Using Model Predictive Control on a Steer-by-Wire Bicycle for Performance Assistance

Simonas Draukšas

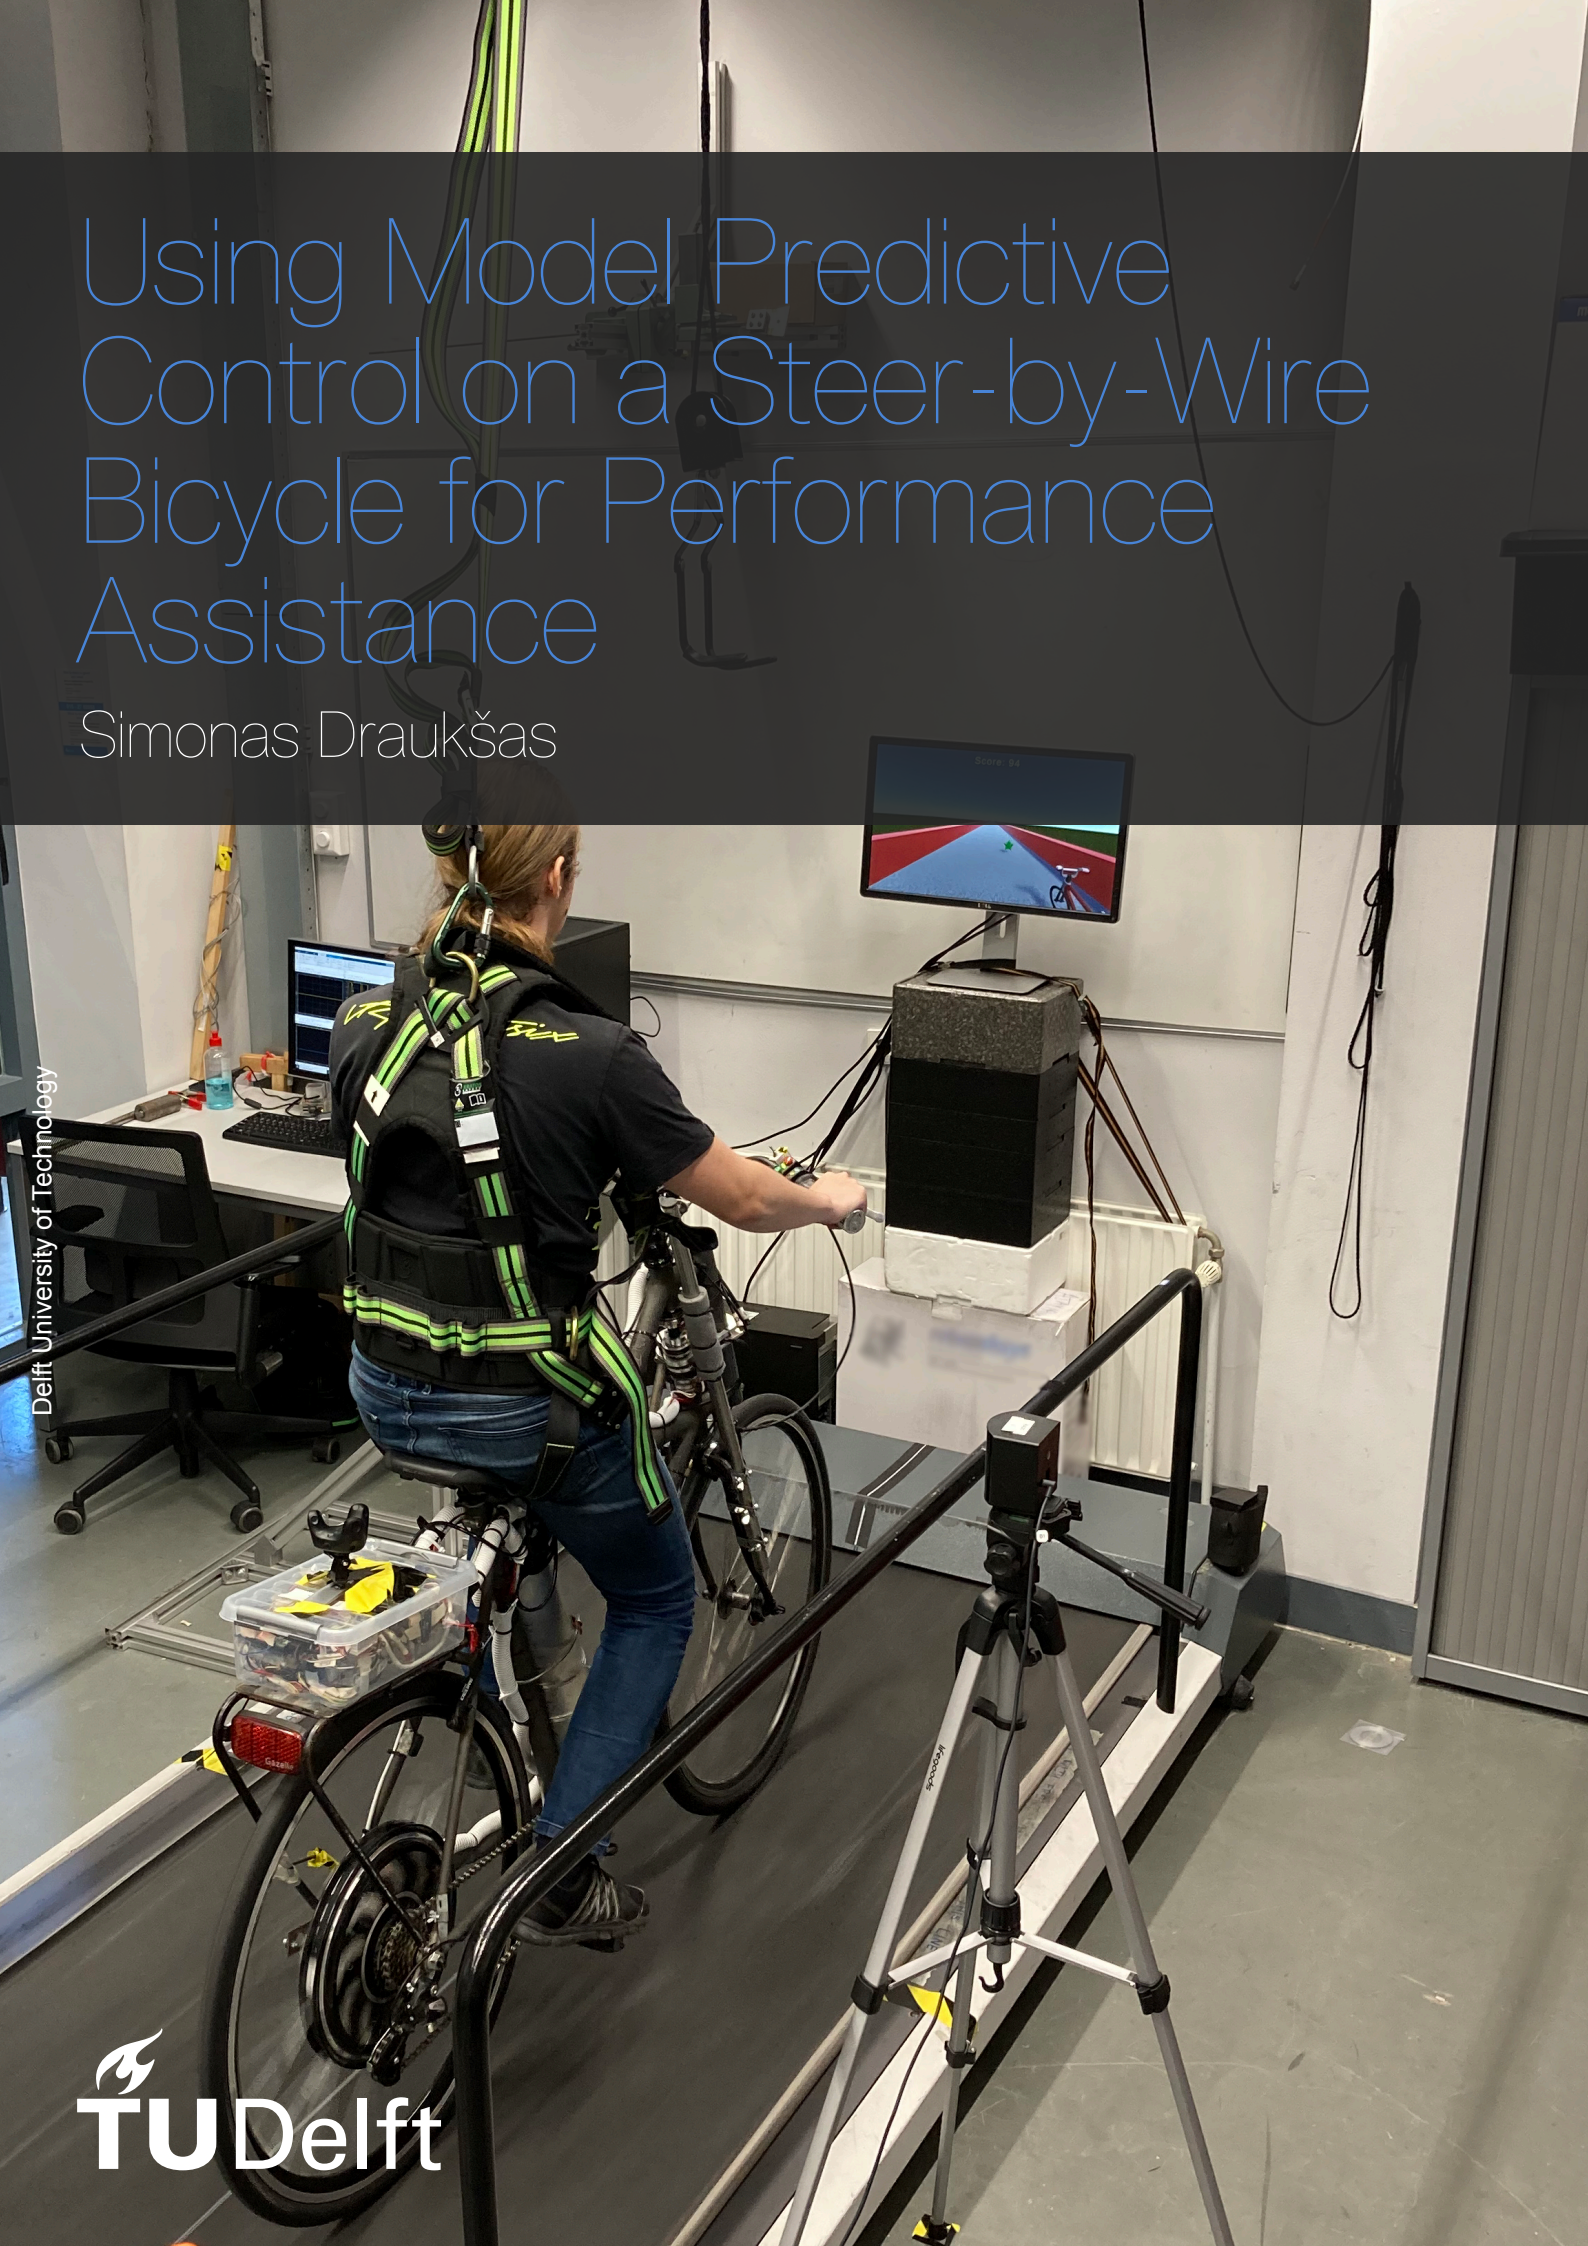

Delft University of Technology

TU Delft

Using Model Predictive Control on a Steer-by-Wire Bicycle for Performance Assistance

by

Simonas Draukšas

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Monday September 26th, 2022 at 14:30.

Student number:	5384559	
Project duration:	March 1, 2022 – September 26, 2022	
Thesis committee:	Dr.ing. L. Marchal Crespo,	TU Delft, chair, supervisor
	Dr. J. K. Moore,	TU Delft, supervisor
	Prof.Dr.ir. R. Happee,	TU Delft, supervisor
	Dr. L. Alizadehsaravi,	TU Delft, daily supervisor
	Dr. B. Shyrokau,	TU Delft

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

Using Model Predictive Control Steer-by-Wire Bicycle for Performance Assistance

Simonas Draukšas
TU Delft, MSc Candidate

September 26th, 2022

Abstract

Bicycle safety is quickly becoming an increasingly important field as the number of electric bicycles on the streets grows faster each year. E-bikes are able to accelerate quicker and travel at faster speeds than conventional bicycles, increasing the severity of injuries in case of an accident. There are a number of ways to improve safety, such as building better infrastructure, implementing active safety systems, improving bicycling skills, or better protective gear. In this thesis, a controller based on Model Predictive Control has been designed to explore whether it could be used as a haptic guidance system that would improve motor learning of a cycling task, and whether it could assist the cyclist during cycling manoeuvres. The cycling task of lane change manoeuvres performed at a constant forward velocity was investigated. The controller was implemented on a desktop PC, which wirelessly controlled the TU Delft's Steer-by-Wire bicycle – a bicycle where the steering is enabled by the use of electric motors instead of mechanical coupling between the front fork and the handlebars. To test the controller, ten participants took part in a pilot study. The study was designed following a counterbalanced measures design and the participants were split into two groups that experienced the controller's haptic guidance in different order. During the study, the participants were asked to ride the bicycle on a treadmill and hit virtual targets, shown on a display mounted in front of the treadmill, by carrying out lane change manoeuvres. The hypothesis stated that the controller improves performance while it is assisting the participants. It was found that the controller did not significantly improve immediate performance and this result is likely caused by too low of the task difficulty. However, it is likely that the controller was more effective at improving motor learning of lower skilled participants compared to higher skilled participants, but a small number of lower skilled participants limited the analysis. A short post-hoc no-hands riding test of the same cycling task was carried out to investigate whether the controller is able to carry out lane change manoeuvres with minimal rider input. A significant performance improvement was found during the no-hands test. In conclusion, due to limitations of the study, no performance or motor learning improvement caused by the controller was found. Yet the controller showed promising results in a no-hands riding test, which suggests that the controller could be used as a starting point for advanced safety systems for bicycles.

1 Introduction

Bicycles are one of the simplest and most accessible modes of transport. They are cheap, do not require much maintenance and, most often, do not require a license to ride on the street. In countries like the Netherlands, bicycle travel even accounts for a quarter of all trips [12]. For the past decade, the number of electric bikes sold each year has been quickly increasing [34]. E-bikes are attractive as they can travel at higher speeds than conventional bicycles for the same rider effort. However, higher speeds also lead to injuries of higher severity in case of an accident [26, 43]. As such, bicycling safety is of utmost importance. There are multiple ways to improve bicycling safety: infrastructure, bicycle design, bicycling skill, and protective gear.

Bicycle designs have barely changed since the so-called safety bicycle – a bicycle with two equally-sized wheels – was introduced in late 19th century. While there have been some developments between the 19th century and modern bicycles, those developments mostly came from trial and error, rather than rigorous analysis of the physics of bicycles. In recent years, there has been a resurgence in interest of bicycle dynamics. In 2007, J. P. Meijaard et al. [36] thoroughly reviewed previous literature and chose the Whipple-Carvallo bicycle model as a canonical model that describes the linearised bicycle dynamics. With the help of nonlinear computer simulations [36] and experimental data [20], they were able to validate that this linearized model is able to describe bicycle behaviour well. Additionally, they provided a set of benchmark bicycle parameters that allows different models and software implementations to be compared and validated.

