ (achieved in BO configuration S3) and $(N_{\text{const}} = 45,\ w_{\text{avoid}} = 5,\ R_0 = 0.5,\ w_{\dot{\psi}} = 0.1)$ (achieved in BO configurations S3 and S6) respectively.

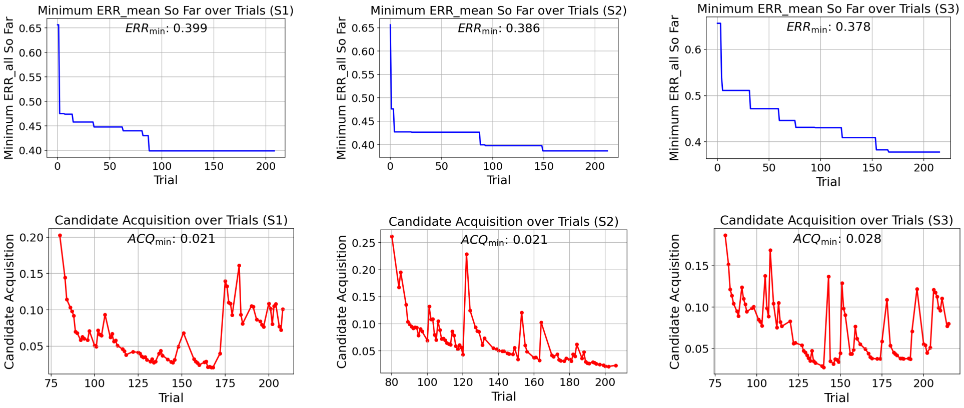

(1). S1: (ν=1.5, seed=3)

(2). S2: (ν=1.5, seed=6)

(3). S3: (ν=1.5, seed=9)

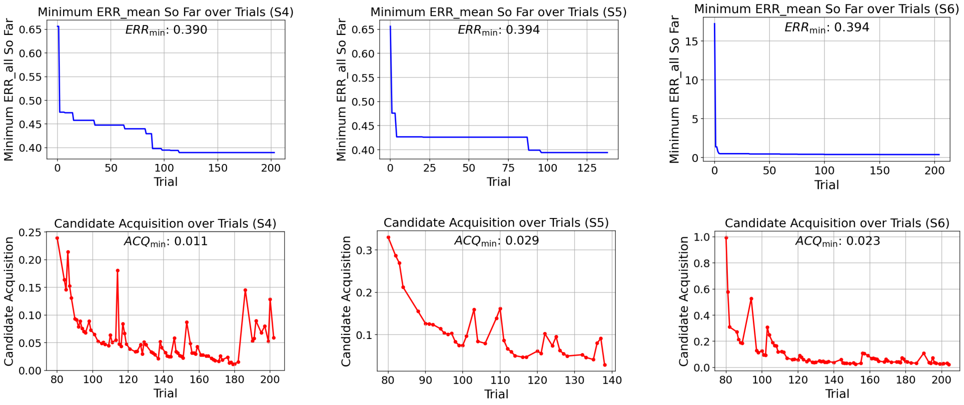

(4). S4: (ν=2.5, seed=3)

(5). S5: (ν=2.5, seed=6)

(6). S6: (ν=2.5, seed=9)

Figure 3.1: Pair 1–18, step 1 (background parameters calibration): evolution of the minimum ERR_mean so far and acquisition value over trials for BO configurations S1–S6.

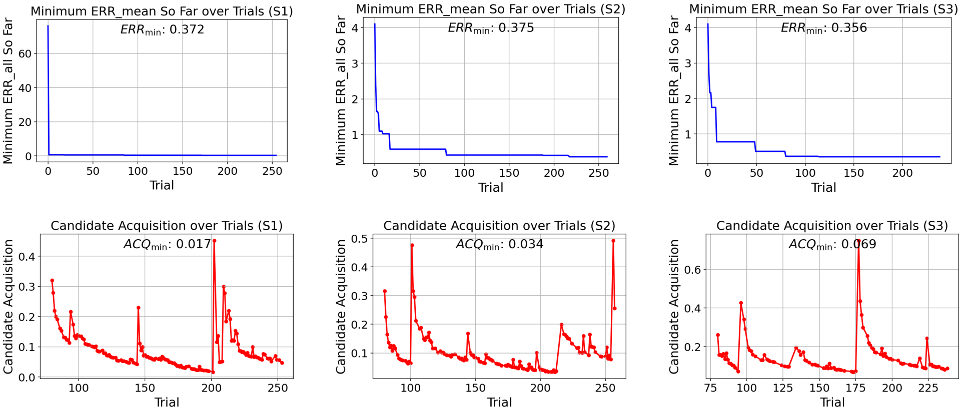

(1). S1: (ν=1.5, seed=3) (2). S2: (ν=1.5, seed=6) (3). S3: (ν=1.5, seed=9)

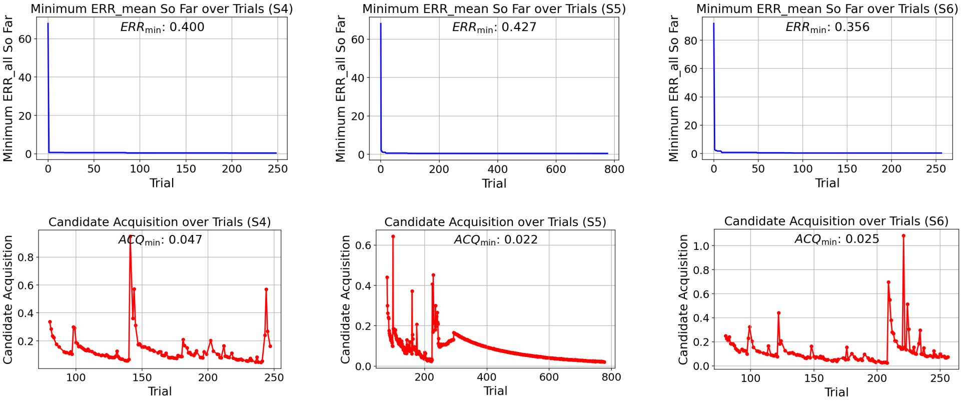

(4). S4: (ν=2.5, seed=3) (5). S5: (ν=2.5, seed=6) (6). S6: (ν=2.5, seed=9)

Figure 3.2: Pair 19–30, step 1 (background parameters calibration): evolution of the minimum ERR_mean so far and acquisition value over trials for BO configurations S1–S6.

3.2 Results of BO Step 2: Horizon-varying Parameter Calibration

Keeping the optimal background parameters (w^*_{goal}, R^*_0, w^*_ψ) obtained in BO step 1 fixed, we then optimize the horizon-varying parameters ($N_{\min}, N_{\max}, d_0$) for each individual pair using the 6 BO configurations S1 to S6, as specified in Table 2.5.

Figures 3.3 to 3.8 illustrate the measured trajectories, the trajectories simulated by the constant-horizon MPC with the optimal parameters calibrated in step 1, and the trajectories simulated by the varying-horizon MPC with the optimal parameters calibrated in step 2. They also show the evolution of the horizon N_{const} or N_{vary} over the simulation time for all 30 pairs, where star markers indicate the overtaking point.

Table 3.1 summarizes the ERR values for the 30 pairs in step 1 (constant-horizon MPC) and step 2 (varying-horizon MPC), together with the corresponding improvement achieved in step 2 and the calibrated optimal horizon-varying parameters. The average ERR decreases substantially from $0.36912m$ in step 1 to $0.16821m$ in step 2, corresponding to a reduction of approximately 54.429%.

After obtaining the optimal horizon-varying parameters in step 2, we further investigate how anticipation-related factors are associated with different interaction scenarios. As summarized in Table 3.2, we report anticipation-related factors and trajectory features for all 30 pairs. Here, $N_{\text{use,min}}$ denotes the minimum prediction horizon actually used by cyclist A during the entire interaction, and $N_{\text{use,overtaking}}$ denotes the prediction horizon of cyclist A at the overtaking point. Note that $N_{\min}$ and $N_{\text{use,min}}$ are conceptually different: $N_{\min}$ represents the theoretical lower bound of the prediction horizon, whereas $N_{\text{use,min}}$ represents the minimum prediction horizon actually used in the simulation, which depends on the inter-cyclist distance according to Equation 2.22. The trajectory features considered are $v_{A,0}$, the initial velocity of cyclist A; $v_{A,\text{overtaking}}$, the velocity of the cyclist A at the overtaking point; $d_{AB,\min}$, the minimum distance between the two cyclists. Although we also examined some other factors (e.g., relative velocity of the two cyclists), we only include the trajectory features that exhibit at least moderate correlation with the anticipation-related factors here.

