Bicycle-rider control identification

Thesis Report

K.D. Wendel

Student ID: 4170555
Department: Cognitive Robotics
Specialisation: Dynamics & Control
Supervisors: Dr.ir. A.L. Schwab & ir. M.M. Reijne
Date: May 26, 2020

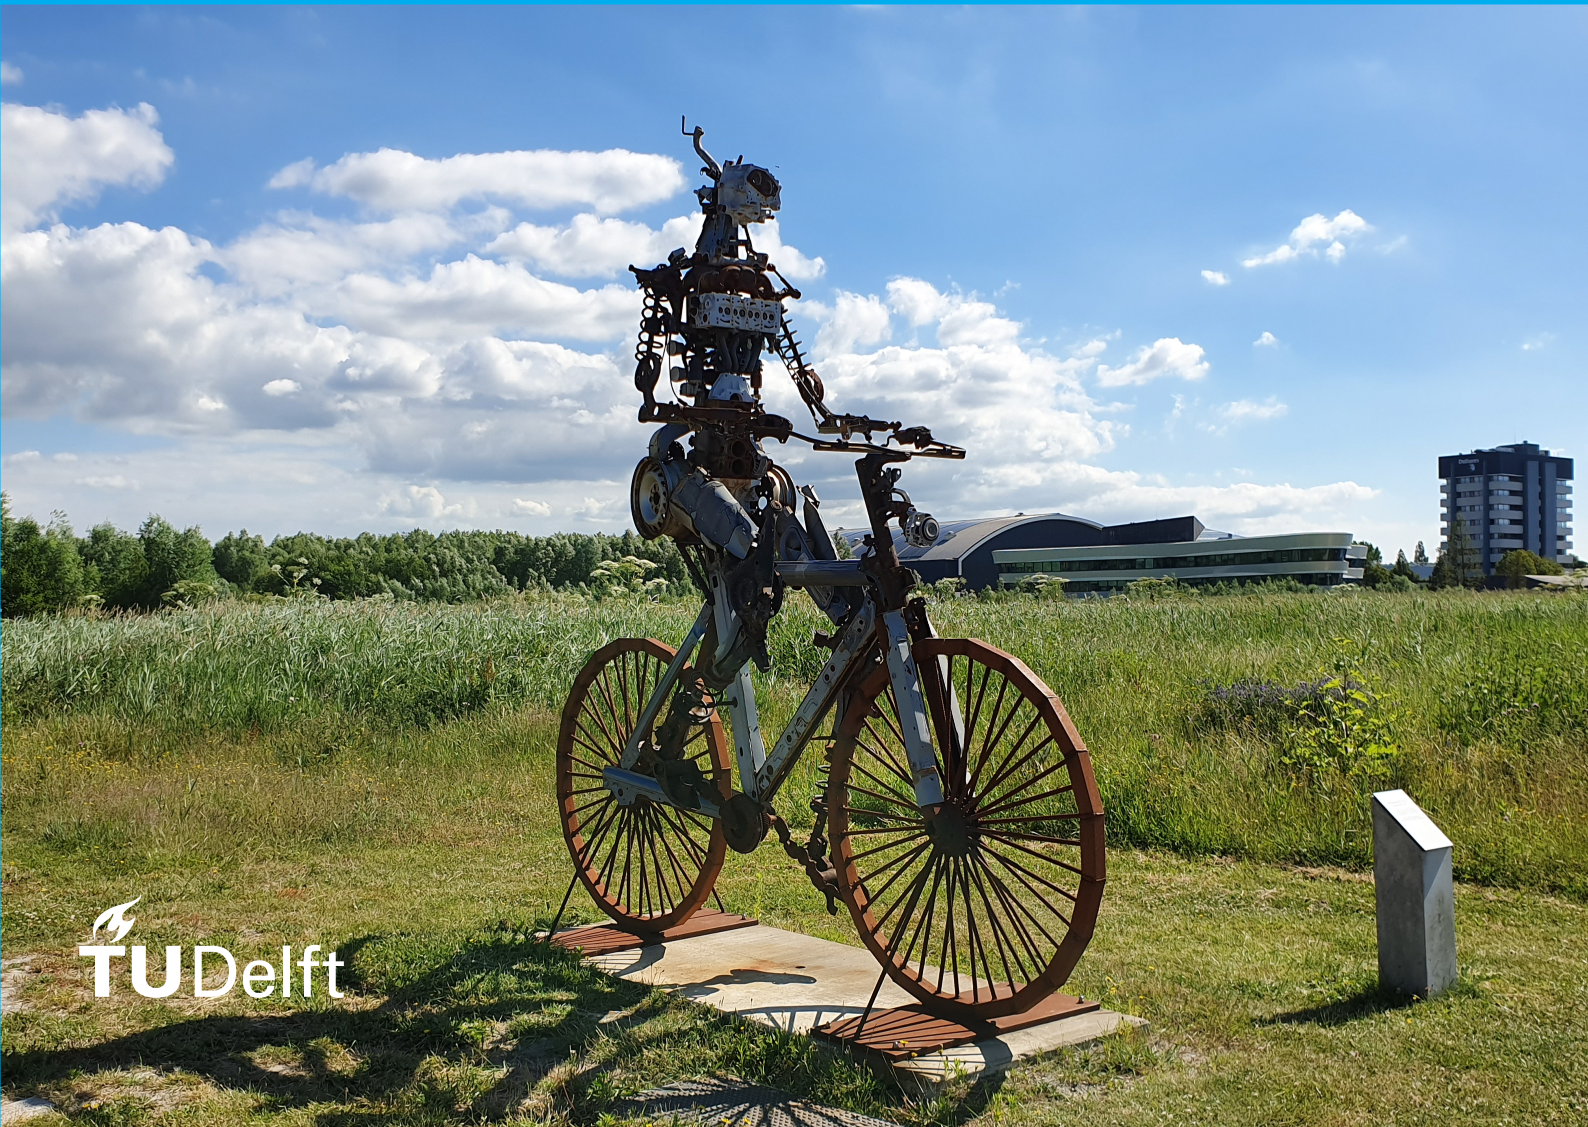

Contents

Glossary

Glossary

Bayes' theorem Bayes' theorem is a statistical rule that can be applied when a known probability **P(A)** of event **A** occurring, is related to the occurrence of event **B**. This is called the conditional probability: **P(A|B)**. Bayes' theorem then states:

$$\mathbf{P(B|A)} = \frac{\mathbf{P(A|B)P(B)}}{\mathbf{P(A)}}$$

23–25

Central-Nervous-System The central-nervous-system consists of the brain and the spinal cord. In the Central-Nervous-System, sensory information is processed and reactive (both voluntary and involuntary) neural control outputs are given. vi, 15

cost function In the context of mathematical optimization, a cost function provides a measure of wrongness of a model. Minimizing a cost function reveals the variables that are associated with the optimal (lowest) cost. 4, 28

feedback In a control theory context, feedback is when the output of a system is, usually manipulated in some way, *fed back* in to the system. vi, 3, 9, 26–28, 52

Finite-Impulse-Response Finite-Impulse-Response is a performance metric that shows how much two signals are alike. A VAF of 100% indicates two signals that are the same. It can be calculated with the following formula: vi, viii, 36, 52, 76

fuzzy control Fuzzy control has been developed in order to convert complex and poorly defined control tasks done by humans, to a digital controller, by allowing linguistically expressed heuristic control policies to be converted to fuzzy rules [105] [106]. Fuzzy rules are rules which can be partially true (as opposed to Boolean rules, which are either true or false), and can therefore be used to describe uncertainty. 28

group delay The group delay describes the delay of the amplitude envelope of a signal as a function of frequency. 36, 37

Kalman filter A Kalman filter can provide, using knowledge about the system, real-time state estimates that are usually better than the single measurements alone. 23–25

Linear-Quadratic-Regulator The linear-quadratic-regulator uses the description of the system and weighting factors to minimize a cost function. If the system description has sufficient fidelity, the algorithm provides feedback gains that are optimal to the weighting factors, and that stabilize the system. vi–viii, 4, 28, 52

Linear-Time-Invariant systems A dynamical system is called linear when the output of the system scales linearly with the input of the system. A system is time invariant when the output is not dependent on when in time the input is applied. Systems with both characteristics are called Linear-Time-Invariant systems. vi, 14

manifold In mathematics, a manifold is a type of topological space. It is used to describe complex structures in the more intuitive and simpler Euclidean space. An often used example of this is the mapping of the surface of the earth, which is a two-dimensional surface manifold known as a sphere, to a topological chart, which is a two-dimensional Euclidean space. This can't be done for the entire surface of the earth at once, but for smaller area's of the earth the transformation is mathematically valid. Hence, for proper charts of the entire planet, we need an atlas. 4

Model-Predictive-Control Model-predictive-control uses a model to predict and anticipate to future events for a finite time-horizon. It then optimizes the current control action by minimizing a control strategy cost function over a receding prediction horizon. vi, 28

neural delay The time delay caused by the transport of signals across the neurons and the time to process or generate those signals in the CNS. 4, 6

open-loop When a system is in an open-loop configuration, the output of that system is not used to provide any feedback. 13, 14, 25, 27

optimal control Optimal control is a method to control a system in such a way that an objective function is optimized. Usually this is done to weigh a certain result or gain against the effort to get there. A well known optimal control method is the linear-quadratic-regulator. 4, 6

state State variables are variables that evolve through time and describe the ”state” of a system at a certain time instant. v, 15, 16, 23–28, 42, 52

Variance-Accounted-For Variance-Accounted-For is a performance metric that shows how much two signals are alike. A VAF of 100% indicates two signals that are the same. It can be calculated with the following formula:

$$VAF = \left(1 - \frac{var(y - \hat{y})}{var(y)}\right) \cdot 100\%$$

vi, viii, 36, 53, 76

Acronyms

CNS Central-Nervous-System vi, 15, *Glossary:* Central-Nervous-System

FIR Finite-Impulse-Response viii, 36–39, 52, 54, 59–61, 63, 64, 67, 68, 76, 78, 79, 82, *Glossary:* Finite-Impulse-Response

LQR Linear-Quadratic-Regulator vi–viii, 4, 5, 28, 40–42, 45, 48, 52–59, 62–64, 66, 67, 70, 71, 75–77, 81, *Glossary:* Linear-Quadratic-Regulator

LTI Linear-Time-Invariant 14, *Glossary:* Linear-Time-Invariant systems

MPC Model-Predictive-Control 28, *Glossary:* Model-Predictive-Control

VAF Variance-Accounted-For viii, 36–39, 44, 47, 53, 55, 57, 59–65, 67–69, 71, 76, 79, *Glossary:* Variance-Accounted-For

Notation and Symbols

$A(i,:)$	The ith row of matrix A
$A(:,i)$	The ith column of matrix A
G^*	Human internal model of true model G
$\tilde{x}$	(Human) Estimate of x
$\hat{x}$	(Human) Optimal estimate of x

Summary

Statistics show that cycling accidents have the biggest (and increasing) contribution to the overall number of hospital visits related to traffic accidents in the Netherlands. In the majority of these cycling accidents, no other road users are involved [97].

Because the majority of the cycling accidents are so-called single vehicle accidents, understanding the control behaviour of the cyclist can spark new insight in to effective preventive measures. This thesis aims to achieve that by reviewing the current state-of-the-art in bicycle-rider control research and subsequently proposing a new rider control model that is developed using system identification techniques.

The first quantitative research in human bicycle control behaviour was done by van Lunteren and Stassen in the 1970s [54], [87] at the TU Delft[1]. In the following decades numerous attempts had been made, primarily to scratch an academic itch.

The literature study into these attempts reveals that no clear consensus is reached on which controller type, control inputs given by the controllers, and the sensory inputs used by the controller result in the best description of human bicycle control. Furthermore, it is found that sensory dynamics and integration are only modelled very rudimentary, and only a select few models are fitted using experimental data. Even less models are validated.

This provides two clear avenues for the advancement of the state-of-the-art:

1. The performance of different controller types and structures in predicting human control behaviour can be compared using a single experimental data set.
2. More realistic sensory dynamics, integration and delays can be modelled in order to get a deeper understanding on what sensory channels are used primarily by bicycle-riders.

Given these avenues for advancement, the scope of the literature review is expanded. For all components that exist in a generic structure that describes human control of a bicycle, literature is gathered that could describe the functionality of these components.

The results of the literature review is used to propose a novel bicycle-rider control model structure of which the block diagram is shown in figure 1. The components of the model structure are:

- **Balancing controller**: A Linear-Quadratic-Regulator (LQR)-type controller that acts on the estimated states $\hat{y}$.
- **Neuromuscular dynamics**: The relation between the steer torque desired by the rider and the realized steer torque is described by a neuromuscular model which takes the form of a second order filter.
- **Bicycle dynamics**: The bicycle dynamics are described with the linear Whipple bicycle model [59]. Passive rider dynamics are also included here.
- **Sensory dynamics**: In this block the dynamics of the various human sensory organs are described.
- **State estimator**: The human sensory fusion and state estimation process is approximated with a Kalman filter.
- **Predictor**: Human predictive capabilities, including the feed-forward effect of the own control actions, are described with a tapped-delay-line predictor.

---

[1] Then called the *Technische Hogeschool Delft*

- **Internal model**: The notion of an internal model is used to describe the knowledge the bicycle-rider has of the dynamics and delays present in the system.

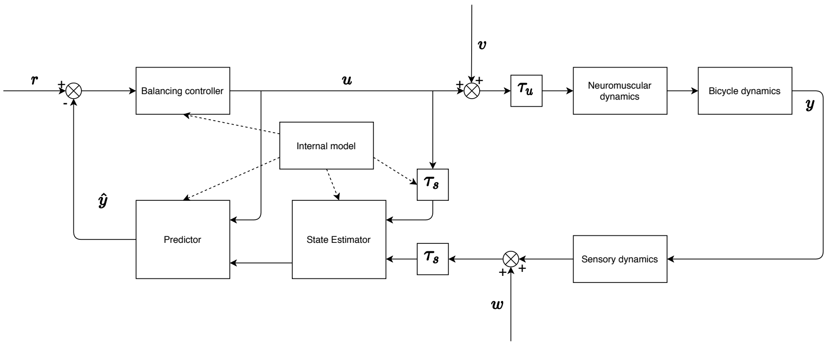

Figure 1: This is the system block diagram of the bicycle-rider system. There are several symbols present in the figure that originate from control theory: the reference roll angle is indicated by the r, the control input by u, process noise by v, the input delay by τ_u, measured outputs by y, measurement noise by w, measurement delay by τ_s and the estimated states by $\hat{y}$.

For validation, two different experimental data sets during which cyclists are laterally (roll) perturbed on instrumented bicycles are used: one collected at the UC Davis on both a horse treadmill and in a sports pavilion and the other collected at the TU Delft on a public cycling path (without other road users).

The human steer response to the perturbation is separated from other effects caused by unknown disturbances and noise with a non-parametric Finite-Impulse-Response (FIR) model. A structured system identification approach is used to determine the final model structure components and free parameter set that approximates the non-parametric model in the best way possible. The data sets are split in a train set and a test set. The model structure is fitted to the train set. The predictive performance is verified by applying the resulting model to the test set.

For the UC Davis data, the resulting model has one rider dependent parameter that describes human bicycle-balancing control behaviour for the entire evaluated forward velocity range (2-8 m/s). This parameter is the weight placed on the roll angle in the LQR control algorithm. The single-run training performance in terms of Variance-Accounted-For (VAF) with the non-parametric model reaches up to 97 %. Test performance incurred a VAF drop, on average, of around 10 %.

The steer responses of the TU Delft experiment can be predicted with a model that has two rider dependent parameters: the weights placed on the roll angle and roll rate in the Linear-Quadratic-Regulator control algorithm. Training performance reaches up till 83 % and test performance incurred a VAF drop, on average, of around 7 %.

The major findings are:

- Roll rate is the most important state to be measured. Without a direct roll rate measurement, no results are found that match the experimental data. This roll rate measurement is likely provided by the vestibular organ.
- The resulting models from the UC Davis treadmill and pavilion experiments are similar and can be interchanged at a VAF penalty up to 30 %. For two of the three participants the average VAF penalty ranged between 0.1 and 7 %.
- The resulting models of the UC Davis and TU Delft experiments are different and can't be interchanged. This can be attributed to the different experimental conditions: upper-body movement at the TU Delft experiment was unrestricted, whereas upper-body movement was restricted in the UC Davis experiment.
- Leaving the neural time delay as a free parameter, usually resulted in a fitted delay close to 0s. By adding a simplistic passive rider model this can be remedied. In that case, the resultant fitted time delay ranges between 40 and 100 ms.

- The individual participants to the studies show distinct control behaviour, which is captured with the proposed model.

1

Introduction

Cycling remains an important mode of transport in the Netherlands. Even though the car is the main mode of transport in the Netherlands, 36 % of the Dutch list the bicycle as the mode of transport most often used on a typical day [30]. The average Dutch resident cycles in excess of 3 kilometers daily, which is about 8 % of the total distance covered with any means of transportation [17].

On face value, this love of cycling of the Dutch has only upsides: it is healthy, puts a lower burden on infrastructure capacity and it has a negligible environmental impact. Hence, many countries try to take measures to cultivate a similar cycling culture as the one that exists in the Netherlands.

However, in recent years, the discussion on how healthy cycling actually is has intensified in the Netherlands. If we look at the total number of hospital visits related to traffic accidents, in 70 % of the cases the mode of transport was a bicycle. Furthermore, this number has increased by 31 % (corrected for population growth) over the last ten years, which is 13 % higher than overall trend in the number of hospital visits related to traffic accidents [97].

Up till now, the Dutch government tried the reduce the number of accidents with different measures mainly focusing on improving the infrastructure and implementing additional legislature such as restricting the use of mobile phones while cycling.

There are two interesting statistics present behind the data on hospital admissions:

1. More than 40 % of the victims was more than 55 years old, which is more than 10 % higher than what one would expect based on the amount of distance traveled [17].
2. in 68 % of the cases the accident was a single vehicle accident, meaning that no other road users were involved. When asked about the cause of their accident, almost half of the victims attributed the cause to own behaviour such as a lack of attention or a steering error [97].

These two statistics (the relativity high number of victims aged higher than 55 years old, and the high percentage of single vehicle accidents) open up more trains of thought on where to seek for solutions for the overall problem of the high injury rate among cyclists. Being a (aspiring) mechanical engineer, it makes sense to approach this problem from its origin; the rider of the bicycle, and to try to understand the control behaviour of the bicycle rider with the tools known to mechanical engineers.

There have been numerous attempts to identify the control behaviour of the bicycle rider, both by mechanical and non-mechanical engineers. Two of those attempts ([63] and [79]) included the extensive validation with an experimental data set that is needed to reach conclusive results. The control structure used by [63] finds its origin in pilot control models, while in [79] a classical control approach is used with a PID-like controller.

Both these studies don't include known human characteristics such as neural time delays and they do not model the characteristics and integration of the sensory pathways available to humans. Modeling these things is important as they affect the stability margin of the rider-bicycle system and can help explain if the inflated injury rate among elderly might be caused by cognitive or sensory decline.

As such, the research objective of the master thesis is to develop a model that includes realistic human neural and cognitive characteristics to predict control behaviour for stabilizing a bicycle.

Research questions that are related to the research objective are:

1. What bicycle states does the rider use as feedback for the stabilization task?
2. What sensors does the rider use to measure the states of the bicycle?
3. What is the relative importance of these sensors?

Outline of the thesis

The content of this master thesis is roughly divided in three parts:

1. A **literature review**, where the current-state-of-the-art of bicycle-control modelling (chapter 2) is determined and the theoretical foundation is laid upon which the realistic human-bicycle model can be built (chapter 3).
2. The **theoretical outline of the proposed approach**. In this chapter (4) the procedure which is used to identify the human-bicycle controller is described.
3. **Results**. Here (chapter 5), both the intermediate results and the final identified model are presented.

As an advice to the time-restricted reader, it is possible to largely skip the literature review and only read the parts of interest. In the following chapters the conclusions of the review are re-iterated with references back to the relevant section in the review itself.

2

Current state-of-the-art of human bicycle control modelling

In this chapter the major research contributions to human bicycle control modelling are listed in chronological order. The chapter leans heavily on the review paper by Schwab and Meijaard [80], were the then current state-of-the-art of human bicycle control modelling is evaluated. Afterwards, the available literature is discussed and the ”gap” in the scientific literature on human bicycle control modelling is elaborated on.

2.1. Chronology

1970-1980

The earliest work done on bicycle rider control modelling and identification is done by van Lunteren and Stassen ([54], [87]). They developed a bicycle simulator capable of simulating roll, pedalling, and steering. Furthermore, the body lean angle of the rider was a measured input. They fitted transfer functions with delays to the data. With this method, they found a 74 ms feedback delay on the steer angle and a 26 ms feedback delay on the roll angle. For his master thesis research, de Lange [22] implemented the transfer functions and found the result to be stable without delays (after correcting some sign errors), but unstable with time delays.

1980-1990

In the eighties, Doyle [27] investigated bicycle riding in a broad sense. He did experiments at different forward velocities, with blinded and non-blinded participants, and both on self-stable and destabilized bicycles[1]. Both based on literature and observations during experiments done with an instrumented bicycle, he proposed a control algorithm describing human bicycle control. The control law provides continuous feedback based on roll rate and roll acceleration, in combination with intermittent roll angle feedback. The control feedback is organised in nested loops with feedback to lateral displacement in the outer loop, feedback to the heading angle in the middle loop, and feedback to the roll rate in the inner loop. The inner loop ensures the roll stability of the system. He investigated feedback to the roll angle, but concluded that there is no evidence to support the proposition that the angle term in the algorithm is used continuously in order to provide steer angle feedback. Instead, when the roll angle accumulates beyond a certain threshold, a (discrete) fairly strong pulse is applied to the handlebar in order to bring the roll angle to zero.

Doyle found early in his research that depriving a rider of vision did not have an influence on riding ability. Therefore, his experiments were done with the participants being blindfolded. In order to make sure that bicycle stabilization was solely achieved with control actions from the bicycle rider, he tried to remove any self stabilizing effects of the bicycle.

---

[1] He validated this by comparing the open-loop behaviour of both bicycles, but his destabilized bicycle probably still has some self-stable characteristics

He concluded by inferring a delay between 60 and 120 ms (different between subjects, but constant for each subject) between the roll rate and steer angle, that the output from the vestibular system must go almost directly to the controlling muscles, making little or no demand on higher cognitive processes.

Doyle [27] also included neural delays in his model, and investigated the stability region as a function of input delay, input gain, forward velocity, and rider height and weight characteristics. Generally speaking, higher delays and higher gains destabilize dynamical systems. This also applies to the rider-bicycle system. Doyle [27] found that the gain margin did not change significantly as a function of forward velocity. However, at lower velocities the stability margin is lower. The reason for this is, that due to the nature of bicycle dynamics, a higher input gain is required at lower speeds to counter a given lean angle. Therefore, disturbance rejection becomes more difficult at lower forward velocities. He found that, by including a delay of 120 ms in the model, the results were comparable with the gathered experimental data.

He concluded that upper-body lean only played a minor part in bicycle stabilization, because even at small bicycle lean angles, upper-body movements cannot bring the centre of mass onto the correcting side of the support point.

1990-2000

For his doctoral thesis, Getz [34] developed an internal equilibrium controller. Dynamic inversion is incorporated into this controller to provide a continuous estimate of the internal equilibrium manifold location. The internal equilibrium controller makes a neighbourhood of the manifold attractive and invariant. This results in approximate tracking of time parametrized paths while retaining balance. Getz applied this method to path-tracking with balancing for a bicycle in [33]. He uses steer angle and rear-wheel torque as inputs. The downside of this method, is that it is strictly path based. For determining rider control actions based on bicycle-rider states, this method is probably less suited [80].

2000-2010

Cook [20] wrote an engaging article where he presents a two-neuron network that can path-track and balance a bicycle. The first neuron outputs a desired lean angle, dependent on the current and desired heading angles. The second neuron outputs the desired steer torque, with as inputs the lean angle and rate, the desired lean angle and its own output. The controller works well for a range of forward velocities, but did not dampen the instabilities that occur in sharp turns or at low forward velocities.

Sharp [81] implemented a Linear-Quadratic-Regulator (LQR) controller with preview in order to model human stabilization and path-tracking of a bicycle. Sharp [81] found that tight control required about 2.5 seconds of preview time, whereas loose control extended that requirement without any real limit. The control tightness was determined by the relative weightings of tracking error and control power in the cost function of the optimal control computations. Based on literature, Sharp [81] considers steer torque to be the primary control input. Reviewing that literature, which was primarily focussed on motorcycles, two arguments were found. The first being that due to the fact that handlebar angles usually are small, handlebar position control is difficult [83]. The other argument [52] is that steer torque is a control input to the (assumed to be correct [59]) Whipple bicycle model [100]. It must be noted that all the references given in the article by Sharp [81] on this subject are circle citations.