Having this canonical model made it easier to conduct further research into bicycles. Kooijman et al. [22] used the model to show that gyroscopic or caster effects are not necessary for bicycle stability; Moore and Hubbard [38] used it to design an optimally handling bicycle; TU Delft, together with Koninklijke Gazelle, [6] used it to design a bicycle that provides steering assistance; and others ([21, 40, 47] and many more) used it to look into human control of a bicycle. The canonical Whipple-Carvallo model is also the model used in this thesis.

In 2012, Appelman [1] designed and built an early prototype of a steer-by-wire bicycle. On a steer-by-wire bicy-

cle, the handlebars are mechanically decoupled from the fork and the front wheel of the bicycle. Instead, two motors and two angular position sensors are installed on the front assembly of the bicycle. One motor, in tandem with one position sensor, controls the steer angle of the fork and the front wheel, while the other motor, combined with the second position sensor, controls the steer angle of the handlebars. These two motors enable the steer-by-wire system designer to change how the bicycle handles. For example: the system could multiply the rider's steering input without a need for a gearbox; the mentioned multiplication could be changed to a nonlinear function; the system could remove the need for the rider to countersteer during cornering [35]; the system could stabilise the bicycle by itself; and a creative mind could come up with many more possibilities. To continue research on steer-by-wire bicycles, a new prototype was built in 2018 by Dialynas [7], which has been used in this thesis.

While a lot of research has been put into bicycle design and human control of the bicycle, not much attention is paid to how humans learn to ride a bicycle. A significant fraction of bicycle riders learnt how to ride a bicycle using training wheels. However, it can be argued that this method is sub-optimal as it masks the real bicycle dynamics, since the bicycle abruptly becomes a tricycle the moment a training wheel touches the ground. Some riders might also develop bad habits and rely on the training wheel to help regain balance instead of learning the correct method of "steer into the fall", which becomes an issue when trying to finally get rid of the training wheels.

Klein et al. [17] came up with a better system than training wheels. Their solution is to replace the bicycle wheels with rollers of varying radii. A new rider starts by riding a bicycle equipped with a roller of infinite radius (a cylinder). The radius is slowly reduced as the riders learn, until the radius of a conventional tyre is reached. Rollers with high radius make the bicycle more stable and easier to balance. The rider, therefore, feels more comfortable during training, during which they build up experience and muscle memory needed to ride a bicycle. This method has been trialed with children with autism spectrum disorder and Down syndrome in a couple of different experiments [13, 17, 31, 48] and was found to be successful at teaching children how to ride a bicycle.

While not much has been explored regarding human motor learning specifically in bicycles, research has been carried out in other motor tasks. Recent advancements in robotics meant that researchers could explore the role of robots in facilitating motor learning. Experiments were carried out on motor tasks of differing difficulty, with and without robotic assistance on healthy subjects. The simplest tasks involved reaching [2, 29, 30] or shape tracing [25, 28, 44, 49]. Other tasks, like driving [14, 24, 32] or tennis swings [16, 33], could be considered more complex.

Literature review [8] carried out before the thesis project found that results in the literature vary from robotic assistance being detrimental to learning, to not having a significant effect, to being effective at improving learning. In general, it was found that robotic training is more effective when the difficulty of the task is matched to the skill of the participant [3]. For example, the robot could assist a participant to make the task easier and more manageable. Looking back on Klein et al. [17], the same mechanism is seen - the complex task of cycling is made easier for novices by replacing the wheels with rollers.

Two of the human motor learning papers were instrumental in this thesis - Özen et al. [42] and Özen, Buetler and Marchal-Crespo [41]. These papers deal with a pendulum-swinging task. In the first paper, Model Predictive Control (MPC) was introduced to motor learning research. The reasoning behind MPC is to allow the participant to experience more variability and increase the sense of agency by using a control method that can adapt and recalculate optimal trajectories on-the-fly, compared to classical control methods, which constrain the participant to a pre-determined trajectory. There, it was found that training with MPC did improve participants' performance during training, measured by absolute error between the participants' location and the targets, when compared to performance without MPC. In the second paper, the study was expanded to include a transfer task and compared MPC to a conventional Proportional-Derivative (PD) controller. It was found that only the participants that trained with an MPC controller acting on the end-effector of the pendulum were able to significantly increase their deviation from the natural frequency of the pendulum, which shows a better control of pendulum's dynamics.

MPC has also been successfully used in other fields. Since its creation in 1970s, it has been an attractive control methodology in petrochemical industry, where it was able to control non-linear systems and outperform classical control in terms of product quality metrics [46]. The ability of MPC to work within specified constraints, as well as the aforementioned ability to control non-linear systems, attracted attention from robotics and automotive fields. For example, MPC was used by Yoshida et al. [50] to model a lane change manoeuvre of an automobile, and Yu et al. [51] have implemented a non-linear MPC model for a path-following task on a simulated vehicle and found that MPC outperformed Stanley and Linear Quadratic Regulator controllers.

This thesis is a successor of the two papers regarding MPC applied to a pendulum swinging task [41, 42]. In this thesis, a complex task of cycling is looked into. An MPC controller is designed and implemented on a steer-by-wire bicycle that applies haptic guidance to assist the rider during lane change manoeuvres. To measure the effectiveness of the controller, a pilot study was conducted. The pilot study involved 10 participants riding a steer-by-wire bicycle on a treadmill and hitting virtual targets shown on a display in front of the treadmill. The performance of the rider was measured by a score, which was calculated according to rider's deviation from the target's centre. The participants were split into two equally-sized groups with a counter-balanced repeated measures design of two training methods: Controller On; Controller Off. The main hypothesis is that the controller improves performance while it is acting on the bicycle.

The thesis is structured as follows. Section 2 describes

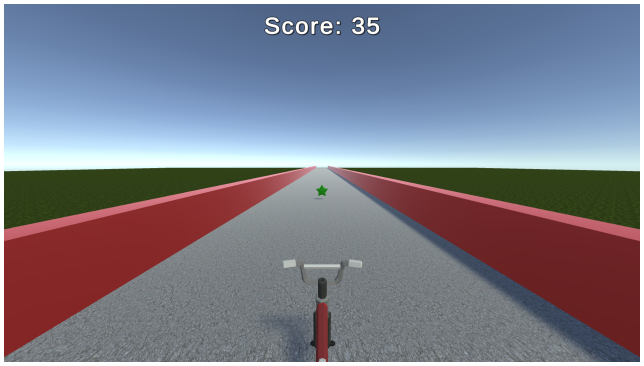

Figure 1: The game shown to the participants. The participants control the lateral position of the virtual bicycle shown in the middle and need to collect green stars that are approaching them. After passing each star, the participants are shown a score at the top of the display that corresponds to their distance to that star's centre. The red walls correspond to the edges of the treadmill.

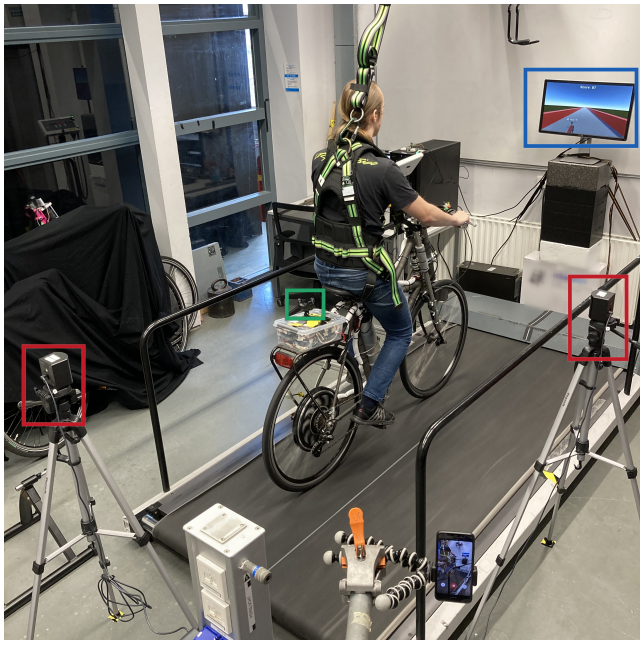

Figure 2: Rear view of the setup, a participant can be seen riding the bicycle. The location of the SteamVR Base Stations 2.0 are shown in red. The location of the HTC Vive Tracker 3.0 is shown in green. The display showing the game is shown in blue. The width of the usable space of the treadmill is 1.1 meters.

the experimental setup, the task and the controller. In Section 3, statistical analysis is carried out on the data collected during the experiment. The results are discussed in Section 4. The thesis is concluded in Section 5. Some suggestions for future work are given in Section 6.

2 Materials and Methods

2.1 Experimental Setup

The experiment involves riding a steer-by-wire bicycle [7] on a treadmill. On this bicycle, the handlebars and the front fork are mechanically separated. A steer-by-wire connection is created using two electric steer motors and two absolute angular encoders. A PD controller is implemented to minimise the steer angle difference between the handlebars and the front fork, mimicking a rigid coupling. The control loop of the PD controller is 1 kHz and the controller runs on a Teensy 4.1 (PJRC, US) microcontroller. Dialynas et al. [7] found that this setup provides good tracking for steering inputs up to 2.5 Hz.

Controllers based on MPC require a lot of computational power due to the need for online optimisation. Thus, it was decided to use a significantly more powerful desktop PC for the controller's implementation, instead of the onboard Teensy microcontroller. The MPC controller is implemented on a Windows 10 (Microsoft, US) desktop computer equipped with Intel i7-7700K 4.2 GHz processor (Intel, US), running Simulink Desktop Real-Time (MathWorks, US). The desktop computer and the bicycle communicate wirelessly using Bluetooth at 200 Hz for the bicycle-to-computer communication, and 75 Hz for the computer-to-bicycle communication.

During the experiment, the bicycle acts as a Human Interface Device for the game, in which the virtual targets are hit. By moving the bicycle on the treadmill laterally, the participant controls the lateral position of the virtual bicycle shown in the game, which is displayed on a 24-inch computer monitor, placed around 2 metres in front of the participant. The game is implemented using Unity (Unity Technologies, US) and can be seen in Figure 1. The lateral position, yaw angle and roll angle of the real bicycle is obtained using an HTC Vive Tracker 3.0 (HTC, Taiwan) installed above the rear wheel. Two SteamVR Base Stations 2.0 (HTC, Taiwan) are used to enable tracking - one installed directly behind the treadmill, and the second one installed on the side of the treadmill. The rear view of the setup showing the treadmill and the locations of the Base Stations can be seen in Figure 2.

The data from the HTC Vive Tracker is sent to the supplied USB dongle, which is connected to a Raspberry Pi 4 Model B 4GB (Raspberry Pi Foundation, UK). This computer runs a 32-bit version of Raspberry Pi OS Lite in headless mode. `Libsurvive`'s [27] Simple API is used to read the raw data coming from the Tracker, calculate the position and the pose of the tracker, and send the data over an Ethernet cable using UDP to the Windows 10 computer at the rate of 220 Hz.

To reduce the risk of injury, a safety harness is used, which is fixed to a point on the ceiling above the treadmill. Due to the harness, the participant does not need to pedal, reducing fatigue and allowing the participant to focus solely on the steering task.

2.2 Target-hitting Task

To analyse the performance of the MPC controller as an assistance system, a target-hitting task is designed. The participant is asked to collect stars that are approaching at a constant velocity of 15 km/h. The interval be-

Table 1: Survey results. The cycling frequency is measured by number of days cycled in a typical 28 day month.

Participant	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Age bracket	25-29	30-34	25-29	60-64	30-34	35-39	25-29	25-29	25-29	25-29
Cycling frequency	28	2	28	24	14	0	25	25	28	16

tween the stars is 6 seconds. The stars' lateral positions can be between -0.2 and $+0.2$ m from the centre of the treadmill. Seven equally spaced-out points are defined spanning this whole region (possible star locations: -0.2, -0.133, -0.067, 0.0, $+0.067$, $+0.133$, $+0.2$ m). The sequence of 10 stars is then randomised from these 7 points before the experiment. Three out of 10 stars were located on the right side of the treadmill, five were on the left, and the remaining two were in the middle. All participants experienced the same sequence.

The stars are hit by placing the virtual bicycle in front of the star and passing through it. To give feedback to the user, a score is shown on the display. The score is calculated using Equation 1, where y_P is the lateral position of the bicycle's rear wheel contact point in metres and y_S is the lateral position of the star's centre in metres. If the participant is within ± 0.02 m of the centre of the star as they pass it, they are given a score of 100. And if the participant is not within the distance of ± 0.02 m, but within the distance of ± 0.22 m, a linear equation is used to calculate the score. Otherwise, the score of 0 is given.

$$\Delta y = |y_P - y_S|$$
$$\text{score} = \begin{cases} 100, & \text{if } \Delta y < 0.02. \\ 500 \times (0.22 - \Delta y), & \text{if } 0.02 \leq \Delta y \leq 0.22. \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

2.3 MPC Controller

MPC is an advanced control method that uses a mathematical model of the controlled system to predict the system's behaviour throughout a specified control horizon. A control signal is chosen such that the predicted system state follows a given reference state. The controller is able to work within the specified constraints put on the system.

What makes MPC different from classical control is that the controller is implicit in the sense that the designer of the controller does not explicitly specify the control law. Rather, the designer specifies a cost function (and its weights) which is then used by the controller to determine the control law at each time step. The control law is obtained through online optimisation, which has a downside of being computationally costly.

The linear MPC problem, used in this application, is stated in Equation 2, where t is the current time, N is the number of steps in the receding horizon, x is the bicycle state, r is the reference state, u is the control input, and Q and R are designer-defined weighing matrices. Subscripts *lb* and *ub* stand for *lower bound* and *upper bound*, respectively, and are used to set the previously mentioned constraints. Matrices A and B are linear time-invariant state-space matrices.

$$\begin{aligned} J = \sum_{k=t}^{t+N} (x_k - r_k)^T Q_k (x_k - r_k) + \sum_{k=t}^{t+N-1} u_k^T R_k u_k \\ \text{subject to } x_{k+1} = Ax_k + Bu_k \\ x_{lb,k} \leq x_k \leq x_{ub,k} \\ u_{lb,k} \leq u_k \leq u_{ub,k} \end{aligned} \quad (2)$$

In this experiment, N is set to 150, which is equal to the time horizon of 2 seconds with 75 Hz sample rate. The control input u is steering torque. The bicycle and reference states, x and r, consist of the lateral position of the rear wheel of the bicycle y_P, the yaw angle ψ, the roll angle ϕ, the steering angle δ, the roll rate $\dot{\phi}$ and the steering rate $\dot{\delta}$. The state-space matrices A and B are obtained using `HumanControl` software [15], which can convert the equations of motion of a linear Whipple-Carvallo bicycle model to state-space representation. Bicycle parameters of the *Davis Instrumented Bicycle* (specified under *Rigid* in pages 91-92 of [40]) are used due to its physical similarity to the TU Delft's Steer-by-Wire bicycle. A forward speed of 15 km/h (4.17 m/s), equal to the treadmill's speed, is chosen. The torque calculated by the MPC controller is applied to both the handlebar and front fork motors of the steer-by-wire bicycle, meaning that the controller not only applies haptic guidance to the handlebars, but also steers the bicycle.

The bounds on the input u for the steering torque are set to ± 3 Nm to allow the participant to easily overpower the motors if needed and to avoid over-current of the motors. The state x has bounds implemented on lateral position, roll and steering angles. The position is limited to ± 0.5 m due to the limited width of the treadmill. The steering angle is limited to ± 40 degrees to avoid hitting the physical stops installed on the bicycle, and the roll angle is limited to ± 20 degrees to avoid unnecessarily big roll angles.