For the following analysis, samples with $ERR > 0.25\,\text{m}$ are treated as poor matches and excluded from the correlation analysis. This applied to pairs 8, 9, 20 and 22, leaving 26 samples for the subsequent investigation. To examine linear associations between variables, we use Pearson's correlation coefficient, as it is appropriate for assessing the strength and direction of a linear relationship between two variables [60]. In interpreting the results, both the correlation coefficient r and the corresponding p-value were considered. Following Cohen's conventional guidelines for interpreting correlation magnitude in behavioral research [61], absolute values of r of 0.1, 0.3, and 0.5 are commonly interpreted as indicating weak, moderate, and strong relationships, respectively. In addition, we assess statistical significance using the conventional threshold of $p < 0.05$, where p-value indicates the probability of obtaining a correlation as large as the observed one purely by chance assuming that no true correlation exists. And $p < 0.05$ is taken to indicate evidence against the null hypothesis at the conventional 5% significance level [62]. Accordingly, only relationships that satisfy both $|r| > 0.3$ and $p < 0.05$ are regarded as meaningful and are therefore discussed below.

As shown in Figure 3.9, among the anticipation-related factors, we find the strongest correlation between $N_{\text{use,overtaking}}$ and $N_{\text{use,min}}$, with a perfect positive linear relation-

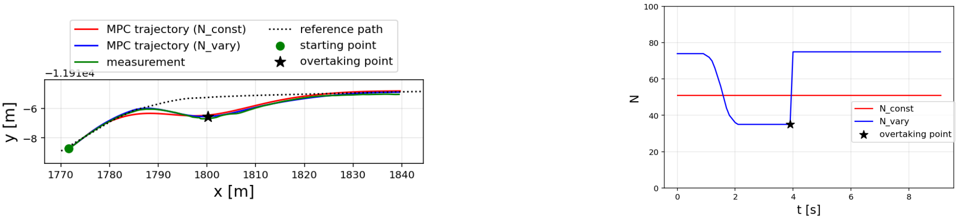

(1). Pair 1: Trajectory Comparison and Prediction Horizon Evolution

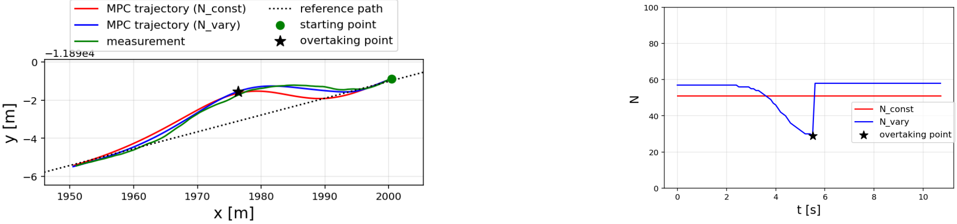

(2). Pair 2: Trajectory Comparison and Prediction Horizon Evolution

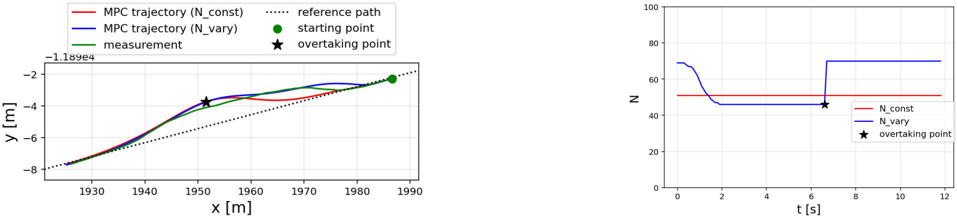

(3). Pair 3: Trajectory Comparison and Prediction Horizon Evolution

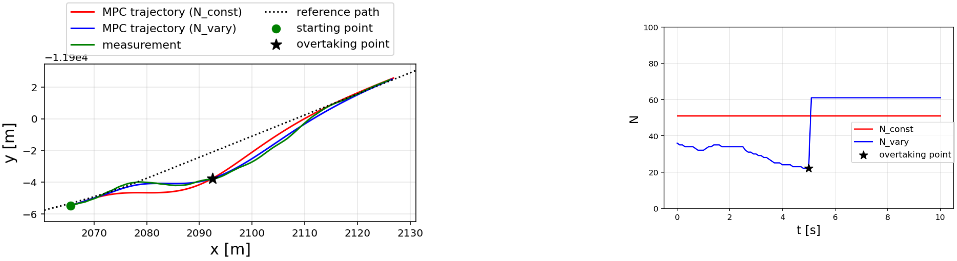

(4). Pair 4: Trajectory Comparison and Prediction Horizon Evolution

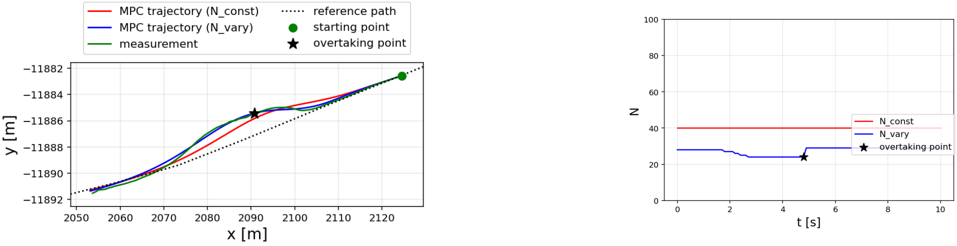

(5). Pair 5: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.3: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 1–5. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

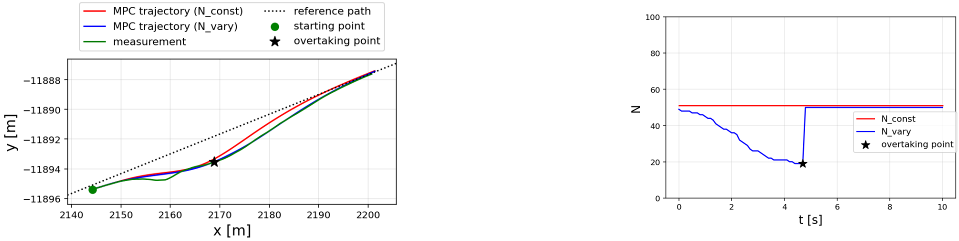

(1). Pair 6: Trajectory Comparison and Prediction Horizon Evolution

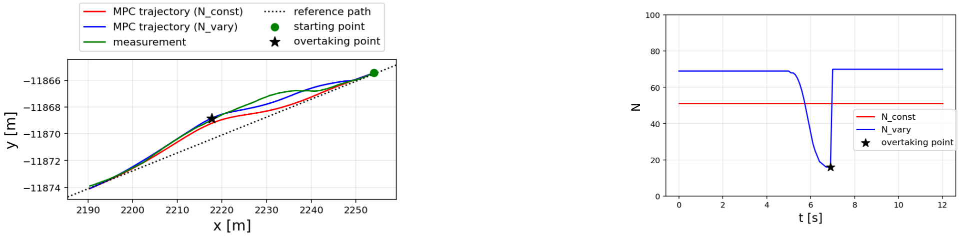

(2). Pair 7: Trajectory Comparison and Prediction Horizon Evolution

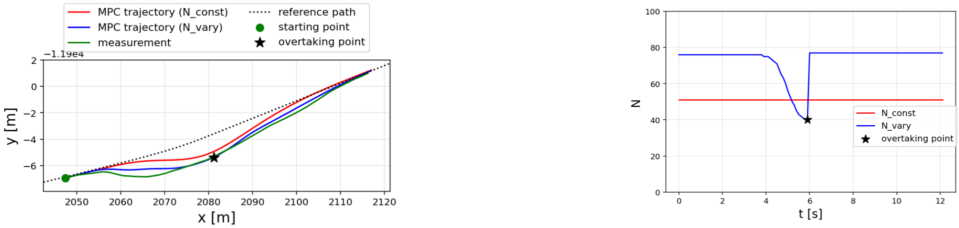

(3). Pair 8: Trajectory Comparison and Prediction Horizon Evolution

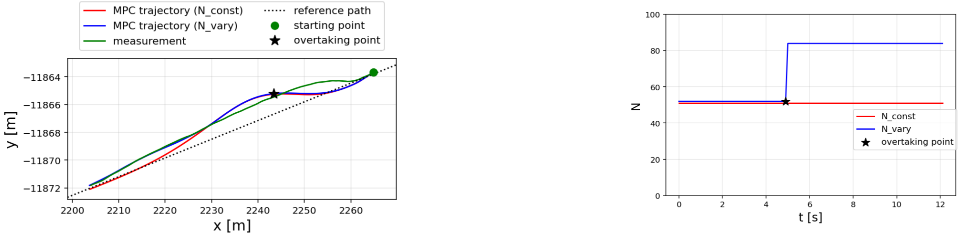

(4). Pair 9: Trajectory Comparison and Prediction Horizon Evolution

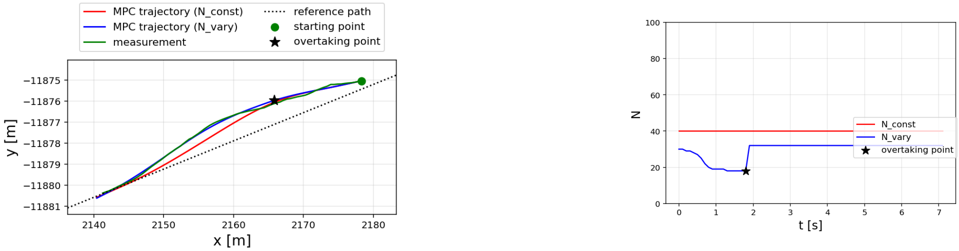

(5). Pair 10: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.4: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 6–10. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