Both an optimal LQR controller and a, so-called, intuitive controller were presented in the article by Schwab et al. [78]. The LQR controller minimizes a cost function were the objective (or state) is weighed against the control cost. The intuitive controller steers the bicycle into the undesired fall using a simple control law. At low speeds, feedback proportional to the roll rate, decreasing with increasing forward velocity, is applied. Above the stable forward velocity range, proportional feedback of the roll angle, increasing with with forward velocity, is applied. Both controllers used steer torque as in input to a linear bicycle model. The rider is attached rigidly to the bicycle frame. Additionally, both control algorithms were extended in order to apply rider upper body control as an additional roll torque input. The authors found no significant difference in the dynamics between the rigid rider configuration and the added upper body control configuration. Furthermore, they deemed the feedback gains found with the LQR controller to be unrealistic because they were outside the range of human capabilities.

2010-2019

Hess et al. [39] developed a bicycle rider model based on pilot models developed in the aerospace research field. The resulting model is parametrized by five gains, two fixed second-order filters, and a preview time. The gains are on the bicycle states roll angle and rate, steer angle, heading angle and lateral displacement. One second-order filter acts on the rider output and as such represents the neuromuscular dynamics. The other on the handling quality performance criterion in order to reduce the magnitudes at frequencies where there would be little power in the visual feedback signal.

The model has multiple nested feedback loops, each for different states of the bicycle model. These loops were associated with sensory mechanisms present in the rider, and are organised hierarchically in a way that low-level control tasks such as balancing take preference above higher-level control tasks such as path-tracking. While the loops were associated with sensory mechanisms, no sensory dynamics or delays were included in the model. Moore [63] gathered experimental data using an instrumented bicycle, and applied grey-box system identification techniques to identify the parameters of a reduced Hess et al. [39] model, where path following is omitted as an control objective. The fundamental, remnant-free, control response of the rigid-rider under lateral roll perturbations can be described reasonably well by this model.

A control structure in which the corrective control inputs (steer angle and upper-body lean torque) are generated proportionally from four different mechanisms is presented in the article by Soudbakhsh et al. [85]. Three of these mechanism provide reflexive control inputs with increasing delays based on proprioceptive feedback. The other mechanism (dubbed the neural controller) integrates vestibular, visual, and proprioceptive signals using static signal weights. Control action is provided by a PID controller. They combined this model with the model of a stationary bicycle on rollers and did an stability analysis. They showed that only upper-body lean is insufficient to stabilize the bicycle. The stability of the bicycle-rider system was found to be greatly affected by variations of the different time delays present in the model. Based on the stability analysis and literature, they set the short latency delay to 21 ms, the medium latency delay to 131 ms and the long latency (neural controller) delay to 288 ms.

Using experimental data gathered at different forward velocities on a narrow treadmill while laterally perturbing the bicycle frame with impulsive forces [63], Schwab et al. [79] presented a rider control model with muscle dynamics and with steer torque as the control input. The reduced control structure, which stabilizes the bicycle, has four gains: a gain on, respectively, the lean angle and rate, a gain on the steer rate, and a gain on the integral of the steer angle. The authors concluded that the gains on the lean angle and rate corresponds with vestibular and/or visual feedback, and the gain on the steer rate corresponds with proprioceptive feedback. The gain on the integral of the steer angle is explained to be similar to a gain on the heading angle, and was deemed to be caused by the narrow treadmill. Time delays were not included. To understand what the rider was optimizing, a LQR method was applied in an inverse manner, resulting in the weight factors (instead of feedback gains). These weight factors indicated that the rider is minimizing control effort at forward velocities below 3 m/s. At higher forward velocities the rider is minimizing heading error.

Wang and Yi [99] presented a human balancing control model for stabilizing a bicycle using two proportional-derivative controllers with time delays acting on the roll angle. One providing steer angle as a control input, and the other upper body movement torque in order for stabilizing the bicycle. A time delay of 235 ms for steer input and a time delay of 110 ms for upper body lean input are used. The time delays are based on literature on human spinal stabilization. They validated their model with experiments done on an instrumented bicycle. While, to the eye, their model matches the results quite decently, no perturbations were given during the experiments. What constituted as perturbations to the controller, remains vague.

2.2. Discussion on the current state-of-the-art

Of the models in the literature that are evaluated, all are able to balance a bicycle. Additionally, a range of models (Getz [33], Cook [20], Sharp [81]), Hess et al. [39]) implemented path tracking as an additional objective. The research by Doyle [27] is a bit difficult to place here, because the modelling structure he proposes includes higher level tasks, but only the balancing part seems to be actually modelled and validated.

A big differentiator amongst the presented models are the control inputs used. Steer torque is used most often (Doyle [27], Cook [20], Sharp [81], Schwab et al. [78], Hess et al. [39], Schwab et al. [79]). Steer angle as a control input is used less often (Lunteren and Stassen [54], Getz [33], Soudbakhsh et al. [85], Wang and Yi [99]). Based on the literature found during this review, no strong case can be made for either steer torque or steer angle as an input. A number of models included the torque generated by upper-body lean as an control input (Lunteren and Stassen [54], Schwab et al. [78], Soudbakhsh et al. [85], and Wang and Yi [99]). However, multiple arguments ([27], [78], [85]) are made that upper-body lean does not provide a significant contribution to the control of a bicycle.

The control strategies used can be roughly divided in four different categories: nested multi-loop, classical state feedback, optimal full state feedback, and other control type controllers. Nested multi-loop controllers (an example is given in figure 2.1 [39]) are presented by Doyle [27] and Hess et al. [39]. Nested multi-loop controllers find their origin in the research done by McRuer and Jex [58], and are applied often for the modelling of aircraft pilots. Hence, they are proven to be able to provide realistic results when used to model human control behaviour. Classical state feedback controllers give feedback on the output of models. Provided that riders can sense all the outputs of the bicycle model used, this could be a realistic method of modelling bicycle rider control behaviour. Classical feedback control was used in the articles by Lunteren and Stassen [54], Schwab et al. [78], Schwab et al. [79] and Wang and Yi [99]. Lastly, Getz [33] and Cook [20] used control methods that aren't covered in the previous categories. In the review paper by [80], the usefulness for determining rider control actions based on bicycle-rider states was questioned. The controller developed by Cook [20] is developed in order to understand the human learning process more. But the author concludes that more research was necessary, as unsupervised learning yielded unsatisfactory results.

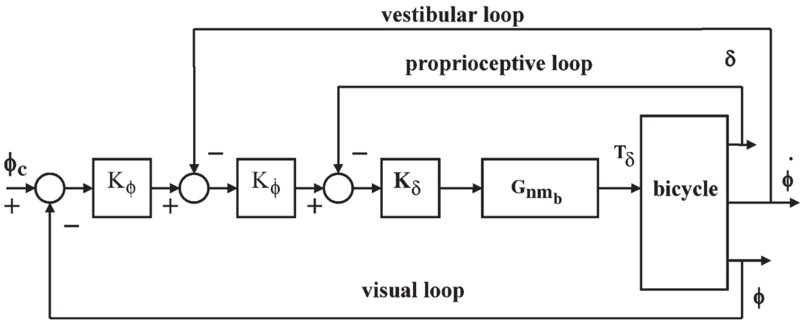

Figure 2.1: Nested multi-loop controller. Here the feedback is given hierarchically: proprioceptive bicycle steer angle feedback takes preference above vestibular bicycle roll-rate feedback, which again has preference above visual bicycle roll feedback (source: [39]).

None of the models incorporate sensory dynamics, and only a few models incorporate neural delays (Lunteren and Stassen [54], Doyle [27], Soudbakhsh et al. [85], and Wang and Yi [99]). In only one model, preview time for path tracking, is implemented (Sharp [81]), and only one model was suitable for non-linear control of non-linear bicycle plants (Getz [33]).

Of the models presented in literature, only a select few were validated with, or based on experimental data of humans controlling a bicycle (Doyle [27], Moore [63], Schwab et al. [79], Wang and Yi [99]). None of these experimentally validated models use an optimal control approach, and, as mentioned before, only the model presented by Wang and Yi [99] incorporated time delays in their complete model structure.

For the benefit of clarity, the sources are categorised in table 2.1[2].

2.3. Avenues for the advancement of the current state-of-the-art

In the previous section, the characteristics of the current bicycle-rider models are discussed. When these models are compared to other contemporary human control models (e.g. [94] for human stance control, or [66] for human driver steering control), it is clear that the modelling of the human sensory system and the integration of sensor signals can be expanded on. Doing this could give a clearer picture on the relative importance of the different sensor signals, and can result in more realistic control behaviour. Furthermore, certain insights and contributions of the different authors can be combined. A good example is the way how Soudbakhsh et al. [85] modelled the reflexive feedback, or how Sharp [81] implemented a path-tracking preview time. Moore [63] made the dataset he gathered using his

---

[2] Inspired in the table by Kooijman [51] in his PhD thesis

Table 2.1: Bicycle-rider models found in literature. An asterisk (*) indicates an experimentally validated model.

Model characteristic	Balancing	Path following
No neural delays	Schwab et al. [78], Moore [63]*, Schwab et al. [79]*	Getz [34], Cook [20], Sharp [81], Hess et al. [39]
Neural delays	Lunteren and Stassen [54]*, Doyle [27]*, Soudbakhsh et al. [85]	Wang and Yi [99]*
Steer torque	Doyle [27]*, Schwab et al. [78], Moore [63]*, Schwab et al. [79]*	Cook [20], Sharp [81], Hess et al. [39]
Steer angle	Lunteren and Stassen [54]*, Soudbakhsh et al. [85]	Getz [34], Wang and Yi [99]*
Body lean	Lunteren and Stassen [54]*, Schwab et al. [78]	Wang and Yi [99]*

instrumented bicycle publicly available, opening opportunities to validate and improve the controllers that have not yet been validated with experimental data.

3

Theoretical foundation for bicycle-rider control modelling

In this chapter a broad literature review is conducted in order to provide a solid theoretical foundation, that can be used to model a sufficiently complete and realistic human-bicycle system from a control theory perspective. This foundation can subsequently be used for the purpose of human control identification.

To illustrate the complexity of human control identification, picture an ordinary control engineer. When this control engineer is tasked to design a controller to control a certain system, the engineer quickly sets about to define the system in terms of inputs and outputs. Sensors are selected to measure outputs, and actuators are selected to give corrective inputs. When the underlying dynamics of the system are not fully understood, a common method is to give known inputs to the system, and based on the measured outputs, characterize a system which predicts the transfer from input to output.

When a human does control, the process is similar. Let's take the relevant example of riding a bicycle. People are often not capable of outright stabilizing a bicycle when riding one for the first time. Also, trying to learn by looking at someone else riding a bicycle is not of much use. As is asking a experienced cyclist. The only option left to learn how, is to hop on one, give some inputs, and by learning from the response, develop an model of the dynamics of the bicycle which enables us to ride one. This process happens subconsciously for the most part.

This is a problem: we can't simply ask someone how he controls a bicycle. But, still we want to understand how. A method of achieving this is the break the problem up into smaller parts and sub-processes that we do understand. These can be organised in such a way that together they form a complete picture of the human-bicycle system we are looking for.

This literature review is partitioned in sections according to these processes and components. The processes and components are organized in a general feedback structure. How information flows between them is illustrated in figure 3.1. They can be separated in two groups: the Human-Bicycle plant, which describes the ”passive” behaviour of the rider-bicycle combination. And the ”active” part, were cognition and control takes place.

Contemporary research on human control identification, primarily in the aviation and, more recent, in the automotive industry, focusses on the active part. For human bicycle control this is not necessary valid, because the rider is a big part of the system. Simply by sitting on a bicycle or holding the handlebars, the rider could passively change the dynamics of the system.

The passive part also includes the dynamics of the actuators and sensors the human has access to. Next to the ability of our brain to do control, the human actuation and sensory system will define the limits of human control.

In each section the possibilities are explored, and their suitability for bicycle-rider identification is evaluated.

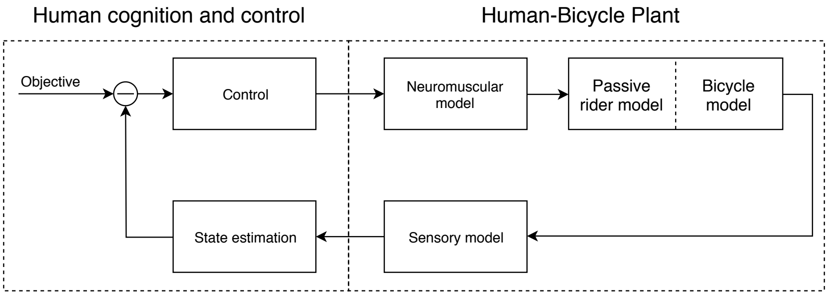

Figure 3.1: Basic control structure.

3.1. Bicycle model

According to the Oxford Dictionary, a bicycle is a ”vehicle consisting of two wheels held in a frame one behind the other, propelled by pedals and steered with handlebars attached to the front wheel”. A bicycle model is a set of equations that describe the dynamic behaviour of this vehicle. The number and complexity of these equations vary depending on the assumptions they are based on, the modelling methods used, and the type of bicycle to be modelled.

An interesting characteristic of a lot of bicycles, is that they are self-stable for some conditions (mainly bicycle geometry and velocity dependent). With self-stable is meant that an uncontrolled bike does not fall over when exposed to limited perturbations: it ”self-corrects”. An indication if a bicycle model describes the dynamics of a bicycle correctly, is that the model predicts self-stable behaviour.

Therefore, we can already discern two categories of bicycle models: those that can, and those that can not predict self-stable behaviour. Extended bicycle models than can predict self-stable behaviour, but that don’t assume rigid bodies and ideal ”knife-edge” wheel surface contacts, are placed in a third category.

1. Linear and non-linear bicycle models that don’t describe self-stability correctly.
2. Linear and non-linear rigid-body bicycle models that describe self-stability.
3. Bicycle models that extend the models in category 2 with any type of complexities such as non-ideal road contact, frame deformations, non-level roads, and so on.

This section on bicycle models is mainly based on four articles which themselves contain literature reviews: the article by Limebeer and Sharp [52], the article by Meijaard et al. [59], the article by Basu-Mandal et al. [6] and the review study by Schwab and Meijaard [80].

3.1.1 Linear and non-linear rigid-body bicycle models that don’t describe self-stability correctly

Bicycle models in this category are not necessary suitable to use as a plant model, as they describe the stability characteristics of the bicycle incorrectly. However, they can be used as an internal model[1].

Meijaard et al. [59] identifies three causes as why models in this category are too limited to describe self-stability: simplified geometry and/or mass distribution, no steer dynamics because steer is fully controlled by the rider, or mathematically simplified models.

One of the, if not the first, scientific publication which includes a mathematical description of a bicycle, is a series of articles by Carlo Bourlet ([8], [9], [10], [11]). These only include a lean equation (and therefore don’t describe self-stable behaviour), the influence of the gyroscopic moment from steer rate is omitted, and the model has front-contact geometry issues in the non-linear equations described in detail in [59]. But other than that, Meijaard et al. [59] concluded that the final linearised lean equation is correct.

Another type of models in this category are non-linear and linear point-mass bicycle models with non-holonomic constraints. The origin of these interesting, but simple models is attributed by Meijaard to two articles from Boussinesq ([12], [13]). In the book on dynamics by Timoshenko and Young

---

[1] Internal models are introduced in section 3.4

[89] the Boussinesq model is presented, and in another (easy accessible) literature review paper by Limebeer and Sharp [52], an analysis of a Boussinesq like model is included. The non-linear and the linearised constant velocity roll equations of that model are presented in equations 3.1 and equation 3.2 respectively. The definition of the parameters and variables used is visualised in figure 3.2 [52].

$$h\ddot{\phi} = gsin\phi - tan\delta\left(\frac{v^2}{w} + \frac{b\dot{v}}{w} + tan\phi\left(\frac{vb}{w}\dot{\phi} - \frac{hv^2}{w^2}tan\delta\right)\right) - \frac{bv\dot{\delta}}{wcos^2\delta} \tag{3.1}$$

$$\ddot{\phi} = \frac{g}{h}\phi - \frac{v^2}{hw}\delta - \frac{bv}{wh}\dot{\delta} \tag{3.2}$$

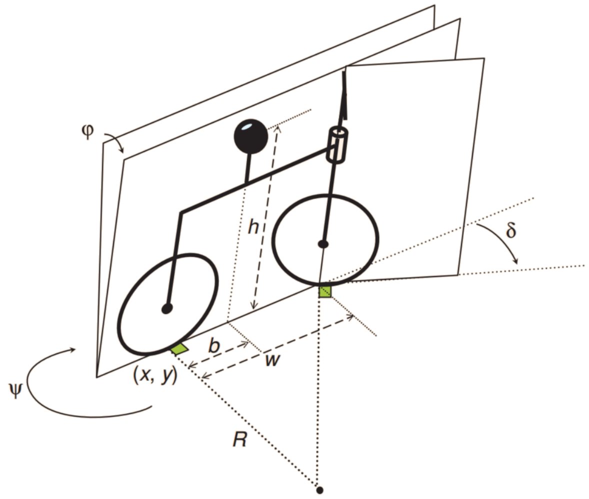

Figure 3.2: Inverted pendulum bicycle model (source: [52]).

3.1.2 Linear and non-linear rigid-body bicycle models that describe self-stability

In this category are models which resemble the so-called Whipple model, of which the linear equations are proven to be correct by Meijaard et al. [59].

The Whipple model was actually derived by two persons, seemingly independent from each other, in the same time period. According to Meijaard, Emmanuel Carvallo was the first who described the dynamics of an uncontrolled bicycle correctly. He identified the four standard eigenmodes, and presented the equations for the lower and upper velocity bounds of the self-stable region.

Slightly later than Carvallo, Francis Whipple derived the correct linear equations of motion for a slightly more elaborated bicycle model than Carvallo, who neglected the mass and moments of inertia of the front frame.

The model consists of four rigid bodies: two wheels, the rear frame, and the front frame. The wheels are connected to the rear frame and front frame through ideal pinned connections that constrict all forces, except moments in the rotation direction of the wheel. The wheels have ideal knife-edge rolling point contacts with the ground. The frame and front assembly themselves are connected through an ideal hinge which defines the steering axis. The four parts of the model and important model parameters are shown in figure 3.3 [59].

Over the last century a lot of linearised bicycle models are developed in the aforementioned framework, some correct and some incorrect. Widely different methods and coordinate systems are used, leading the different sets of equations. These sets of equations are hard to compare, except through eigenvalue analysis. Meijaard et al. [59] presented the ”canonical” linearised equations which describe the lateral dynamics for the Whipple bicycle model. These equations can be considered as canonical because two different methods to derive them are used and they also are checked with two non-linear dynamics simulations. Furthermore, Meijaard provides two benchmark bicycles which can be used for test cases. Finally, the model is experimentally validated by Kooijman et al. [48], and is shown to be in fairly good agreement in the forward velocity range of 0 to 6 m/s.

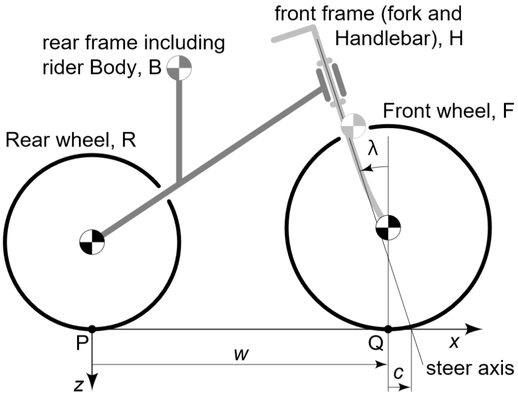

Figure 3.3: Whipple model naming and parameter convention (source: [59]).

The equations can be written in the same form as general mechanical systems: the forces, or torques in this case, are proportional to the inertia terms multiplied with accelerations, the damping terms multiplied with velocities and the stiffness terms multiplied with displacements. Hence, they can be written as:

$$\mathbf{M}\ddot{\mathbf{q}} + v\mathbf{C}_1\dot{\mathbf{q}} + [g\mathbf{K}_0 + v^2\mathbf{K}_2]\mathbf{q} = \mathbf{T} \tag{3.3}$$

Where $\mathbf{q} = [\phi, \delta]^\mathsf{T}$ and $\mathbf{T} = [T_\phi, T_\delta]^\mathsf{T}$. The constant terms ($\mathbf{M}$, $\mathbf{C}_1$, $\mathbf{K}_0$, $\mathbf{K}_2$) follow from the parameters of the bicycle. $\mathbf{C}_1$ is linear with forward velocity, $\mathbf{K}_2$ is quadratic with forward velocity v, and $\mathbf{K}_0$ is proportional to the gravitational acceleration g.

This linear model has four different modes, of which two are a complex oscillatory eigenvalue pair. The modes can be found through an eigenvalue analysis and are called capsize, weave and castering. The capsize mode can be envisioned as the bicycle falling over. The weave mode is an oscillatory motion of steering lagged by leaning. The last mode is the castering mode, which is always stable if the forward velocity is positive. Its motion can be compared to the straightening of a wheel on a shopping cart.

In the paper by Basu-Mandal et al. [6] the literature on bicycle models that describe circular motions is reviewed. That literature has roughly the same pattern as Meijaard found: a lot of the research done is either partly complete, nor correct or cross-checked. As a consequence, Basu-Mandal wrote a complete description of hands-free circular motions of an idealized benchmark bicycle (the same one as Meijaard used) using two different methods (Lagrange and Newton-Euler). These two sets of equations are both checked with each other, and match with machine precision. The linearised Lagrange equations are checked with the equations as derived by Meijaard et al. [59]. The eigenvalues obtained from the two systems matched to 14 decimal spaces.

The set of equations are quite extensive (34 equations for the Newton-Euler derived model) or very long (the Matlab m-file is 3.5 MB for the Lagrange's equations of motion), and are therefore not included in this literature review. They can be found in the electronic supplementary of the paper by Basu-Mandal et al. [6].

Turnwald and Liu [90] derived a non-linear bicycle model in the Lagrangian and Hamiltonian framework, yielding a pseudo-Hamiltonian system that also describes self-stable behaviour. This model is then simplified in structure preserving manner with constant tensors. In this way passivity-based controller and observers can be designed for this non-linear bicycle model. This model was also compared to the Meijaard benchmark bicycle, to give an indication of validity.

3.1.3 Extended bicycle models

Category 3 models are models that extend upon the previous category bicycle models. A first subcategory amongst these extensions is the modelling of ever more complex tyre behaviour.

A simple approach is to model toroidal shaped wheels instead of the knife-edge contacts used in the models in the previous two categories. This can both be done as an extension to linear [60] and to non-linear [98] models. Meijaard and Schwab [60] also considers effects of pneumatic trail and damping due normal spin at the tyre contact patch.

The toroidal shape reduces the arm of the gravity force, and thereby slows down the dynamics, easing the balance task. The added tyre contact force effects did not have a big influence on the dynamics of a bicycle [80].

Complex tyre models describe tyres that generate forces dependent on normal force, tyre slip and camber with time lags dependent on wheel velocity and tyre relaxation length. These complex tyre models can be roughly separated in two categories: physics based models and empirical models. For bicycle tyres the choice is quite limited and we have to rely on empirical formula's. Schwab and Meijaard [80] give a nice overview of the research done on tyre modelling until then. They cite four sources for the gathered experimental data sets. In the seventies Roland and Lynch [75] measured side-slip and camber forces. In the last two decades Cole and Khoo [19] measured side-slip forces, Dressel and Rahman [28] measured side-slip and camber forces, and so did Doria et al. [26]. Dressel and Rahman [28] also measured tyre relaxation lengths, which was found to be around 0.1m.

The data sets were gathered with different measurement methods, but all measurements were done dependent on tyre load and tyre inflation pressure. They have the same correlations between them: with increasing loads, decreasing coefficients and with increasing inflation pressure, increasing coefficients. Also, side slip coefficients are found to be in rough agreement with each other and where in the range of 0.8-1.5. This can not be said for the camber coefficients. Roland and Lynch [75] found camber force coefficients in the range 0.15-0.6, while Doria et al. [26] found coefficients between 1.3-1.5.

Sharp [82] developed a mathematical model in order to show typical stability characteristics and how these depend on various parameters values. Sharp found that including a side slip tyre model introduces a new mode: wobble. Wobble is a fast oscillation of the front wheel around the steering axis. He found that the natural frequency of wobble to be almost independent of forward velocity and well damped at low and medium velocities (till approximately 5 m/s). At higher velocities the damping drops rapidly. The wobble mode is primarily affected by the steer damping coefficient and the tyre relaxation length. Sharp [82] found that weave stability decreases when a tyre slip is modelled.

Bulsink et al. [16] developed an open-loop bicycle-rider model in the commercial multibody dynamics software ADAMS. They analysed the effect of tyre and rider properties on bicycle stability. The tyres are modelled based on the data gathered by Doria et al. [26]. Just like Sharp [82], they found that extending simplified models with a realistic tyre model leads to a notable decrease in the weave stability and a stabilization of the capsize mode. Furthermore they found that different tyre types and tyre inflation pressures had little effect on the bicycle's stability, in the case of riding straight at a constant forward velocity. Tyre load did have a large effect on bicycle stability.