The weighing matrix R_k is always kept constant at the value of 1 to minimise the control inputs, while Q_k is time-dependent. The weights in Q_k are always kept at zero, except for the 4 second period before a star reaches the bicycle. During that 4 second period, the weights are linearly increased to the specified values: 10 for y_P and ψ; 3 for ϕ; and 0 for the remaining states. The weight on y_P instructs the controller to carry out lane change manoeuvres, while the weight on ψ encourages the controller to keep the bicycle parallel to the treadmill's centreline as much as possible. Lastly, ϕ's weight ensures that the bicycle is kept upright. The values were obtained through trial and error.

Since the MPC cost function is quadratic and con-

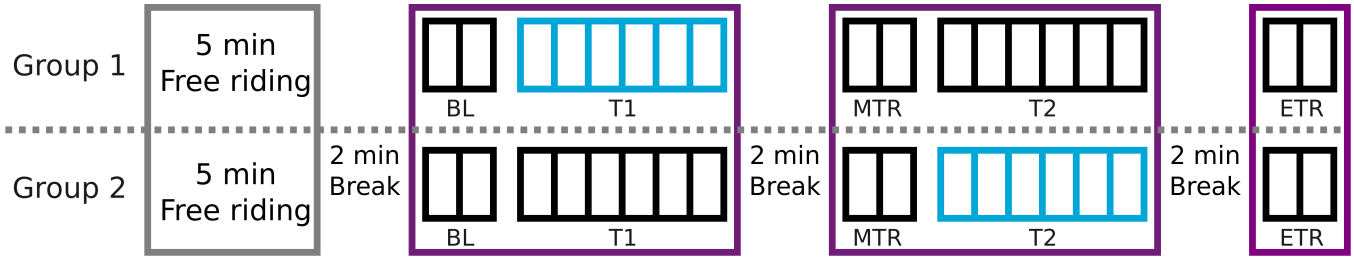

Figure 3: Study protocol. Participants were assigned to one of two groups. Blue rectangles represent blocks where the controller is assisting the participant, black represents the blocks where the controller is not assisting the participants. Each rectangle is 1 minute long and contains 10 stars. *BL* stands for Baseline, *MTR* stands for Mid-training retention, *ETR* stands for End-of-training retention, and *T1* and *T2* stand for Training 1 and Training 2, respectively.

vex, a Quadratic Programming (QP) solver is used to find the optimal control input. The solver of choice here is `qpOASES` [11] and, more accurately, its Simulink interface for solving QP problems with varying matrices `qpOASES_SQProblem`. The QP problem is solved at every time step (at 75 Hz).

The state of the bicycle is measured using absolute encoders on the bicycle and the HTC Vive Tracker. The measurements are then passed through a real-time second-order low-pass Butterworth filter with a cutoff frequency of 8 Hz (adapted from [39]). The Butterworth filter also calculates the angular rates $\dot{\phi}$ and $\dot{\delta}$ from their respective angles.

The details on how the MPC problem is translated into a QP problem and a visualisation of a typical reference path can be found in Appendix B.

2.4 Study Protocol

Ten healthy adult participants (7 male, 3 female) took part in the study. The participants were informed of the risks and consented to the study. Before the study, the participants were asked for their age bracket and how many days they cycle during a typical month (a typical month was considered to be 28 days long, out of which 8 are weekends). The answers can be seen in Table 1.

The study is split into five trials: Baseline (*BL*), Training 1 (*T1*), Mid-Training Retention (*MTR*), Training 2 (*T2*), and End-of-Training Retention (*ETR*). Baseline, Mid-Training Retention and End-of-Training Retention are considered to be "no-intervention" trials, as the controller was always turned off during these trials. The MPC controller is only turned on during Training 1 or Training 2. All trials are further separated into blocks, where each block is 1 minute long and contains a sequence of 10 stars. The participants were randomly assigned to either Group 1 or Group 2. The participants were informed that the MPC controller will assist them during the experiment, but were not informed when the assistance will be provided.

Before the study, the participants were given 5 minutes to familiarise with riding a bicycle on a treadmill. During these 5 minutes, the MPC controller was turned off, the stars did not appear on the display, and the participants were encouraged to carry out lane change manoeuvres of varying amplitudes.

The experiment starts with 2 minute long no-intervention Baseline trial. Immediately after Baseline, the 6 minute long Training 1 trial begins. During Training 1, Group 1 has the controller assisting them, while Group 2 trains without the controller. After Training 1, a very short (less than 2 minutes) break is needed to prepare the upcoming trials. The participants are asked whether they would like to extend this break. While all participants wanted to continue immediately, the break was extended (to approximately 5 minutes) for two participants due to technical issues.

After the break, another 2 minute long no-intervention Mid-Training Retention trial takes place, immediately followed by a six minute long Training 2 trial. During Training 2, Group 2 has the controller assisting them, while Group 1 trains without the controller. As before, a very short break is required after Training 2 and the participants are given a chance to extend the break. All participants wanted to continue without extending the break.

The study is completed with the last 2 minute long no-intervention End-of-Training Retention trial. The visualisation of the study protocol can be seen in Figure 3.

2.5 Riding Without Hands

After the study was completed, an additional short test was conducted to check whether the controller is able to complete lane change manoeuvres by itself. The test followed the structure of one of the 8 minute long sections of the study – 2 minutes of Baseline with the controller turned off, immediately followed by 6 minutes of riding with the controller turned on. However, this time the controller had the full control authority over the bicycle. Throughout this test, the author was riding the bicycle and used as little as possible active steering by hovering his hands above the handlebars. Additionally, he was ignoring the game shown on the screen and was watching the handlebars instead. The video of this test can be seen in [9].

2.6 Data Processing and Statistical Analysis

The sensor data from the experiment is recorded in a Matlab `.mat` file. After the experiment, the scores are calculated for each star from the sensor data using Equation 1.

Participants receive a score between 0 and 100 for each star. Therefore, the performance of the participant throughout each trial is measured by the **mean score**. A within-subject **score variance** of the trial is calculated as well. A low score variance can be indicative of high precision and participant's skill to repeatedly place the bicycle in the correct location. The calculated values are exported as `.csv` files and imported into `R` [45] for statistical analysis.

Normality of the data is checked by visually inspecting Q-Q plots and using Shapiro-Wilk tests. A Welch two sample t-test is used to compare the **mean scores** and **score variances** of the two groups during Baseline to see if there were any significant differences in performance between two groups.

The controller's ability to improve performance while the controller is assisting the participant is investigated by calculating relative **difference of the mean score** and **difference of the score variance** between the Training trial and the preceding no-intervention trial. To calculate the differences during Training 1, equation $PM_{T1} - PM_{BL}$ is used, where PM stands for "performance metric" and can be either the **difference of the mean score** or **difference of the score variance**. Similarly, Training 2 uses the equation $PM_{T2} - PM_{MTR}$. The obtained values are compared between groups (*T1* Group 1 – *T1* Group 2; *T2* Group 1 – *T2* Group 2) using two Welch two sample t-tests.

The α-level is set to .05 and any result with a p value lower than this is considered significant.

3 Results

All 10 participants completed all five trials without falling. However, one participant fell during the 5 minutes of familiarisation, but was not injured and chose to continue the study. No data was excluded from the analysis.

The **mean score** and **score variance** evolution by block for all participants can be seen in Appendix A, as well as the scores of the groups as a whole. No significant differences were found ($F(1,8) = 0.216, p = .655$ and $F(1,8) = 0.003, p = .958$) between the groups during *BL*, therefore is it assumed that the differences between groups in other trials are due to the controller's effect.

3.1 Improvement in Performance with MPC Assistance

To test whether the controller had an effect on immediate performance, the **difference of the mean score** and **difference of the score variance** between *BL* and *T1* or *MTR* and *T2* were compared using two Welch two sample t-tests. The graph showing the variables can be seen in Figure 4.

During *T1*, there was no difference between the groups in mean score change ($t(6.92) = 0.042, p = .968$) and score variance change ($t(7.96) = -0.190, p = .854$). The same can be said about *T2* – no significant difference in mean scores ($t(7.35) = -1.981, p = .086$) and variances ($t(8.00) = 2.069, p = .072$). However, both p-values in *T2* show that there might be a trend, therefore, the previously mentioned figures need to be inspected.

Looking at the score mean differences (Figure 4a), an interesting phenomenon can be seen. During both *T1* and *T2*, Group 2 had a negative median (performance reduction). However, Group 1 managed to have a positive median (performance increase) during *T1* (when they had the controller assisting them), but the change became negative (performance decreased) during *T2* (when the controller was not assisting them).

The story is similar in the amount of difference of score variance (Figure 4b). Since this graph regards the difference in variance, negative values are considered as performance improvement, while positive values show the reduction of performance. In this case, Group 1 exhibited a slightly negative median during *T1*, which then turned into a positive median during *T2*, mirroring the mean score change. Group 2, on the other hand, were able to keep their variance at a relatively same level (illustrated by the near-zero median) during *T2*.

3.2 Riding Without Hands

During the Baseline of the no-hands trial, since there was no active steering towards the stars, the author stayed in the middle of the treadmill. The **mean score** of this Baseline trial was 51.1, with a **score variance** of 31.6. During the following 6 minute long Controller trial, during which MPC was controlling the bicycle, the **mean score** was 87.0 and the **score variance** was 17.9. The bicycle's lateral position without and with the controller can be seen in Figure 5. Additional time domain graphs can be found in Appendix D.

To check whether this difference is statistically significant, a Wilcoxon Paired Signed-Ranks test is carried out. A Wilcoxon test was used in place of a t-test to err on the side of caution, as Shapiro-Wilk normality test returned a non-significant, but close to significance, result ($W = 0.905, p = .052$). Since paired tests require equal sample sizes, the 6 minute long trial was split into three 2 minute long segments - Controller 1, Controller 2, Controller 3. Each of the three segments is then compared to the Baseline. The boxplot showing the distribution of the scores during the Baseline (shown under *NH_NC* scenario) and Controller (shown under *NH_C* scenario) can be seen in Figure 6. The Wilcoxon test shows significance in the differences between Baseline and Controller 1 ($Z = -2.94, p = .003$), Baseline - Controller 2 ($Z = -2.01, p = .045$), and Baseline - Controller 3 ($Z = -2.99, p = .003$). No significant differences found between the different Controller parts. The results are corrected with Bonferroni correction for multiple comparisons.

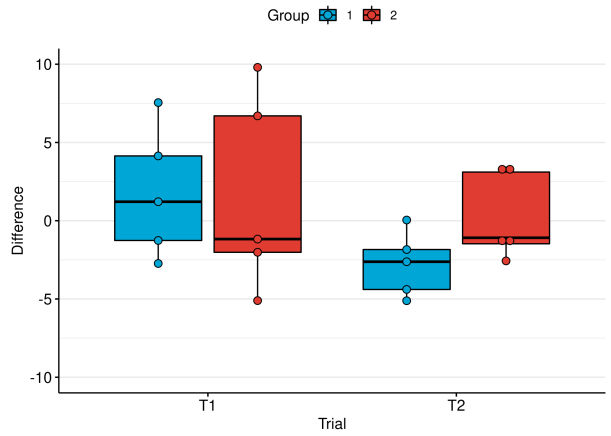

(a) Score. Relative to the previous no-intervention trial (between *BL* and *T1*, or between *MTR* and *T2*), positive difference indicates performance improvement.

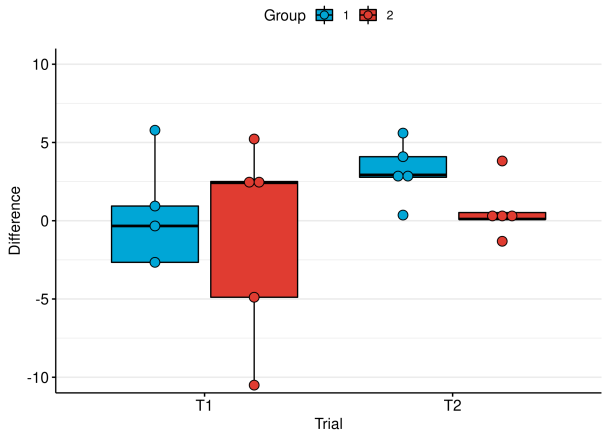

(b) Variance. Relative to the previous no-intervention trial (between *BL* and *T1*, or between *MTR* and *T2*), negative difference indicates performance improvement.

Figure 4: Boxplot of the **difference of the mean score** and the **difference of the score variance** for both groups during *T1* and *T2*.

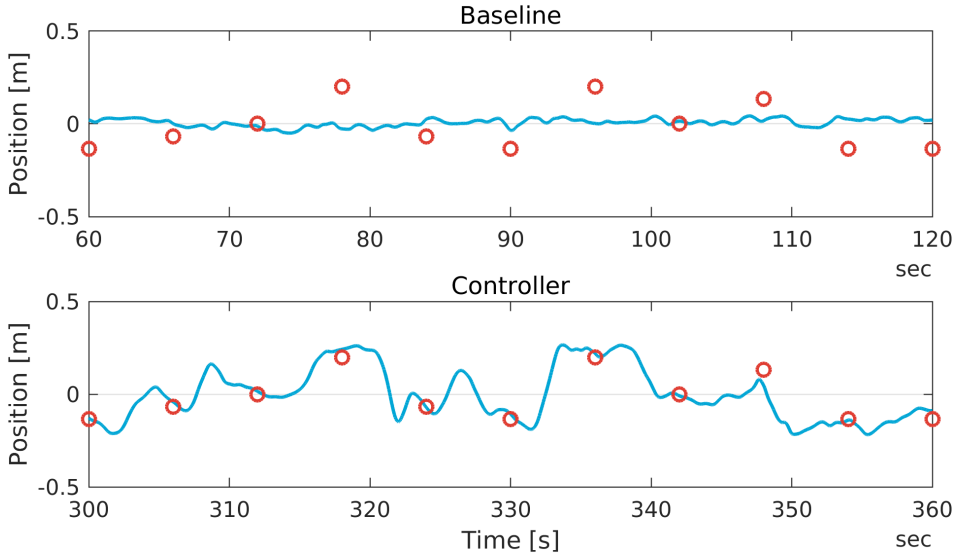

Figure 5: Bicycle's lateral position during the no-hand test. The top subfigure shows the lateral position of the bicycle throughout the last minute of the *Baseline* part of the test. The bottom subfigure shows the lateral position of the bicycle throughout the last minute of the *Controller* part of the test. The red circles are the star locations.

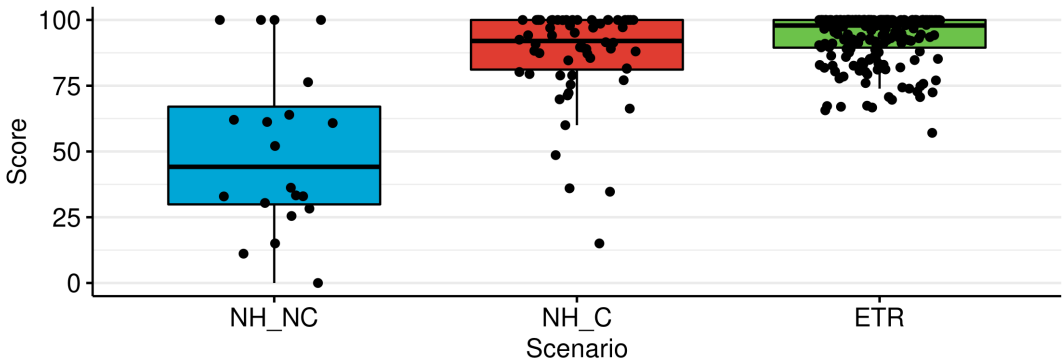

Figure 6: Score distribution during the no-hands test and during the main study's last trial. Scenario *NH_NC* represents the *Baseline* part of the no-hands test, when the controller was not active (20 scores). Scenario *NH_C* represents the *Controller* part of the no-hands test, when the controller was steering the bicycle (60 scores). Scenario *ETR* contains all 200 scores (10 participants, 20 stars each) of the main study's End-of-Training Retention trial.