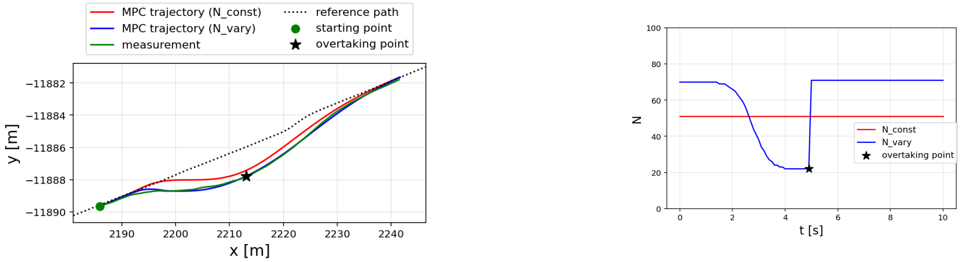

(1). Pair 11: Trajectory Comparison and Prediction Horizon Evolution

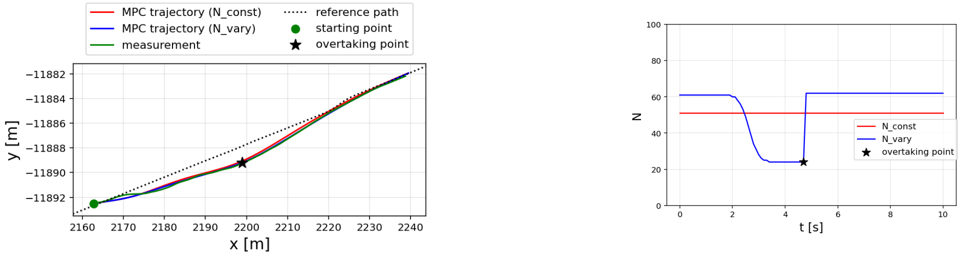

(2). Pair 12: Trajectory Comparison and Prediction Horizon Evolution

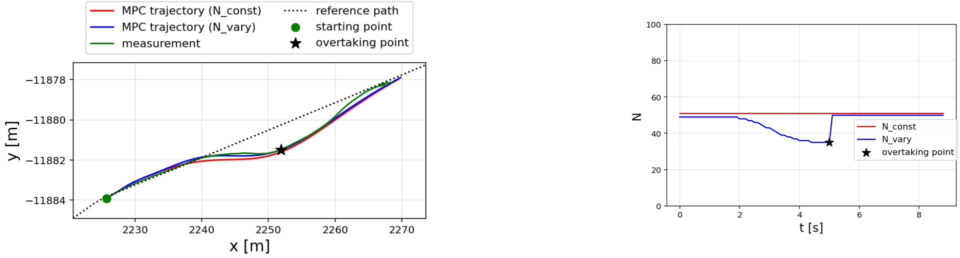

(3). Pair 13: Trajectory Comparison and Prediction Horizon Evolution

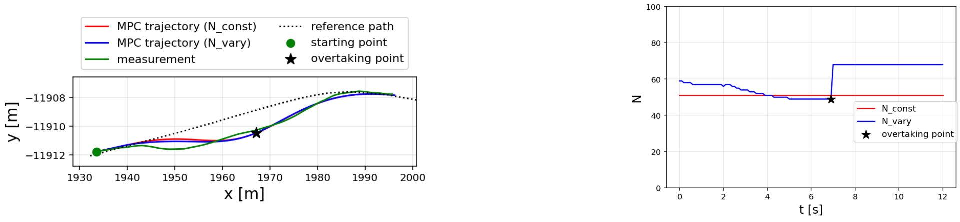

(4). Pair 14: Trajectory Comparison and Prediction Horizon Evolution

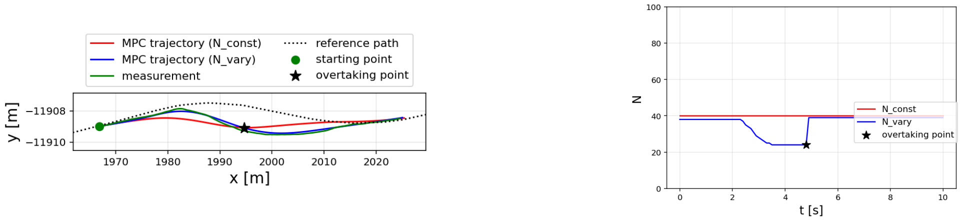

(5). Pair 15: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.5: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 11–15. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

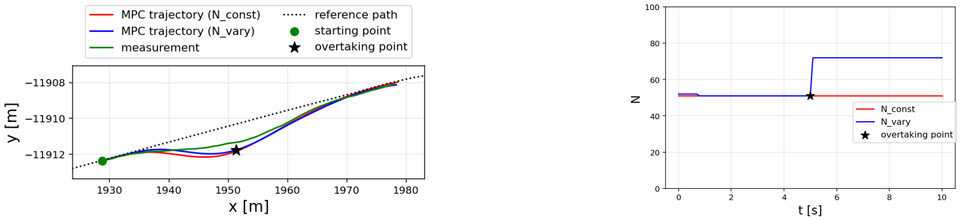

(1). Pair 16: Trajectory Comparison and Prediction Horizon Evolution

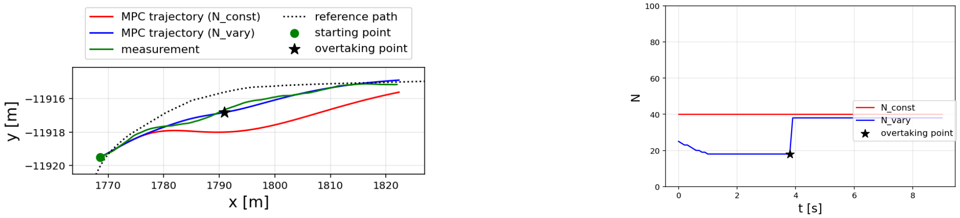

(2). Pair 17: Trajectory Comparison and Prediction Horizon Evolution

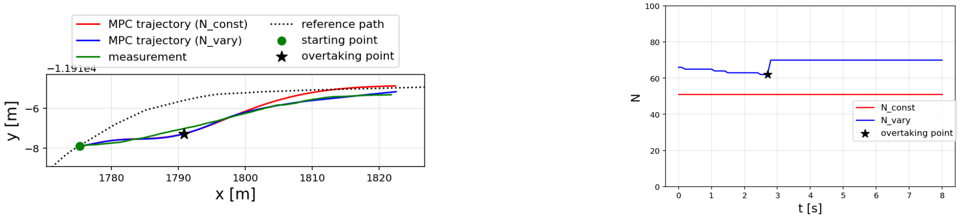

(3). Pair 18: Trajectory Comparison and Prediction Horizon Evolution

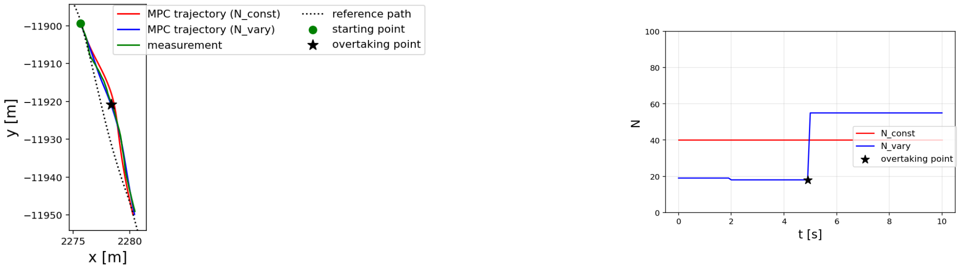

(4). Pair 19: Trajectory Comparison and Prediction Horizon Evolution

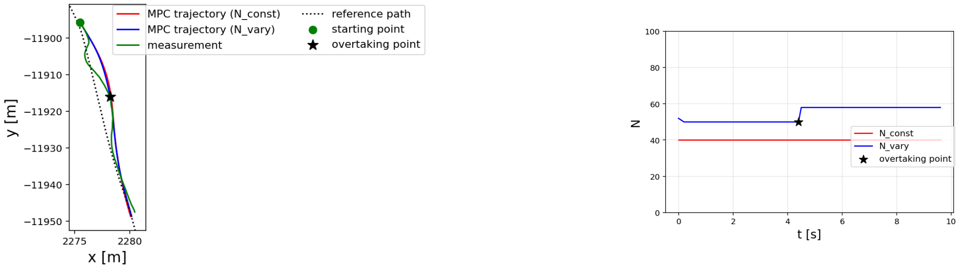

(5). Pair 20: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.6: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 16–20. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