Frame compliance can also be modelled. In the context of bicycles this is done by adding an extra degree of freedom in the connection between the rear and front frames which allows lateral bending [69]. Frame compliance is then simulated by a spring and damper that act on that displacement. Including frame stiffness introduces, just like adding tyre compliance, the wobble mode. Therefore the stability of a non-rigid frame bicycle is reduced when compared to a rigid bicycle. Plöchl et al. [69] also tried to validate the inclusion of frame compliance and lateral-slipping tyres to the linear Whipple model, and found some agreement. However, the results depended greatly on rider posture, and results indicated that riders were able to damp out, or prevent, wobble easily.

Finally, [60] extended the linear Whipple model to include cycling over a road with a gradient, aerodynamic effects and driving and braking torques.

3.1.4 Discussion on the different bicycle models

Both validity and usability are important criteria when selecting a model. Considering validity, Kooijman et al. [48] found that the linear Whipple bicycle model, as derived by Meijaard et al. [59], agreed with reality up till a forward velocity of 6 m/s, but was difficult to validate at forward velocities below 3 m/s. This indicates that extensions to that model are unnecessary within that velocity range. At least for cycling on level surfaces, with no additional wind effects, with bicycles that are similar to the benchmark bicycle and for small roll and steer angles. This is proven further by Sharp [82], who found that the unstable wobble mode can be caused by tyre dynamics, but is well damped below velocities of 5 m/s. Plöchl et al. [69] found that that riders can prevent or damp out any wobble oscillations quite easily. At a higher velocity range, a lot of modelling possibilities are available. Unfortunately, validation is lacking.

Usability is a bit of a vague term that indicates if the model is useful in a system identification and control context. Here it is true that most system identification and control methods are developed for Linear-Time-Invariant (LTI) systems. For any model that falls in that category, it will be easier to develop tailored control and identification algorithms for.

Reviewing the experimental data sets which are to be used for this research, the assumptions that tie in with the use of the Whipple bicycle model are valid. Namely, the experiments are carried out with normal city bikes, on level ground, without (to much interference) from the environment and the measured angles rarely exceed 15 degrees (approximately the point where truncation errors start to exceed 1%). Also, all the experiments are carried out at constant forward velocities, so a model which only describes the lateral dynamics is sufficient. Therefore, the Whipple model is the best candidate to use as a bicycle model.

3.2. Passive rider model

Passive rider models describe behaviour that is not active control behaviour, but does effect the dynamics of the bicycle model. This includes control strategies that bypass the usual sensor-control loop, and by doing that have little to zero delay. Examples of this are posture and muscle co-activation.

Schwab et al. [77] added a passive rider model to the linear Whipple model, without adding any degrees of freedom. Therefore, only the indices of the constant tensors change. They evaluated two different postures: a forward leaning posture and a upright posture (see figure 3.4, source: [77]). In the case of a forward leaned rider posture with stretched arms and hands on the handlebar, the open-loop dynamics only change marginally. However, an upright rider posture with flexed arms and hands on the handlebar cause the open-loop dynamics to become unstable.

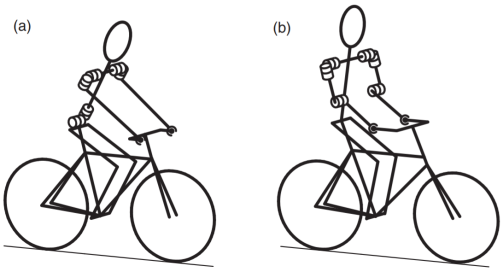

Figure 3.4: Rider postures: (a) forward leaning and (b) upright (source: [77]).

A remark to the study by Schwab et al. [77], is that they modelled the rider body as a system of cylinders connected with ideal joints. Therefore, any damping and stiffness effects caused by, for example, the muscle spindles are not included.

In his PhD dissertation [63], Moore presents an arm model similar to the one proposed by Schwab et al. [77]. He found it to perform better than the regular Whipple bicycle model.

In a more recent article by Doria and Tognazzo [25], these issues are addressed. They identified stiffness and damping parameters that are relevant for cycling and lumped them in a biomechanical model using experimental results. This model was then coupled with the benchmark bicycle model. They found that uncontrolled damping reduces the self-stability region and that uncontrolled arm stiffness causes the open-loop self stable behaviour of the benchmark bicycle to disappear entirely.

The study done by Schouten et al. [76] provides an outlook on modelling passive human behaviour that is based on physiological characteristics of the human body. They investigated the dynamic behaviour of the neuromusculoskeletal system. Their model consists of transfer functions describing input-output behaviour and includes afferent feedback of the muscle spindles and Golgi tendon organs, as well as muscle visco-elasticity. The muscle visco-elasticity is of primary interest here, because it provides a passive contribution to the dynamic behaviour of the neuromusculoskeletal system. The afferent feedback of the muscle spindles and Golgi tendon organs less so, because this tends to be control guided behaviour. They fitted parameters to the model, using an experimental set-up, where the human was force-perturbed with a manipulator held in the right hand.

3.3. Sensory model

The purpose of the sensory model is the supply the Central-Nervous-System (CNS), where cognition and control takes place, with information. In the context of the system overview given in figure 3.1, the sensory model consist of all the sensors a human has access to, that provide the information necessary to stabilize a bicycle. It also consist of the neurons and neural pathways connecting the sensors to the CNS, and any organs that post-process the neural data in order to make it suitable for state estimation.

In this chapter, a short overview of the sensor types, their sensitivity, and the neurons that transmit the information is given. After that, the sensors that are relevant to bicycling are presented. Finally, the perception windows and dynamics of the relevant sensors are quantified.

A sensor can be placed in one of three different categories: exteroreceptors, interoreceptors and homeostasis. The exteroreceptors give information of the state of the outside world. Think of gathering information on the weather, but also of the velocity of an approaching car. Exteroreceptors also give information of the state of the human in relation to the outside world. So, our own orientation and velocity are measured with respect to our surroundings.

Interoreceptors gather information on the posture and movement of the human body. Practically speaking, they tell us the whereabouts of our limbs. The accelerations acting on us are also measured by interoreceptors (by the vestibular organ in this case).

The final sensor subset is homoeostasis. They supply internal control loops, of which we are largely unaware, with sensory information. For example, to manage the pH level in the stomach.

A research question is what sensors are used and their relative importance. In table 3.1, the states that can exist in any bicycle model are summarised. In table 3.2, all the sensors organs which can measure these states are listed.

Table 3.1: Bicycle states.

Name	Symbols
Bicycle roll angle, rate and acceleration	ϕ, $\dot{\phi}$, $\ddot{\phi}$
Bicycle steer angle, rate and acceleration	δ, $\dot{\delta}$, $\ddot{\delta}$
Bicycle forward position, velocity and acceleration	x, v, a
Bicycle lateral position, velocity and acceleration	y, $\dot{y}$, $\ddot{y}$
Bicycle heading angle, rate and acceleration	ψ, $\dot{\psi}$, $\ddot{\psi}$
Roll torque	T_ϕ
Steer torque	T_δ

Table 3.2: Sensors relevant for bicycle control.

Type	Name	Measured states	Symbols
Exteroreceptors	Vision	Position, velocity, heading	ϕ, $\dot{\phi}$, δ, $\dot{\delta}$, x, v, y, $\dot{y}$, ψ, $\dot{\psi}$
Interoreceptors	Vestibular - Semi circular channels	Roll rate, heading rate	$\ddot{\phi}$, $\ddot{\psi}$
	Vestibular - Otoliths	Linear accelerations	a, $\ddot{y}$
	Muscle sensors - Golgi tendon organs	Muscle force	T_ϕ, T_δ
	Muscle sensors - Muscle spindles	Muscle length and velocity	ϕ, $\dot{\phi}$, δ, $\dot{\delta}$

The sensitivity of the different sensor organs is dependent on three different things: intensity, frequency, and the duration of the signal. Intensity is the magnitude of the signal. E.g., brightness or loudness. If the intensity is too low, the sensor won't perceive anything. Too high, and the sensor gets damaged. The frequency of the signal has to be in a perceivable range. Think of the wavelength of light, or the frequency of sounds. Lastly, the duration indicates how long a signal has to be measured in order for it to be perceived. After the maximum duration for a certain intensity, the sensor gets damaged. These things can be viewed as thresholds and together can be thought of as a perception window, with each threshold range on a different axis. If the signal is within the bounds of the perception window, we can measure it. However, it is still possible that we don't perceive a certain signal, because our attention is directed somewhere else, or limited by, for example, sleep deprivation. Generally speaking, the size of our perception window decreases with age. It is also possible to influence the size of our perception window. An example is a infra-camera in order to increase our frequency range for vision. Painkillers would lower the intensity range of our pain receptors. The perception window is visualized in figure 3.5.

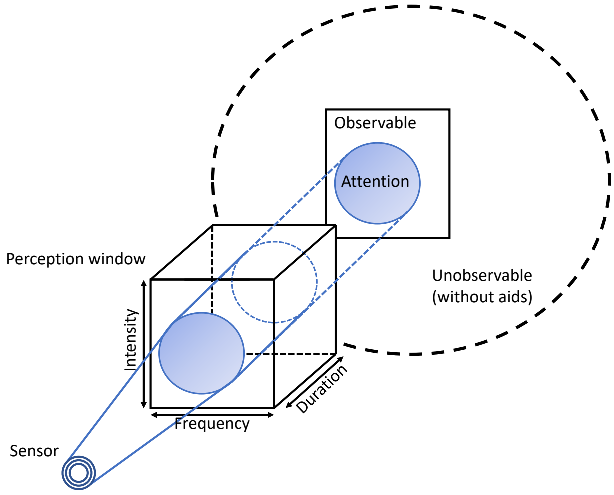

Figure 3.5: Sensory system perception window.

The collected information is transmitted from the sensors to the CNS through the afferent nerve fibres by means of ”spikes”. A spike is generated by the triggering of a sensor or another neuron. In simple terms, it is a short reversal of the potential of a nerve cell. The information is frequency encoded, this means that the number of spikes per second is an indication of sensor activation. The spikes propagate along the nerve cells with a neural transmission speed as low as 2 m/s and up to 100 m/s. Considering that the longest nerves can have a length in excess of 1 m, the neural transmission speed can introduce significant delays to the system.

In order to implement these sensors in our model, we need a description of the transfer between the to be perceived signal and the afferent neural response to the CNS. Also, we need to know the perception window and delay characteristics of each sensor as they can influence the control behaviour. For each sensor listed in table 3.2, such a representation is given in the subsequent sections.

3.3.1 Visual system

Table 3.2 shows that the visual system is the most versatile sensor, as it can measure the most states. In the context of the human control structure used in this thesis (figure 3.1), the visual system is considered to be the entire pipeline from eyes to processed data (i.e., position, velocity and heading information) which can be used by CNS. Therefore, it consists of the eyes, the parts of the brain that processes the information, and the nerves connecting them with each other and to the CNS. An in-depth neurological description of the visual system can be found in the book by Purves et al. [72]. Here, we focus on the perception mechanism and the associated dynamics, thresholds and delays.

The exact mechanism(s) humans use to derive information using the visual system is not fully understood yet. Generally, it is accepted that we use ”visual-flow” cues of some sort to infer self-motion.

This both takes place at the central visual field, where our attention is directed, or at the surrounding peripheral visual field. In the central visual field precise information of the motion of objects is attained, whereas in the peripheral visual field, more general (non-cognitive) information on ego-motion is gathered [14].

One of the first (if not the first) source to describe visual flows and how humans could use this, is found in the book by Gibson [35]. Gibson notes that towards the end of World War II[2] the hypothesis was formed, which still to this day has implications. In simple words, Gibson phrased it like this: ”the visual space should be conceived not as an object or an array of objects, but as a continuous surface or an array of adjoining surfaces”[35]. In his book, Gibson drew visual flow-fields humans could use to derive depth and motion. However, they lacked general mathematical descriptions which were provided by a great number of papers sparked by the book by Gibson, and more recent, by the developments for autonomous driving. A fairly recent review of the methods developed to estimate ego-motion using a visual observer until then is given by Raudies and Neumann [73]. A nice visual representation of visual flow under the influence of ego-motion can be found in the article by Britten [15], which can be seen here in figure 3.6.

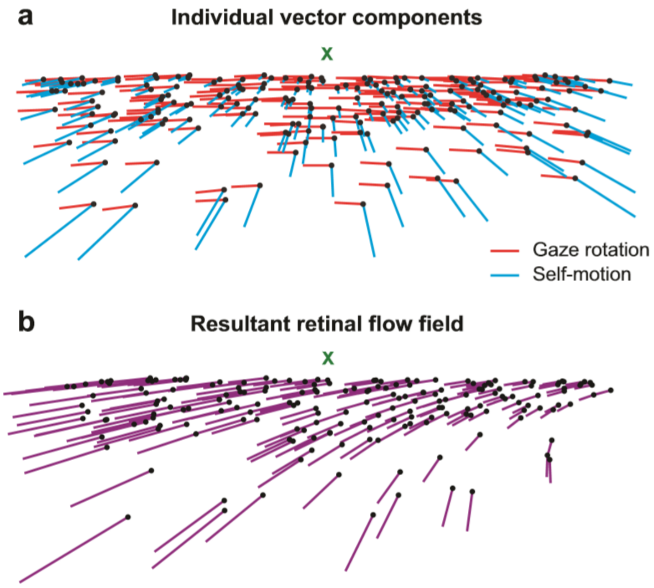

Figure 3.6: a) Visual flow vectors resultant from movement towards the green X and horizontal eye movement to the right. b) The resultant visual (or retinal) flow field. (source: [15]).

Another mechanism mentioned in Nash et al. [65], which is slightly different than the other visual-flow based methods, is described by Zacharias et al. [104]. Here, instead of measuring how the visual scene flows, the rates of change of vectors between ourselves and certain objects in our visual field are measured. These can be used to calculate the time before we cross a certain boundary, or collide with a certain object.

Using the mechanisms above, it can readily be explained that it is possible to determine distances, translational and rotational velocities, and heading. However, it is more questionable if we measure accelerations with our visual system, or just infer it from successive attentive visual velocity measurements. Gordon [36] makes the case that the acceleration visual field looks unnatural, and therefore that acceleration is not a primary visual input. Unfortunately, no research has been found that investigated time-bound visual acceleration perception thresholds, as delays could be an indication of the perception mechanism used.

Due to the complexity (and differing possibilities) of the organs and mechanisms involved with visual perception, it is difficult to give a dynamic representation of the system. Often a unity gain between the signal and the resulting neural afferent response to the CNS, is assumed.

---

[2] During World War 2 the understanding of visual perception was of sudden interest in order to aid aircraft pilots, primarily with landing when the estimation of the correct distance to the ground was of great importance

Hosman [40] described the velocity measurement mechanism of humans, based on the work by Van de Grind [91], as a simple bilocal motion detector (figure 3.7). Each bilocal motion detector is only sensitive to a single velocity. Therefore, a very large number of motion detectors, each sensitive to a different velocity, are located in the visual cortex. Furthermore, Van de Grind et al. [92] found that the bilocal motion detectors could be roughly divided in two groups: one sensitive for higher delays, which have a constant time delay (T in figure 3.7) and a varying span between the retinal receptors (S in the figure 3.7). The other group of bilocal motion detectors is sensitive to lower velocities, of which the time delays vary, but the retinal receptors have a constant span. The time delay of the detector varies between 65 ms for higher velocities, and between 1 to 5 seconds for lower velocities. The critical velocity that divides these two groups cannot be given precisely. Other than the time delay caused by the characteristics of the bilocal motion detector, an additional delay is introduced by the transmission time between the visual organs and the CNS and the processing time needed for perception.

Summarizing the above, the unity gain, denoted by K, extended with the bilocal motion detector delay τ_1, and the neural transmission and processing delay τ_2, results in a transfer function of the following form:

$$H_{vis}(s) = \frac{\text{Afferent neural response}}{\text{Visual stimulus}} = Ke^{-s(\tau_1+\tau_2)} \tag{3.4}$$

It is difficult, if not impossible, to measure the additional delay τ_2 directly with some certainty. Usually the entire delay from stimulus to motor command is measured, and even then it is difficult the discern between neuromuscular lags and neural processing delays.

An early, well known, source to visual time delays, is the article by McRuer and Jex [58]. He found by fitting his well known crossover model to simple tracking tasks that visual time delays decreased with increasing forcing function bandwidth. McRuer and Jex [58] inferred that this was done by tightening up the neuromuscular loop. In the case of their experiment, that meant that subjects increased their grip on the control stick. Depending on the dynamics of the controlled element and the forcing function bandwidth, delays were found to be between 0.1 and 0.5 seconds.

An article that provides some insight in the make-up of the visual perception time delay is written by Kawakami et al. [46]. They measured the latency of magnetoencephalograph (MEG) response of the cortical regions involved in visual motion detection to real motion in a wide speed range, and compared it with the human reaction time to pressing a button for the same objective. MEG response latency was found to be between 0.1 and 0.2 seconds. Response times were found to be between 0.25 and 0.5 seconds. Both the MEG response latency and reaction time decreased with increasing motion stimulus speed. One should note that the delay could also depend on the stimulus.

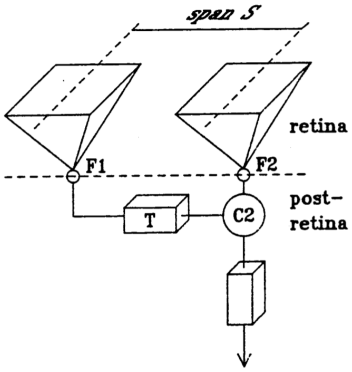

Figure 3.7: Bilocal motion detector, which only activates if the T seconds delayed response measured at $F1$ coincides at the comparator $C2$ with the response of $F2$. (source: [91])

Another interesting avenue to evaluate the visual perception of motion by humans, is through vection. Previc and Mullen [71] compared the latencies of visually induced postural change and self-motion perception under identical visual conditions. They found an increase in postural instability several seconds before a subject begins experience vection. The deemed that delay of several seconds in the postural response to visual field shifts would render it ineffective in helping to prevent most naturally occurring falls.

In the literature review article by Nash et al. [65], the doctoral thesis of Bigler [7][3] is mentioned. Bigler used a driving simulator simulator to measure yaw angle and lateral displacement thresholds. He found the visual system to have a low-pass characteristic, and could be described for both yaw and sway motions with the following low-pass filter:

$$H_{SCC}(s) = \frac{\text{Afferent neural response}}{\text{Visual yaw/sway stimulus}} = \frac{0.810}{s + 0.810} \tag{3.5}$$

3.3.2 Semi-circular-canals

The semi-circular-canals (SCC) are part of the vestibular organ and are sensitive to rotational accelerations. In total there are three canals, which located in each inner-ear approximately perpendicular to each other. Two types of hair cells in the organ measure deflections of the canals. The afferent neural response of type I hair cells correlate with the deflection and rate of deflection of the hairs (ciliae) on the cell. Type II hair cells generate an afferent response which correlates only with the deflection of the ciliae.

Fernandez and Goldberg [31] formulated a transfer function describing the relationship between rotational accelerations and the afferent neural response based on the dynamical properties of the SCC's and the measurements done on monkeys:

$$H_{SCC}(s) = \frac{\text{Afferent neural response}}{\text{Rotational acceleration stimulus}} = \frac{\tau_a s}{1 + \tau_a s} \cdot \frac{K(1 + \tau_L s)}{(1 + \tau_1 s)(1 + \tau_2 s)} \tag{3.6}$$

In his doctoral thesis, Hosman [40] presents parameters for the formula found by **?**] that the describe the human SCC's. The parameters were based both on literature review and experimental data he gathered. He omitted the first part of the transfer function dependent on τ_a (called the adaptation term), because that part is only important for long duration or very low frequency stimuli. These parameters are:

- Adaptation time constant $\tau_a = 80$ s
- Neural lead term $\tau_L = 5.92$ s
- Time constant $\tau_1 = 0.11$ s
- Time constant $\tau_2 = 0.005$ s
- Gain $K = 5.73$

Plotting transfer function 3.6 and the limited transfer function without the adaptation term using the above parameters, results in the bode diagram shown in figure 3.8. It can be seen that, when the stimulus frequency is higher than 0.1 rad/s, the SCC's approximately acts as an integrator. Hence, the afferent neural response correlates with the rotational velocity instead of rotational acceleration in this region.

Numerous studies have been done to establish SCC measurement thresholds. An overview is given by Nash et al. [65]. Thresholds were found to be frequency dependent, and between 0.02 to 0.1 rad/s.

In literature, there is a lot of discussion on the vestibular time delay. Some studies find significantly longer vestibular delays than visual delays (a review of these studies: [5]). While in other studies, the latency to the onset of the Visual-Occular-Reflex is found to be only 5-9 ms [4]. In a closed-loop pilot-control experiment, vestibular delays have been found to be lower than visual delays (a 150 ms difference) [103].

3.3.3 Otoliths

The otoliths measure transversal accelerations[4] in three orthogonal axes. They are similar to the SCC's, as they also generate a afferent neural response as a function of the deflection and rate of

---

[3] Not freely accessible, so the findings of Nash et al. [65] are used

[4] More correct, but less clear, is using specific force instead of acceleration. The specific force can be found by subtracting the gravity vector from the acceleration vector acting on the otoliths. In this thesis, when talking about accelerations, the gravity contribution is never included unless specified otherwise

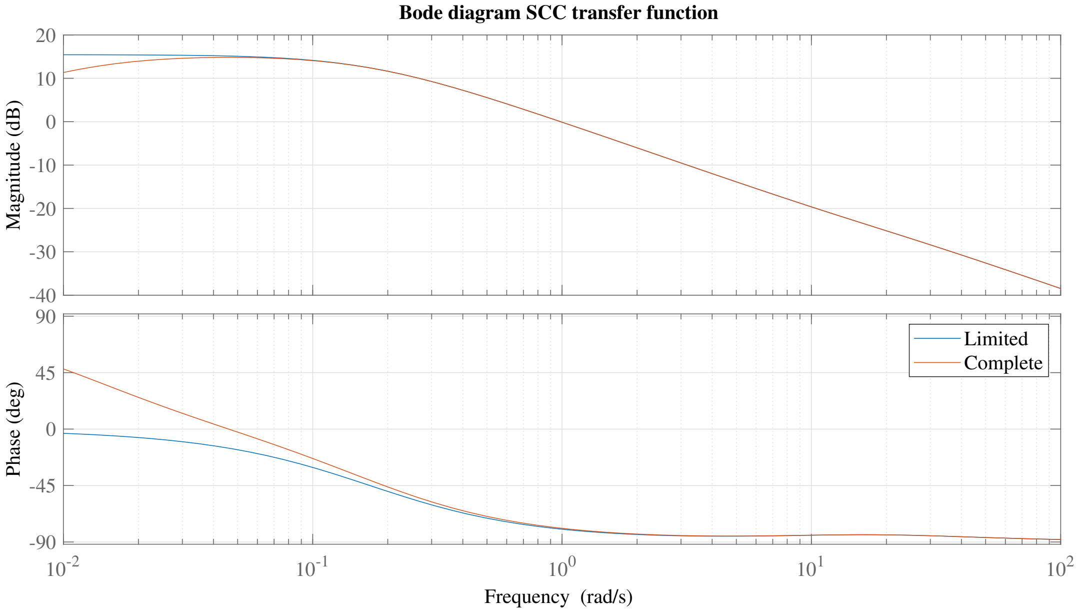

Figure 3.8: The bode diagram of the transfer function describing the relation between rotational acceleration stimulus and afferent neural response in red as derived by Fernandez and Goldberg [31] using the parameters found by Hosman [40]. In blue the transfer function without the adaptation term is used.

deflection of the ciliae on type I and II hair cells. The ciliae extend in a jelly-like layer which contains calcium-carbonate stones. This layer moves under the influence of transversal accelerations.

Fernandez and Goldberg [32] found, using similar methods as for the SCC's, the transfer function describing the transferral velocity stimulus and afferent neural response of monkeys:

$$H_{oto}(s) = \frac{\text{Afferent neural response}}{\text{Transversal acceleration stimulus}} = K\frac{1+K_a\tau_a s}{1+\tau_a s}\cdot\frac{1+K_v(\tau_a s)^{K_v}}{1+\tau_m s} \tag{3.7}$$

However, the term $(\tau_a s)^{K_v}$ makes it difficult to use this transfer function for modelling purposes. Therefore, often a simplified transfer function is used (first derived by Young and Meiry [101]):

$$H_{oto}(s) = K\frac{1+\tau_n s}{(1+\tau_1 s)(1+\tau_2 s)} \tag{3.8}$$

Over the years numerous research papers were written ([101], [68], [40], [88])[5] that derive the constants of this transfer function.

In figure 3.9 the bode diagram of transfer functions using the different derived parameters 3.8 is shown. It can be seen that the the relation between transversal acceleration stimulus and afferent neural response is quite non-linear.

In similar fashion as with the SCC thresholds, Nash et al. [65] compiled the otolith thresholds of different studies done on them. There is a less clear frequency-threshold relation as with the SCC's. This is attributed by Nash et al. [65] to different experimental methods. The thresholds found are between 0.02 and 0.09 m/s^2.