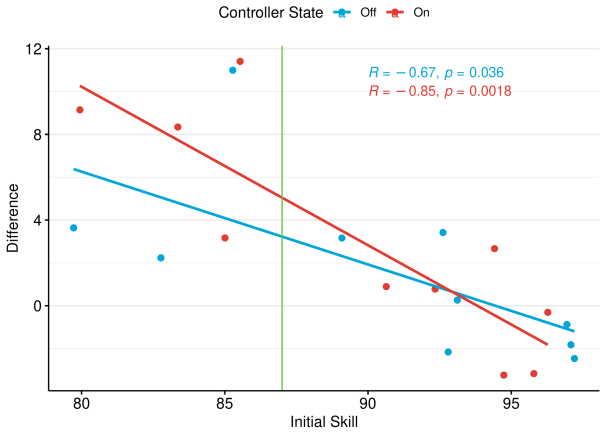

(a) Scatter-plot showing the initial score versus the score difference.

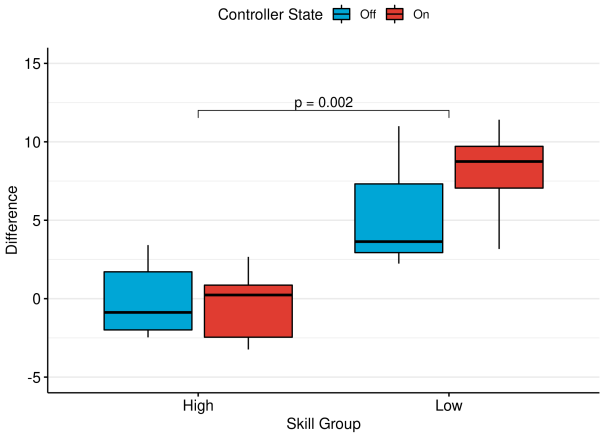

(b) Boxplot showing the score difference versus the skill level grouped by k-means.

Figure 7: Skill dependency. The score difference is calculated by taking the score of the no-intervention trial following the Training trial and subtracting the score of the no-intervention trial preceding the Training trial.

3.3 Skill Dependency

Literature regarding human motor learning has shown that the efficacy of robotic assistance depends on functional task difficulty, which depends on participant's skill [3]. Therefore, a post-hoc analysis is carried out to explore whether there is an interaction between initial skill level of the participant and the controller. While this analysis was not planned before the experiment took place, the results found here could be useful for future studies.

The score of the participant during the preceding no-intervention trials (*BL* or *MTR*) is assumed to be a good estimate of that participant's skill at that moment in time. To estimate the skill increase due to the controller during the intervention trials (*T1* or *T2*), the **difference in mean score** between the preceding and following no-intervention trials is used (*MTR* - *BL* or *ETR* - *MTR*). A scatter-plot showing the initial score versus the score difference can be seen in Figure 7a. Two regression lines fitted to the scenarios of training with the controller (*Controller On*) and training without the controller (*Controller Off*) can be seen. K-means clustering is applied to the data, and the score threshold represented by the vertical green line is found. The results to the left of the green line are considered as "low skilled", while the results to the right are considered as "high skilled". A boxplot using these groupings is shown in 7b. A Welch two sample t-test shows a significant difference between the groups in score differences ($t(8.39) = -4.50, p = .002$).

To further see if there is an effect of the controller, a Linear Mixed Effects (LME) model with the Equation 3 is used, where $Difference$ is the score difference as shown on the y-axis of the scatter plot, $Skill$ is the initial score (x-axis of the scatter plot), $Controller$ is either On or Off (colour of the data points in the scatter plot), and ID is the participant ID. The LME model is then used in ANOVA using `lmerTest` [23].

$$Difference \sim Skill * Controller + (1|ID) \quad (3)$$

The ANOVA results can be seen in Table 2. The results show that the amount that the participant improved by strongly depends on the initial score. Additionally, a likely effect of the controller can be seen, as there is the one-side significant main effect of *Controller* and the one-side significant interaction effect of the skill and the controller.

4 Discussion

The results have shown that the controller is able to carry out lane change manoeuvres by itself, which can be seen from the no-hands test. However, in general, it did not seem like the controller was able to improve performance of the participants. Yet, when taking into account the skill of the participants, it is possible that the controller was more effective for lower skilled participants, while even being detrimental for higher skilled participants. The main results will be discussed here. For additional observations, such as participant feedback, refer to Appendix C.

4.1 Performance Improvement

4.1.1 Task Difficulty

It is thought that the difficulty of the task was not properly matched to the skill level of the subjects. While designing the study protocol, due to author's own experience testing the experimental setup, the assumption was made that the participants will need more than 5 minutes to fully familiarise with riding on a treadmill. Thus, it was expected that the scores during Baseline trials would be somewhere around 70, which would provide the participants with a lot more room for improvement and the effect of the controller might have been more pronounced. However, it was found that the mean score of the participants during Baseline was 88.6. A higher-than-expected score meant that the amount of improvement throughout the experiment was limited. Additionally, during the

Table 2: ANOVA analysis of the LMM model shown in 3

	DF Numerator	DF Denominator	F	p
Skill	1	19.6544	47.7045	.000005***
Controller	1	9.7537	5.3920	.07837*
Skill * Controller	1	9.7669	5.1671	.08428*

Significance codes: * - 0.1, *** - 0.0

no-hand test, the controller achieved a mean score of 87, which means that the participants were able to match the controller's performance from the very beginning and possibly limited the amount of assistance that the controller was able to supply.

As such, the task should be made more difficult in the follow up studies. The difficulty can be increased either by reducing the time between the stars (here it was set to 6 seconds), by increasing the amplitude of the lane changes (here, the maximum possible amplitude was 0.4 metres), or by combining both of these changes.

4.1.2 Low Torques

During the no-hands trial, the controller rarely ever exceeded the absolute torque of 0.5 Nm, with the mean absolute torque over the trial being equal to 0.16 Nm. During the study itself, the mean absolute torque over all participants was 0.22 Nm. The participants' hands were placed roughly 0.3 m away from the axis of rotation, which means that the mean absolute torque of 0.22 Nm translates to an absolute force of only 0.73 N.

After the experiment, majority of the participants reported that they did not feel the controller acting on the handlebars. A follow-up question revealed that those participants have had their arm muscles tensed up during the experiment due to trying to stay as straight with the bicycle as they can. Only two participants reported feeling some external forces. One participant said she was relaxed and therefore could feel the low amplitude torques. The other participant could only feel the controller during big amplitude lane changes, which required higher torques to complete.

However, it is possible that some participants did not notice the controller as they were in agreement with controller's applied torques. In order to prove this, the rider's torques would need to be measured, which was not done during this study. For a future study, the measurement of the rider's torques or rider's intentions should be considered as they can also be a measure of how much the rider trusts the controller.

To address the low torques, the torques applied to the handlebars should be increased in magnitude to make them more obvious for the participant. Yet, as shown by the no-hands trial, high torques are not needed to carry out lane change manoeuvres. The solution could be to fully exploit the steer-by-wire system of the bicycle and apply torques of different amplitudes to the front fork and the handlebars. On the other hand, care must be taken not to increase the torques felt on the handlebars too much, as too high torques can lead to the loss of sense of agency for the participant, which can be detrimental to learning [10].

Another solution would be to instruct the participants to relax. However, this approach was not taken due to the concern that the participants would then rely too much on the controller and would not learn how to control the bicycle themselves.

4.2 Riding Without Hands

4.2.1 Controller's performance

The no-handed test showed that the MPC controller is able to control the bicycle and hit the stars. In fact, throughout the whole trial, the controller managed to reach a median score of 92 (and the median of 97 during the last two minutes of the test), while all participants of the main study combined were able to achieve a median score of 97 during *ETR* (Figure 6). By reordering Equation 1, a median absolute distance to the star's centre can be obtained. In this case, a score of 92 is equal to 3.6 cm error from the centre of the star, while the score of 97 is equal to a distance of 2.6 cm. This means, that the controller was able to get within 1 cm of human performance if the whole trial is taken into account.

4.2.2 Trust in the Controller

Since the test was carried out with the experienced author riding the bicycle, it is possible that, during this test, the rider trusted the controller more than the ten participants who experienced the controller for the first time. However, even with the additional experience, at the start of the test the author did not completely trust the controller and interfered during a few manoeuvres by taking control, as he felt that the controller was about to lose stability. By the end of the test, the author was able to fully trust the controller, which can also be seen from higher scores during the last 2 minutes of the test.

4.3 Skill-dependency

Post-hoc analysis of the scores led to Figure 7a, in which two regression lines of the two controller states can be seen crossing. At lower initial scores, the *Controller On* line suggests that training with the controller's support led to higher improvements. Interestingly, at the initial score of around 92, the *Controller On* line crosses the *Controller Off* line, which can suggest that for participants that initially scored high, controller's support was detrimental

and training without the controller is better in that situation. This seems to match the observations seen in motor learning literature, where robotic assistance is more effective for lower skilled subjects, while higher skilled subjects improve more under robotic disturbance [3]. The scatter plot also shows that lower skilled participants were able to improve more than higher skilled participants, which is expected, as the room for improvement reduces as the score increases, due to the upper limit of 100.

The further analysis of the data using the LME model showed that there might be a trend that the initial skill and the controller state are interacting. However, the analysis is limited due to a small number of low skilled participants and that a lot of data points are clustered around the initial score of 92. Therefore, a future study should aim to have a wider range of participant skills, as well as a higher number of participants, in order to explore this phenomena further.

5 Conclusions

During this thesis a controller, based on the Model Predictive Control framework, capable of controlling a bicycle was designed and a pilot study investigating controller's effectiveness at improving performance was conducted.

A test, during which the author rode a bicycle without actively steering it, showed that the controller is able to carry out lane change manoeuvres if the rider trusts the controller and does not interfere with it. Therefore, the controller could be used as a starting point for future projects regarding self-riding bicycles or advanced bicycle safety systems.

However, no significant effect of the controller on performance improvement during the main study was found. Two main reasons are attributed to this observation – too low task difficulty, and controller's torques being too low for the participants to notice.

Data analysis after the study showed that the initial baseline scores were very high, which suggests an easy task. As previous human motor learning literature has shown, haptic guidance is not effective when the relative task difficulty is low. Therefore, the controller was not able to significantly improve participants' performance, as the participants were already quite proficient at the task. In fact, when taking into account the initial skill level of the participant, a trend appears that suggests that the controller could have had an effect on lower skilled participants, for whom the task was more difficult, however, a low number of participants prevents a more concrete conclusion. For future experiments, the functional difficulty of the task should be increased either by making the task more difficult, or by recruiting beginner cyclists as participants.

The no-hands trial showed that the torque needed for the lane change manoeuvres was low, as the mean absolute torque was only 0.16 Nm. Additionally, the maximum absolute torque applied by the controller rarely exceeded 0.5 Nm. During the main study, the mean absolute torque increased only slightly to 0.22 Nm and, therefore, was hard to notice for the participants, especially when they tensed up their arm muscles while concentrating on the task. For future experiments, the Steer-by-Wire mechanism of the bicycle could be exploited to increase the haptic guidance torques felt by the participants, while applying low torques to the front fork only.

6 Future Work

Future work could take a number of different paths: human motor learning research; smart bicycles and assistance systems; assessment of bicycle handling. This section will look at them and provide some suggestions.

6.1 Human Motor Learning

As the idea for this thesis came about from motor learning research, it is natural to continue on this path further. During this project, a task and an experimental setup have been designed. The pilot study has shown that motor learning has occurred throughout the experiment, but no significant effect of the controller has been found. Therefore, some changes are proposed for future studies:

- **Increase the sample size.** There were only 10 participants in this pilot study, which were then split into two groups of 5. Due to a small sample size, statistical analyses did not have a lot of statistical power. Additionally, a wider spread of skill levels should be sought out. However, a bigger sample size for this thesis was not achievable due to time constraints.
- **Increase task difficulty.** In this study, the participants were skilled bicyclists with several years of experience and the task proved to be easier than anticipated. Task difficulty could be increased by reducing the time between stars, or by recruiting bicyclists without a lot of experience.
- **Increase guiding torques.** Participants noted that they could not feel controller's guiding torques. Increasing the torques applied to the handlebars might address this issue.

6.2 Smart Bicycles and Assistance Systems

The controller designed in this thesis was able to control the bicycle by itself sufficiently to get a comparable score to experienced participants. Thus, the author believes that the controller could be used on bicycles or motorcycles as a part of Advanced Driver Assistance Systems. Possible applications could be Emergency Obstacle Avoidance, Lane Change Assist, or Lane Keeping Assist. However, before that, some steps should be taken:

- **Integrate the controller into the bicycle itself.** The current iteration of the controller had to be ran on a desktop PC, which limits the application to

only a laboratory setting. In order to integrate the controller into the bicycle, more powerful embedded hardware, such as Jetson (NVIDIA, US), needs to be installed on the bicycle. The controller itself should be parallelised to take advantage of Jetson's GPU cores. Paper by Kogel and Findeisen [18] could be used as a starting point.

- **Estimate bicycle's state with on-board sensors.** In this thesis, only the steering angle was measured using the bicycle's sensors. The yaw and roll angles, with their corresponding angular rates, were obtained from an external HTC Vive Tracker, constraining the controller's use to laboratory settings.
- **Estimate bicycle's location with on-board sensors.** For features like obstacle avoidance or lane keeping assist, the bicycle first needs to know its location relative to the obstacle or lane boundaries. The need of localisation and computer vision could be fulfilled by using Intel RealSense Technology (Intel, US) or similar.

6.3 Bicycle Handling Assessment

Currently, there are no standardised handling quality metrics or their tests for bicycles. This thesis has designed a lane change task, which produces a quantitative result in terms of a score. Different bicycles could be used to complete the task and the score could be used as a metric that describes how manoeuvrable that bicycle is. Additionally, the test setup is easy to put together, as only an HTC Vive Tracker is needed for sensing, which can be quickly transferred between different bicycles. Some improvements are recommended:

- **Star sequence.** Currently, a random star sequence with constant distance between stars is used. Randomisation is great when dealing with human subjects, but might not be suitable if a standardised metric needs to be measured repeatedly. The distance between stars could also vary with time so that the test could capture bicycle's handling at different levels of manoeuvre aggressiveness.
- **In-game perspective.** In order to get rid of as many human factors as possible from the test, the rider should have as much information about the current situation as possible. This means, that the complaints from the pilot study participants regarding the in-game perspective in the visualisation need to be addressed, for example, by providing lines that lead towards the centre of the star.
- **Calibration.** The test needs to be repeatable, therefore a repeatable and reliable calibration of the HTC Vive Tracker method should be thought of.

Software developed for this thesis is available at `https://github.com/mechmotum/TUDelft-SbW-Bicycle`

References

[1] N. Appelman. "Dynamics and Control of a Steer-by-Wire Bicycle". In: (2012). URL: https://repository.tudelft.nl/islandora/object/uuid%3A373f1f52-f149-4a47-b744-3050a2608f0d (visited on 07/27/2022).