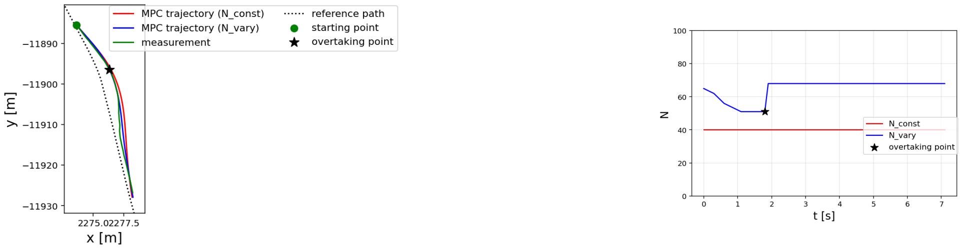

(1). Pair 21: Trajectory Comparison and Prediction Horizon Evolution

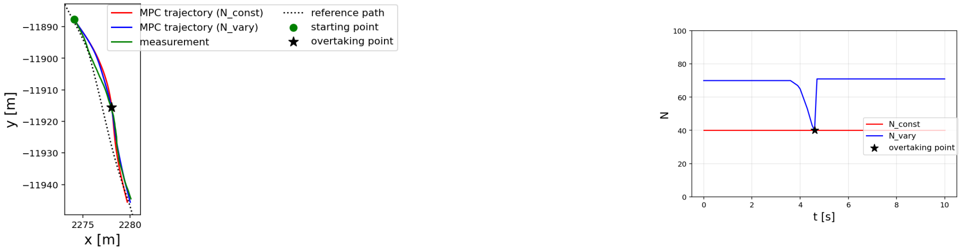

(2). Pair 22: Trajectory Comparison and Prediction Horizon Evolution

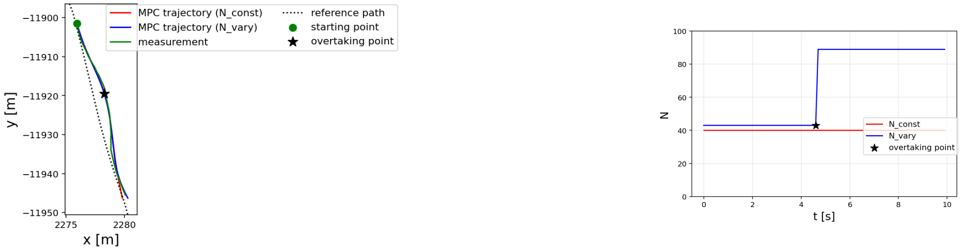

(3). Pair 23: Trajectory Comparison and Prediction Horizon Evolution

(4). Pair 24: Trajectory Comparison and Prediction Horizon Evolution

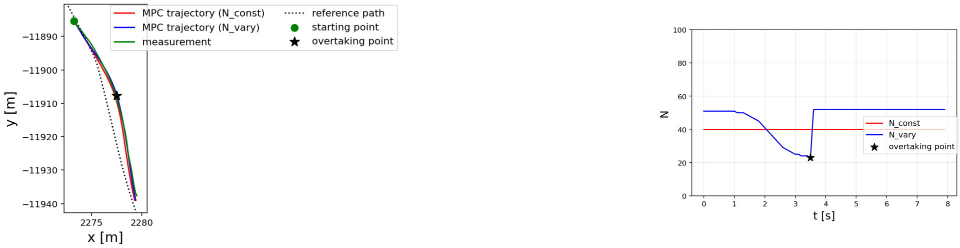

(5). Pair 25: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.7: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 1–5. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

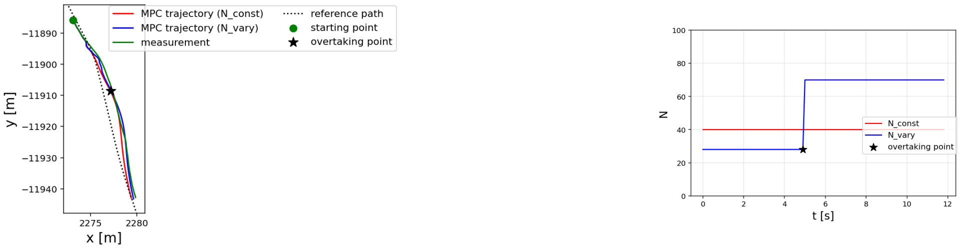

(1). Pair 26: Trajectory Comparison and Prediction Horizon Evolution

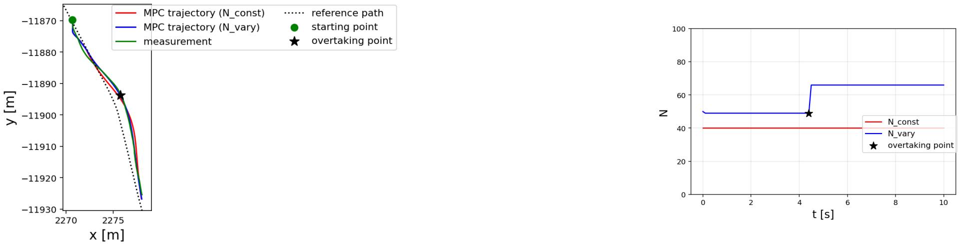

(2). Pair 27: Trajectory Comparison and Prediction Horizon Evolution

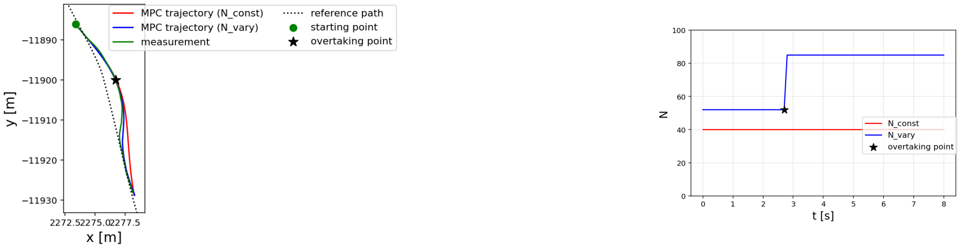

(3). Pair 28: Trajectory Comparison and Prediction Horizon Evolution

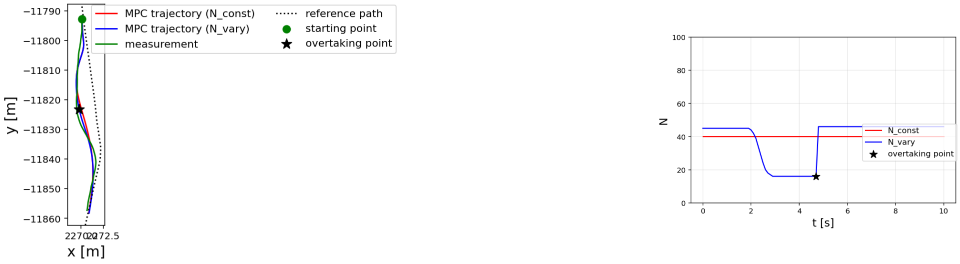

(4). Pair 29: Trajectory Comparison and Prediction Horizon Evolution

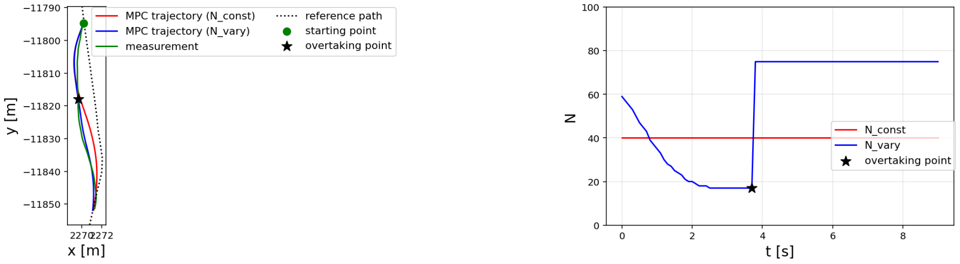

(5). Pair 30: Trajectory Comparison and Prediction Horizon Evolution

Figure 3.8: Comparison between the measured trajectory and the trajectories simulated by MPC with a constant prediction horizon and a varying prediction horizon for pairs 26–30. The left column shows the trajectory comparison, while the right column shows the corresponding prediction horizon N over simulation time.