3.3.4 Golgi tendon organs

Deformation of a muscle tendon correlates with the force generated along that tendon. This deformation is measured by the Golgi Tendon Organs (GTO's). Therefore, the afferent neural signal generated by the GTO's provides the CNS with information on muscle forces. The first article to model this relation was by Houk and Simon [42]. Their linear model was of the following form for the measured afferent responses of the GTO's located in the legs muscles of cats:

$$H_{GTO}(s) = \frac{\text{Afferent neural response}}{\text{Muscle tendon force}} = K\frac{(s+n_1)(s+n_2)(s+n_3)}{(s+d_1)(s+d_2)(s+d_3)} \tag{3.9}$$

---

[5] Based on a recent review paper by Asadi et al. [3]

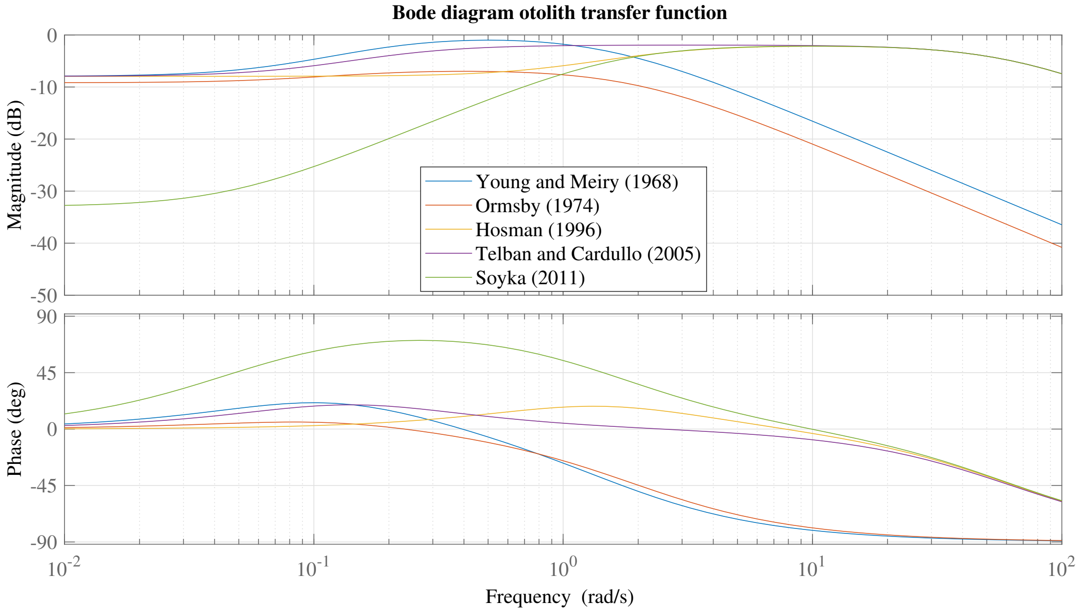

Figure 3.9: The bode diagram of the transfer function describing the relation between transversal acceleration stimulus and afferent neural response (equation 3.8) using the parameters found by the different studies. The gains K have been scaled just as in [65] in order to give comparable outputs.

The Houk and Simon [42] model is still used, however a more complicated non-linear model which captures the effects of self- and cross-adaptation was developed by Mileusnic and Loeb [62].

In a literature review on the GTO's by Jami [43], the characteristics of the GTO's are described in depth. Summarizing, the sensitivity of the GTO's is found to be as low as 5 to 20 μN for small force amplitudes, and independent of frequency in a frequency range between 10^{-1} and 10^{2} Hz. Furthermore, the sensitivity was found to be independent of the tendon dynamics itself. Delays under 10 ms between tendon stretch and GTO activation were measured.

The bode diagram transfer function described in equation 3.11, using parameters describing the triceps surae (calf muscle) found by Houk and Simon [42], is shown in figure 3.10.

3.3.5 Muscle spindles

Just like the GTO's, the muscle spindles are sensors attached to the muscle fibers. There are two types of afferent nerves connected to the muscles spindles: the Ia and II sensory nerves. The Ia afferent response correlates with muscle length and velocity. The II afferent response correlates only with length. A visualization of the frequency encoded afferent signals of both nerves is shown in figure 3.11. The first to derive a model, describing the behaviour of the afferent sensory nerves of the muscle spindles under isolated conditions, were Poppele and Bowman [70]. They concluded that the behaviour was linear, and formulated the following transfer functions:

$$H_{Ia}(s) = \frac{\text{Afferent neural response}}{\text{Muscle length and velocity}} = \frac{s(s+0.44)(s+11.3)(s+44)}{(s+0.04)(s+0.816)} \tag{3.10}$$

$$H_{II}(s) = \frac{\text{Afferent neural response}}{\text{Muscle length}} = \frac{(s+0.44)(s+11.3)(s+44)}{(s+0.816)} \tag{3.11}$$

These transfer functions are based on a data fit, which in this case resulted in improper transfer functions. These can't be used as is in a modelling context. The authors of [95] conclude that the above transfer functions are often to detailed to use in an integrated model. They advice that, for good modelling, it is more important to have realistic time delays and noise levels. For good measure, the bode diagrams of the Ia and II sensory nerves are included in figure 3.12.

Matthews [57] investigated the delays of the muscle spindles. He found a short latency delay of 25-30 ms for the Ia afferent nerve, and a longer latency delay of 40 ms for the II afferent nerve. This does not include any neural processing time.

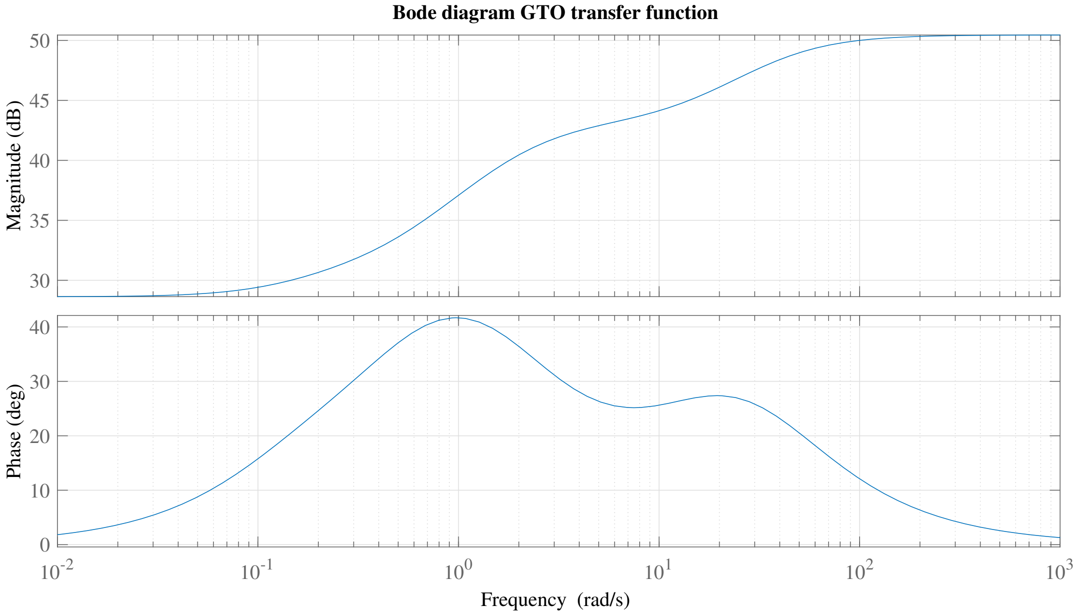

Figure 3.10: The bode diagram of the transfer function describing the relation between rotational acceleration stimulus and afferent neural response of the triceps surae as derived by Houk and Simon [42].

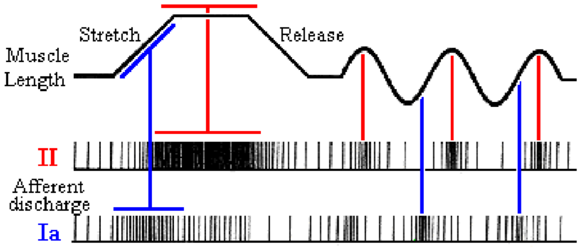

Figure 3.11: The frequency encoded response of the II and Ia afferent nerves. The II afferent nerve is mainly sensitive to length, while the Ia afferent nerve is sensitive to muscle velocity (source: [95]).

Nash et al. [65] reported on the findings by Newberry et al. [67] and Bigler [7]. They found, when investigating the perception of the steering wheel angle in cars, that the average displacement threshold was between 0.005 and 0.03 rad.

3.3.6 The effect of multimodality

Up till now, all the sensors and their characteristics have been discussed as separate entities. However, for the cycling task, it is very likely that different sensors are used together to estimate all the relevant states. This can have a profound effect on the overall perception window of the rider.

When investigating these effects in more detail, it became clear that there many compounding effects that influence all dimensions of the perception window (intensity, frequency and duration).

In the research by [86], (optimal) sensory integration of the Otoliths and the SCC's is proven, which results in a lowered eccentric rotation perception threshold. Similar findings, but for visual-vestibular integration are presented in [23].

Using natural motion cues, Karmali et al. [45] found that visual and vestibular perceptual thresholds each demonstrate better precision at specific frequencies and also exhibit optimal integration. Precisely, they found that the visual and vestibular thresholds were indistinguishable at 0.05 Hz and 2 Hz, vision was more precise between 0.1 and 1 Hz and less precise than the vestibular system above 2 Hz.

Zaal et al. [102] found that in a combined pitch tracking and disturbance-rejection task, pitch motion in addition to visual cues significantly improved tracking performance due to a lower effective time delay. However, they also noted that this effect is likely task-dependent.

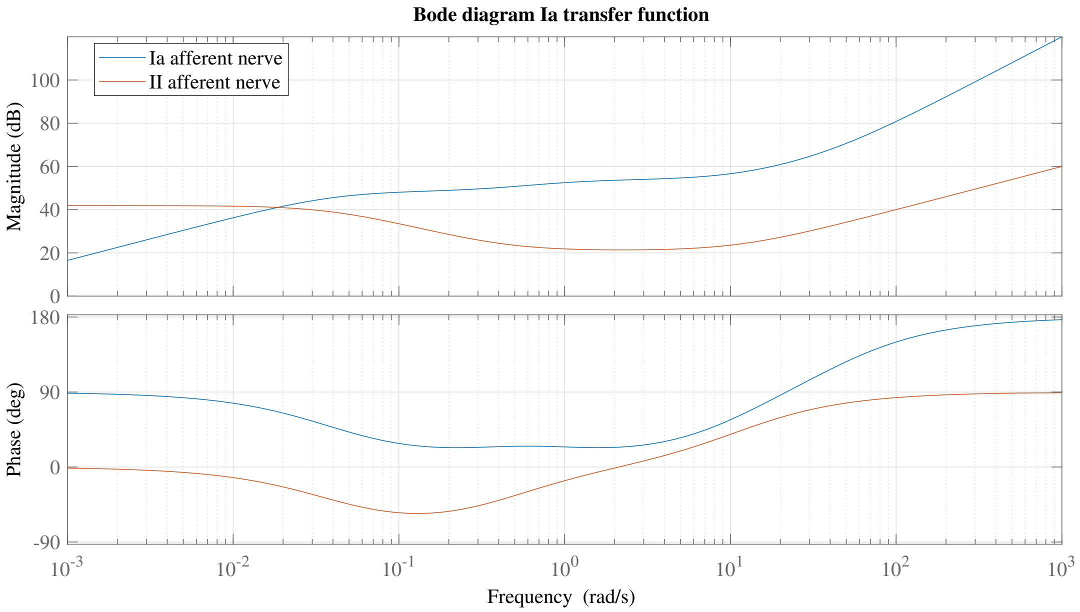

Figure 3.12: The bode diagram of the transfer functions of the II and Ia afferent nerves. The parameters used are from [70].

3.3.7 Discussion on sensory model selection

A main topic of research should be, to what level of detail a sensory model needs to be implemented. As found in section 2, were the current state-of-the-art was evaluated, sensory models are not included in any bicycle model until now.

For all the sensory organs in table 3.2, mathematical descriptions of the sensory dynamics, perception windows, and neural delays are available. To what degree these descriptions need to be implemented in order to yield realistic overall results, is debatable according to some authors, but at least the neural delays should be implemented [95]. For bicycle-rider control modelling it depends on how big the effect is of the sensory dynamics on the closed-loop behaviour of the system. For the control of a unstable bicycle, for which an unrecoverable fall can occur within seconds, it makes sense that the neural delays have a large impact on the closed-loop behaviour.

In general, the vestibular delays were found to be lower than visual delays.

3.4. State estimation

As discussed in the previous section, humans have a wide suit of sensors available that can be used to determine the relevant states of a bicycle-rider system (refer to table 3.2 for an overview of these states and sensors). In the context of this research, the process of deriving these states of the bicycle model (e.g., roll, steer, and heading), from measurements done by the human senses, in a way that mimics human behaviour is called state estimation.

It covers methods to integrate the sensor signals and predict effects of motor actions to improve the estimation of the states of the bicycle. There are several options, which may be combined, to achieve this. They are categorised as followed:

- No sensory integration
- Linear weighing of sensory signals
- Bayes' theorem based
- Kalman filtering based
- Prediction

In this chapter, these five state estimation methods are shortly introduced.

3.4.1 No sensory integration

The simplest option is to skip integration or estimation steps, and to just provide the states of the bicycle model directly to the controller. This has some degree of realism, because as found in section 3.3, all the bicycle states can in theory be measured by the rider. This is the method that is used primarily when modelling bicycle rider control behaviour. Validated examples are given by Moore [63] and Schwab et al. [79]).

3.4.2 Linear weighing of sensory signals

Depending on how the human senses are modelled, certain states can be measured by multiple senses (e.g. roll rate by the vestibular and by the visual system). A simple way of fusing multiple sensory signals, is by assigning weights to them in order to add them up to a weighted sum. Hosman and Stassen [41] uses this approach to model the perception of a pilot controlling an aircraft. Vestibular and visual cues are combined using linear weights, which are found using experimental data. The model Hosman and Stassen [41] uses is shown in figure 3.13.

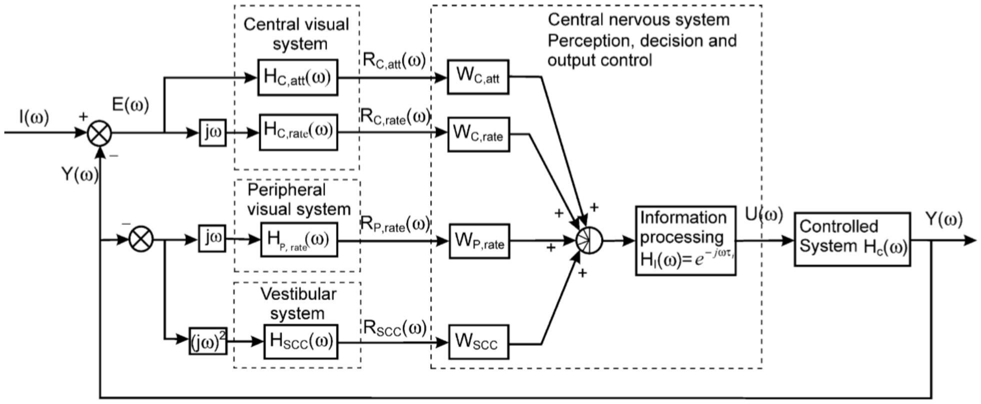

Figure 3.13: A pilot perception model using linear weights $\mathrm{W}_{\mathrm{c,att}}$, $\mathrm{W}_{\mathrm{c,rate}}$, $\mathrm{W}_{\mathrm{P,rate}}$, and $\mathrm{W}_{\mathrm{SCC}}$ (source: [41]).

In the context of bicycle rider control modelling, linear weighting is used by Soudbakhsh et al. [85].

3.4.3 Bayes' theorem based

Bayes' theorem enables you to calculate the probability of an event, using prior knowledge related to that event. This relation can be described with the following equation:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{3.12}$$

Here, $P(A|B)$ is the probability of event A occurring, given that event B already occurred. $P(B|A)$ is the reverse of that: the probability that event B occurs, given that event A already occurred. $P(A)$ and $P(B)$ are the probabilities of event A or B occurring respectively. Ma et al. [55] wrote a article, describing the principle of using Bayesian inference, to derive an optimally estimated mean (equation 3.13) and variance (equation 3.14) from multiple observations in a human context.

$$\mu = \frac{\sigma_1^2}{\sigma_2^2 + \sigma_2^2}\mu_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}\mu_2 \tag{3.13}$$

$$\frac{1}{\sigma^2} = \frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2} \tag{3.14}$$

3.4.4 Kalman filtering based

A Kalman filter uses knowledge about the process and measurement dynamics of the system, inputs to the system, and the measured outputs of the system, to get a better estimate[6] of the internal states.

---

[6] A minimum mean-square error estimate, to be precise

This makes sense if not all states are (accurately) measured, or multiple measurements of the same state can be fused.

The Kalman filter is divided in two steps: the predict step and the update step. In the predict step, the previous state estimate is used to predict the state at the current time-step. In the update step, this *a priori* estimate is then combined with the current observed state, in order to yield an improved *a posteriori* state estimate.

Assuming that a well-motivated and well-trained human operator behaves in a near optimal manner, Kleinman [47] developed a model that uses a cascade combination of a Kalman filter, a least mean-squared predictor, and a set of gains acting on the estimated state. This model provides excellent agreement between theoretical and measured quantities.

A similar approach was used by van der Kooij et al. [93] in order to model multi-sensory integration for human stance control. This model was later extended with an adaptive Kalman filter to mimic human adaptation to noisy measurements [94].

In a review article [29] a detailed case is made that humans integrate information optimally. The remaining problem is, is that for a Kalman filter to function, process and sensory noise information has to be provided. In [29], it is argued that it is unlikely that this is prior only information, because that the sensors can be used in a infinite number combinations of tasks and environments, which causes the noise characteristics to change. Therefore, it is likely that the noise characteristics are updated on-line.

3.4.5 Prediction

It can be illustrated with a simple example that the human uses some sort of prediction. Let's consider someone holding a bottle at a constant position. If that person self-generates a perturbation of the bottle, this perturbation can be counteracted easily: the grip force is matched and timed with the perturbation. If the perturbation has an external origin, this is more difficult. A higher baseline grip force is required to prevent slippage, and the grip force lags behind the perturbation due to neural time delays.

A good overview of the common interpretation of this mechanism, is given by Marc Jeannerod in [44]: The notion is that each time the motor centers generate an outflow signal for producing movement, a ”copy” of this command (the ”efference copy”) is retained. The reafferent inflow signals (from the sensory system) generated by the movement are compared with the copy. If a mismatch between the two types of signals is recorded, new commands are generated until the actual outcome of the movement corresponds to the desired movement. In order to give this mechanism a predictive role, one must assume the existence of an ”internal model”. This internal model is used to forward simulate the effects of the action, without waiting for sensory reafference.

A prediction model that aims to mimic human behaviour, is described by Miall et al. [61]. They suggested that the cerebellum forms two types of internal models. One is a rapid forward predictive model, predicting the consequences of motor movement. The other model is of the time delays that exist within the control loop. An adaptation of their model, which mimics a Smith predictor [84], is shown in figure 3.14. If the internal model is a close match to the plant model, the output of the internal plant model likely matches the output of the plant. Any mismatch between the two outputs will be corrected by the delay model. However, a Smith predictor requires the plant to be stable in open-loop configuration [37]. This limits it's usefulness for modelling bicycle rider control.

An other way to model human prediction, is the Tapped Delay Line (TDL). When used together with an Kalman filter, the TDL provides an optimal state estimate [47]. [94] uses a TDL to model human postural balancing. It should be noted that a TDL only can be implemented in discrete time. The TDL is illustrated in figure 3.15. The A* and B* matrices in the figure denote the knowledge the rider has about bicycle dynamics and are used to forward model the impact of the control actions towards the prediction horizon. The optimal state estimate $\hat{x}[k]$, as a function optimal state estimate $\hat{x}[k-\tau_s]$ from a Kalman filter, the (discrete) delay τ_s ,and the delayed inputs, can be calculated using:

$$\hat{x}[k] = A^{\tau_s} x[k-\tau_s] + \sum_{n=1}^{\tau_s} A^{n-1} B u[k-n] \tag{3.15}$$

3.4.6 Discussion on the methods of state estimation

Based on the evidence that humans have an optimal and adaptive nature when integrating sensory information, both Kalman filtering and Bayes' theorem based integration can be used for modelling

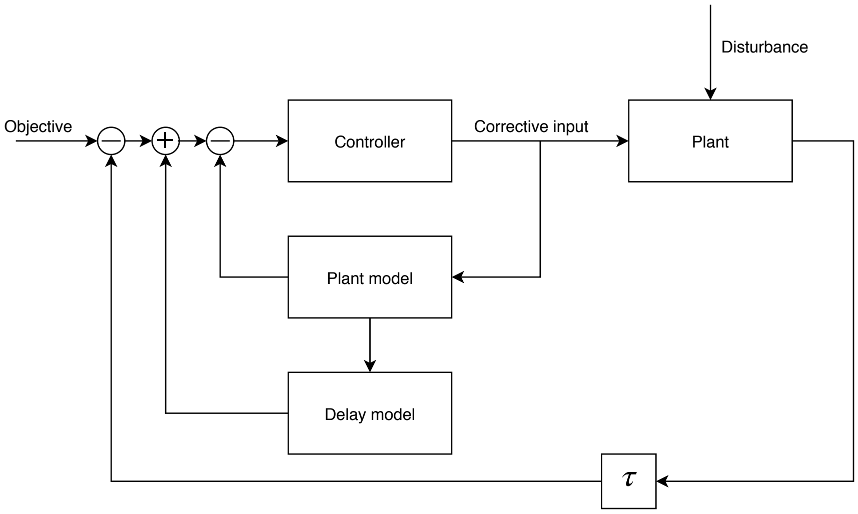

Figure 3.14: The Smith predictor architecture for a system with delay τ (adapted from [61]).

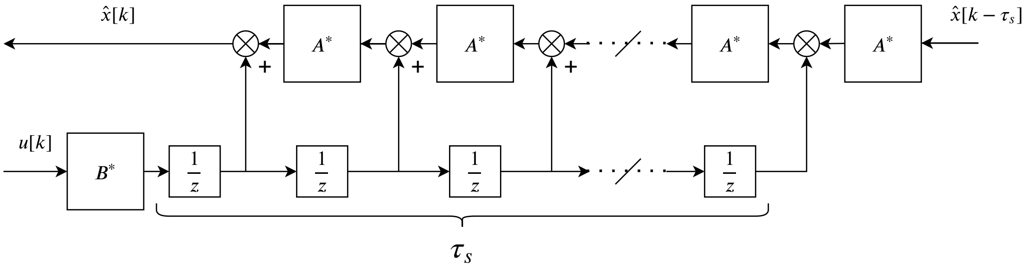

Figure 3.15: Visual representation of a tapped delay line with delay τ_s. A* and B* denote the knowledge the bicycle-rider has of the dynamics of the bicycle and the effect of the control input $u[k]$ on these dynamics. In this figure, the delayed state estimate $\hat{x}[k-\tau_s]$ which is optimal if provided by a Kalman filter, is used together with the inputs given during the delay period τ_s, to forward calculate the current state $\hat{x}[k]$.

human state estimation. Linear weighing of the sensory channels also is proven to work, but ties in less well with the notion of an internal model. Linear weighing also requires some assumptions on if a state is directly measured by a sensor, or derived from other measurements. The Kalman filter combines these aspects. Furthermore, the approach to gauge the relative importance of the different sensory channels laid out by [93], where a Kalman filter is used, could also be applied to this research. This results in the Kalman filter to be the best candidate for modelling sensory integration. On the use of a prediction algorithm: it is established that humans do some manner of prediction on the effects of their control actions. Based on literature, the best way of modelling this, is trough the TDL predictor.

3.5. Control

In what is defined as the control part of human cognition, the states derived in the state estimation block are compared to the control objective (usually balancing of the bicycle), and corrective control action is applied when necessary. In this section, the aspects of the controller are divided in three sections. First, the possible structures in which a controller can placed are discussed. After that, the possible control methods, where the corrective control action is calculated, are discussed. Lastly, the types of inputs the controller could use for corrective action to a bicycle model, are discussed.

3.5.1 Control structure

Most often, control is applied through a feedback structure, of which the basic representation is given in figure 3.16. In layman's terms: the output of the plant is compared to a certain objective. The controller then provides a corrective input to reduce any error between the output and objective signals.

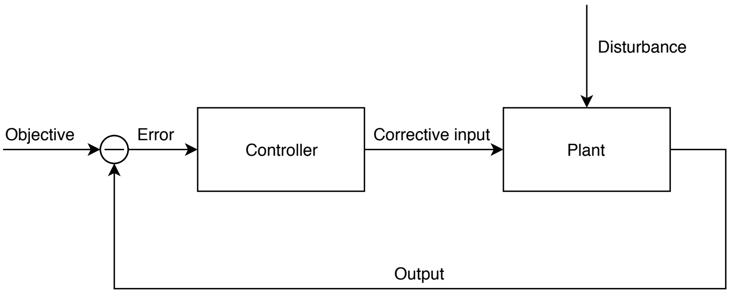

Figure 3.16: Visual representation of a basic negative feedback control structure.