[2] Edwin H. F. van Asseldonk et al. "Influence of haptic guidance in learning a novel visuomotor task". en. In: *Journal of Physiology-Paris*. Neurorobotics 103.3 (May 2009), pp. 276–285. ISSN: 0928-4257. DOI: 10.1016/j.jphysparis.2009.08.010. URL: https://www.sciencedirect.com/science/article/pii/S0928425709000515 (visited on 10/06/2021).

[3] Ekin Basalp, Peter Wolf, and Laura Marchal-Crespo. "Haptic training: Which types facilitate (re)learning of which motor task and for whom Answers by a review". In: *IEEE Transactions on Haptics* (2021). Conference Name: IEEE Transactions on Haptics, pp. 1–1. ISSN: 2329-4051. DOI: 10.1109/TOH.2021.3104518.

[4] Miguel Borges et al. "HTC vive: Analysis and accuracy improvement". In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018, pp. 2610–2615.

[5] Sylvain Daronnat et al. "Inferring Trust from users' behaviours; agents' predictability positively affects trust, task performance and cognitive load in human-agent real-time collaboration". In: *Frontiers in Robotics and AI* 8 (2021). DOI: 10.3389/frobt.2021.642201.

[6] TU Delft. *Smart motor in handlebars prevents bicycles from falling over*. 2019. URL: https://www.tudelft.nl/en/2019/tu-delft/smart-motor-in-handlebars-prevents-bicycles-from-falling-over (visited on 07/27/2022).

[7] George Dialynas, Riender Happee, and AL Schwab. "Design and implementation of a steer-by-wire bicycle". In: *International Cycling Safety Conference*. 2018. URL: https://www.researchgate.net/publication/328808185.

[8] Simonas Draukšas. "Augmented haptic feedback system to promote motor learning of a bicycle task: A literature review". In: (Sept. 2022). DOI: 10.6084/m9.figshare.21069973.v1. URL: https://figshare.com/articles/online_resource/Augmented_haptic_feedback_system_to_promote_motor_learning_of_a_bicycle_task_A_literature_review/21069973.

[9] Simonas Draukšas. "Using MPC on a SbW Bicycle for Performance Improvement, No-Hands Test Video". In: (Sept. 2022). DOI: 10.6084/m9.figshare.21029197.v1. URL: https://figshare.com/articles/media/Using_MPC_on_a_SbW_Bicycle_for_Performance_Improvement_No-Hands_Test_Video/21029197.

[10] Satoshi Endo et al. "Effect of external force on agency in physical human-machine interaction". In: *Frontiers in Human Neuroscience* 14 (2020). DOI: 10.3389/fnhum.2020.00114.

[11] H.J. Ferreau et al. "qpOASES: A parametric active-set algorithm for quadratic programming". In: *Mathematical Programming Computation* 6.4 (2014), pp. 327–363.

[12] Lucas Harms and Maarten Kansen. *Cycling Facts 2018*. 2018. URL: https://www.government.nl/topics/bicycles/documents/reports/2018/04/01/scycling-facts-2018 (visited on 07/27/2022).

[13] Zoë Hawks et al. "Accelerating motor skill acquisition for bicycle riding in children with ASD: A pilot study". In: *Journal of autism and developmental disorders* 50.1 (2020), pp. 342–348.

[14] Masakazu Hirokawa et al. "Effect of Haptic Assistance on Learning Vehicle Reverse Parking Skills". In: *IEEE Transactions on Haptics* 7.3 (July 2014), pp. 334–344. ISSN: 2329-4051. DOI: 10.1109/TOH.2014.2309135.

[15] Jason K. Moore. *HumanControl*. Version 37107d228e502ff940efa3dcbea1c8430e6af310. June 22, 2019. URL: https://github.com/moorepants/HumanControl.

[16] Julius Klein, Steven J. Spencer, and David J. Reinkensmeyer. "Breaking It Down Is Better: Haptic Decomposition of Complex Movements Aids in Robot-Assisted Motor Learning". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20.3 (May 2012), pp. 268–275. ISSN: 1558-0210. DOI: 10.1109/TNSRE.2012.2195202.

[17] Richard E. Klein et al. "Adapted Bicycles for Teaching Riding Skills". In: *TEACHING Exceptional Children* 37.6 (July 2005), pp. 50–56. ISSN: 0040-0599. DOI: 10.1177/004005990503700606. URL: https://doi.org/10.1177/004005990503700606 (visited on 11/16/2021).

[18] Markus Kögel and Rolf Findeisen. "Parallel solution of model predictive control using the alternating direction multiplier method". In: *IFAC Proceedings Volumes* 45.17 (2012). 4th IFAC Conference on Nonlinear Model Predictive Control, pp. 369–374. ISSN: 1474-6670. DOI: https://doi.org/10.3182/20120823-5-NL-3013.00081. URL: https://www.sciencedirect.com/science/article/pii/S1474667016314719.

[19] *Experimental Validation of the Lateral Dynamics of a Bicycle on a Treadmill*. Vol. Volume 4: 7th International Conference on Multibody Systems, Nonlinear Dynamics, and Control, Parts A, B and C. International Design Engineering Technical Conferences and Computers and Information in Engineering Conference. Aug. 2009, pp. 2029–2034. DOI: 10.1115/DETC2009-86965. eprint: https://asmedigitalcollection.asme.org/IDETC-CIE/proceedings-pdf/IDETC-CIE2009/49019/2029/4598337/2029_1.pdf. URL: https://doi.org/10.1115/DETC2009-86965.

[20] J. D. G. Kooijman, A. L. Schwab, and J. P. Meijaard. "Experimental validation of a model of an uncontrolled bicycle". In: *Multibody System Dynamics* 19.1 (Feb. 1, 2008), pp. 115–132. ISSN: 1573-272X. DOI: 10.1007/s11044-007-9050-x. URL: https://doi.org/10.1007/s11044-007-9050-x (visited on 02/23/2021).

[21] J. D. G. Kooijman, A. L. Schwab, and Jason K. Moore. "Some Observations on Human Control of a Bicycle". In: *Proceedings of the ASME 2009 International Design and Engineering Technical Conferences & Computers and Information in Engineering Conference*. ASME 2009 International Design and Engineering Technical Conferences & Computers and Information in Engineering Conference. ASME, Aug. 30, 2009. DOI: 10.1115/DETC2009-86959. URL: http://proceedings.asmedigitalcollection.asme.org/proceeding.aspx?articleid=1649894.

[22] J. D. G. Kooijman et al. "A Bicycle Can Be Self-Stable Without Gyroscopic or Caster Effects". In: *Science* 332.6027 (Apr. 15, 2011), pp. 339–342. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1201959. URL: https://science.sciencemag.org/content/332/6027/339 (visited on 10/03/2019).

[23] Alexandra Kuznetsova, Per B. Brockhoff, and Rune H. B. Christensen. "lmerTest Package: Tests in Linear Mixed Effects Models". In: *Journal of Statistical Software* 82.13 (2017), pp. 1–26. DOI: 10.18637/jss.v082.i13.

[24] Hojin Lee and Seungmoon Choi. "Combining haptic guidance and haptic disturbance: an initial study of hybrid haptic assistance for virtual steering task". In: *2014 IEEE Haptics Symposium (HAPTICS)*. ISSN: 2324-7355. Feb. 2014, pp. 159–165. DOI: 10.1109/HAPTICS.2014.6775449.

[25] Jaebong Lee and Seungmoon Choi. "Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance". In: *2010 IEEE Haptics Symposium*. ISSN: 2324-7355. Mar. 2010, pp. 335–342. DOI: 10.1109/HAPTIC.2010.5444635.

[26] T. L. Lefarth et al. "[Pedelec users get more severely injured compared to conventional cyclists]". In: *Der Unfallchirurg* 124.12 (Dec. 2021), pp. 1000–1006. ISSN: 1433-044X. DOI: 10.1007/s00113-021-00976-x.

[27] *libsurvive*. Version 1.0. Feb. 6, 2022. URL: https://github.com/cntools/libsurvive.

[28] Rakshith Lokesh and Rajiv Ranganathan. "Haptic Assistance That Restricts the Use of Redundant Solutions is Detrimental to Motor Learning". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 28.6 (June 2020), pp. 1373–1380. ISSN: 1558-0210. DOI: 10.1109/TNSRE.2020.2990129.

[29] Dylan P Losey et al. "Improving short-term retention after robotic training by leveraging fixed-gain controllers". In: *Journal of Rehabilitation and Assistive Technologies Engineering* 6 (Jan. 2019), p. 2055668319866311. ISSN: 2055-6683. DOI: 10.1177/2055668319866311. URL: https://doi.org/10.1177/2055668319866311 (visited on 09/29/2021).

[30] Dylan P. Losey, Laura H. Blumenschein, and Marcia K. O'Malley. "Improving the retention of motor skills after reward-based reinforcement by incorporating haptic guidance and error augmentation". In: *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. ISSN: 2155-1782. June 2016, pp. 857–863. DOI: 10.1109/BIOROB.2016.7523735.

[31] Megan MacDonald et al. "Bicycle training for youth with Down syndrome and autism spectrum disorders". In: *Focus on Autism and Other Developmental Disabilities* 27.1 (2012), pp. 12–21.

[32] Laura Marchal-Crespo and David J. Reinkensmeyer. "Haptic Guidance Can Enhance Motor Learning of a Steering Task". en. In: *Journal of Motor Behavior* 40.6 (Nov. 2008), pp. 545–557. ISSN: 0022-2895, 1940-1027. DOI: 10.3200/JMBR.40.6.545-557. URL: http://www.tandfonline.com/doi/abs/10.3200/JMBR.40.6.545-557 (visited on 09/07/2021).

[33] Laura Marchal-Crespo et al. "The effect of haptic guidance and visual feedback on learning a complex tennis task". en. In: *Experimental Brain Research* 231.3 (Nov. 2013), pp. 277–291. ISSN: 0014-4819, 1432-1106. DOI: 10.1007/s00221-013-3690-2. URL: http://link.springer.com/10.1007/s00221-013-3690-2 (visited on 09/24/2021).

[34] Manuel Marsilio. "A New Era for Cycling in the Post COVID-19 Outbreak". In: *WFSGI Magazine* (2021). URL: https://wfsgi.org/wfsgi-magazine/.

[35] Y Marumo and M Nagai. "Steering control of motorcycles using steer-by-wire system". In: *Vehicle System Dynamics* 45.5 (2007), pp. 445–458.

[36] J.P Meijaard et al. "Linearized dynamics equations for the balance and steer of a bicycle: a benchmark and review". en. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 463.2084 (Aug. 2007), pp. 1955–1982. ISSN: 1364-5021, 1471-2946. DOI: 10.1098/rspa.2007.1857. URL: https://royalsocietypublishing.org/doi/10.1098/rspa.2007.1857 (visited on 09/06/2021).

[37] JP Meijaard and AL Schwab. "Linearized equations for an extended bicycle model". In: *III European Conference on Computational Mechanics*. Springer. 2006, pp. 772–772.

[38] Jason K Moore and Mont Hubbard. "Expanded Optimization for Discovering Optimal Lateral Handling Bicycles". In: Bicycle and Motorcycle Dynamics: Symposium for Dynamics and Control of Single Track Vehicles. Padova, Italy, 2019, p. 12.

[39] Jason K. Moore. *row_filter*. Version 1.0.0. Aug. 2019. DOI: 10.5281/zenodo.3378965. URL: https://doi.org/10.5281/zenodo.3378965.

[40] Jason Keith Moore. *Human control of a bicycle*. University of California, Davis, 2012.

[41] Özhan Özen, Karin A. Buetler, and Laura Marchal-Crespo. "Promoting Motor Variability During Robotic Assistance Enhances Motor Learning of Dynamic Tasks". In: *Frontiers in Neuroscience* 14 (Feb. 2, 2021), p. 600059. ISSN: 1662-4548. DOI: 10.3389/fnins.2020.600059. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7884323/ (visited on 09/20/2021).

[42] Özhan Özen et al. "Multi-purpose Robotic Training Strategies for Neurorehabilitation with Model Predictive Controllers". In: *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*. ISSN: 1945-7901. June 2019, pp. 754–759. DOI: 10.1109/ICORR.2019.8779396.

[43] H. P. a. M. Poos et al. "[E-bikers are more often seriously injured in bicycle accidents: results from the Groningen bicycle accident database]". In: *Nederlands Tijdschrift Voor Geneeskunde* 161 (2017), p. D1520. ISSN: 1876-8784.

[44] Dane Powell and Marcia K. O'Malley. "The Task-Dependent Efficacy of Shared-Control Haptic Guidance Paradigms". In: *IEEE Transactions on Haptics* 5.3 (2012), pp. 208–219. ISSN: 2329-4051. DOI: 10.1109/TOH.2012.40.

[45] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria, 2022. URL: https://www.R-project.org/.

[46] Max Schwenzer et al. "Review on model predictive control: an engineering perspective". In: *The International Journal of Advanced Manufacturing Technology* 117.5 (Nov. 1, 2021), pp. 1327–1349. ISSN: 1433-3015. DOI: 10.1007/s00170-021-07682-3. URL: https://doi.org/10.1007/s00170-021-07682-3.

[47] Robin S. Sharp. "On the Stability and Control of the Bicycle". In: *Applied Mechanics Reviews* 61.6 (Oct. 2008). 060803. ISSN: 0003-6900. DOI: 10.1115/1.2983014. eprint: https://asmedigitalcollection.asme.org/appliedmechanicsreviews/article-pdf/61/6/060803/5442322/060803_1.pdf. URL: https://doi.org/10.1115/1.2983014.

[48] Dale A Ulrich et al. "Physical activity benefits of learning to ride a two-wheel bicycle for children with Down syndrome: A randomized trial". In: *Physical therapy* 91.10 (2011), pp. 1463–1477.

[49] Xing-Dong Yang, Walter F. Bischof, and Pierre Boulanger. “Validating the Performance of Haptic Motor Skill Training”. In: *2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. ISSN: 2324-7355. Mar. 2008, pp. 129–135. DOI: 10.1109/HAPTICS.2008.4479929.

[50] Hidehisa Yoshida, Shuntaro Shinohara, and Masao Nagai. “Lane change steering manoeuvre using model predictive control theory”. In: *Vehicle System Dynamics* 46.sup1 (2008), pp. 669–681. DOI: 10.1080/00423110802033072. eprint: https://doi.org/10.1080/00423110802033072. URL: https://doi.org/10.1080/00423110802033072.

[51] Chuanyang Yu et al. “MPC-based Path Following Design for Automated Vehicles with Rear Wheel Steering”. In: *2021 IEEE International Conference on Mechatronics (ICM)*. 2021, pp. 1–6. DOI: 10.1109/ICM46511.2021.9385606.

A Score and Variance Evolution Graphs

This appendix shows the **mean scores** (Figure 8) and **score variance** (Figure 9) achieved by each participant during each block (1 minute long, 10 stars) of the pilot study. The participants are grouped by the group they were assigned to during the study. Additionally, Figures 10 and 11 show the whole group's performance during each trial.