Table 3.1: Best objective ERR_{min} obtained from Step 1 (background parameters calibration) and that from Step 2 (horizon-varying parameters calibration) for all 30 pairs, and the N^*_{min}, N^*_{max}, and d^*_0 values corresponding to ERR_{min} in Step 2.

pair	ERR_{min} [m] (step 1)	ERR_{min} [m] (step 2)	improve	N^*_{min}	N^*_{max}	d_0 [m]
1	0.17238	0.07297	0.57666	35	75	8.0
2	0.32678	0.12957	0.60349	26	58	2.5
3	0.29196	0.19027	0.34830	46	70	9.5
4	0.37369	0.12243	0.67238	21	61	3.0
5	0.40793	0.17599	0.56859	24	29	5.5
6	0.32514	0.11427	0.64854	17	50	3.0
7	0.49716	0.23056	0.53624	15	70	3.5
8	0.62598	0.26353	0.57901	31	77	2.5
9	0.47204	0.34731	0.26424	52	84	10.0
10	0.27537	0.08428	0.69396	18	32	3.0
11	0.37961	0.08601	0.77343	22	71	5.0
12	0.17277	0.08960	0.48138	24	62	6.5
13	0.18839	0.12453	0.33897	35	50	3.0
14	0.29475	0.24118	0.18175	49	68	4.0
15	0.38226	0.09636	0.74792	24	39	5.0
16	0.26605	0.20889	0.21483	51	72	5.0
17	1.08696	0.16001	0.85279	18	38	6.5
18	0.26423	0.11296	0.57247	40	74	10.0
19	0.20257	0.08842	0.56353	18	55	3.5
20	0.52016	0.47834	0.08040	50	58	7.5
21	0.25286	0.12542	0.50400	51	68	3.5
22	0.35786	0.25543	0.28624	10	71	1.0
23	0.12971	0.10252	0.20962	43	89	10.0
24	0.82226	0.14286	0.82625	12	30	2.5
25	0.21669	0.11947	0.44864	22	52	2.5
26	0.36074	0.19551	0.45803	28	70	7.0
27	0.25179	0.13896	0.44812	49	66	9.5
28	0.33869	0.11589	0.65783	52	85	6.0
29	0.29188	0.23265	0.20295	16	46	10.0
30	0.52494	0.20015	0.61873	17	75	5.0
average	0.36912	0.16821	0.54429			

Table 3.2: Selected trajectory features and anticipation-related parameters

pair	$d_{AB,\text{min}}$ $[m]$	$v_{A,\text{overtaking}}$ $[m/s]$	$v_{A,0}$ $[m/s]$	N_{min}	$N_{\text{use,min}}$	$N_{\text{use,overtaking}}$	N_{max}
1	1.4304	7.0287	7.0129	35	35	35	75
2	1.2744	5.2545	3.1694	26	29	29	58
3	1.1577	5.1099	5.8385	46	46	46	70
4	1.4332	6.3175	4.8993	21	22	22	61
5	1.6937	7.4288	6.6534	24	24	24	29
6	1.6778	5.2960	5.1238	17	19	19	50
7	4.4684	4.2322	4.2322	15	16	16	70
8	3.2763	4.7269	4.7269	31	40	40	77
9	3.7109	3.3399	3.3399	52	52	52	84
10	1.2822	4.7238	5.9610	18	18	18	32
11	1.9197	5.0425	5.0425	22	22	22	71
12	4.6690	6.7875	6.7875	24	24	24	62
13	2.4855	5.1047	5.7149	35	35	35	50
14	2.1824	4.2997	4.3222	49	49	49	68
15	4.1028	6.1890	6.0358	24	24	24	39
16	3.2456	4.1108	4.1108	51	51	51	72
17	5.1292	5.0945	5.0945	18	18	18	38
18	2.3276	4.8541	4.8541	40	62	62	74
19	1.0004	4.1280	4.1280	18	18	18	55
20	2.5910	6.0448	6.0463	50	50	50	58
21	4.7751	5.3775	5.3775	51	51	51	68
22	1.0351	6.4872	6.4872	10	10	10	71
23	4.1223	2.8733	2.8733	43	43	43	89
24	5.0057	6.2910	6.2910	12	12	12	30
25	5.9962	5.6941	5.6941	22	22	22	52
26	1.9705	4.4835	4.4835	28	28	28	70
27	4.8880	5.4817	5.4817	49	49	49	66
28	4.4924	4.4353	4.4353	52	52	52	85
29	1.9971	6.5583	5.8599	16	16	16	46
30	1.6931	6.5617	5.6104	17	17	17	75

ship ($r = 1.000$). This indicates that, across all 26 samples, the cyclist's minimum prediction horizon in practice is always reached at the overtaking point. In addition, we observe a very strong positive correlation between $N_{\text{use,min}}$ and N_{min} ($r = 0.955$), indicating that the lower bound of prediction horizon N_{min} is not merely a theoretical lower bound. Rather, overtaking cyclists in almost all samples (except pair 18) have reduced their prediction horizon to this lower-bound level at the overtaking point in practice.

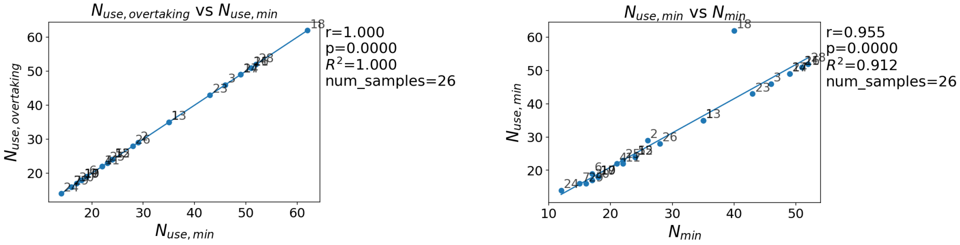

(1). $N_{\text{use,overtaking}} - N_{\text{use,min}}$ Correlation (2). $N_{\text{use,min}} - N_{\text{min}}$ Correlation

Figure 3.9: Correlations between $N_{\text{use,min}}$ and $N_{\text{use,overtaking}}$, N_{min} (blue dots indicate pair index)

As shown in Figure 3.10, we observe a moderate negative correlation between N_{min} and $N_{\text{use,overtaking}}$ ($r = -0.409$), as well as between $v_{A,\text{overtaking}}$ and $N_{\text{use,overtaking}}$ ($r = -0.396$). This suggests that, at the overtaking moment, a higher overtaking speed is associated with a more reactive behavior, reflected by both a smaller theoretical lower bound of the prediction horizon N_{min}, and a shorter practically used prediction horizon $N_{\text{use,overtaking}}$.

In addition, both N_{min} and $N_{\text{use,overtaking}}$ are negatively correlated with $d_{AB,\text{min}}$, with correlation coefficients of $r = -0.490$ and $r = -0.460$, respectively. This trend may seem counterintuitive if a smaller minimum distance represents a more urgent situation that should result in more reactive behavior. However, a larger $d_{AB,\text{min}}$ may also indicate that the overtaking cyclist starts adapting earlier and maintains a larger safety margin throughout the maneuver. This is consistent with the findings reported in [1], where a larger prediction horizon lead to safer interactions. It is also in agreement with the observations of large-scale bicycle experiment in [17], where cyclists anticipated a bottleneck earlier and adapted accordingly, ultimately maintaining a larger inter-cyclist distance before the bottleneck. Under this interpretation, a larger minimum inter-cyclist distance is associated with a more anticipatory rather than more reactive strategy, which in turn requires a larger theoretical lower bound N_{min} and a larger practically used prediction horizon at the overtaking point $N_{\text{use,overtaking}}$. This interpretation remains tentative, but it suggests that the calibrated horizon-varying parameters may also capture how early cyclists start anticipating and adapting to the interaction, rather than only reflecting the urgency of the situation. Also, this does not contradict the assumed distance-dependent horizon function (2.22) in the methodology chapter. The model assumption concerns how the prediction horizon varies with the instantaneous inter-cyclist distance within a single interaction pair, whereas the correlation analysis here concerns how calibrated horizon parameters vary across different cyclist pairs as a function of the minimum distance reached during the whole interaction $d_{AB,\mathrm{min}}$. These two levels of analysis are different and can coexist.

These results may point to an additional mechanism that is not explicitly captured in the current model representing the variation of prediction horizon as in Equation (2.22), where a larger inter-cyclist distance leads to a larger prediction horizon. The observed correlation between larger calibrated prediction horizon parameters and larger minimum inter-cyclist distances suggests that the horizon may also reflect a higher-level tendency (desire) to maintain a safer interaction margin. In this sense, the prediction horizon may depend not only on the actual distance at a given moment, but also on the cyclist's anticipatory regulation toward a desired or preferred spacing that the cyclist tends to maintain during the interaction. A larger calibrated horizon may therefore reflect a stronger tendency toward safer, more anticipatory overtaking behavior.

Furthermore, N_{max} is negatively correlated with the speed of cyclist A at the overtaking point $v_{A,\mathrm{overtaking}}$ ($r = -0.400$) and the initial speed of cyclist A, $v_{A,0}$ ($r = -0.480$). In other words, when the overtaking cyclist rides faster (either at the overtaking moment or already at the beginning of the interaction), the theoretical upper bound of the prediction horizon N_{max} tends to be smaller. This trend suggests that cyclists riding faster rely less on long-horizon anticipation and instead exhibit more reactive behavior with shorter prediction horizons.