The other, more specialized, scheme that is used often for the modelling of human control, are multiple-loop controllers. An example is given in figure 3.17. Multiple-loop controllers enable the use of control methods that were originally developed for single-input-single-output systems to multiple-input-multiple-output systems. A downside is that multiple-loop controller structures often result in elaborate and over-determined models, which as a consequence, are difficult to identify [64].

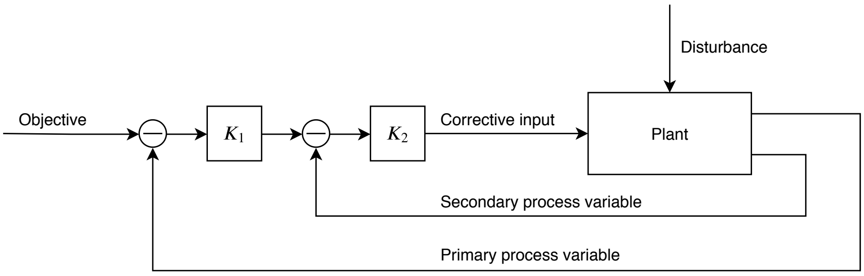

Figure 3.17: Visual representation of a multi-loop controller. Here two loops are included, where the inner loop governs the feedback to the primary control variable (e.g. roll) and the outer loop governs the feedback to a secondary control variable (e.g. heading).

3.5.2 Control method

Still to this day, one of the leading works that gives insight on human compensatory control, is given by McRuer and Jex [58]. They found that a human adopts sufficient lead or lag equalization, so that the slope of the open-loop response of the combined human and controlled element system is approximately -20 dB/decade in the crossover region. Based on this, they proposed a two-parameter crossover model:

$$Y_{OL}(s) = Y_P Y_C \doteq \frac{\omega_c e^{-\tau_e s}}{s} \tag{3.16}$$

Here, Y_P is the transfer function describing the human controller, Y_C is the transfer function of the controlled element, ω_c is the crossover frequency, and τ_e is the effective time delay (from human controller input to output of the controlled element). The simplest model that can describe human compensatory behaviour, is given by McRuer and Jex [58] to be:

$$Y_P = K_P \frac{T_L s + 1}{T_I s + 1} e^{-\tau_e s} \tag{3.17}$$

In the above formula, T_L and T_I are respectively the lead and lag time constants. K_P is the static gain of the human controller, and τ_e again is the effective time delay. Depending on the controlled element, a human adjusts the parameters in such a way that the crossover model is valid in the crossover region. However, it should be noted that bicycle balancing is not a pure compensatory task.

Another way to model human control, that is more suitable for multiple-input-multiple-output models, is by using proportional state feedback gains. As noted in section 2, the majority of the models developed in order to mimic human bicycle control use proportional state feedback gains. There are usually chosen based on a data fit to experimental data.

There is evidence that humans optimize their control effort in some way [74]. This can be mimicked by using a Linear-Quadratic-Regulator (LQR). The LQR algorithm uses a plant model to minimize a cost function, that weighs undesired state deviations against control effort, to find a state-feedback controller. A desirable characteristic of the LQR algorithm, is that if the cost function is specified and the system model is provided, the feedback gains are automatically defined by the algorithm and do not have to be chosen individually.

Model-Predictive-Control (MPC) is another control method that might be suitable to mimic human control behaviour. MPC uses a model to predict and anticipate to future events for a finite time-horizon. It then optimizes the current control action by minimizing a control strategy cost function over a receding prediction horizon. The same argument as for the LQR controller is valid: humans are known to have an gain-to-effort control strategy. MPC has been used for human control modelling in motorcycles ([56], [18]). Both articles used MPC to track a reference roll angle, while Massaro et al. [56] also considered forward velocity as a reference. MPC has the same desirable characteristic that the LQR algorithm also has: the feedback gains follow from the algorithm. However, now there are slightly more tunable parameters, namely additional weighting coefficients in the cost function and the length of the finite time horizon. Also, because MPC does not calculate a single optimal solution for the whole time horizon (like LQR), it may reach a suboptimal solution. Of course, for human control modelling this is not necessarily a downside.

Yet another control method that can be used to model human control behaviour, is fuzzy control. Fuzzy control has been developed in order to convert complex and poorly defined control tasks done by humans, to a digital controller, by allowing linguistically expressed heuristic control policies to be converted to fuzzy rules [105] [106]. Although, the usefulness of fuzzy control for modelling bicycle-rider control is debatable, since riders of bicycles are not capable of describing their control strategy in-depth [27].

3.5.3 Control inputs

As revealed in section 2, all the models had one or two different control inputs. Either steer angle or steer torque was used, sometimes extended with input generated by upper body movement. There is quite some disagreement between different bicycle control modelling efforts if steer torque or steer angle should be considered as an control input. Sharp [81] argued, that because the (validated) Whipple bicycle model has torques as an input, it is logical that the rider uses torque based control inputs. However, a simple inverse model can be used to make a angle controller suitable for the Whipple bicycle model, without affecting the results.

It also is shown to be unlikely that body lean is a control input that is used for the balancing task at all, except the very lowest cycling speeds [77]. This also is shown qualitatively by Kooijman et al. [50], who filmed bicycle-riders while cycling.

3.5.4 Discussion on control

A controller that follows the principles established by McRuer and Jex [58], would be very powerful in it's predictive capabilities. However, it is likely that McRuer and Jex [58] principles do not hold true, because bicycle balancing is not a pure compensatory task.

LQR or MPC algorithms also provide enticing control methods, as they don't rely on determining individual feedback gains, but on weighing control effort versus task performance. This leaves less free parameters, and can therefore yield better predictive capabilities. These algorithms can only be placed in the basic feedback structure, so picking these would exclude a multiple-loop structure. This is not necessarily a downside due to the aforementioned issues that arise when such a controller structure is used for system identification purposes.

Fuzzy control is difficult the implement, because it is required to express the control law linguistically. This is difficult, if not impossible, for rider-control tasks.

It is not important, from a modeling perspective, to differentiate between position or torque control. Body lean as a control input can be excluded, unless cycling at very low velocities is modeled.

3.6. Neuromuscular dynamics

The use of neuromuscular dynamics in human control models is widespread (Mulder et al. [64] gives a nice overview). They are usually used to translate a control command given by a human to the actual realized control input to the controlled system. Essentially, it is a limitation on the bandwidth of the human control input and it can described by a model which lumps the neuromuscular and control inceptor systems together [21]. The resulting model has the form of a second order low pass filter.

In the case of bicycle control, the neuromuscular model describes the relation between the neural torque command given by the rider ($\hat{T}_\delta^c$) and the actual output steer torque (T_δ^c). This results in the following description of the neuromuscular dynamics:

$$H_{nm}(s) = \frac{T_\delta^c(s)}{\hat{T}_\delta^c(s)} = \frac{\omega_c^2}{s^2 + 2\zeta\omega_c s + \omega_c^2} \tag{3.18}$$

For bicycle-rider identification, Schwab et al. [79] used shoulder muscle parameters from Happee et al. [38]: the cut-off frequency $\omega_c = 2.17 \cdot 2\pi$ rad/s and the damping coefficient $\zeta = \sqrt{2}$. Moore [63] found that by adding neuromuscular dynamics, his model gave a better approximation of the experimental data.

4

Theoretical outline of the proposed approach

To reach the research objective, and to answer the secondary research questions, a model has to be developed that is realistic enough in order to conclusively say which sensor channels are used and what their relative importance is, but also is general enough to be applied beyond the data it is modelled after. In this chapter the system identification procedure used to reach the research objective is explained. This results in a range of model structure choices and free parameters. In the second section, the modelling aspects, detailing how the model is implemented, are elaborated on.

4.1. System identification

System identification is, in principle, quite a straightforward procedure. It can be divided in three steps:

1. Acquiring suitable data.
2. Model structure selection.
3. Fitting the model structure to the data.

The first step is to acquire data that can be used to fit a model to: an experiment has to be designed and data has to be collected. Afterwards the data often needs to be pre-processed. For example to filter out noise, or to remove measurement bias.

The second step is to select a model structure. This starts at an elementary level with the choice between a between a white-box, grey-box, or a black-box model. White-box models are models where all information is available a priori. So, governing equations (i.e., the mathematical relationship between the parameters) and their parameters are fully given based on the known nature of the system that has to be identified. If this yields good results, system identification stops after this step. If not, another approach is to use a grey-box model. Here, the governing equations are prescribed, but (some of) the parameters are left free. The equations and their parameters still have a physical interpretation. The free parameters can then be fitted to the experimental data. The last option is to use a black box model. Now, only the number of parameters (or order) of the model is prescribed. This results in the parameters losing all physical meaning. For some applications this is fine, but for the purpose of this research, only white-box or grey box model structures are applicable, because we are interested in the underlying phenomena (e.g., the bicycle states that are used by the rider), that result in human control behaviour. However, this still leaves a lot of more detailed model structure choices open. These options are covered in more detail in section 4.1.3.

The third step is to fit the model to the data and to validate the model. Fitting is done by defining a criteria that tells how well the model predicts the actual data. Usually this criterion takes the form of an error that has to be minimized. The minimization itself can be done with a lot of different methods, both iterative and non-iterative. Validation is usually done by splitting the experimental dataset in two parts: a part which is used to fit the parameters, and a part which is used to test the fitted model.

If the performance of the model remains good (i.e., the performance criterion remains above some predefined limit), the system identification procedure is complete.

What needs to be considered, is that in practice, the results at each step influence the final result. So, if, for some reason, the performance criterion is not met, the procedure becomes iterative. Then some, if not all, of the previous steps have to be done again using an alternative approach. The steps, and the iterative nature of the system identification cycle, are visualized in figure 4.1.

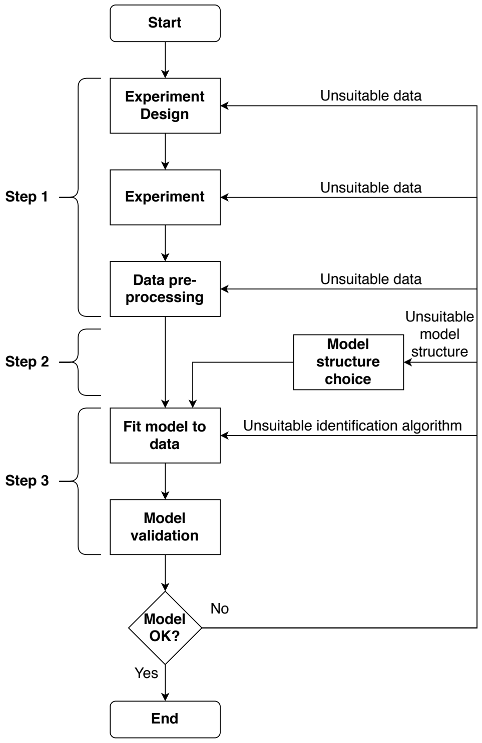

Figure 4.1: The system identification cycle (adapted from Ljung [53]).

4.1.1. Experiment design and description

Two experiments where conducted with cycling on instrumented bicycles that can be used for system identification of human balancing of a bicycle. One by Jason Moore (and associates) at the University of California in Davis. The other by George and Christos at the TU Delft. A good overview of the experiment done at the UC Davis can be found in the dissertation of Moore [63]. Dialynas et al. [24] describes the experiment done at the TU Delft. In the following two sections, a summary and an evaluation of the validity of both experiments are included.

UC Davis experiment

The experiment done at the UC Davis was designed to excite the rider/bicycle system in such a way that the parameters of the rider control system could be identified. Three different subjects of similar age, mass and (good) bicycling ability where used. The riders where attached rigidly to the bicycle frame by means of a harness which restricted upper-body movement, and magnets attached to the knees of the rider, to prohibit leg movement[1]. This ensured that the only means to stabilize the bicycle was through control of the handlebars.

There are two reasons for this. In a study on human control of a bicycle [49], visual observation showed that very little upper-body lean occurred, and that stabilization is done by steering control actions. This is an indication that a high fidelity model can be identified with just handlebar control as an input.

---

[1] While [49] found that very little upper-body lean occurred during normal cycling, at speeds below 1.4 m/s (5 km/h) knee movements became very large. Hence the restriction of knee movement.)

Table 4.1: Overview of the number of runs done in different categories during the UC Davis experiment. A portion of the runs done are not suitable for identification purposes for various reasons such as corrupted data and large forward velocity changes. The numbers in brackets indicate the total number of runs done in a certain forward velocity range, while the unbracketed number indicates runs that are selected for use in the further identification procedures. In the ”Other” category are runs with other done with other task descriptions.

Environment	**Treadmill**		**Pavillion**	
Task description	Balance with disturbance	Other	Balance with disturbance	Other
1-2 m/s	0	0	3 (5)	20
2-3 m/s	7 (8)	19	14 (17)	40
3-4 m/s	3 (3)	13	18 (25)	39
4-5 m/s	7 (7)	19	29 (50)	72
5-6 m/s	1 (1)	0	4 (16)	52
6-7 m/s	2 (2)	4	0	2
7-8 m/s	6 (8)	19	0	0
8+ m/s	0 (3)	3	0	0
Total	**26 (32)**	**77**	**68 (113)**	**225**

Table 4.2: The sensor suite of the Davis instrumented bicycle

Measurement	**Sensor type**
Roll angle	Potentiometer
Steer angle	Potentiometer
Roll rate	IMU
Yaw rate	IMU
Pitch rate	IMU
Steer rate	Single axis rate gyro
Rear wheel rate	Rotary encoder
Steer rate	Gyro
Linear accelerations	IMU
Lateral perturbation force	Load cell
Steer torque	Magnetostrictive torque sensor

The second reason is that the Whipple bicycle model does not have upper-body lean as a degree of freedom. Allowing upper body lean as an additional degree of freedom would probably require an extension to the Whipple model equations in order to reach matching behaviour of both the predicted bicycle roll and steer angle. This opens up a whole new can of worms: how the model a free moving bicycle rider?[2]

The experimental runs where performed in two different environments: on a horse treadmill and in a large gymnasium. Both environments had their up and downsides. The horse treadmill was only 1 m wide and the environment was static. This reduces the fidelity of the environment and could influence the riders' control strategy. The gymnasium was limited in space, and runs could only be performed at a straight sections which was 30 m long. The maximum travel speed was limited to 7 m/s.

The runs where conducted at different forward velocities between 1.4 and 9 m/s (5 and 32 km/h), during which the bicycle was impulse perturbed by applying a lateral force at the seat post. The direction of the perturbation was randomized. In table 4.1 an overview is provided of the number of runs done. The numbers of the runs, according to the labeling system used by Moore, is included in appendix A.5.

The sensor suite on the bicycle was quite extensive. In table 4.2, all the directly measured states and their respective sensors are listed. The data was collected with a sampling interval of 1/200 s. Furthermore, the bicycle was equipped with motor, which could keep the bicycle at a constant forward velocity.

---

[2] *Realistically* is the magic word here. The modelling itself is not that difficult and is not unprecedented. But, to do this properly, it is likely that a separate experiment is required in order to validate any assumptions made in such a model

Figure 4.2: A participant of the UC Davis study riding the instrumented bicycle on a horse treadmill (source: [63]).

TU Delft experiment

At the TU Delft, an experiment to identify the importance of torque feedback on the control task was conducted. In total, fifteen different subjects conducted two types of runs during which the riders were laterally perturbed. For the first run type, steering feedback was enabled. In the second type of runs steering feedback was disabled. The riders where not rigidly attached to the bicycle frame. This has the disadvantages described in the section on the UC Davis experiment. An advantage of not attaching the riders rigidly to the bicycle frame, is that this is more realistic. In a third type of runs, the steer motor gave a impulse perturbation. These runs are not used in this research, because they can't be compared to the UC Davis runs. Furthermore, the posture of rider on the bicycle influenced the results greatly. This increases the inter subject variability, which makes fitting a general model even harder than it already is.

The trials where performed at a public cycling path (without interference from other road users). This has as a benefit, that this is the actual environment we are interested in, so the fidelity of the environment is high. However, it also introduces a lot of uncontrolled (and unmeasured) variables, such as weather and cycling path conditions. These variables could influence the control behaviour of the cyclist.

The runs where performed at four different forward velocities (2.6, 3.7, 4.5, and 5.6 m/s). During the runs, the subjects where laterally impulse perturbed by means of a rope. This only allowed perturbations in one direction, but the authors found that this did not influence the results. In table 4.3 an overview is provided of the number of runs done.

The bicycle itself is not only equipped with the mechanism necessary to influence the steer torque feedback, but also with additional sensors which measure the states of the bicycle. The sensors are listed in table 4.4. The sensor suite is less extensive than the one on the Davis instrumented bicycle. Relevant signals, such as roll angle and steer torque, are estimated using the available sensor channels. The data was collected with a sampling interval of 1/1000 s. Just as the Davis bicycle, the Delft bicycle is equipped with a motor and cruise control system, which keeps the bicycle at a constant forward velocity.

Table 4.3: Overview of the number of runs done during the TU Delft experiment. In the ”Other” category are runs done without steer feedback or are done with steer perturbations.

	Cycling path	
Task description	Balance with disturbance	Other
2.6 m/s	15	30
3.7 m/s	15	30
4.5 m/s	15	30
5.6 m/s	15	30
Total	**60**	**480**

Table 4.4: TU Delft instrumented bicycle sensor suite

Measurement	**Sensor type**
Steer angle	Rotary encoder
Roll rate	IMU
Yaw rate	IMU
Pitch rate	IMU
Rear wheel rate	Rotary encoder
Linear accelerations	IMU
Lateral perturbation force	Load cell

Figure 4.3: A participant of the TU Delft study riding the instrumented bicycle while being perturbed by George (source: [24]).

4.1.2. Data pre-processing

The general aim of data pre-processing is to create a data set on which system identification methods can be applied and of which the individual trials can be compared fairly. Three different stages of data pre-processing are defined:

1. Preliminary data pre-processing
2. Extracting the relevant human response
3. Labelling outliers

Preliminary data pre-processing

In the first stage the raw measurement data is processed in order to have signals that can be used for further system identification. So, if necessary, the raw measurement data is synchronized, calibrated, and filtered. Furthermore, unmeasured states that are required for system identification are estimated. The processing procedures used are described in Moore [63] and Dialynas et al. [24] for the UC Davis and the TU Delft experiments, respectively. In this research, these procedures are unaltered.

Extracting the relevant human response

In the second stage of data pre-processing, the aim is to separate the human response to the perturbation from other effects caused by unknown disturbances and noise.

As explained in section 3.5, the rider controls the bicycle primarily through the handlebars, the handlebar steer angle is selected is the relevant human response. Handlebar steer torque is not selected for three reasons:

1. The steer torque generated by the bicycle rider was not directly measured in the TU Delft experiment.
2. There does not have to be a clear relation between steer torque and steer angle. I.e., it is possible to end up with a model that predicts the steer torque correctly, but the steer angle incorrectly. Moore also found this in his research and mentioned two possible causes: the knife-edge, no side-slip wheel contact assumptions in the Whipple bicycle model and the un-modeled passive rider contribution. It is easy to get lost in trying the figure out the actual relationship between the steer torque and the steer angle.
3. The handlebar angle has a very big influence on the heading of the bicycle. [79] found that bicycle heading might be a control objective of the bicycle rider.

For this research, approximately the same method as in [22] is used, which was fitting a Finite-Impulse-Response (FIR) model to the experimental bicycle-rider data in order to describe the relationship between the perturbation and the human response. This FIR model is then subsequently filtered with a 8th order low-pass Butterworth filter with a cut-off frequency of 10 Hz.

The validity of this approach is based on three assumptions:

1. The bicycle-rider system is asymptotically stable.
2. The process noise is uncorrelated with the measurement noise.
3. The human response has a linear relationship with the perturbation.

The last two assumptions are not likely to be strictly true, but in practice the FIR model has a decent Variance-Accounted-For (VAF) with the raw measurement data. This indicates that only little potentially useful information is removed from the data.

The approach described by [22] is slightly altered because two problems were encountered when using it for this research:

1. Upon visual inspection, the low-pass filtering of the FIR model caused smearing.
2. For some of the pavilion runs in the UC Davis data, the perturbations followed each other too closely. This caused one of the principles of the FIR model to become invalidated: the oscillations caused by the perturbation do not die out after the number of samples which was prescribed by [22], because a second perturbation was applied within that sample range (which was 768).

The cause of the first problem was traced to the properties of the Butterworh filter. For the cut-off frequency used in [22], the filter induces a significant group delays (frequency dependent delays) that can be between 0.08 and 0.15 seconds, which is in the same order of magnitude as human control delays. The group delay is plotted in figure 4.4. It can be seen that the 20 Hz cut-off frequency filter has the same, albeit reduced, issue. Removing the filtering step improves the results in terms of VAF with the experimental data for all runs in the data-set used by [22], leaving questions to why it is used. One explanation might be that when the FIR data pre-processing method is applied to the measured roll angle data from the UC Davis experiment, the unfiltered result is quite noisy. However, also [22] only used the steer angle to determine the fitting criterion. In this research the filtering step is skipped entirely.

The second problem can be solved by reducing the number of samples to 600. This did not have a big effect (<1% VAF difference for most data) on the VAF of the FIR model with the experimental data.

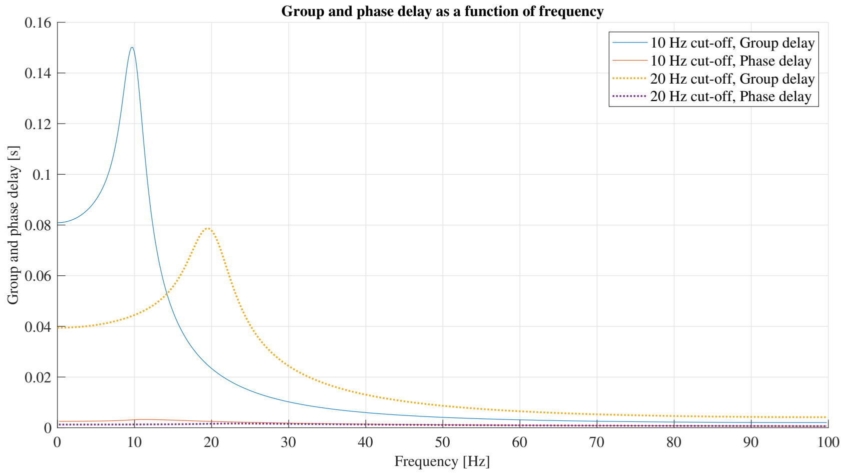

Figure 4.4: The group and phase delay are shown, in seconds, as a function of signal frequency for a 10 Hz and a 20 Hz low-pass filter cut-off frequency. The signal sampling frequency is 200 Hz.

In figure 4.5, a FIR model fitted to a run is shown for three different conditions (filter cut-off frequency = 10 Hz, filter cut-off frequency = 10 Hz, and no filtering) and for three different FIR sample sizes (200, 400, 600). A run from the UC Davis data set is used, which is sampled at 200 Hz.

This figure gives insight in several things:

- It shows the effect of the group delay induced by the low-pass Butterworth filter.
- There is little difference between the 400 and 600 sample FIR models. This indicates that using larger sample sizes when calculating the FIR don't have a big influence on the final FIR model form.
- The FIR model is unaffected by artifact present in the experimental data (the spike slightly after 2.5 seconds).

The VAF with the experimental data of an unfiltered 600 sample size FIR model for the entire UC Davis data-set is shown in figure 4.6. Several interesting things can be noted:

- The VAF is, on average, lower for the treadmill data. This is because the run durations are longer for the treadmill data: a generic pavilion run has between 2 and 5 perturbations, while the treadmill runs usually have more than 9 perturbations. Therefore the intrasubject variability is filtered out better for the treadmill data, but the VAF with the experimental data is lower.
- For both environments, the FIR-models fitted to the data generated by Charlie have a variance that is significantly higher than those of Luke and Jason. This is of course strange, as it essentially means that for some runs the FIR model can approximate Charlie better than for other runs. This could indicate that Charlie might not be a proficient bicycle rider, but this is not noted by [63], as he describes the participants to all have similar cycling ability. Maybe Charlie just has a more variable way of cycling.