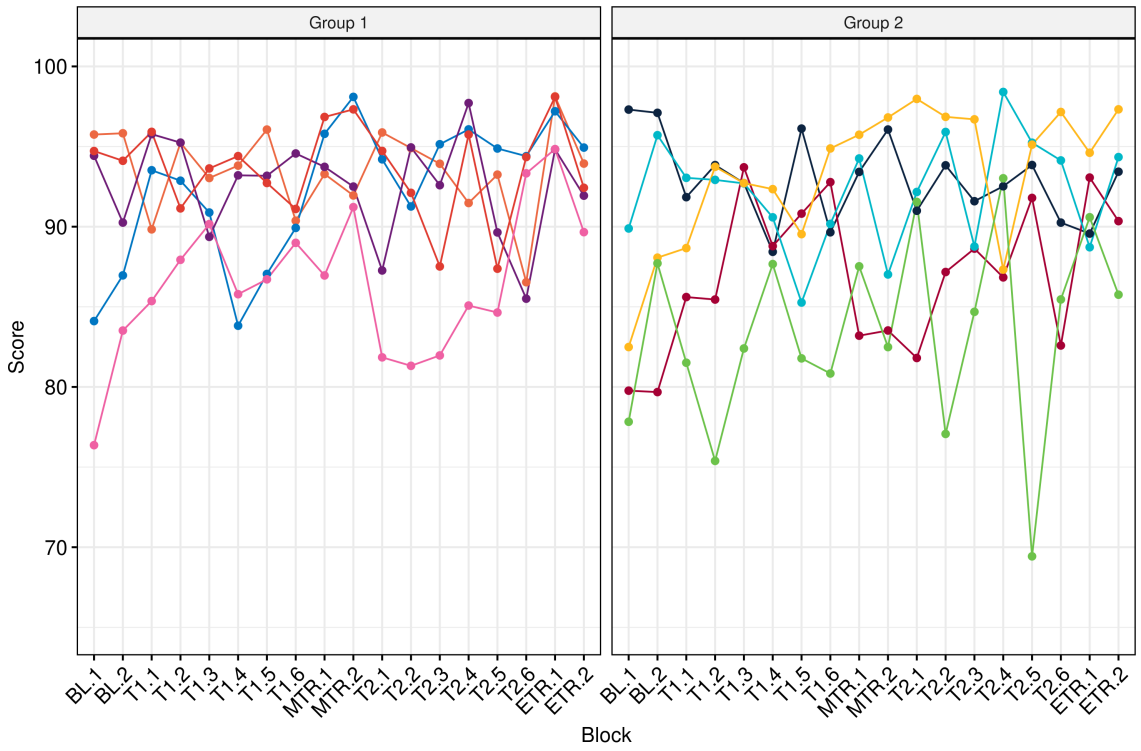

Figure 8: Mean score evolution throughout the experiment, separated by groups. Higher score is better.

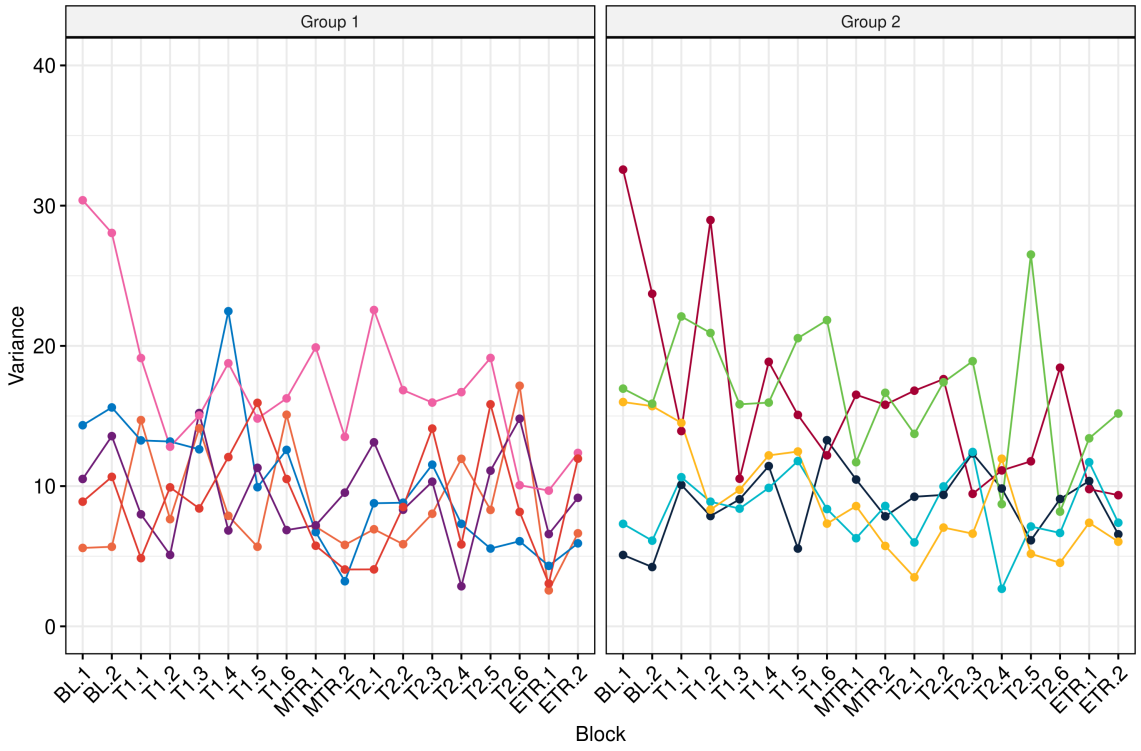

Figure 9: Score variance evolution throughout the experiment, separated by groups. Lower variance is better.

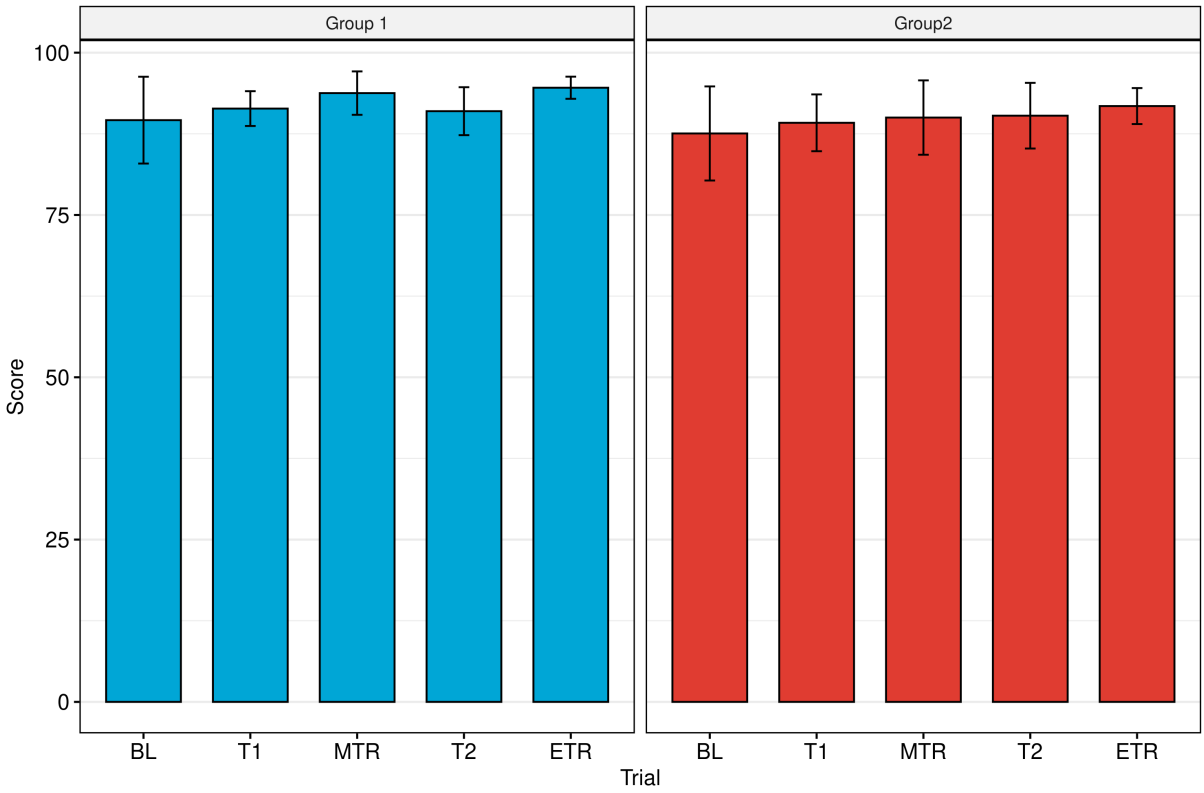

Figure 10: Group mean scores during each trial. Higher score is better. The error bars show standard deviation between the subjects.

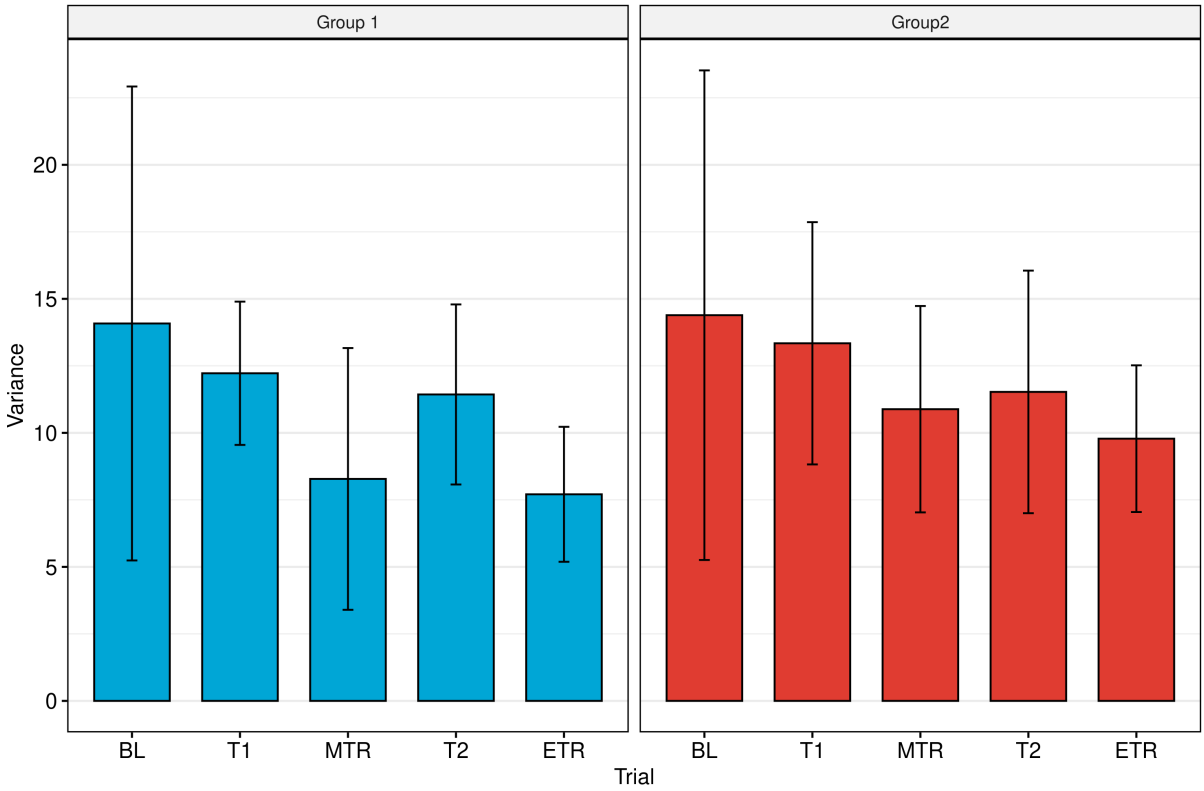

Figure 11: Group score variances during each trial. Lower variance is better. The error bars show standard deviation between the subjects.

B More details on MPC implementation

B.1 Model Predictive Control Problem

The MPC problem is stated in Equation 4, where N is the number of steps in the receding horizon, x is the bicycle state, r is the reference state, u is the control input, and Q and R are user-defined weighing matrices. Subscripts *lb* and *ub* stand for *lower bound* and *upper bound*, respectively, and are used in designer-defined constraints on the system states and control inputs. Matrices A and B are linear time-invariant state-space matrices.

$$\begin{aligned} J = \sum_{k=0}^{N} (x_k - r_k)^T Q_k (x_k - r_k) + \sum_{k=0}^{N-1} u^T R_k u \\ \text{subject to } x_{k+1} = Ax_k + Bu_k \\ x_{lb,k} \leq x_k \leq x_{ub,k} \\ u_{lb,k} \leq u_k \leq u_{ub,k} \end{aligned} \tag{4}$$

In order to calculate the cost, all states of the system need to be known, but at each moment in time only x_0 – the current state of the system – is known. In this instance, using the linear time-invariant state-space representation makes overcoming this issue easy, as it is possible to obtain all the future states while only knowing the initial state and a sequence of future control inputs. Equations in 5 show the principle behind it. By stacking these states into one vector $\bar{x} = \begin{bmatrix} x_0 & x_1 & \cdots & x_N \end{bmatrix}^T$, the control inputs into $\bar{u} = \begin{bmatrix} u_0 & u_1 & \cdots & u_{N-1} \end{bmatrix}^T$, the sequence of equations can be simplified into one equation 6. The reference states, and lower and upper bounds are combined into their own vectors as well, as shown in Equation 7. The weighing matrices are joined into diagonal matrices Equation 8. All of these vectors and matrices lead to a simplified MPC problem, seen in 9.

$$\begin{aligned} x_{i+1} &= Ax_i + Bu_i \\ x_{i+2} &= Ax_{i+1} + Bu_{i+1} = A^2 x_i + ABu_i + Bu_{i+1} \\ x_{i+3} &= Ax_{i+2} + Bu_{i+2} = A^3 x_i + A^2 Bu_i + ABu_{i+1} + Bu_{i+2} \\ x_{i+N} &= Ax_{i+N-1} + Bu_{i+N-1} = A^N x_i + A^{N-1} Bu_i + \cdots + Bu_{i+N-1} \end{aligned} \tag{5}$$

$$\bar{x} = \begin{bmatrix} I \\ A \\ A^2 \\ \vdots \\ A^N \end{bmatrix} x_0 + \begin{bmatrix} 0 & 0 & \cdots & 0 \\ B & 0 & \cdots & 0 \\ AB & B & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A^{N-1}B & A^{N-2}B & \cdots & B \end{bmatrix} \bar{u} = \bar{A}x_0 + \bar{B}\bar{u} \tag{6}$$

$$\bar{r} = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_N \end{bmatrix} \quad \bar{x}_{ub} = \begin{bmatrix} x_{ub,1} \\ x_{ub,2} \\ \vdots \\ x_{ub,N} \end{bmatrix} \quad \bar{x}_{lb} = \begin{bmatrix} x_{lb,1} \\ x_{lb,2} \\ \vdots \\ x_{lb,N} \end{bmatrix} \quad \bar{u}_{ub} = \begin{bmatrix} u_{ub,1} \\ u_{ub,2} \\ \vdots \\ u_{ub,N} \end{bmatrix} \quad \bar{u}_{lb} = \begin{bmatrix} u_{lb,1} \\ u_{lb,2} \\ \vdots \\ u_{lb,N} \end{bmatrix} \tag{7}$$

$$\bar{Q} = \begin{bmatrix} Q_1 & 0 & \cdots & 0 \\ 0 & Q_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & Q_N \end{bmatrix} \quad \bar{R} = \begin{bmatrix} R_1 & 0 & \cdots & 0 \\ 0 & R_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & R_N \end{bmatrix} \tag{8}$$

$$\begin{aligned} J = (\bar{x} - \bar{r})^T \bar{Q} (\bar{x} - \bar{r}) + \bar{u}^T \bar{R} \bar{u} \\ \text{subject to } \bar{x} = \bar{A}x_0 + \bar{B}\bar{u} \\ \bar{x}_{lb} \leq \bar{x} \leq \bar{x}_{ub} \\ \bar{u}_{lb} \leq \bar{u} \leq \bar{u}_{ub} \end{aligned} \tag{9}$$

B.2 Optimisation Problem

The only unknown variable left in the simplified MPC problem is the vector of control inputs, which is where optimisation comes in. The goal is to minimise the cost J by using a sequence of minimal control inputs $\bar{u}$ that

minimises the difference between the system states and the reference states $\bar{x} - \bar{r}$. The relative importance of the magnitude of state errors and the magnitude of the control inputs can be adjusted by changing the weights in the matrices $\bar{Q}$ and $\bar{R}$.