Taken together, these correlations suggest that the prediction horizon of the overtaking cyclist is not fixed, but varies systematically with both the cyclist's speed and the inter-cyclist outcome of the interaction. Higher speeds, both at the beginning of the interaction $v_{A,0}$ and at the overtaking moment $v_{A,\mathrm{use,overtaking}}$, are associated with smaller bounds of prediction horizon N_{min} N_{max} and the prediction horizon adopted at the overtaking point $N_{\mathrm{use,overtaking}}$, indicating a shift toward more reactive behavior with shorter prediction horizons. In contrast, larger minimum inter-cyclist distances $d_{AB,\mathrm{min}}$ are associated with larger lower bound N_{min} and practically used horizons at the overtaking point $N_{\mathrm{use,overtaking}}$, suggesting that more anticipatory planning contributes to maintaining a larger safety margin during overtaking. This implies that the calibrated horizon-varying parameters may reflect a balance between two behavioral tendencies: reactive behavior under high-speed conditions and anticipative behavior when aiming to maintain safer inter-cyclist spacing.

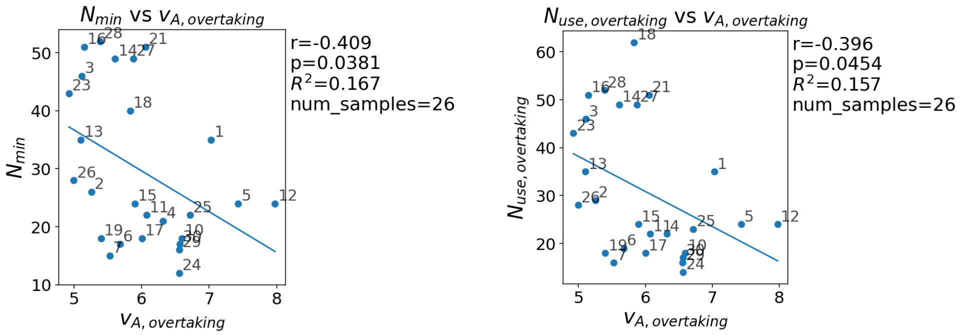

(1). $N_{\text{min}} - v_{A,\text{overtaking}}$ Correlation

(2). $N_{\text{use,overtaking}} - v_{A,\text{overtaking}}$ Correlation

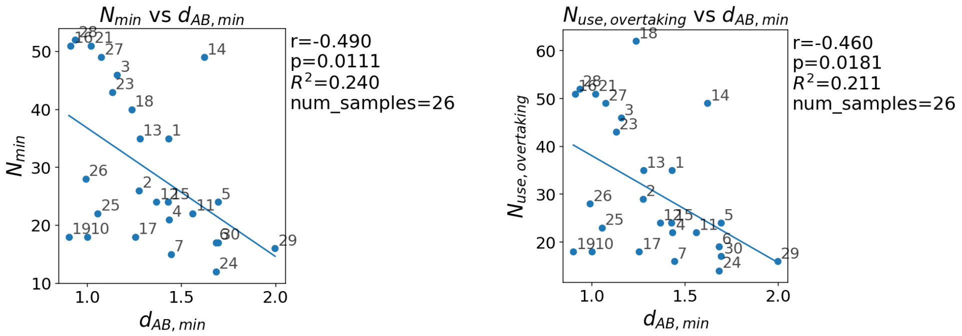

(3). $N_{\text{min}} - d_{AB,\text{min}}$ Correlation

(4). $N_{\text{use,overtaking}} - d_{AB,\text{min}}$ Correlation

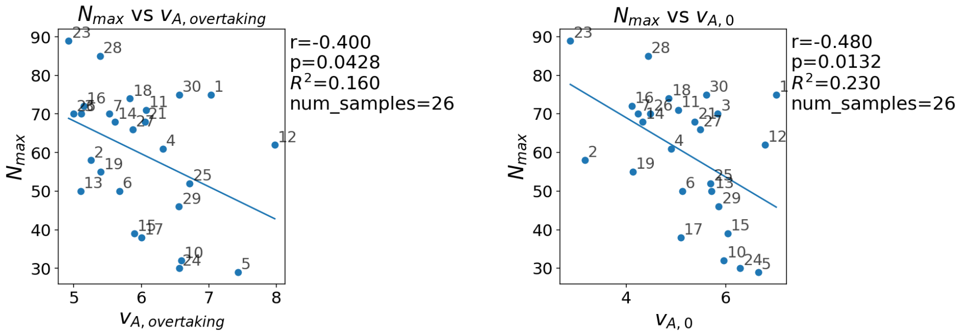

(5). $N_{\text{max}} - v_{A,\text{overtaking}}$ Correlation

(6).$N_{\text{max}} - v_{A,0}$ Correlation

Figure 3.10: Correlations between interaction features and anticipation-related parameters (blue dots indicate pair index).

4

CONCLUSION AND DISCUSSION

4.1 Conclusion

This thesis investigates how cyclists regulate anticipation during pairwise overtaking interactions and proposes an adaptive-horizon Model Predictive Control (MPC) framework to explicitly represent this mechanism. The prediction horizon was modeled as a smooth and bounded function of the inter-cyclist distance, such that the horizon could continuously vary between a lower bound N_{min} and an upper bound N_{max}. We calibrate the model parameters against measured overtaking cyclist trajectories from the TUM-DOT dataset using a two-step BO procedure: first, background MPC parameters are calibrated under a constant-horizon MPC assumption, and second, the horizon-varying parameters are calibrated for each cyclist pair individually. In this way, anticipation is not treated as a fixed parameter, but as a context-dependent variable that changes during the interaction.

The results show that this adaptive-horizon formulation works effectively. Across the 30 calibrated overtaking pairs, introducing a varying prediction horizon substantially improves trajectory reproduction compared with the constant-horizon model. The average trajectory error decreased from 0.3691 in Step 1 to 0.16821 in Step 2, corresponding to an error reduction of approximately 54.43%. This indicates that allowing the prediction horizon to vary with interaction conditions provides a better representation of cyclists' observed overtaking behavior than keeping the horizon fixed throughout the maneuver. Rather than merely improving the fit, the calibrated horizon-varying parameters also revealed interpretable and systematic behavioral patterns.

A first important finding is about the internal relationships among the anticipation-related quantities themselves. The practically used minimum prediction horizon $N_{\text{use,min}}$ shows a perfect positive correlation with the prediction horizon at the overtaking point $N_{\text{use,overtaking}}$ ($r = 1.000$), indicating that in all samples, the cyclist's minimum horizon in practice is reached at the overtaking point. In addition, $N_{\text{use,min}}$ is very strongly positively correlated with the theoretical lower bound of prediction horizon N_{min} ($r = 0.955$). This shows that the lower bound is not merely a formal parameter in the model: in almost all cases, cyclists reduce their effective prediction horizon down to this lower-bound level during the interaction. Together, these results suggest that the overtaking point is the moment at which the cyclist applies the shortest planning horizon, and that the calibrated lower bound has practical behavioral meaning rather than being only a mathematical bound.

A second group of findings is about the relationship between the horizon-related parameters and trajectory features. At the overtaking point, higher overtaking speed $v_{A,\text{overtaking}}$ is associated with both a smaller theoretical lower bound N_{min} ($r = -0.409$) and a shorter practically used horizon $N_{\text{use,overtaking}}$ ($r = -0.396$). In addition, the theoretical upper bound N_{max} is negatively correlated with both the speed at the overtaking point $v_{A,\text{overtaking}}$ ($r = -0.400$) and the initial speed of the overtaking cyclist $v_{A,0}$ ($r = -0.480$). Taken together, these results suggest that when the overtaking cyclist rides faster, both the lower and upper bounds of the calibrated prediction horizon become smaller, and the cyclist also uses a shorter horizon at the overtaking moment. This points to a shift toward more reactive behavior under higher-speed conditions, presumably because less effective time is available for longer-term planning.

The correlations with the minimum inter-cyclist distance $d_{AB,\text{min}}$ provides an additional and more nuanced insight. Both N_{min} and $N_{\text{use,overtaking}}$ are negatively correlated with $d_{AB,\text{min}}$, with coefficients $r = -0.490$ and $r = -0.460$, respectively. Although this may initially seem counterintuitive, the interpretation is that a larger minimum spacing may reflect earlier adaptation and a stronger desire to maintain a safety margin throughout the maneuver. Under this interpretation, larger calibrated horizon-related values correspond to more anticipatory overtaking behavior. This means that the horizon may reflect not only the cyclist's response to the instantaneous distance, but also a higher-level behavioral tendency to regulate the interaction in a safer manner to keep a larger distance from the opponent cyclist.

Overall, this thesis shows that cyclists' anticipation during overtaking can be meaningfully represented as a dynamically adapted prediction horizon within an MPC framework. The proposed adaptive-horizon model not only improves agreement with measured trajectories, but also yields interpretable anticipation-related parameters that capture how cyclists adjust their planning depth in response to different interaction situations.

4.2 Limitations and Future Work

MPC-related limitations and future work:

- The surrounding environment is only represented in a limited way, while other factors such as road width may also influence how far cyclists look ahead [22]. In our current model, only virtual infrastructure is introduced to enforce the desired overtaking side, and in future work we should include more realistic and more complete modeling of nearby infrastructure.