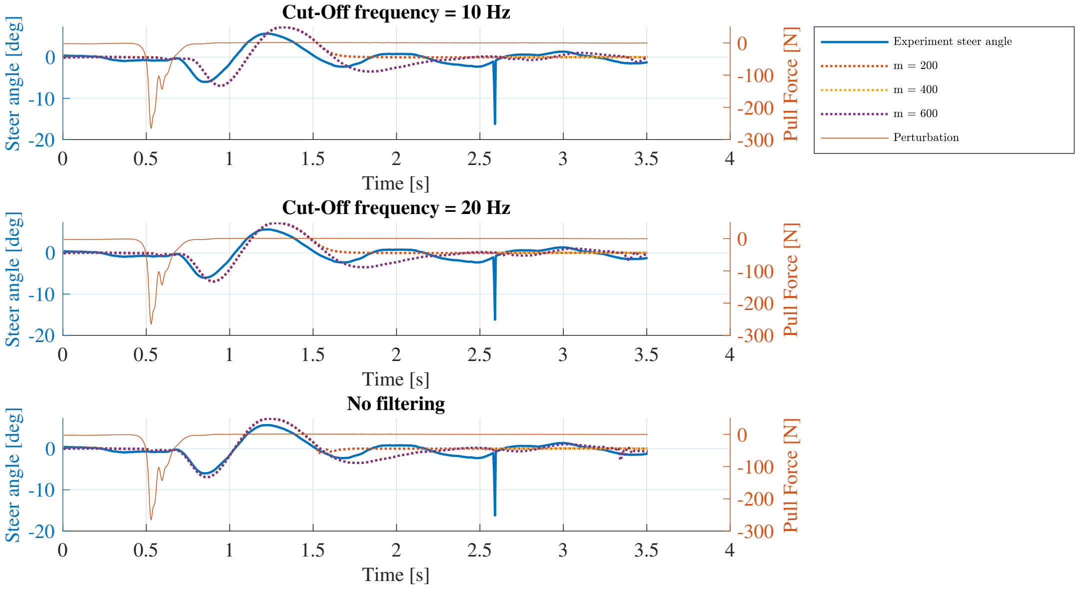

Figure 4.5: The FIR model fitted to a run is shown for three different conditions (filter cut-off frequency = 10 Hz, filter cut-off frequency = 10 Hz, and no filtering) and for three different FIR sample sizes (200, 400, 600). The original data is sampled at 200 Hz. Note that the evaluated run has a total length of 60 seconds and contains a total of 7 lateral roll perturbations perturbations which all contribute to the final FIR model. For visualization purposes, only a small slice of 3.5 seconds of the run is shown in the figure.

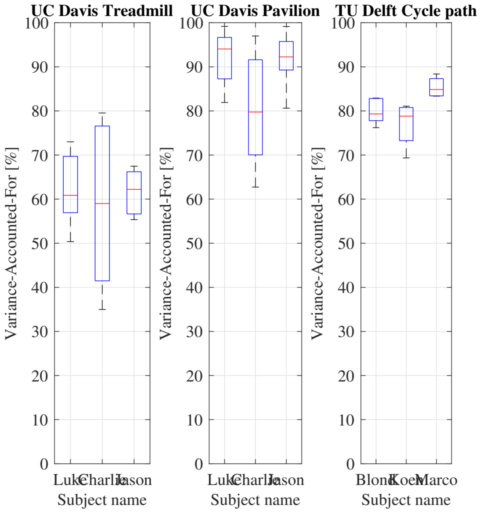

Figure 4.6: Boxplots of the VAF values of the FIR model with the experimental data.

Outliers

An outlier is a data point that differs significantly from the set of observations it belongs to. A difficult aspect of data analysis are classifying outliers and how to deal with them. In this case, we want to find experimental runs that are outliers; in the FIR models that are fitted to these runs to be precise. This is done by calculating the median FIR model from all runs in a data set. The outliers are then classified using the VAF's the individual responses have with this median response.

The classification itself is usually done with methods based on the statistical distribution of the data set. An example of such a method is removing outliers that are farther away than three times the data set variance from the mean of a normally distributed dataset. Unfortunately, in this case, the data set does not have a clear statistical distribution. Therefore, a lot of outlier detection algorithms can't be applied. In the end, the following (robust) outlier detection methods are used and compared:

1. Detection based on the Median-Absolute-Deviation (MAD)

2. Detection based on the Interquartile Range (IQR)
3. Naive detection based on a minimum VAF threshold (TH)

The MAD is the median of the distances between the data points and the median of the data set. It is formally defined in equation 4.1 for data set X. A data point is classified as an outlier if it is more than three Median-Absolute-Deviations away from the median. The IQR is the distance between the 25th and 75th percentiles. Using the IQR, outliers are defined as data points that are farther than 1.5 times the IQR away from the 25th or 75th percentile[3]. Extreme outliers lie farther than 3 times the IQR away. The IQR, and how it can be used to detect outliers, is visualized in figure 4.7. Both the MAD and IQR outlier detection methods are robust in the sense that they are not greatly affected by either very high or very low values. TH outliers are FIR models that have a VAF with the median response that is lower than a certain threshold. This means that this classifier does not depend on any specific distribution or spread of the data, but simply marks a data point as an outlier if it has a too low (or none if the VAF is zero) resemblance to the median. This is an advantage of using the VAF as a criterion.

$$MAD = median(|X_i - median(X)|) \tag{4.1}$$

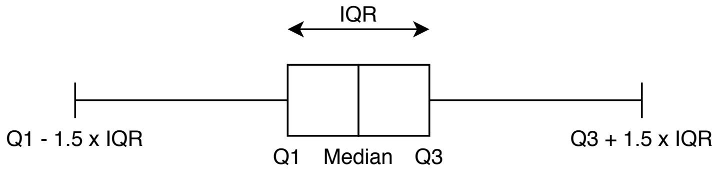

Figure 4.7: The IQR outlier detection method visualized. An outlier is a data point that lies beyond either the *Q1 - 1.5 x IQR* or *Q3 + 1.5 x IQR* bound

Because the outlier detection methods are a bit questionable statistically speaking, no experimental data is removed from the test data set. However, extreme IQR outliers are removed from the training set.

4.1.3. Model structure choice

As mentioned before, only white-box and grey-box model structures are applicable for this research. The reason for this is straightforward: I am not only interested in input-output behaviour which mimics a human bicycle rider, but I am also interested in what underlying sensory pathways the human uses. It is not possible to gain these insights using a black-box model structure.

Furthermore, it is quite naive to assume that a white-box model will be sufficient. There is simply not enough known about the internal processes of a human to arrive at a fully parametrised model that accurately predicts human bicycle balancing behaviour. This leaves a grey-box model structure as a suitable choice for parameter fitting.

Based on the literature review, the model structure can be divided in several components in a system block diagram, visualized in figure 4.9. Usually, there are many solutions for each block and these solutions often have multiple free parameters that can be fitted. Based on the research question and/or literature review, several choices already can be made:

- The Whipple bicycle model is the best available model to describe the bicycle dynamics (section 3.1)
- At minimum, the sensory delays τ_s should be included in the sensory model (section 3.3)
- Sensory integration and state estimation is best modeled with a Kalman filter (section 3.4)
- Neuromuscular dynamics ought to be included (section 3.6)

---

[3] For normally distributed data, the $1.5 * IQR$ bounds envelop 99.3 % of the data.

The system block diagram is a direct representation of the rider-bicycle system that has to be identified. The ”locations” of these components are shown in figure 4.8. How the components in figure 4.9 are implemented and what free parameters are associated with them, is discussed in the following sections.

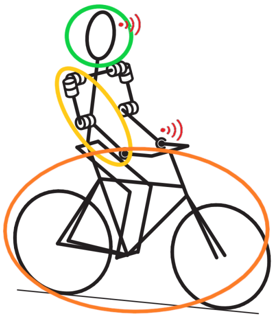

Figure 4.8: Graphical representation of the bicycle-rider system (adapted from Schwab et al. [77]). The state estimation, internal model, predictor and controller blocks are processes that originate from the CNS, which consists of the brain and spinal cord. This is represented by the green circle. The neuromuscular (arm) dynamics and input delay are tied to the yellow circle. The sensory dynamics and sensory delay that originate from different locations in the human body are indicated by the red signal signs. Lastly, and straightforwardly, the bicycle dynamics describe the bicycle, which is in the orange circle.

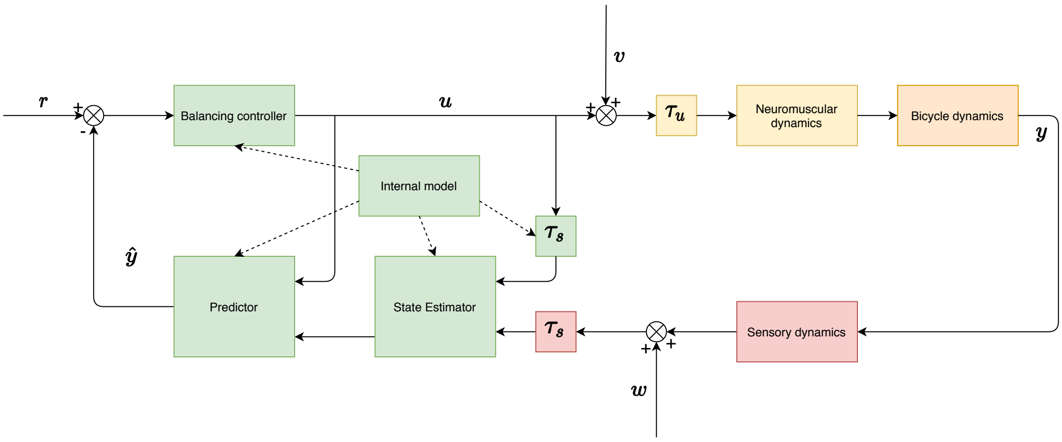

Figure 4.9: This is the system block diagram of the bicycle-rider system. The colors of the blocks corresponds with the colors in figure 4.8. There are several symbols present in the figure that originate from control theory. The reference is indicated by the r, the control input by u, process noise by v, the input delay by τ_u, measured outputs by y, measurement noise by w, measurement delay by τ_s and the estimated states by $\hat{y}$.

Controller

For the reasons mentioned in section 3.5, the controller will be structured in a single-loop[4] structure where feedback is applied with steer torque by (non-scheduling) gains on the estimated state $\hat{y}$. Two methods of determining the feedback gains are considered:

1. Directly fitting the feedback gains to the experimental data.
2. Calculating the gains with the LQR algorithm of which the weights are fitted to the experimental data.

When feedback gains are fitted directly, the free parameters are the states that exist in the internal model. The number of free parameters might be reduced evaluating their impact on the performance criterion.

---

[4] While there are multiple loops present in the system block diagram (figure 4.9), the actual feedback is calculated and provided in a single loop.

The parameters that define the LQR algorithm are the state and control weighting matrices **Q** and **R**. The weights are on the diagonal of these matrices and penalize state deviations from the reference and control effort, respectively. The dimension of these matrices depends on the dimensions of the state-space matrices that define the internal model. As minimum the internal model only contains simplified bicycle dynamics. If we consider the inverted pendulum bicycle model (section 3.1) as an example of such a model, the **A** internal model belongs to $\mathbb{R}^{2x2}$. With the steer angle as an input the **B** matrix is then $\mathbb{R}^{2x1}$. In this case the LQR **Q** and **R** matrices have dimension 2 and 1, associated respectively.

Because the rider is assumed to have only one control input, namely through manipulation of the handlebars, the dimension of the **R** weighting matrix will always be one. This enables us to set it to a constant value (1), as its effect is only relative to the weights in the **Q** matrix. Therefore, the free parameters when the LQR algorithm is used are weighting factors that penalize the states of the internal model. The number of free parameters might be reduced by setting weighting factors to zero. In that case, the state that is associated with that weighting factor is not penalized. In theory, one non-zero weighting factor is sufficient. The only requirement is that the system remains controllable.

Predictor

As discussed in section 3.4, the TDL (Tapped-Delay-Line) algorithm is the algorithm most suitable to model the predictive capabilities of the bicycle-rider. There are no free parameters to be fitted, only the selected internal model matters. In terms of model structure choice, the question is whether a dedicated prediction algorithm improves the results.

Bicycle dynamics - Passive rider model

The passive rider model describes the behaviour that is not active control behaviour, but does effect the dynamics of the bicycle model. This includes control strategies that bypass the usual sensor-control loop, and by doing that have little to zero delay. Hence, it is included in the bicycle dynamics block in figure 4.9. Examples of this are posture and muscle co-activation (section 3.2). Note that the the rider also increases the overall mass, mass distribution and inertia of the bicycle by sitting on it. In the case of the UC Davis experiment, where the rider is attached rigidly to the bicycle frame, this is control strategy independent and therefore only depends on the body characteristics of the participants. Practically speaking, this means that this contribution to the bicycle dynamics does not have to be fitted, but can simply be calculated based on some measurements of the participants.

The passive rider model is implemented inspired by the papers by Schwab et al. [77] and Doria and Tognazzo [25]. Rider upper-body lean is not considered for two reasons. The first reason is that it is unnecessary to include it for the UC Davis experimental dataset, because there upper-body lean was restricted. The second reason is that, while the TU Delft experiment did not restrict upper-body lean, the upper-body lean was not measured during the experiment. This means that the active nor the passive control input given by the upper-body can't be fitted to the experimental dataset.

Knowing this results in the following addition to the steer equation of the Whipple bicycle model:

$$\begin{gathered}\bar{\mathbf{M}}\begin{bmatrix}\ddot{\phi}\\ \ddot{\delta}\end{bmatrix} + v\bar{\mathbf{C}}_1\begin{bmatrix}\dot{\phi}\\ \dot{\delta}\end{bmatrix} + [g\bar{\mathbf{K}}_0 + v^2\bar{\mathbf{K}}_2]\begin{bmatrix}\phi\\ \delta\end{bmatrix} = T_\delta + I_r\ddot{\delta} + c_r\dot{\delta} + k_r\delta \\ \bar{\mathbf{M}}, \bar{\mathbf{C}}_1, \bar{\mathbf{K}}_0, \bar{\mathbf{K}}_2 \in \mathbb{R}^{1x2}\end{gathered} \tag{4.2}$$

The free parameters are: I_r as the inertial contribution, c_r as the damping contribution and k_r as the stiffness contribution of the passive rider holding the handlebars.

Neuromuscular dynamics

Neuromuscular dynamics are often included when modelling human control behaviour [64]. In the model, the neuromuscular dynamics describe the relationship between the efferent motor signal and the actual output torque generated by both arms on the handlebars[5]. Effectively, it limits the bandwidth of the efferent motor signal (and thus can be compared to a filter). It is described by the following transfer function:

---

[5] And because of this, the neuromuscular dynamics do not directly influence the bicycle dynamics. Therefore it's contribution is seen as separate to the passive rider model, whereas the passive rider model is included in the bicycle dynamics block in figure 4.9.

$$H_{nm}(s) = \frac{T_\delta^c(s)}{\hat{T}_\delta^c(s)} = \frac{\omega_c^2}{s^2 + 2\zeta\omega_c s + \omega_c^2} \tag{4.3}$$

For bicycle-rider identification, Schwab et al. [79] used shoulder muscle parameters from Happee et al. [38]: the cut-off frequency $\omega_c = 2.17 \cdot 2\pi$ rad/s and the damping coefficient $\zeta = \sqrt{2}$. Moore [63] found that by adding neuromuscular dynamics, his model gave a better approximation of the experimental data. These parameters are also used for this research.

Internal model

Both a LQR controller type, the Kalman filter state estimator and the prediction algorithm require knowledge of the system dynamics. In a human control context, this knowledge is called the *internal model*. Usually, all system dynamics are included in the internal model, which in this case consist of all "dynamics" blocks and delays in the system block diagram (figure 4.9). One could question the validity of this approach: does the human really has such detailed knowledge about its own dynamics and the bicycle dynamics? As a part of the model structure choice, the components of the internal model are either removed or dumbed down. The following (combinations of) possibilities are considered:

- Include all the dynamics present in the plant model (so, the neural, bicycle and sensory dynamics) in the internal model. This assumes that the human rider has perfect knowledge of his own bandwidth, the bicycle dynamics and sensor organ characteristics.
- Remove the neuromuscular dynamics from the internal model. Now we assume that the human does not have knowledge of his own bandwidth.
- Remove the sensory dynamics.
- Simplify the bicycle dynamics.

In order to have an internal model that can be used by the LQR controller, the TDL prediction algorithm and the Kalman filter state estimator, the internal model is compiled in a state-space representation.

Sensory dynamics

For all states that exist in the bicycle model, there exist appropriate sensors that can measure them. In table 4.5, the bicycle model states are summarized. In table 4.6, all the sensors organs which can measure these states are listed.

It must be noted, that while these sensors might measure the bicycle states directly, it could be possible that a derivative measurement actually provides the neural signal used by the human. An example: a human uses his vestibular organ to level his head, and (probably) also does this while cycling. This enables the possibility of the human to use the measurement of the deflection or movement of the spinal cord as a measure of roll or roll rate of the bicycle. However, in such cases, the importance of the primary sensor (the vestibular organ in this case) still stands.

Any combination of these sensors, measuring any kind of the states associated with them, may yield a well performing bicycle model structure. For all of these sensors, there are transfer functions available that describe their dynamics (section 3.3 provides a more detailed insight in the origins of these transfer functions).

However, just as the authors of [94] concluded, these transfer functions are in some cases to detailed to be implemented (primarily the ones describing the muscle spindles, as they are improper, excluding the possibility to convert them to a state-space notation). It does not make sense to "cherry pick" certain sensory dynamics that are convenient to implement in the framework provided by the chosen model structure.

Also, it is known that when different sensors are used in junction, the overall perception window of a human changes. This affect is not incorporated in the descriptions found in section 3.3.

Therefore, it is decided to not use the descriptions of any of the sensory dynamics during the identification procedure, but restrict ourselves to the sensory delays which are assumed to have the biggest impact.

Table 4.5: Bicycle states.

Name	Symbols
Bicycle roll angle, rate and acceleration	ϕ, $\dot{\phi}$, $\ddot{\phi}$
Bicycle steer angle, rate and acceleration	δ, $\dot{\delta}$, $\ddot{\delta}$
Bicycle forward position, velocity and acceleration	x, v, a
Bicycle lateral position, velocity and acceleration	y, $\dot{y}$, $\ddot{y}$
Bicycle heading angle, rate and acceleration	ψ, $\dot{\psi}$, $\ddot{\psi}$
Roll torque	T_ϕ
Steer torque	T_δ

Table 4.6: Sensors relevant for bicycle control.

Type	Name	Measured states	Symbols
Exteroreceptors	Vision	Position, velocity, heading	ϕ, $\dot{\phi}$, δ, $\dot{\delta}$, x, v, y, $\dot{y}$, ψ, $\dot{\psi}$
Interoreceptors	Vestibular - Semi circular channels	Roll rate, heading rate	$\ddot{\phi}$, $\ddot{\psi}$
	Vestibular - Otoliths	Linear accelerations	a, $\ddot{y}$
	Muscle sensors - Golgi tendon organs	Muscle force	T_ϕ, T_δ
	Muscle sensors - Muscle spindles	Muscle length and velocity	ϕ, $\dot{\phi}$, δ, $\dot{\delta}$

Sensory and input delays

Time delays can have a significant impact on the response of a closed loop system. As described in chapter 3.3, they are roughly two types of delays in the human controller: the afferent neural- and processing delays, and the efferent neural (control/input) delay. Combined they are the effective delay in the control loop. Depending on what sensory channels are used, the effective delay can vary a lot. Generally speaking, states are estimated with a higher delay by the visual system than states estimated with interoreceptor type sensory organs such as the vestibular system.

For bicycling, the states of the bicycle may both be measured by visual system or by the interoreceptors listed in table 4.6. The challenge is knowing which sensor type takes preference.

A problem is, is that different available sensory channels can be combined. This not only affects the sensing accuracy, but also the effective delay in the control loop. In other words: when a human controls something, he apparently does so with a single control delay. And this control delay is affected by which combination of sensors are used for the control task.

It is difficult, and might not be realistic, to mimic this fusing effect with a multiple-delay Kalman filter state estimator (I briefly go into my attempt at this in section 6.4).

A better solution is to approach the problem differently: instead of trying to model the time-fusion of all different sensory channels realistically, we can try to figure out the control delay of the bicycle rider and the states that are used by the bicycle rider.

By establishing the delay of the control loop, we can already learn which *type* of sensor organs are the most important. An effective delay lower than 0.15 s indicates that interoreceptors are the primary source of information. An effective delay higher than that does not definitely tells us that the visual system is used as the primary source of information, but at least tells us that it *could* be used in that manner.

Furthermore, we can establish which bicycle states are the most important. For example: if the model is only supplied with a roll-angle measurement, and that yields good results when compared to the experimental data, we can note that the vestibular organ is not required for bicycle stabilization (the vestibular organ provides rate and acceleration information).

State Estimator

To model human state estimation, a Kalman filter will be used. As mentioned in section 3.4, this method has precedence in human control modelling. A Kalman filter uses knowledge about the process and measurement dynamics of the system, inputs to the system, and the measured outputs of the system to give an estimate the internal model states (usually more accurate than single measurements alone). For it to work, a model of the plant dynamics, control input dynamics, the measurement dynamics (in state-space **A**, **B**+**D** and **C**, respectively) and the covariance matrices of the process and measurement noise have to be provided: **Q** and **R**, respectively.

The models of the various dynamics are present in the internal model of the human, and as such don't have to be specified separately. The noise covariance matrices are unique to the Kalman filter. The process noise covariance matrix **Q** contains variance and covariance information about the noise that is caused by unmeasured and unmodelled disturbances (e.g., unevenness in the road surface). The measurement noise covariance matrix **R** contains information about the noise caused by the measurement devices (organs in our case). It is difficult to provide this information as a lot these noise characteristics are unknown and have to be estimated. And then it becomes questionable: does the human rider estimate these noise characteristics in the same manner? To circumvent this problem, it is decided to assume the human has no access to noise covariance information, but only to noise variance information. This causes the **Q** and **R** matrices to only have non-zero values on the diagonal.

The size of the **Q** and **R** matrices depends on the dimensionality of he provided state-space matrices. The **Q** matrix has the same dimension as the **A** matrix of the internal model and the **R** matrix has the same dimension as the number of rows in the **C** matrix of the internal model. The number of rows in the **C** matrix is the same as the amount of sensory pathways included in the internal model.

The free parameters are the values on the diagonals of the **Q** and **R** matrices.

4.1.4. Model structure and free parameter overview and selection procedure

The structure choices and the available free parameters are listed in tables 4.7 and 4.8, respectively. Due to the many possibilities, the first step is to reduce the number of model structure choices. We want to do this while retaining a high number of free parameters, because this helps to ensure that the search space still encapsulates the global optimum. The global optimum is defined as the model that describes the entire experimental data set in the best way possible. To account for inter and intra subject variability, we also want to test as much model structure options as possible on an as large as possible fraction of the experimental data set. Due to computational constraints[6] these two objectives; trying all model structure combinations with all the parameters left free versus trying them on as much experimental data as possible, contradict.

By assuming that the effective delay and the controller have, next to the model structure choice, the most impact on performance[7] and that the state-estimation and passive rider model components have a relatively minor impact, we can limit the number of free parameters to those associated with the effective delay and the controller.

Now, by following a sequential approach to the structure selection in the order of assumed importance of the components, we can evaluate the importance of the separate components in terms of performance. In theory, if the structure is viable in terms of stability, this should result in the highest possible VAF value for the candidate model structure. When the possibilities are exhausted, the best candidate model structure is selected as a final model structure. Of this model structure, the remaining free parameters can be optimized and the free parameters of the controller can be reduced in order to yield a more general model. The entire procedure is visualized in figure 4.10. The optimization workflow and how it interacts with the model components in the system block diagram (figure 4.9) is visualized in figure 4.11.

How large of a fraction of the data set has to be used in order to conclusively fix any model structure

---

[6] The total number of model structure choices are, strictly speaking, not the faculty of the total number of possibilities, but there are certainly thousands of combinations to be evaluated before the entire search space is covered. Even when one function evaluation takes less than a second, the time needed to evaluate all possibilities is unpractical because a non-convex optimization algorithm for a problem with more than 10 free parameters quickly requires 1000+ function evaluations.

[7] In this research, performance is a measure of how well the model approximates the post-processed data. When an other type of performance is discussed, such as computational speed, this is stated explicitly.

option or exclude any free parameter can't be said in advance as it depends on the variability of the results.

A downside of the entire approach is, is that it is possible that the global optimum is excluded from the search space. It is often easier to prove that a certain solution is *not* the global optimum than that it *is* the global optimum. Therefore, we should not exclude common sense and heuristic search techniques in our search for the best bicycle-rider model.

An already identified common sense based technique, is the possibility to estimate the effective delay directly from the experimental data or with black-box type models which include delays. This can be done because both the bicycle states and the rider input were measured during the experiments. This means that a separate system identification cycle will be conducted in order to determine the effective time delay, which in turn gives information on what type of sensors are the most important. These results can then be compared to the results of the grey-box identification described earlier in this section.

Table 4.7: Model structure choices

Component	**Options**
Controller	Feedback gains / LQR
Prediction	Included / not included
Internal model	Same as plant model / without sensory dynamics / without neuromuscular dynamics / Simplified bicycle dynamics
Sensor channels	Options listed in table 4.5

Table 4.8: Model structure free parameters.