Since the cost function is quadratic and is constrained, the optimization method of choice is Quadratic Programming (QP). In order to implement MPC, the QP problem needs to be solved online at each time step, which requires a fast solver. In this thesis, `qpOASES`[11] solver was chosen, as it features a Simulink interface that is easy to implement.

`qpOASES` solves the QPs of the form given in equation 10, where y is the variable to optimise, H is the Hessian, ω_0 is a parameter, g is a parameter-dependent gradient vector, E is a matrix, and lbE, ubE, lb and ub are parameter-dependent lower and upper constraint bounds.

$$\begin{aligned} \min_{y} \quad & \frac{1}{2} y^T H y + y^T g(\omega_0) \\ \text{subject to } & lbE(\omega_0) \leq Ey \leq ubE(\omega_0) \\ & lb(\omega_0) \leq y \leq ub(\omega_0) \end{aligned} \tag{10}$$

The MPC problem can be translated to this QP problem by substituting the stacked state-space equation 6 into the cost function. By expanding and collecting terms, equation 11 is obtained. The last three terms can be excluded as x_0 and $\bar{r}$ are known beforehand, are constant during optimisation, and act as scaling factors. Dividing the equation by two and equating $\bar{u}$ with y, the equation takes the form of the `qpOASES` minimisation function, with $H = \bar{B}^T \bar{Q} \bar{B} + \bar{R}$ and $g(\omega_0) = \begin{bmatrix} \bar{B}^T \bar{Q} \bar{A} & -\bar{B}^T \bar{Q} \end{bmatrix} \begin{bmatrix} x_0 \\ \bar{r} \end{bmatrix}$.

$$J = \bar{u}^T (\bar{B}^T \bar{Q} \bar{B} + \bar{R}) \bar{u} + 2\bar{u}^T (\bar{B}^T \bar{Q} \bar{A}) x_0 + 2\bar{u}^T (-\bar{B}^T \bar{Q}) \bar{r} + x_0^T (\bar{A}^T \bar{Q} \bar{A}) x_0 + 2x_0^T (-\bar{A}^T \bar{Q}) \bar{r} + \bar{r}^T \bar{Q} \bar{r} \tag{11}$$

Substituting 6 into the constraints of the MPC problem 9 and reordering results in 12. Which gives that $lb(\omega_0) = \bar{u}_{lb}$, $ub(\omega_0) = \bar{u}_{ub}$, $lbE(\omega_0) = \bar{x}_{lb} - \bar{A}x_0$, $ubE(\omega_0) = \bar{x}_{ub} - \bar{A}x_0$, and $E = \bar{B}$.

$$\begin{aligned} \bar{x}_{lb} - \bar{A}x_0 \leq \bar{B}\bar{u} \leq \bar{x}_{ub} - \bar{A}x_0 \\ \bar{u}_{lb} \leq \bar{u} \leq \bar{u}_{ub} \end{aligned} \tag{12}$$

B.3 Time-variance

As mentioned in subsection 2.3, the weight matrix Q_k varies with time and can take on different values for each sampling time. This means that $\bar{Q}$ is different at each sampling time and QP problem's H matrix needs to be recalculated at each sampling time.

`qpOASES` provides a standard class `QProblem`, but it requires the H matrix to be constant. Luckily, the toolbox also provides `SQProblem`, which allows for a varying H (or E) matrix.

A visualisation of how the matrix Q_k varies can be seen in Figure 12. The three coloured lines represent the three varying weights - lateral position y_P, yaw angle ψ and roll angle ϕ. All weights start at the value of 0. When 4 seconds are left before the star, the weights start to increase. The weight for y_P starts at 0.1, while the weights for ψ and ϕ start at 1. The weights reach their final values of 10 for y_P and ψ and 3 for ϕ at the same moment as the star is passed. Afterwards, the weights are set back to 0.

B.4 Reference

In this thesis, the controller is able to preview the reference, meaning that it knows exactly what is waiting for it in the future. The reference contains the desired lateral location of the bicycle's rear wheel contact point y_P at each time step, while all other state references are set to 0. Setting the yaw and steering rate reference to 0 acts like a damper. Setting the yaw, roll, and steering angle reference to 0 balances the bicycle, as the controller is going to try to keep these values at zero.

The reference for y_P is the only one that is varying between time steps. A visualisation of a part of it is shown in Figure 13. A time snippet of 16 seconds are shown. At the current time step, the star in position 2 is being hit. The next star is located 6 seconds away in position 1. The MPC controller does not "see" the stars (shown in red) themselves, but "sees" the reference lines (shown in green) that passes through the centre of the stars.

The reference line begins 4 seconds before the star and corresponds to the period during which the weights are increased, as mentioned in subsection 2.3. The line extends past the star for 2 seconds to ensure that the controller passes through the star. If this extension did not exist and the line immediately shifted to another star, the controller would start turning towards the other star and miss the current star. The length of the extension was chosen to be the same length as the time horizon of the controller, to make sure that at the moment in time when a star is passed, all the controller sees is the location of that exact star and not the location of a future star.

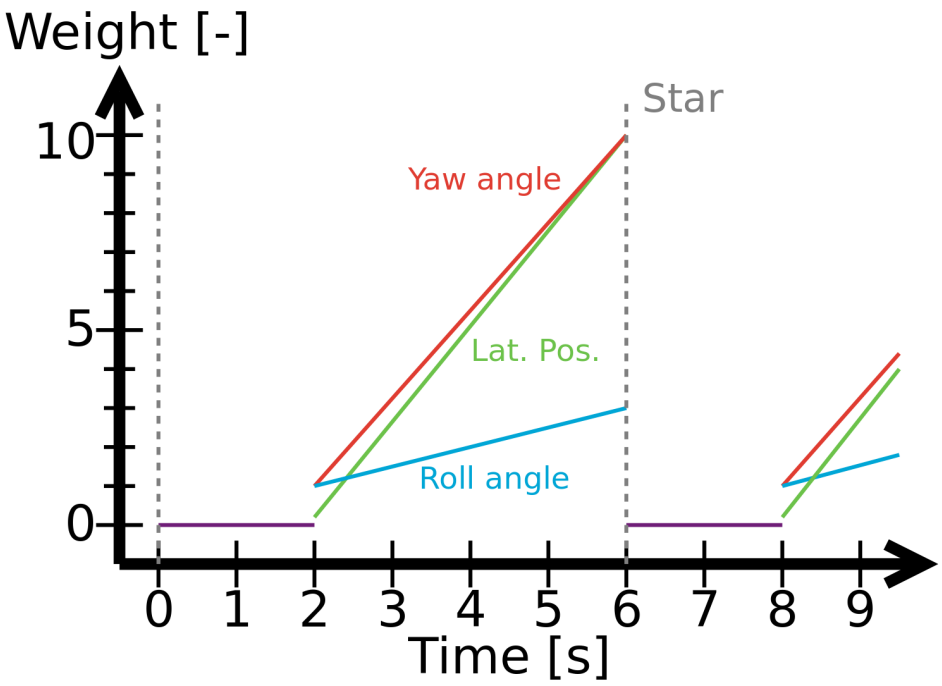

Figure 12: Visualisation of the controller weights.

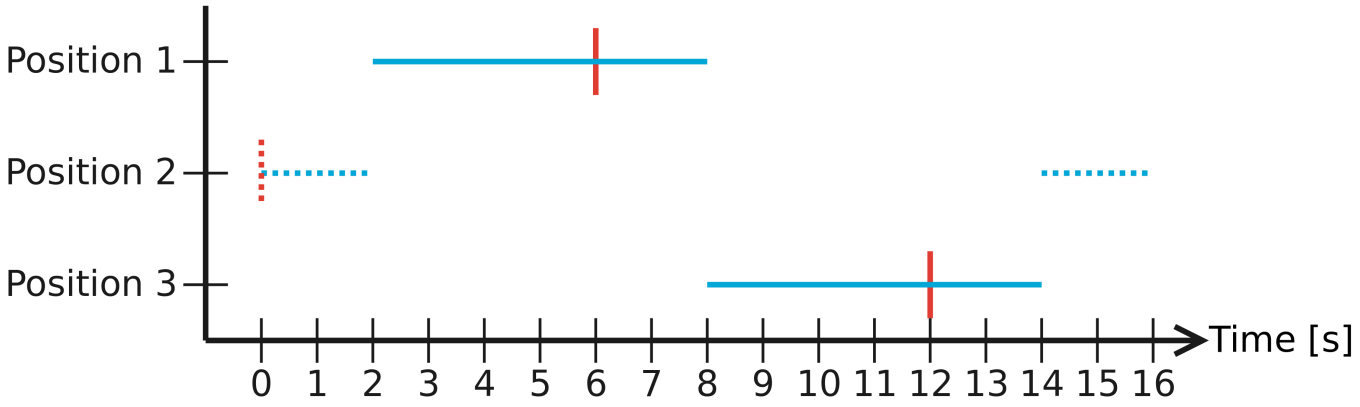

Figure 13: Reference line of the rear wheel contact point y_P supplied to the MPC controller.
Red vertical lines represent the stars, with the centre of the line representing the centre of the star that would produce a score of 100. Blue horizontal lines represent the reference position of y_P, shown to the MPC controller.

C Additional Observations

This appendix contains additional observations that could be useful for future projects and experiments.

C.1 MPC Controller

C.1.1 Controller Weights and Trust in the Controller

The reasoning behind implementing varying controller weights was to provide the participants with some freedom of exploration during the early stages of the manoeuvre [42]. However, in practice this led to one of two things - the participant had already completed the manoeuvre before the controller had a chance to act, or, if the participant had not completed the manoeuvre, the controller would act too late to have any meaningful performance improvement. Additionally, varying weights reduce the ability of the participant to predict what the controller is going to do next. Now the participant needs to figure out a three way mapping between the bicycle's location relative to the star, time to the star and the controller's torque. With constant weights, only a two way mapping would be needed – between the location of the bicycle relative to the star and the controller's torque. This can be an issue, since inability to predict the controller might lead to reduced trust [5].

And trust in the controller is important in ensuring performance improvements. The participants need to trust the controller to be able to follow its guidance. If the participants do not follow the guidance, they will be fighting against it, reducing the performance.

C.1.2 Sensors

Better state estimation and cleaner sensor data would lead to better controller performance and lower deviations. While the current setup with the HTC Vive Tracker 3.0 works well, some possible areas of improvement are found.

A more rigid tracker mount should be designed to avoid as much of relative motion between the tracker and the bicycle as possible. Better physical alignment between the tracker's coordinate system and the bicycle's system should be sought out as well. Currently the misalignment is solved using a calibration procedure before each participant. However, this approach is not 100% reliable and a bias in the roll angle was found for one participant, which led to a bias in controller's torques.

Additionally, it would be beneficial to know how accurate the tracker is. A paper by Borges et al. [4] found that in a static condition, the precision of the tracker is under a millimetre. However, under dynamic conditions, they found that the precision decreases into the range of centimetres. The precision was measured using earlier Tracker versions, as well as earlier versions of the Base Stations, and they used Vive's proprietary algorithms. Therefore, it is recommended to carry out new precision tests with the new hardware and using `libsurvive`'s algorithms. To be able to do that, another, well-tested, localisation-capable sensor is needed to obtain a ground truth, such as a motion capture camera setup.

C.1.3 Autonomous Bicycle

In this thesis, the controller was designed to play a role of an assistant for the human rider, rather than a system that would ride the bicycle by itself. However, the no hands trial showed that the controller, with some improvements, could hypothetically be used for a self-riding bicycle. To test this hypothesis, the rider could be replaced with a rigidly-attached mass and the MPC weights should be adjusted. After a satisfactory performance is achieved, different masses could be fitted without adjusting the MPC weights to investigate how sensitive the controller is to unmodelled changes in dynamics.

C.2 The Game

After the study, participants expressed some annoyances regarding the game. Two main complaints were about the in-game perspective and the delay between the movement of the real and virtual bicycles.

C.2.1 In-Game Perspective

The game is displayed on a 24-inch display placed around 2 metres in front of the rider. The camera in the game is stationary and fixed to the middle of the treadmill.

Participants noted that, as the star appeared in the horizon, it was hard to accurately predict where the star was going to end up in a couple of seconds. Additionally, as the star came closer, it was difficult to tell how far away from the centre of the star the virtual bicycle was. A couple of possible improvements are suggested.

One of them is to constantly show the participants their relative distance to the centre of the star. This can be achieved in two ways. The simplest one is to display a longitudinal line starting at the bottom of the screen and

ending at the star's centre, which would give the participant a line that they could follow. The second way is to show a "ruler" that laterally extends from the virtual bicycle's current position, to the lateral position of the star's centre. The length of this ruler represents the lateral distance to the star and changes length according to the movement of the bicycle.

Another improvement is to increase the size of the monitor or replace it with a projector or a virtual reality headset. The use of a relatively small monitor compared to the treadmill led to the participants needing to turn their heads towards the display, which can interfere with the perception of the relative positions. The need to rotate the head can be reduced by increasing the display size. By replacing the display with a projector or a virtual reality headset, the real and virtual worlds could be joined seamlessly and lead to better spatial awareness for the participants. Alternatively, the experimental setup could be modified to work outside the laboratory setting and the experiment could be carried out on the road.

Some participants suggested to make the camera move sideways, instead of always being fixed to the middle of the treadmill. In fact, this was piloted before the experiment with two participants. The camera was fixed to the virtual bicycle to simulate a "first person view". However, this led to the feeling of confusion as the participants of the pilot test reported that any movement they made felt like it was twice as big. This issue could be overcome by smarter camera movement, such as allowing the camera rotate towards the middle of the virtual road to mimic the rotation of the head of the participant.

C.2.2 Delay

Seven participants commented that they noticed a slight delay between moving with a real bicycle and the virtual bicycle. A part of the delay is introduced by design - the virtual bicycle represents the location of the rear wheel of the real bicycle. As bicycle riders are not actively thinking about the location of the rear wheel while riding a bicycle, but rather the location of the front wheel, this choice was made to make the task slightly more difficult. The participants were not informed about this before the experiment as it was assumed that the participants will be able to adapt to the induced delay by themselves.

Another component of the delay came from the delays in communication and filtering. The raw data from the HTC Vive Tracker 3.0 is sent using Bluetooth to the Raspberry Pi, which then sends the calculated position and pose data to the desktop PC running the controller. The data then passes through a second order Butterworth filter, and, finally, the lateral position data is extracted and sent to the game. Each of those steps has an inherent delay that is likely smaller than what a human could perceive, but in this case, the delays compound and become big enough for human perception. This delay, which happened to be around 100 ms, could be reduced by reducing the number of communication steps, for example, by making the tracker communicate directly with Simulink, or by getting rid of Simulink altogether and implementing the controller on a bicycle or in the game itself.

C.3 Treadmill Riding

C.3.1 Starting From a Standstill

It was found that the most difficult part of riding on the treadmill is the very beginning. On the road, when starting from a standstill, the rider quickly accelerates the bicycle to a more stable and controllable speed range. On a treadmill, the acceleration depends on the treadmill's capabilities and the treadmill might be accelerating too slowly, which forces the rider to spend more time at speeds where it is hard to balance. During testing before the study, some riders were not able to start riding on this particular treadmill from a standstill, therefore, handrails were installed on the sides of the treadmill. The main study's participants were asked to hold on to the handrail at the beginning of the different sections, whilst the treadmill accelerated to the controller's design speed of 15 km/h. However, this introduced another issue, as some participants had trouble letting go of the handrail.