- We assume the overtaking cyclist to predict the motion of the overtaken cyclist perfectly, since the measured trajectory of the overtaken cyclist over the prediction horizon is directly used for calculating the obstacle-avoidance cost for the overtaking cyclist. This assumption neglects the uncertainty in human prediction. Future work could address this limitation by formulating the interaction as a distributed MPC problem, where each cyclist plans based on the prediction of the other cyclist's future trajectories with an explicit uncertainty representation.

- In our model, we activate the obstacle-avoidance cost only when the opponent cyclist is ahead of the ego cyclist. This choice is consistent with a previous study modeling a cyclist's reactive range (i.e., the region within which cyclists respond to other road users [29]), which is primarily oriented toward the front as shown in Figure 1.3a. Accordingly, cyclists behind the ego cyclist are assumed to have a negligible influence on the ego cyclist. However, this simplification may not fully capture the continuous nature of interaction throughout the entire overtaking maneuver. Intuitively, a cyclist may still wish to maintain a certain distance from another cyclist immediately after overtaking. Thus, future work could investigate models in which the obstacle-avoidance cost remains active throughout the full maneuver, while allowing different distance-decay rates for interactions with opponent cyclists ahead of and behind the ego cyclist.

- The horizon-varying model assumes that the prediction horizon recovers immediately to $N_{\max}$ after the overtaking point. This assumption may be too abrupt, since cyclists may still remain aware of and influenced by the overtaken cyclist for some time after passing. Future work could therefore explore smoother recovery formulations in which the prediction horizon gradually returns to $N_{\max}$ after the overtaking point.

- The system dynamics is discretized using Forward Euler integration, which is fast but is less accurate. Given more time budget, we should try higher-order schemes such as the midpoint method or fourth-order Runge–Kutta (RK4) method.

- The obstacle-avoidance cost is not normalized in the same way as the other cost components. Most cost terms in the objective function have relatively simple forms and are normalized using the interquartile ranges (IQRs) of the corresponding measured quantities, ensuring that different cost components have comparable magnitudes. However, the proximity cost has a more complex form and is therefore not normalized. This may cause its contribution to the total MPC cost to depend more strongly on the related weight w_{avoid}, potentially affecting the calibration results. Future work could explore normalization approaches for the obstacle-avoidance cost.

BO-related limitations and future work:

- The dimension of the calibrated MPC parameter space remains limited. Due to the restricted time budget, we simplify the calibration problem by keeping several MPC parameters constant. A future direction would be to examine how these fixed parameters affect cyclist behavior. For instance, the cost function includes an exponential discount factor, $e^{-\eta t}$, which represents the tendency that cyclists place less weight on events predicted further into the future. In this thesis, we fix η at 0.05. A smaller value of η would weaken this discounting effect, thereby increasing the influence of future events and potentially making the varying prediction horizon more important in determining cyclist behavior.

- The current BO objective has limitation in that it is defined only in terms of the difference between the simulated MPC trajectory and the measured trajectory. Although trajectory matching is a meaningful performance measure, it does not account for other relevant aspects of cyclist behavior. For example, agreement in the full bicycle state (e.g., steer angle, steer rate) would provide a more complete assessment of whether the model reproduces realistic cyclist behavior. However, such an extension is not feasible with the TUMDOT dataset, since these variables were not measured. Therefore, a valuable direction for future work would be to consider datasets that include more complete cyclist state measurements, so that the BO objective can be expanded to include agreement in more cyclist states.

- Although we have tested several BO configurations (kernel smoothness combined with different random seeds for initial sampling), the hyperparameter space of BO is still explored to a limited extent. For example, more BO configurations could be examined to better balance exploration and exploitation, such as tuning the β parameter in UCB-based BO [63]. It would also be worthwhile to test other covariance functions, such as the RBF kernel, which is a commonly used alternative to the Matérn kernel in Gaussian process models [56]. In this way, we may find a better local optimum for MPC parameters, resulting in better agreement with the measurement.

- BO does not reach strict convergence under the classical 1% acquisition-based stopping criterion. Due to the limited available time, we use the relaxed 10% criterion instead, under which most BO configurations have converged. In future work, we should consider a larger optimization time budget to assess whether more BO trials can achieve stricter convergence and help us find better local optima for MPC parameters with better agreement with the measurement.

- The current grouping of samples into pair 1–18 and pair 19–30 is based mainly on certain observed differences between cyclist trajectories on the two road segments. A more systematic grouping strategy, for example, based on clustering or feature-based classification, may provide a stronger basis for calibration.

- The two-step calibration procedure could be further refined. In this study, we calibrate the background parameters assuming constant-horizon MPC in step 1, after which we keep the optimal background parameters fixed and calibrate horizon-varying parameters in step 2. However, this means that step 2 depends on background parameters obtained in step 1 under the simplified assumption of constant prediction horizon that is later relaxed to varying prediction horizon. An extension would therefore be to iterate between the two calibration steps. For example, we could try recalibrating the background parameters after obtaining the optimal horizon-varying parameters in step 2, and then optimize the horizon-varying parameters again using the newly calibrated background parameters. Such an iterative procedure may lead to a more realistic MPC parameter set.

BIBLIOGRAPHY

[1] S. Hoogendoorn, A. Gavriilidou, W. Daamen, and D. Duives, “Game theoretical framework for bicycle operations: A multi-strategy framework,” *Transportation Research Part C: Emerging Technologies*, vol. 128, p. 103 175, 2021.

[2] A. Kutsch, M. Margreiter, and K. Bogenberger, “Tumdot–muc: Data collection and processing of multimodal trajectories collected by aerial drones,” *Data Science for Transportation*, vol. 6, no. 2, p. 15, 2024.

[3] J. Pucher and R. Buehler, “Cycling towards a more sustainable transport future,” *Transport reviews*, vol. 37, no. 6, pp. 689–694, 2017.

[4] X. Ma and D. Luo, “Modeling cyclist acceleration process for bicycle traffic simulation using naturalistic data,” *Transportation research part F: traffic psychology and behaviour*, vol. 40, pp. 130–144, 2016.

[5] F. Kchour, S. Cafiso, and G. Pappalardo, “Understanding cyclists’ visual behavior using eye-tracking technology: A systematic review,” *Sensors*, vol. 25, no. 1, p. 22, 2024.

[6] S. Useche, L. Montoro, F. Alonso, and O. Oviedo-Trespalacios, “Infrastructural and human factors affecting safety outcomes of cyclists,” *Sustainability*, vol. 10, no. 2, p. 299, 2018.

[7] T. W. Gaither et al., “Estimated total costs from non-fatal and fatal bicycle crashes in the usa: 1997–2013,” *Injury prevention*, vol. 24, no. 2, pp. 135–141, 2018.

[8] A. C. Scholten, S. Polinder, M. J. Panneman, E. F. Van Beeck, and J. A. Haagsma, “Incidence and costs of bicycle-related traumatic brain injuries in the netherlands,” *Accident Analysis & Prevention*, vol. 81, pp. 51–60, 2015.

[9] A. C. Scholten, J. A. Haagsma, M. J. Panneman, E. F. Van Beeck, and S. Polinder, “Traumatic brain injury in the netherlands: Incidence, costs and disability-adjusted life years,” *PLoS one*, vol. 9, no. 10, e110905, 2014.

[10] City of Copenhagen, *Copenhagen: City of cyclists – bicycle account 2012*, Accessed: 2026-03-11, 2014. [Online]. Available: `https://www.fietsberaad.nl/CROWFietsberaad/media/Kennis/Bestanden/Fietsmonitor%20Kopenhagen%202012.pdf`.

[11] S. Daraei, K. Pelechrinis, and D. Quercia, “A data-driven approach for assessing biking safety in cities,” *EPJ Data Science*, vol. 10, no. 1, p. 11, 2021.

[12] M. Kilani and S. Bennaya, “Environmental impacts of bicycling in urban areas: A micro-simulation approach,” *Transportation Research Part D: Transport and Environment*, vol. 125, p. 103 967, 2023.

[13] F. Spotswood, T. Chatterton, A. Tapp, and D. Williams, “Analysing cycling as a social practice: An empirical grounding for behaviour change,” *Transportation research part F: traffic psychology and behaviour*, vol. 29, pp. 22–33, 2015.

[14] M. V. Butz, O. Sigaud, and P. Gérard, *Anticipatory behavior in adaptive learning systems: Foundations, theories, and systems*. Springer, 2004, vol. 2684.

[15] K. Huesmann and F. Loffing, “Perception-action coupling in anticipation research: A classification and its application to racket sports,” *Frontiers in psychology*, vol. 15, p. 1 396 873, 2024.

[16] R. Poli, “The many aspects of anticipation,” *Foresight*, vol. 12, no. 3, pp. 7–17, 2010.

[17] A. Gavriilidou, M. J. Wierbos, W. Daamen, Y. Yuan, V. L. Knoop, and S. P. Hoogendoorn, “Large-scale bicycle flow experiment: Setup and implementation,” *Transportation research record*, vol. 2673, no. 5, pp. 709–719, 2019.