Component	**Free parameters**
Controller	$\mathbf{Q}_{LQR}$ or $\mathbf{K}$
State Estimator	$\mathbf{Q}_{Kf}$ and $\mathbf{R}_{Kf}$
Passive rider model	I_r, c_r and k_r
Sensory and input delays	τ_{sensor}

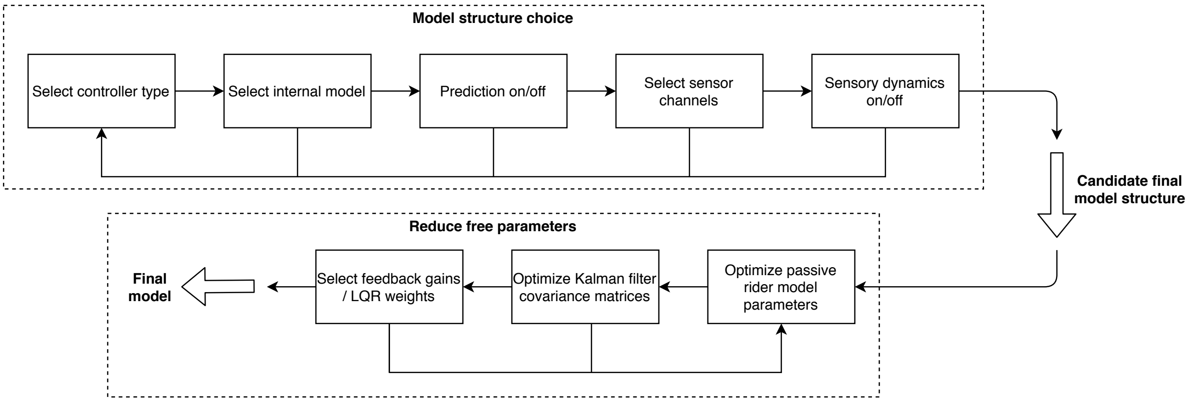

Figure 4.10: Model structure components and free parameter reduction approach. This diagram visualizes the proposed step-by-step method to reduce the number of model structure components and free parameters. First, the model structure is determined iteratively while the parameters associated with the controller are all kept free (so, either all of the LQR weights or the maximum number of feedback gains). After that, the passive rider parameters and the state estimation parameters are optimized. Finally, the free parameters of the controller are reduced. The effective delay is always kept as a free parameter.

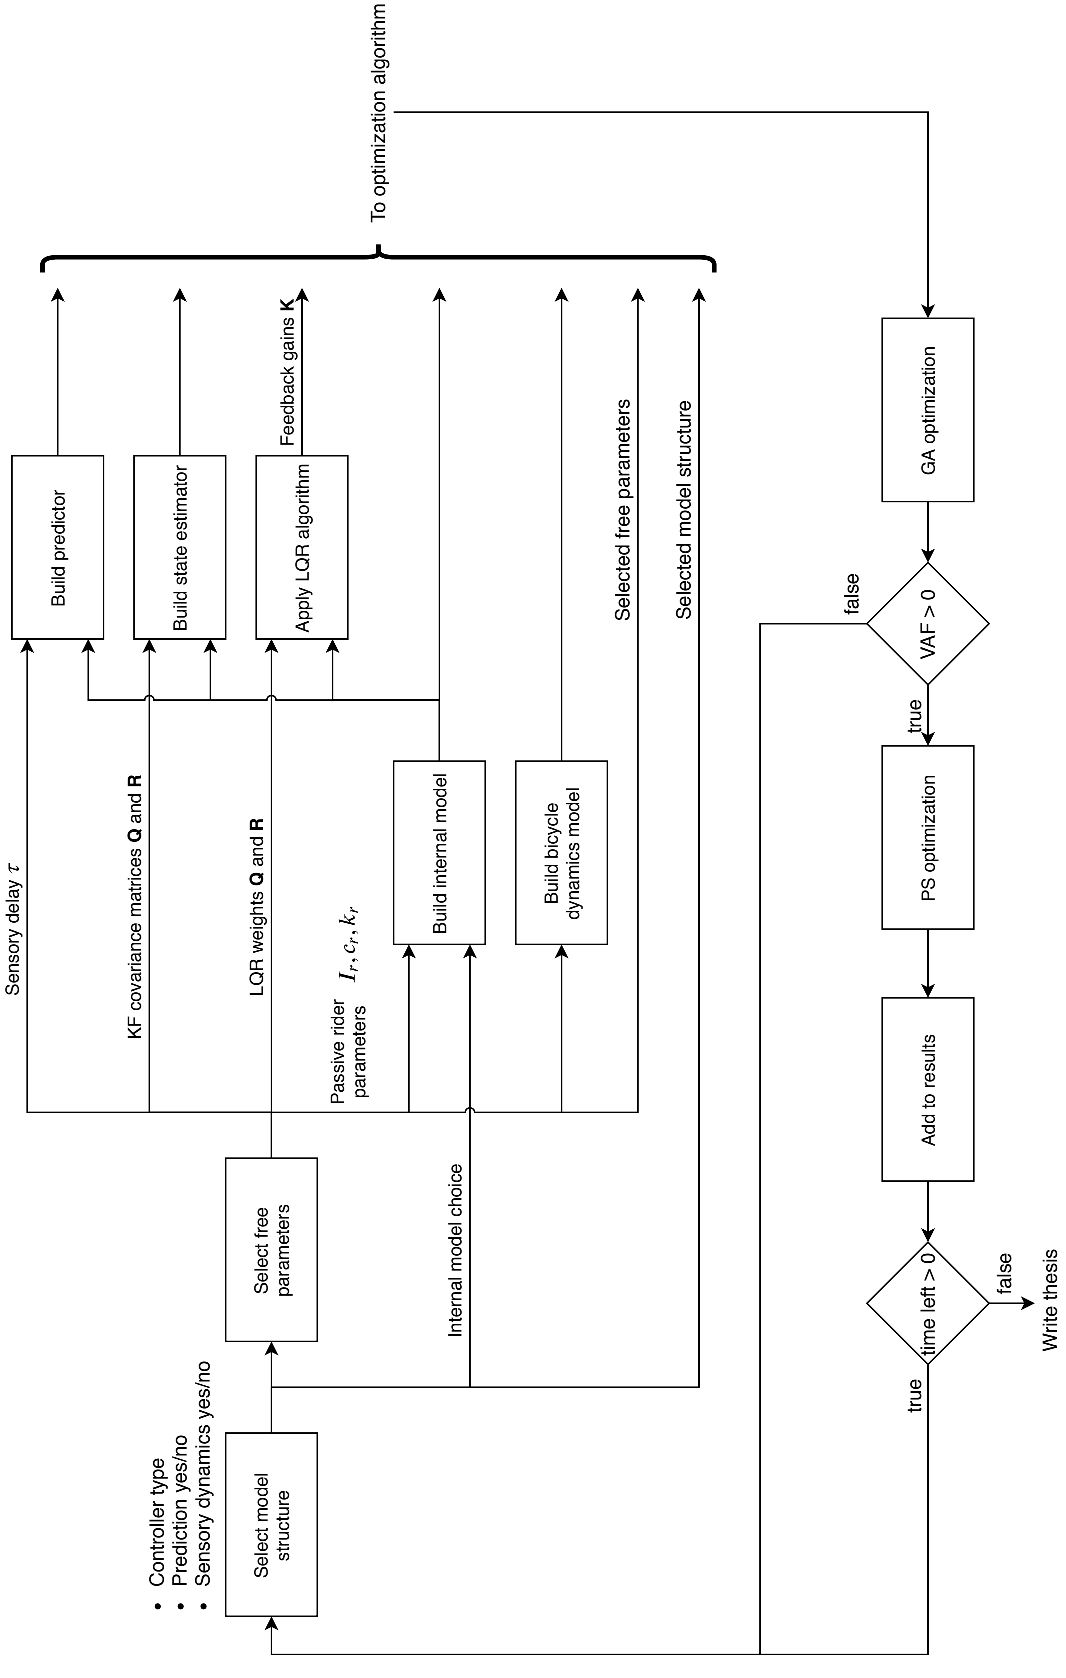

Figure 4.11: Model and optimization workflow diagram. All the possible changes and their impact on the model are visualized in this figure. Also, the iterative nature of the system identification procedure is visible in this block diagram.

4.1.5. Fitting procedure

Once a candidate model structure and the free parameters are selected, the free parameters have to be fitted in such a way that the resulting model mimics the experimental data as good as possible, while preventing overfitting or underfitting[8]. The ”goodness” of the fit is usually evaluated with a fitting criterion, which is calculated using a cost function. Depending on the criterion used, this results in a minimization or maximization problem, which are essentially the same. As mentioned in the introduction, this problem can both be solved using iterative and non-iterative methods. Here, we focus on iterative methods, as they are less dependent on model structure choice and more straightforward to implement.

Optimization method

The free parameters are fitted in two stages. First, due to the unstable nature of the bicycle, a parameter set has to be found which stabilizes the bicycle. This is done by means of a Genetic Algorithm (GA). The GA is an iterative optimization method that is based on natural selection. It generates candidate solutions and, at each iteration, selects the most fit candidates to generate new solutions. Because a certain amount of randomization is allowed by also allowing some less fit candidates to survive and through candidate mutations, a GA can find the solution that corresponds to the global optimum.

In theory, using only a GA is sufficient to find the parameters that yield the best solution. However, the downside of this is, is that this takes an extreme amount of time. Therefore, if the fitness function stabilizes, which means that a solution has been found which already approximates the experimental data quite well, the final parameters are found by using a local optimization method. A short introduction to the various optimization methods can be found in appendix A.3.

Fitting criterion

Each optimization method uses a cost function to evaluate the fitness of a candidate solution. There are several criteria available that can be used to compare different signals. A straightforward and intuitive criterion, is to use the VAF criterion. A higher VAF indicates that two signals are more similar, with a VAF of 100 % indicating two signals that are exactly the same. The VAF can be calculated using equation 4.4.

$$VAF = \left(1 - \frac{var(y-\hat{y})}{var(y)}\right) \cdot 100\% \tag{4.4}$$

Another, less intuitive criterion, is the Root-Mean-Square-Error (RMSE) criterion. The RMSE is related to the VAF, as it also indicates how close two signals are together. A RMSE of zero indicates two signals that are the same. Values larger than zero indicate some measure of difference between the signals. This value is not normalized, so no direct conclusions can be drawn, except for that larger values indicate a higher dissimilarity.

$$RMSE = \sqrt{\frac{\sum_{k=1}^{T}(\hat{y}_k - y_k)^2}{T}} \tag{4.5}$$

4.1.6. Model validation

The aim of model validation is to test the predictive capabilities of the fitted model structure. This is done by splitting the experimental data set in two parts: a part where the model structure is fitted to, and a part which is used to test the found model. To allow for a fair comparison, it is important is that the two parts where gathered under the same experimental conditions. The splitting method used in this thesis is straightforward: the runs are ordered according to the forward velocity. Then, all uneven runs are placed in the train data set. This mixes the runs both in terms of participants and in date they were conducted, but also ensures that the test set has runs included evenly over the entire forward velocity range.

The result is a set of fitted free parameters for each run in the train set. Parameter sets for which the fitting procedure failed are excluded. The goal is to extract a single forward velocity independent parameter set from the remaining parameter sets. If a certain parameter is found to be forward velocity dependent, they can be made forward velocity independent with the method described in [22].

---

[8] Overfitting means that to many parameters are fitted, which causes the model to lose generality, which results in poor predictive performance. Underfitting is the reverse of overfitting.

In a nutshell, this approach is to fit a forward velocity dependent polynomial (preferably a first order polynomial) to the fitted parameters.

If a parameter is found to be forward velocity independent, the median parameter in the train set is selected. Using the median instead of an average reduces the impact of possible outliers.

The model validation approach is visualized in figure 4.12.

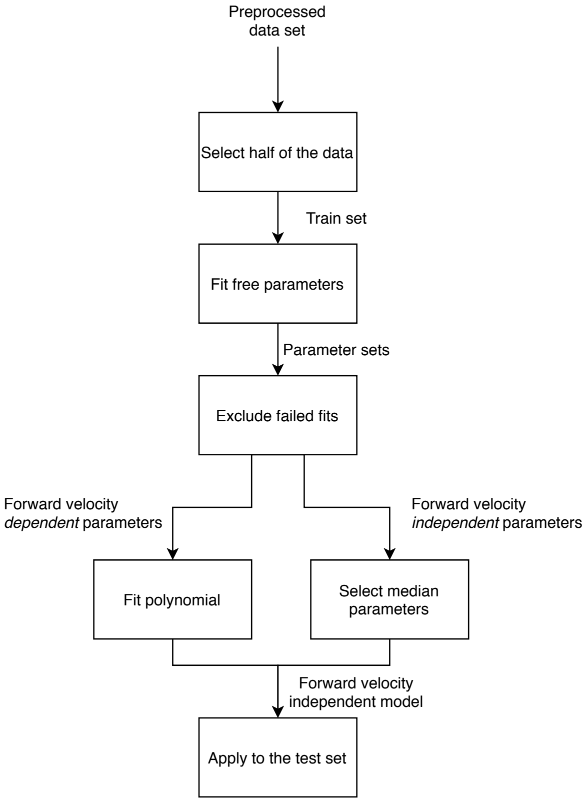

Figure 4.12: Validation approach

4.2. Modelling aspects

In this section the implementation of the model is discussed.

4.2.1. Software used

All modelling was done in Matlab. A complete list of dependencies can be found in table 4.9.

Table 4.9: List of dependencies

Dependency name	Version
MATLAB	9.7
Control System Toolbox	10.7
System Identification Toolbox	9.11
Signal Processing Toolbox	8.3
Symbolic Math Toolbox	8.4
Optimization Toolbox	8.4
Global Optimization Toolbox	4.2
Statistics and Machine Learning Toolbox	11.6
Curve Fitting Toolbox	3.5.10

4.2.2. Mathematical derivation of the model components

Because the TDL predictor, the Kalman filter and the LQR algorithms are all defined using state-space system notations, all the descriptions of the dynamics should also be in this format. After the conversion of the individual components, they are combined in a single state-space description. For the sake of readability, they derivations are not included here, but can be viewed in appendix A.

4.2.3. Discretization

The plant is discretized using zero-order-hold with a sampling frequency of 200 Hz. This sampling frequency is selected because it is the minimum frequency at which the experimental data set is collected. To verify that this approach does not introduce significant truncation or discretization errors, an analysis is done which compares several numeric integration methods and sampling frequencies.

The system is evaluated for a duration of ten seconds with a step input. In figure 4.13 it can be seen that the discretized system has a final error for this situation of 10^{-10} rad when compared with ODE4 and ODE8 numeric integration methods. This is deemed acceptable.

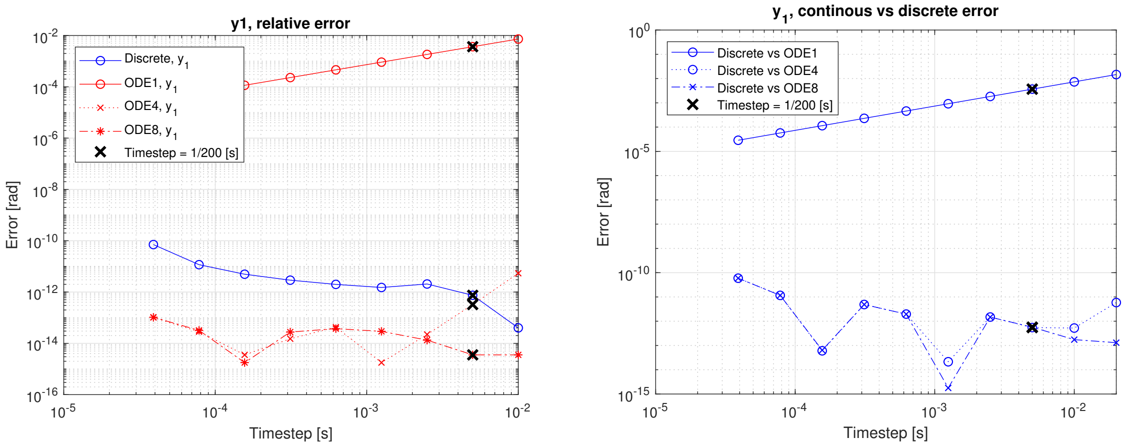

(a) **(b)**

Figure 4.13: Accuracy error

5

Results

As discussed in the previous chapter, the first step to a viable rider control model is to reduce the number of model structure choices. To give insight in this process, intermediate results and their implications are also discussed in this chapter. This is done for each model structure component for which a choice had to be made (refer to table 4.7 for an overview of the choices and to figure 4.10 for the approach). Afterwards, the final model structure is presented. This model structure is then applied to the available data sets, of which the results are presented. For the sake of readability, the results are interpreted as they are presented. The overarching interpretations, implications and limitations are discussed in chapter 6.

5.1. Fitting method

While for the global optimization step the Genetic Algorithm is used, for the subsequent local optimization method two different optimization algorithms were tried and compared: interior-point optimization and pattern-search optimization. In terms of the quality of the optimum found, pattern-search optimization outperformed interior-point optimization. A downside of pattern-search optimization covariance based parameter reduction methods can't be used, because it is a gradient free optimization method. Also, just like other polling based optimization methods, pattern search optimization slows down quite a bit with an increasing number of parameters. Up till five free parameters is manageable, meaning that overnight optimization of an entire data set is possible.

It is important that the optimization method finds a global optimum, defined as the parameters set that yields the best performance score in the entire search space. This is difficult (read: impossible for our problem) to prove, so we limit ourselves to a more practical definition of a ”good” optimum: the optimization method should reliably find the same optimum for multiple random starting conditions. This is verified by running the optimization algorithm multiple times.

The settings used for the optimization algorithms are in table 5.1.

Table 5.1: Genetic algorithm and pattern search settings

	Max generations / iterations	Tolerance	Cross-over fraction [%]	Pop. size	Stall generation
GA	100	0.001	85	150	35
PS	2500	0.001	N.A.	N.A.	N.A.

In figure 5.1, the performance of the GA is illustrated.

5.1.1. Remarks on parameter bounds

While it is possible to leave the optimization problem unbounded, the GA population size can be kept smaller if the parameters are bounded. The reason for this is that both a larger initial population size and bounded parameters increase the chance of finding viable solutions. But, increasing the population size has as downside that it increases the computation time required. The GA generally fails to converge if

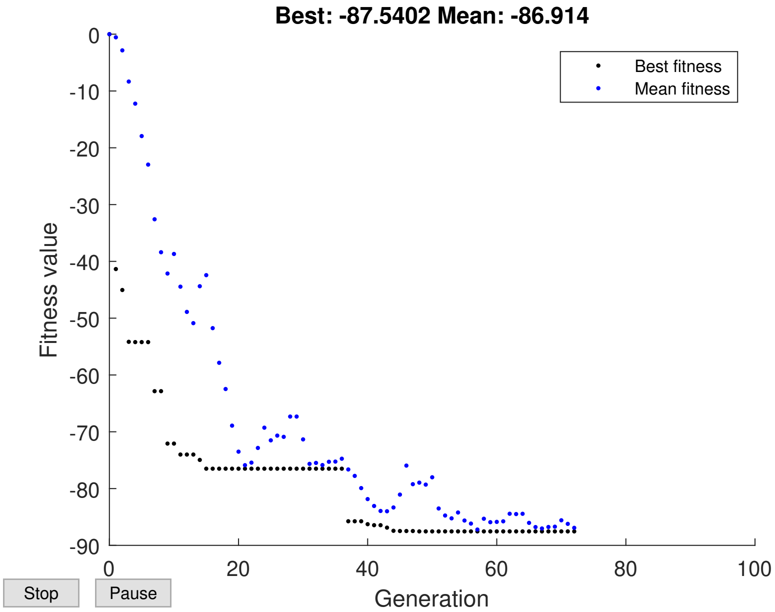

Figure 5.1: This figure shows the evolution of the performance criterion as a function of generations. Note that the optimization problem has been recast as a minimization problem.

no viable solutions are found within the fist few generations. Of course, care has to be taken that the bounds don't influence the optimum found.

Generally, the lower bounds are set at 0.001. The upper bounds are set at 950.000 and 0.5 for the Linear-Quadratic-Regulator (LQR) weights and effective delay, respectively.

5.2. Model structure selection

5.2.1. Controller

Two different controller types are applied to the experimental data sets: fitting the feedback gains directly, or using the LQR algorithm to determine optimal feedback gains based on the weighting factors. Structure wise, both controller types are the same, as they both give state-feedback. However, both conceptually and in practice there are differences:

1. The LQR algorithm has a higher dependency on the correctness of the plant model used, whereas determining individual feedback gains can "capture" the in-correctness of the model within the gains. However, when the plant model is correct, the LQR algorithm requires less free parameters to be fitted (minimally only one to calculate all the feedback gains for all forward velocities), and thus results in a more general model.
2. Due to the issues mentioned above, the LQR algorithm has reduced performance when the plant model is less correct.
3. Gains determined with the LQR algorithm are more likely to stabilize the plant, this simplifies the parameter fitting procedure.
4. The weights in the LQR algorithm are penalties on state deviations. Knowing which states are penalized the most, gives insight in the control strategy the rider uses.

The second point is illustrated in figure 5.2. Here, the two controller types are fitted (in both cases leaving all state weights/gains free) to the Finite-Impulse-Response (FIR) model of a pavilion run in the UC Davis data set for two different scenarios. In the first scenario, the bicycle dynamics model based on the UC Davis experiment is used. For the second scenario, this model is altered by removing the cross-terms in the mass, damping and velocity independent stiffness matrices of the Whipple model equations. The velocity dependent cross-term in the stiffness matrix is halved for good measure. Essentially, this leaves a simplified model, which is somewhat reminiscent of the inverted pendulum bicycle model (3.1). It can be seen that both control methods perform well with an intact bicycle model, but the LQR controller degrades more when the bicycle model is simplified.

This effect also causes problems when fitting model structures to the TU Delft data. The LQR type controller does not perform well for that data. This can be caused because in that experiment the upper-body movement of the riders is not restricted. It is difficult to compare the raw data of both experiments fairly (different perturbation sizes, different riders, different bicycle, different environment), but it looks like that the roll angles are higher for the TU Delft experiment. This is likely due to the roll-motion of the bicycle frame being less stiff.

Concluding, the LQR algorithm will be used as a controller. It should result in better fitting performance and more insight in the control strategy used by the rider. Due to the issues with the TU Delft data and the experimental conditions of the UC Davis experiment match the chosen model structure better, we will focus on the UC Davis data-set for further model structure choices.

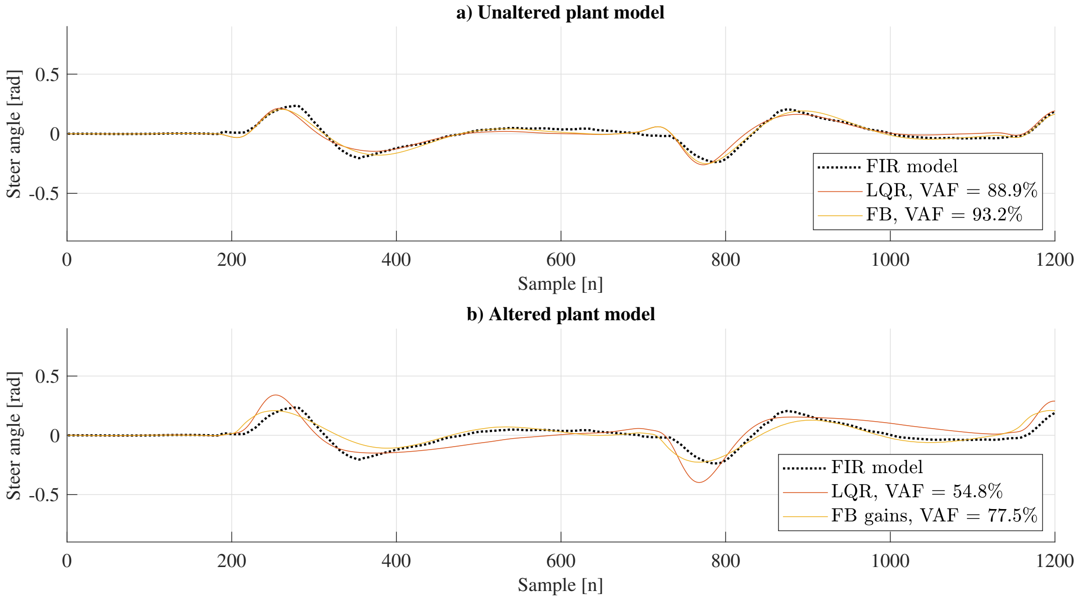

Figure 5.2: The above plots show a comparison between both controller types fitted using **a)**, a known to be reasonably correct plant model, and **b)**, a simplified plant model. The Variance-Accounted-For (VAF) drop for the feedback gain controller is a lot lower (-15.7%), than the VAF drop of the LQR controller (-34.1%).

5.2.2. Internal model

In this section, the results of the various internal model options applied to the experimental data set are discussed. Starting with an internal model were the descriptions of all dynamics present in the plant model (neural, bicycle and sensory dynamics) are included in the internal model, the options are:

- Simplify the bicycle dynamics.
- Remove the neuromuscular dynamics from the internal model.

Simplify the bicycle dynamics

As the previous section illustrated, the performance of the LQR algorithm depends on the correctness of the system model provided. When the plant model is kept the same, and just the fidelity of the internal model is reduced, both the feedback gain controller and the LQR controller quickly become unable to stabilize the plant. This can be attributed to the reduced performance of the Kalman filter. Finding a combination of a simplified plant model and Kalman filter settings is an extensive process which has yet to be described in literature. Therefore, after some initial trials, this is not further pursued.