In order to hold on to the handrail, the participants had to place their bicycle near the edge of the treadmill's belt, which limited the amount of space in which the participants could manoeuvre, compared to starting in the middle of the belt. Additionally, the height of the handrail is not adjustable, therefore, taller participants had to slightly roll with the bicycle to be able to reach the low handrail. As the bicycle is rolling towards the handrail, in order to cancel the roll angle, the rider needs to steer towards the handrail, which brings the bicycle even closer to the edge of the treadmill's belt and the rider grabs onto the handrail again to make sure that they do not ride off the side of the belt. Thus, at the beginning of the study, it generally took a couple of tries for the participant to fully let go of the handrails and become stable. The number of tries needed would then quickly decrease and, by the end of the experiment, the majority of participants were able to let go of the handrail on the first try.

Taking into account the different amounts of time the participants needed to let go of the handrail, the 5 minute timer of the familiarisation trial was started only when the participant was stable and in the middle of the treadmill.

C.3.2 A "Different Feel"

The participants reported that riding on a treadmill felt different to riding on the road. The exact reasons of why the feelings differ are not yet known. The author believes that the bicycle dynamics on the treadmill do not change significantly from those on the road, as shown by Kooijman and Schwab [19, 20]. Therefore, the author suspects that the reasons lie in motion perception and safety perception of the rider.

One of the contributing factors could be that the participants were instructed not to pedal during the experiment, and some people commented that they normally use pedalling to help them balance. However, the author wants to bring attention to coasting. During normal cycling there are many times when the rider does not pedal (for example while slowing down or going downhill), yet can still balance and carry out manoeuvres, such as lane changes. Therefore, the author does not believe that not pedalling has had a big effect on participants' performance.

Other remark from the participants was that they did not feel comfortable riding close to the edge of the treadmill. It is likely due to the manoeuvrability being limited and can be compared to riding on a very narrow bicycle path on the road.

The author thinks, however, that the main reason is the lack of optical flow, which is how humans perceive velocity. The author hypothesises that the internal bicycle dynamic models, that the rider has learned, have velocity dependent gains, just as the bicycle dynamics are velocity dependent. Therefore, to be able to use the correct gains, the rider needs to be able to estimate the velocity at which they are currently moving. On the road, this can be done using the visual cues as the surrounding objects move past the rider. On the treadmill, however, the surroundings are stationary and the rider does not have any reference to estimate the velocity from. While the participants were informed that the treadmill was going to be set to 15 km/h, it is likely that it was not enough. A solution for this is to use a virtual reality headset or a display that surrounds the sides of the treadmill and can show moving objects in the peripheral vision.

Vestibular cues (cues regarding accelerations) are also important for movement perception and the absence of them can make movement feel different and unnatural. As the bicycle on a treadmill is moving at a constant velocity, there is no longitudinal acceleration and is equivalent to riding a bicycle on the road at a constant velocity. Other major vestibular cues come from lateral and angular accelerations, which are expected to exist during lane change manoeuvres while riding on a road. The same accelerations also exist on the treadmill, as the bicycle is able to move from side to side, and is able to change its yaw angle. Therefore, the author thinks that there are no significant vestibular cues missing from this experiment, which could contribute to the different feel that the participants experienced.

C.4 Steer-by-Wire Bicycle

C.4.1 Play in the Steering Assembly

The participants reported feeling a bit of play in the steering assembly that led to the front fork not reacting to small steering inputs (especially during micro-corrections) or reacting to them with a slight delay. This is due to the bicycle not having a mechanical link between the fork and the handlebars, but rather having two electric motors, which have gearboxes, connected to the handlebars or the fork using timing belts. The electric motors are controlled by PD controllers.

The lack of reaction to small steering inputs (approximately < 0.5 degrees) is therefore happening due to a couple of compounded reasons. First is that when a small steering angle is applied to the handlebars, only a small angular error between the handlebars and the fork is observed. A small angular error leads to a small torque request by the PD controller. A big part of this already small torque is then absorbed by the gearbox of the electric motor, either due to backlash or friction losses. The remaining part of the torque reaches the timing belt, where, again, a bit of the torque is lost in the stretching of the belt. Additionally, there are friction losses in the fork bearings as well, which absorb the remaining amount of torque. To overcome these losses, the rider needs to increase their steering angle inputs slightly, compared to a normal bicycle. This issue could be overcome by forgoing the gearbox and the belts and instead opting to connect the fork and the handlebars straight onto the shafts of their respective electric motors. However, it is not an easy improvement to implement as new motors would need to be sourced and the bicycle frame would need to be redesigned. Therefore, a possible improvement could be achieved electronically through PD controller tuning or feed-forward friction compensation. As an alternative, a bicycle with one electric motor connected to a conventional mechanically coupled steering assembly could be used, if advanced features enabled by the steer-by-wire system (such as application of different amplitude torques to the fork and the handlebars) are not needed.

C.4.2 On-board Controller

A significant improvement would be to integrate the whole controller into the bicycle itself, which would get rid of the wireless communication delays and also allow the controller to be used outside of the laboratory setting. To do

this, two major things would need to be changed.

Currently, the bicycle's brain is a Teensy 4.1. It is one of the most powerful affordable microcontrollers. However, it is still likely not powerful enough to run the current iteration of the MPC controller. To make the controller work on a Teensy, it would need to be written in C or C++ instead of being implemented as a Simulink model, and would likely need to be simplified by reducing the time horizon, the sample rate, or by simplifying the dynamic model. Another approach could be to replace the Teensy with a Raspberry Pi, which is a tiny computer that has General-Purpose Input/Output pins that could be used to communicate with various sensors.

The second major change would be to implement a way to accurately estimate the position (mainly lateral) and pose (mainly roll and yaw angles) of the bicycle without the need for the HTC Vive Tracker and its Base Stations. The angles could be estimated using IMU data and Kalman (or similar) filters, while the (relative) lateral position would likely require cameras and machine vision.

D Supplementary Time Domain Analysis

This appendix contains time-domain graphs (Figures 14 and 15) of the excerpt from the data recorded during the no-hands test. This particular excerpt is taken from the last minute of the no-hands test, between the fifth and sixth stars of the sequence. These two stars are interesting due to the lane change manoeuvre being the biggest in amplitude (0.33 metres, from -0.13 m to 0.2 m) in the study.

Figure 14 shows the bicycle's lateral position, yaw angle, roll angle, steering angle and the controller torque in solid blue. The red circles in the lateral position graph represent the locations of the stars. Green dashed lines represent an "ideal" trajectory, obtained from the simulation of the MPC controller with the same weights. Please note that the MPC torque's graph contains two y-axes for easier legibility, as the torques in the simulation were significantly lower.

Figure 15 shows the bicycle's roll and steer rates in solid blue, while dashed green represents the simulation data. Please note the different y-axes for both rates.

It can be seen that the lateral position and the yaw angle of the real bicycle roughly matches the ideal simulation graphs. At the beginning of the manoeuvre, the roll angle, steer angle and the torque try to follow the ideal trajectories. However, as time passes, the real bicycle's data drifts away from the simulation data. The author suspects that this disagreement arises due to a number of reasons.

Approximate bicycle and rider parameters. The steer-by-wire bicycle used in this project has not been measured for its parameters, as the bicycle would need to be disassembled to be measured. Due to the complexity of the steer-by-wire system, disassembly and reassembly would be a very time-consuming task. Instead, the parameters of a physically similar (*Davis Instrumented Bicycle* [40]) bicycle have been used. Both bicycles feature an electric hub motor installed on the rear wheel, a battery fixed to the seat tube, an electronics box placed above the rear wheel, and a modified steering assembly. Not only that, but the rider parameters are important in bicycle dynamics as well. Similarly to bicycle parameters, rider parameters were also approximated, as measuring each participant before an experiment would be very time consuming. Therefore, the controller used during the experiment did not know the actual parameters of the system it was controlling, while the simulated controller had a perfect knowledge of the system.

Linear Whipple-Carvallo model is not a perfect description of bicycle dynamics. It is the simplest model that is able to capture bicycle's self-stability. To achieve that simplicity, several assumptions have been made. For example, the wheels are modelled as an ideal knife-edge contact with the ground, instead of having a toroidal shape and exhibiting tyre forces. The rider is also assumed to be rigidly connected to the bicycle, instead of being free to move relative to the bicycle like in the real world. It is possible to extend this model further to relax some assumptions (like it was done in [37]). However, the canonical Whipple-Carvallo model was chosen due to the need for a real-time controller, and each additional extension of the model introduces additional equations that need to be solved online, requiring more computational power.

Unmodelled disturbances. As mentioned before, the Whipple-Carvallo model assumes that the rider is rigid. That was not the case during the experiment, where the rider could interfere by rolling in different directions and acting as a disturbance. Additionally, the model assumes that there are no mechanical losses in the system. In this situation, the losses in the steer-by-wire system might have played a big part. The author thinks that their influence can be seen in the torque graph, as during the simulation the torques needed to carry out the manoeuvre were nearly a magnitude lower than observed in reality, while the yaw, roll and steer angles are of similar magnitude.

To conclude, the no-hands trial shows that the controller tries to follow an ideal line, but the approximations and unmodelled dynamics and disturbances prevent the controller from performing at its best. Regardless, the controller is still able to carry out lane change manoeuvres, and the performance can be improved over time by addressing the above mentioned limitations.

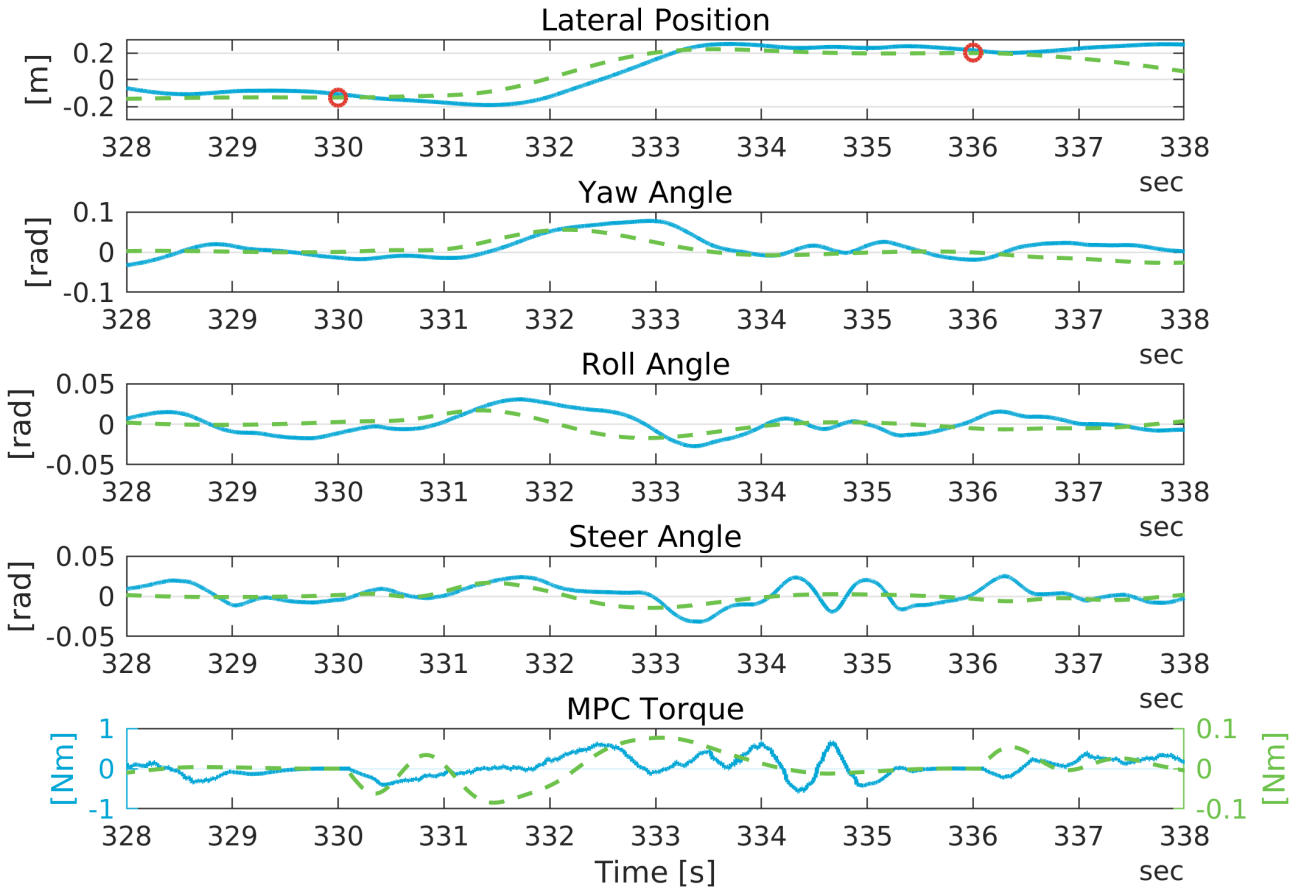

Figure 14: Data excerpt from a no-hands trial, compared to simulation data. Blue solid line is the real data, green dashed line is the simulation data, red circles represent the location of the stars. Please note the different y-axes in the torque graph.

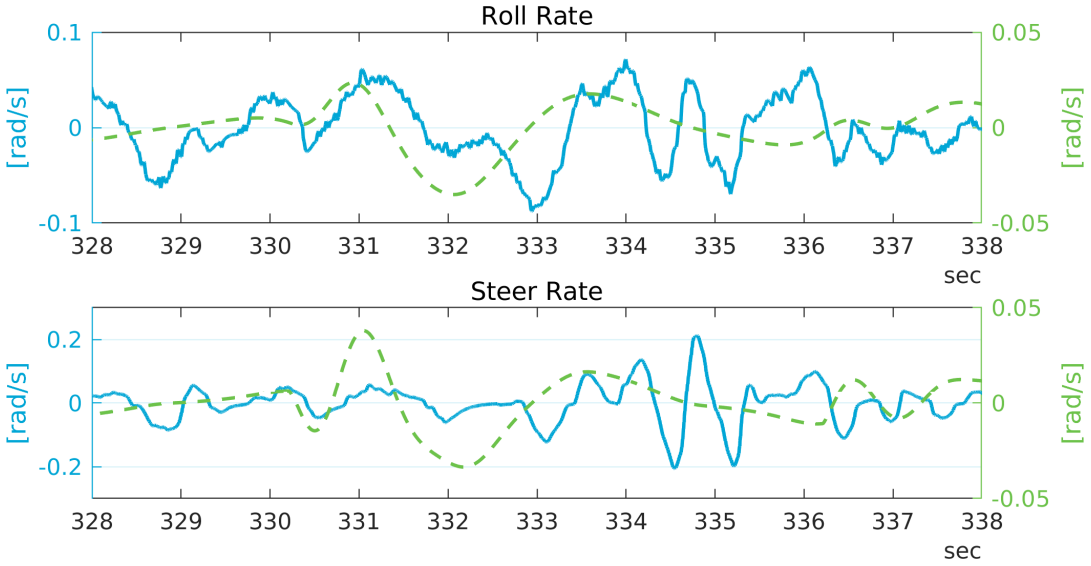

Figure 15: Roll rate and steer rate data from the no-hands trials, compared to simulation data. Blue solid line is the real data, green dashed line is the simulation data. Please note the different y-axes.