[18] G. P. Castro, “Towards microscopic models for bicycle traffic simulation,” Licentiate thesis, Linköping University, Linköping, Sweden, 2023.

[19] H. Twaddle, T. Schendzielorz, and O. Fakler, “Bicycles in urban areas: Review of existing methods for modeling behavior,” *Transportation research record*, vol. 2434, no. 1, pp. 140–146, 2014.

[20] M. Nadin, “Anticipation–the underlying science of sport. report on research in progress,” *International Journal of General Systems*, vol. 44, no. 4, pp. 422–441, 2015.

[21] O. Lee, A. Rasch, A. L. Schwab, and M. Dozza, “Modelling cyclists’ comfort zones from obstacle avoidance manoeuvres,” *Accident Analysis & Prevention*, vol. 144, p. 105 609, 2020.

[22] P. Vansteenkiste, G. Cardon, E. D’Hondt, R. Philippaerts, and M. Lenoir, “The visual control of bicycle steering: The effects of speed and path width,” *Accident Analysis & Prevention*, vol. 51, pp. 222–227, 2013.

[23] X. Liang, M. Xie, and X. Jia, “New microscopic dynamic model for bicyclists’ riding strategies,” *Journal of Transportation Engineering, Part A: Systems*, vol. 144, no. 8, p. 04 018 034, 2018.

[24] M. F. Land and S. Furneaux, “The knowledge base of the oculomotor system,” *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 352, no. 1358, pp. 1231–1239, 1997.

[25] T. Miura, “Behavior oriented vision: Functional field of view and processing resources,” in *Eye movements from physiology to cognition*, Elsevier, 1987, pp. 563–572.

[26] S. I. Khan and W. Raksuntorn, “Characteristics of passing and meeting maneuvers on exclusive bicycle paths,” *Transportation research record*, vol. 1776, no. 1, pp. 220–228, 2001.

[27] Y. Yuan, W. Daamen, B. Goñi-Ros, and S. P. Hoogendoorn, “Investigating cyclist interaction behavior through a controlled laboratory experiment,” *Journal of transport and land use*, vol. 11, no. 1, pp. 833–847, 2018.

[28] M. Huber, Y.-H. Su, M. Krüger, K. Faschian, S. Glasauer, and J. Hermsdörfer, “Adjustments of speed and path when avoiding collisions with another pedestrian,” *PloS one*, vol. 9, no. 2, e89589, 2014.

[29] X. Liang, B. Mao, and Q. Xu, “Perceptual process for bicyclist microcosmic behavior,” *Procedia-social and behavioral sciences*, vol. 43, pp. 540–549, 2012.

[30] J. S. Brunner, Y.-C. Ni, A. Kouvelas, and M. A. Makridis, “Microscopic simulation of bicycle traffic flow incorporating cyclists’ heterogeneous dynamics and non-lane-based movement strategies,” *Simulation Modelling Practice and Theory*, vol. 135, p. 102 986, 2024.

[31] Y. H. Cho and J. Kim, “Stability analysis of the human controlled vehicle moving along a curved path,” *Vehicle system dynamics*, vol. 25, no. 1, pp. 51–69, 1996.

[32] T.-K. Dao and C.-K. Chen, “Path-tracking control of a riderless bicycle via road preview and speed adaptation,” *Asian Journal of Control*, vol. 15, no. 4, pp. 1036–1050, 2013.

[33] C. Dai, C. Zong, and G. Chen, “Path tracking control based on model predictive control with adaptive preview characteristics and speed-assisted constraint,” *IEEE Access*, vol. 8, pp. 184 697–184 709, 2020.

[34] J. Edelmann, M. Haudum, and M. Plöchl, “Bicycle rider control modelling for path tracking,” *Ifac-Papersonline*, vol. 48, no. 1, pp. 55–60, 2015.

[35] J. Li, Y. Ni, and J. Sun, “A two-layer integrated model for cyclist trajectory prediction considering multiple interactions with the environment,” *Transportation research part C: emerging technologies*, vol. 155, p. 104 304, 2023.

[36] J. Li, Y. Ni, and J. Sun, “Considering cyclists’ aggressiveness and bounded rationality: A self-learning motion prediction model for cyclists,” *IEEE Transactions on Intelligent Vehicles*, 2024.

[37] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, 2000.

[38] J. P. Meijaard, J. M. Papadopoulos, A. Ruina, and A. L. Schwab, “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*, vol. 463, no. 2084, pp. 1955–1982, 2007.

[39] C. M. Konrad, R. Happee, J. K. Moore, and A. Dabiri, *Stochastic control behavior of the balancing rider for cycling safety in traffic simulation*, EngrXiv preprint, 2025. DOI: `10.31224/6107`.

[40] A. L. Schwab and J. P. Meijaard, “A review on bicycle dynamics and rider control,” *Vehicle system dynamics*, vol. 51, no. 7, pp. 1059–1090, 2013.

[41] C. M. Konrad, *Cyclistsocialforce: Modified social forces for cyclists with realistic dynamics in sumo*, `https://github.com/chris-konrad/cyclistsocialforce`, GitHub repository, 2023.

[42] C. M. Schmidt, A. Dabiri, F. Schulte, R. Happee, and J. K. Moore, “Essential bicycle dynamics for microscopic traffic simulation: An example using the social force model,” in *Proceedings of the 5th Symposium on the Dynamics and Control of Single-track Vehicles: Bicycle and Motorcycle Dynamics 2023, October 18–20, Delft, The Netherlands*, TU Delft OPEN Publishing, 2024. DOI: `10.59490/65a5124da90ad4aecf0ab147`.

[43] BicycleParameters Authors, *Bicycleparameters documentation (release 1.2.0)*, `https://bicycleparameters.readthedocs.io/_/downloads/en/v1.2.0/pdf/`, Accessed 2026-03-09, 2025.

[44] BicycleParameters Authors, *Browserbenchmark.txt (browser bicycle benchmark parameter set)*, `https://github.com/moorepants/BicycleParameters`, File: data/bicycles/Browser/Parameters/BrowserBenchmark.txt. Accessed 2026-03-09, 2025.

[45] J. K. Moore, "Human control of a bicycle," Ph.D. dissertation, University of California, Davis, 2012.

[46] E. Hairer, G. Wanner, and S. P. Nørsett, *Solving ordinary differential equations I: Nonstiff problems*. Springer, 1993.

[47] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, vol. 106, no. 1, pp. 25–57, 2006.

[48] A. Shahzad and P. J. Goulart, "A new hot-start interior-point method for model predictive control," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 2470–2475, 2011.

[49] J. Li, Y. Ni, and J. Sun, "Considering cyclists' aggressiveness and bounded rationality: A self-learning motion prediction model for cyclists," *IEEE Transactions on Intelligent Vehicles*, 2024.

[50] A. Kutsch, L. Kessler, and K. Bogenberger, "Analyzing bicycle riding characteristics based on naturalistic urban drone observations: Free riding, following behavior, and overtaking maneuvers," *Transportation Research Record*, vol. 2679, no. 9, pp. 900–914, 2025.

[51] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint arXiv:1807.02811*, 2018.

[52] R. Garnett, *Bayesian optimization*. Cambridge University Press, 2023.

[53] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.

[54] C. Hvarfner, D. Eriksson, E. Bakshy, and M. Balandat, "Informed initialization for bayesian optimization and active learning," *arXiv preprint arXiv:2510.23681*, 2025.

[55] S. Ament, S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy, "Unexpected improvements to expected improvement for bayesian optimization," *Advances in neural information processing systems*, vol. 36, pp. 20 577–20 612, 2023.

[56] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2.

[57] Meta Platforms, Inc., *Ax: Adaptive experimentation platform*, `https://github.com/facebook/Ax`, Accessed: 2026-03-05, 2024.

[58] M. Balandat et al., "Botorch: A framework for efficient monte-carlo bayesian optimization," *Advances in neural information processing systems*, vol. 33, pp. 21 524–21 538, 2020.

[59] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient global optimization of expensive black-box functions,” *Journal of Global optimization*, vol. 13, no. 4, pp. 455–492, 1998.

[60] P. Schober, C. Boer, and L. A. Schwarte, “Correlation coefficients: Appropriate use and interpretation,” *Anesthesia & analgesia*, vol. 126, no. 5, pp. 1763–1768, 2018.

[61] J. Cohen, *Statistical power analysis for the behavioral sciences*. routledge, 2013.

[62] R. L. Wasserstein and N. A. Lazar, “The asa statement on p-values: Context, process, and purpose,” *The American Statistician*, vol. 70, no. 2, pp. 129–133, 2016.

[63] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, “Gaussian process optimization in the bandit setting: No-regret and experimental design,” *arXiv*, vol. 0912.3995, 2009.

5

AI STATEMENT

Generative AI (ChatGPT) was used to polish language and help interpret programming errors during debugging. All final text was written by me, and all technical decisions, analyses, and conclusions were developed and verified by me. No AI-generated text was directly inserted into the thesis.