Remove the neuromuscular dynamics from the internal model

The question is if the human rider has ”internal” knowledge of the transfer between the efferent motor signal and the actual output torque. Comparing two model structures, one without the neuromuscular dynamics included in the internal model, and one with the neuromuscular dynamics included, revealed that a description of the neuromuscular dynamics need to be included in the internal model to reach a stabilizing controller.

This is in line with other research into human control which include neuromuscular dynamics (such as the research by McRuer). In their research the human controller is usually mimicked by state feedback gains of some sort are used. It is explained in section 5.2.1 that the feedback gains fitted with this procedure account for the effect of the neuromuscular dynamics if they are included in the model.

While the use of neuromuscular dynamics in models describing human control behaviour is widespread (refer to [64] for an overview) and it’s importance in bicycle control modeling was already noted by [63], it is nice to see those findings verified. In figure 5.3, a quick and dirty comparison for a single experimental run is shown.

The figure features the result of model structure fitted to a single run from the UC Davis data. Free parameters were the weighting factors of the LQR algorithm and the effective neural delay.

It can be seen that the difference between both models is quite significant, telling us that the literature is correct.

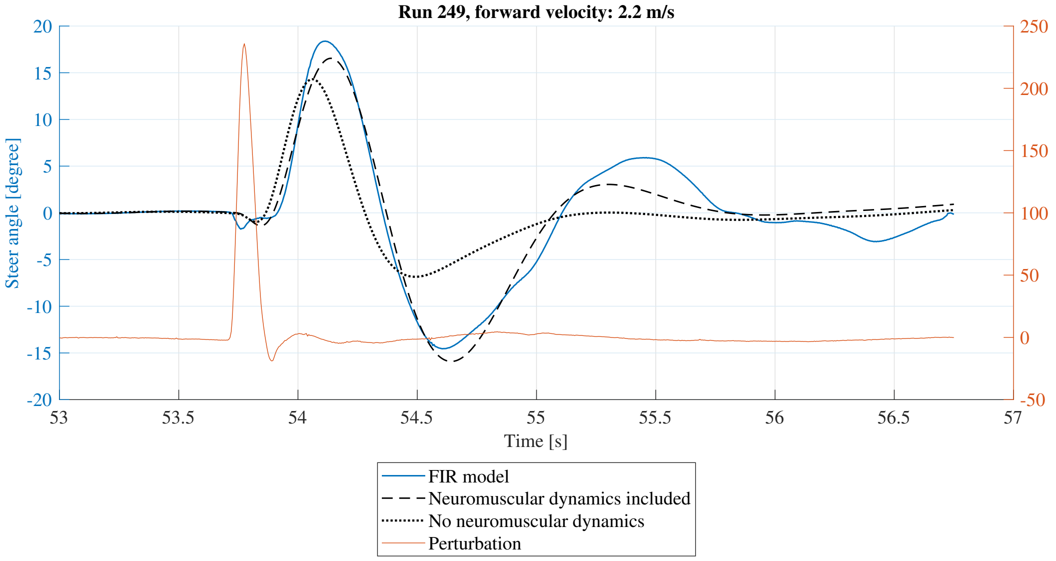

Figure 5.3: A comparison between a model which includes neuromuscular dynamics (both in the internal model and in the plant) and a model which does not include neuromuscular dynamics. Do note that the results are only plotted for a small time range. The original run, to which the FIR model is fitted, is over 80 seconds long and includes nine lateral perturbations.

Final internal model form

The final internal model has the components listed in table 5.2. The states of the neuromuscular dynamics (which is derived in appendix A) correspond with the steer torque and steer torque rate that is given as input to the bicycle. These states can be penalized in the LQR algorithm. This can be interpreted as a penalty on control effort.

Table 5.2: Internal model components

Component	States
Bicycle dynamics	ϕ, δ, $\dot{\phi}$, $\dot{\delta}$
Neuromucsular dynamics	T, $\dot{T}$
Heading	ψ

5.2.3. Prediction

As established in chapter 3.4, humans are capable of predicting the effects of their actions. After some experimentation with turning the prediction algorithm on and off for various runs, the impact of including the prediction algorithm never had a negative effect. Depending on the set effective delay and the forward velocity it either had a marginal impact (around 5 % negative VAF difference for effective delays below approximately 0.06 seconds) or it went off a cliff into instability.

For this reason, the prediction algorithm is included in the final model structure.

5.2.4. Sensory channels used

The bicycle model that is used is described by four states: the roll and steer angle, and their derivatives. By enabling different combinations of sensory channels measuring these states, a combination can be found which gives the best fit to the experimental data. This is done while leaving all the other possible free parameters free: the LQR algorithm can place a weight on all states and the sensory delay is left free to be fitted. A downside of this is, is that it results in a high computational demand, increasing the time before the optimization algorithms find a solution. To reduce the required time to a practical amount, the combinations of sensory channels are only tried for a limited number of runs across the entire forward velocity range.

The resulting best sensory channel combinations are in table 5.3. Using $\dot{\phi}$ and δ as measurement channels results in the best performance in all of the runs for which the combinations where tried. No forward velocity dependency is found. Other things, not shown in the table, can also be learned from the results:

- Simulations without the roll rate as a sensory input failed to reach VAF values higher than 20 %. And only using the steer angle or steer rate as a sensory input never results in a successful fit to the experimental data. Therefore, it is likely that roll-rate is the most important measured quantity.

- Adding sensory channels beyond the roll and steer rates reduced VAF. This is also interesting, because this should only increase the accuracy of the estimated states, indicating that some degree of inaccuracy in state estimation increases fidelity. Note that this inaccuracy does not have to originate from the sensory dynamics, but could also be caused by process or sensor noise. The inaccuracy of the estimate is mainly present in the roll angle estimation, which is overestimated. The effect is illustrated in figure 5.4.

- The fitted sensory delay always approached zero. This is of course quite a shame, and likely is because the passive rider dynamics are not included in the model structure. This is further discussed in section 5.3.4.

Table 5.3: Best sensory channel combinations.

Run ID	Best combination	VAF [%]	v [m/s]	Environment	Rider
249	$\dot{\phi} + \delta$	85	2.2	Treadmill	Jason
489	$\dot{\phi} + \delta$	88	2.2	Pavilion	Jason
253	$\dot{\phi} + \delta$	88	4.2	Treadmill	Jason
661	$\dot{\phi} + \delta$	91	4.2	Pavilion	Luke
670	$\dot{\phi} + \delta$	83	4.5	Pavilion	Luke
TUD 1	$\dot{\phi} + \delta$	87	2.5	Cycling path	Blond

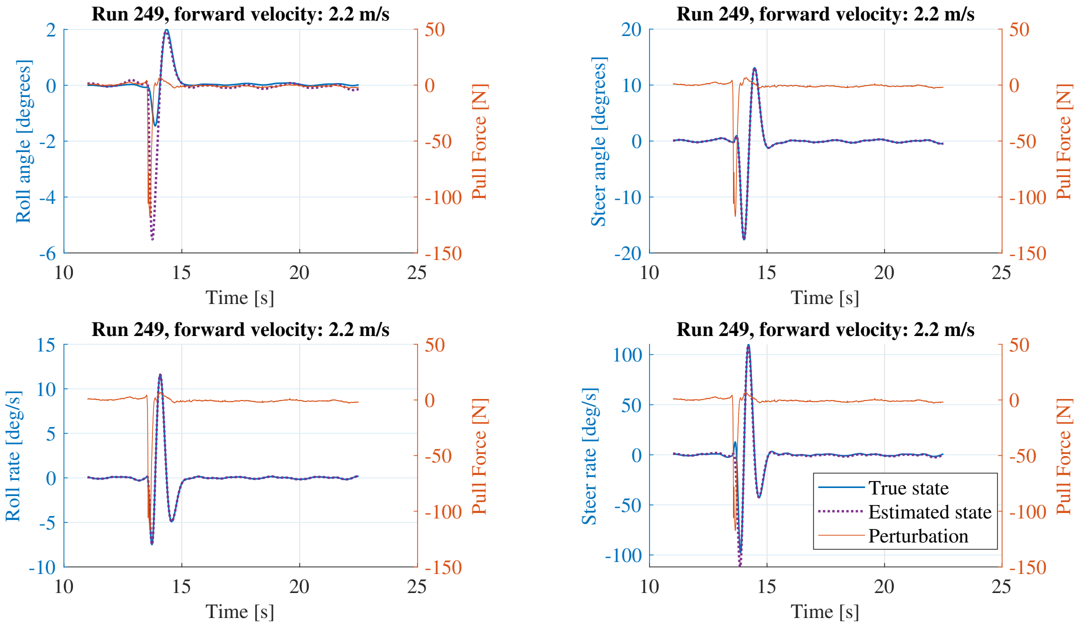

Figure 5.4: In this figure the true bicycle states (blue) and the estimated bicycle states (purple) are plotted together with the lateral perturbation force (red).

5.2.5. Model structure summary

Using the lessons learned in the preceding sections, a model structure with the following components and free parameters was selected:

Table 5.4: Model structure choices

Component	Options
Controller	LQR
Prediction	Included
Internal model	Whipple bicycle model with neuromuscular dynamics

Table 5.5: Model structure free parameters.

Component	Free parameters
Balancing controller	$\mathbf{Q}_{LQR}$ (7)
State estimator	$\mathbf{Q}_{Kf}$ (2) and $\mathbf{R}_{Kf}$ (7)
Neural sensory and input delays	τ_{sensor} (1)
Passive rider model	I_r, c_r and k_r

5.3. Reduce free parameters

To prevent over fitting and to reduce simulation times, the number of free parameters should be reduced. In following section, the importance of the free parameters associated with the passive rider model, the state estimator and the controller are evaluated.

5.3.1. Controller

A property of the LQR algorithm is that it is not required to put weights on all states or inputs. Because the weights are relative, and there is only one control input, the weight on the control input is kept constant. By trying different combinations of LQR weights that are left free to be fitted, we can find the best performer. Good performance is defined as a set of weights that have good predictive capabilities: they represent the entire data set well.

Time-wise, it is not possible to train and test each controller weight combination for each environment (treadmill, pavilion, cycling path) in order to find the best performer. Some experimentation revealed the following things:

- As expected, leaving all weights free yields the best VAF when they are fitted to experimental data, but when the controller defined by those weights is tested with another slice of the experimental data, the results are not satisfactory: the free parameters are overfitted.
- Good performing weight combinations are different for the UC Davis and TU Delft data. This is likely because the Whipple bicycle model that is used in both cases is incorrect for the TU Delft data (due to the upper-body movement not being restricted). This confirms that the data-sets can't be compared directly.
- A single weight on the roll angle or the heading already yields quite good training results for the UC Davis data for both the treadmill and pavilion environments.
- Two weights are required for good results with the TU Delft data gathered on the cycling path.
- A weight on the heading angle generally makes predictive performance worse.
- There seem to be differences in control behaviour between different subjects.

Because one or two LQR weights already yield good training performance, the search is limited to combinations of one or two weights. Well performing combinations are shown in tables 5.6, 5.7 and 5.8.

Across the environments, weights placed on the roll rate and roll angle perform best. However, for the UC Davis data, only placing a weight on roll already yields good results. Because of the enticing prospect of describing human bicycle balancing with only one fitted parameter, the detailed results of that study are presented and compared in the next sections.

The results of the TU Delft data are also presented, but for the case when $Q(\phi)$, $Q(\dot{\phi})$ are selected as free parameters.

Table 5.6: The standard deviation σ and the mean μ for train and test **treadmill** data of different LQR weights left as free parameters. Note that only 1 test subject (Jason) was compared to reduce the effects of inter-subject variability.

	Train VAF		Test VAF	
Q weights on	$\mu[\%]$	$\sigma[\%]$	$\mu[\%]$	$\sigma[\%]$
ϕ	85.0	7.8	77.6	10.9
δ	55.3	11.0	39.9	11.2
$\dot{\phi}$	32.4	9.0	33.2	8.4
$\dot{\delta}$	0	0	0	0
T	15.0	22.6	13.8	22.5
$\dot{T}$	0	0	0	0
ψ	87.0	10.9	63.6	20.8
ϕ, ψ	90.9	5.5	63.6	20.8
$\dot{T}, \psi$	89.8	7.6	69.6	22.4
$\phi, \dot{\phi}$	86.7	7.3	78.6	10.8

5.3.2. State estimation

It is decided to keep the measurement and process noise covariance matrices of the Kalman filter as identity matrices. This is done for the following reasons:

- Because during the model structure evaluations discussed until now, the model performs well without altering them.
- Sensory noise is not modeled.

Table 5.7: The standard deviation σ and the mean μ for train and test **pavilion** data of different LQR weights left as free parameters. Note that only 1 test subject (Luke) was compared to reduce the effects of inter-subject variability.

	Train VAF		Test VAF	
Q weights on	μ[%]	σ[%]	μ[%]	σ[%]
ϕ	65.5	23.7	55.6	27.7
$\dot{\delta}$	25.6	17.4	15.2	4.4
T	16.6	22.7	11.3	21.4
ψ	58.9	34.9	49.9	36.5
$\phi, \dot{\phi}$	74.4	13.2	55.4	27.8

Table 5.8: The standard deviation σ and the mean μ for train and test **cycling path** data of different LQR weights left as free parameters. Note that only 1 test subject (Blond) was compared to reduce the effects of inter-subject variability.

	Train VAF		Test VAF	
Q weights on	μ[%]	σ[%]	μ[%]	σ[%]
$\phi, \dot{\phi}$	79.5	3.1	73.6	8
$\phi, \dot{\delta}$	18.5	35	0	0
ϕ, δ	84.7	5.5	72.1	15.9

- It greatly reduces the number of free parameters (even if only the diagonal values of the noise covariance matrices are left free).

5.3.3. Sensory delay

As discussed in section 4.1.3, it is likely that the sensory delay both has an big impact on the results, and gives an indication of which sensory pathways are the most important for the balancing task.

Unfortunately, as discovered, the fitted sensory delay tends to go to zero in the simulations done up till now. This is likely due to failing to model the passive rider dynamics correctly.

Because of this, it does not make sense to leave the sensory delay as a free parameter.

5.3.4. Passive rider model

The lack of a decent passive rider model is the prime candidate for causing the fitted sensory delay to go to zero. Unfortunately, it is not practical to optimize the passive rider parameters for overall train-test performance. This is because the plant and internal models change for each cost function evaluation, which excludes the possibility to do parallel computations in Matlab. This increases the duration of a optimization, for for example all treadmill runs, to multiple days if not weeks.

The papers by [78] and [25] present some inertia, stiffness and damping values that could be incorporated in the bicycle model. Initial results are promising and are shown in figure 5.5, for which two model structures with a 0.1 s delay are fitted. In one model structure the bicycle model has increased steering inertia and stiffness terms inspired by [78] and [25] (the steer inertia is increased by $1kg \cdot m^2$ and the steer stiffness is increased by $200Nm/rad$. The bicycle model in the other model structure is unaltered.

The figure reveals that the steer angle remains constant as the bicycle is perturbed. However, for the model structure without passive rider dynamics, this is not the case: the handlebar steers out of the fall. This is because the steer and roll angles can influence each other. This movement is prevented by including passive rider dynamics in the model.

Unfortunately, due to time constraints, the prospects of this method could not be explored elaborately. But the model structure with the added passive as described above is applied to a slice of the data set in section 5.4.4.

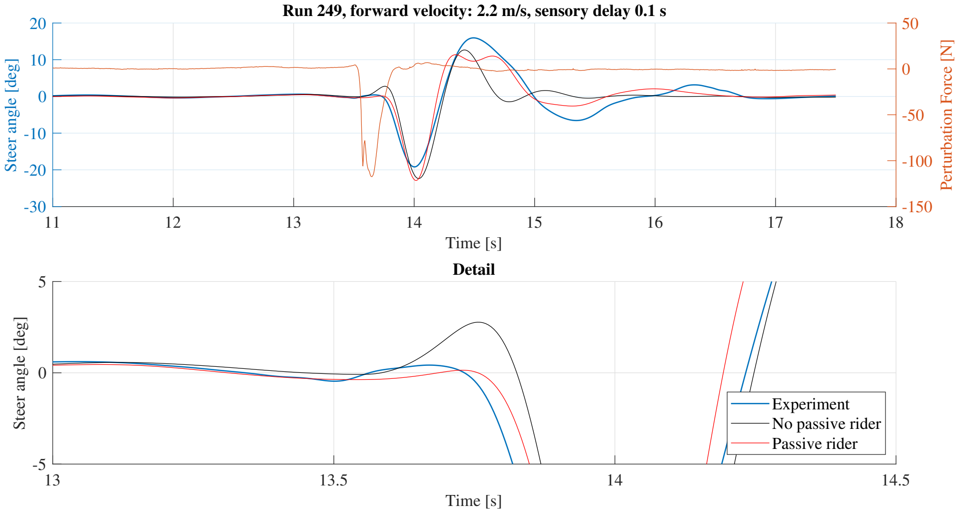

Figure 5.5: The FIR model fitted to the experimental data of run 247 are compared with a fitted model structures with and without added passive rider dynamics. The passive rider dynamics are included in a simplistic fashion: the steer intertia and stiffness terms are increased by $1kg \cdot m^2$ and $200Nm/rad$, respectively. In the top figure the entire response to a perturbation is shown. A detail of the first 0.5 seconds after the perturbation is shown in the bottom figure.

5.4. Validation

The results of the application of the chosen model structure to the experimental data sets are covered in this section. For the sake of brevity the data is presented in four steps:

1. The FIR models fitted to the experimental data are shown and outliers are marked. Also, the variability between the FIR models is evaluated.
2. Train and test results: the predictive performance of four different median controllers and the median effective delay applied to different slices of the data is shown. The four different median controllers and effective delay originate from the following partitions of the data:
    - Test data
    - MAD outlier excluded data
    - IQR outlier excluded data
    - TH outlier excluded data
3. The distribution of the fitted parameters is visualized and the median controllers are presented in a table, with the median controller being the fitted LQR weight.
4. The forward velocity dependent feedback gains that result from the median controller are shown and compared with the results from the trained data.

The first step gives insight in the variability of the experimental data set. The variability between the FIR models is an easy performance criterion: the test-set median controller should have a lower variability and a higher average VAF. At the second step this is verified. If the median controller does not perform better, the parameters are likely to be overfitted. The third and fourth step are used to gain insight in the fitted controller: are there things that can be improved? Is there a forward velocity dependency? How do the individual participants compare?

5.4.1. UC Davis experiment - Treadmill

The model structure is fitted to the data used in [79]. This is a subset of the treadmill data consisting of 14 runs of one participant (Jason). These runs where conducted at forward velocities between 2.1 and 7.4 m/s.

Additionally, the model structure is fitted to runs of the other two participants of the UC Davis experiment (Luke and Charlie). Unfortunately, they have done only a few runs (6 for Luke and 5 for Charlie).

1) FIR model distribution and outliers

In figure 5.6, the results of the application of the different outlier detection algorithm is shown. Both the MAD and IQR algorithms mark the same high forward velocity runs as outliers, but upon visual inspection this does not make a lot of sense.

The naive threshold-based outlier detection algorithm tend to mark runs that are done at forward velocities lower than 3.5 m/s. The only exception is a high speed run done at 7 m/s (for which the optimization procedure failed to find a solution).

While not directly clear in the figure, the amplitude of the steer angle decreases with increasing forward velocity. Also, it can be seen that some of the runs exceed steering angles of 15 degrees. These are the runs done at forward velocities lower than 3.5 m/s. As mentioned before, this is non-ideal for the linear Whipple plant model.

Lastly, the fitted FIR models indicate there is quite a bit of variance between the runs. This is further illustrated in figure 5.7, where the VAF distribution of the fitted FIR models to the median FIR model is shown in a boxplot. The median test-set controller should at the very least perform better than this.

In the end, all runs are retained in the data-set. For obvious reasons, the run for which the fitting procedure failed is excluded from the train data.

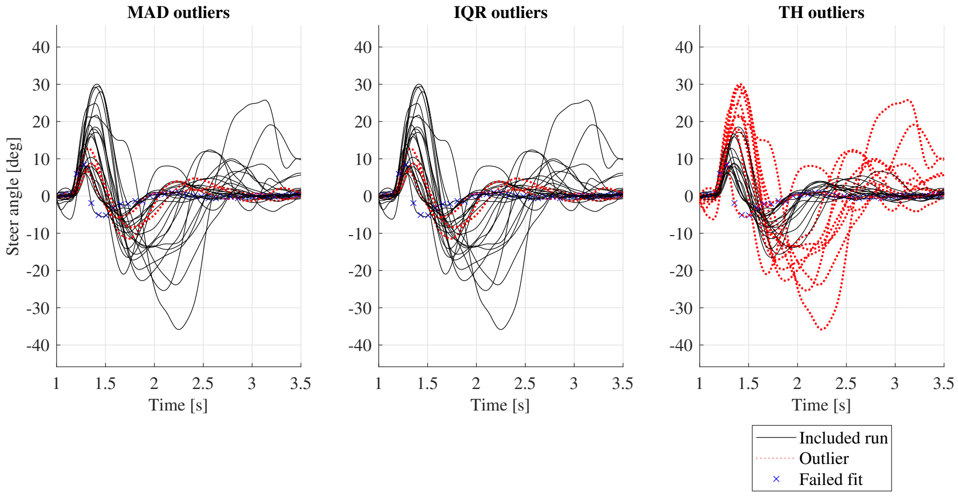

Figure 5.6: A comparison of the FIR models fitted to all raw pavilion data sets. The outliers for the different outlier detection algorithms are marked in red. The runs for which the rider-model failed to find a good fit are marked with blue crosses.

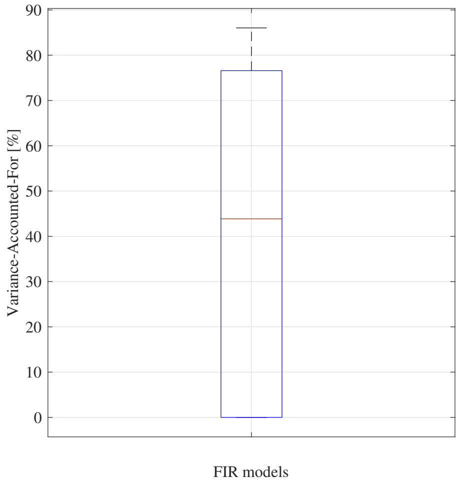

Figure 5.7: This plot shows the VAF distribution of the fitted FIR models to the median FIR model. The variance of this data set is $\sigma^2 = 34.7\%^2$

2) Train and test results

The four different median controllers are extracted from the results and plotted together with the train results in figure 5.8. The spread of the data can be viewed in the boxplot in figure 5.9. The test results are shown in terms of standard deviation and mean in table 5.9 .

A couple things can already be noted:

- There is no clear forward velocity dependency.
- The runs done by Charlie can not be identified well and the quality of the fit is worse than that of the median FIR model. The variability of the runs done by Charlie was already noted in the previous chapter and is likely the cause of this.
- If we don't consider Charlie, the VAF drop between train and test data is quite low, with a drop of 10 % for the data from Jason and even increasing for the data from Luke. This confirms that the fitted model structure can approximate the data for the entire velocity range with only one fitted parameter.
- The outlier detection algorithms are not of a big help. Even when no outliers are excluded, the test and train data are quite close together in terms of VAF.
- The roll and steer angle VAF values are similar for the test results.

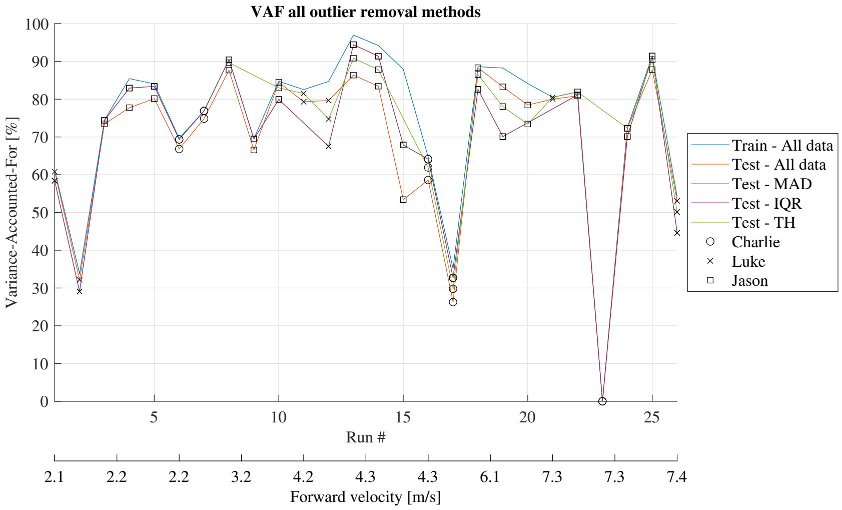

Figure 5.8: The performance of the train and test procedures in terms of the VAF . On the second x-axis, the velocity distribution associated with the run numbers is shown. For each test result, the participant that generated the data is marked.

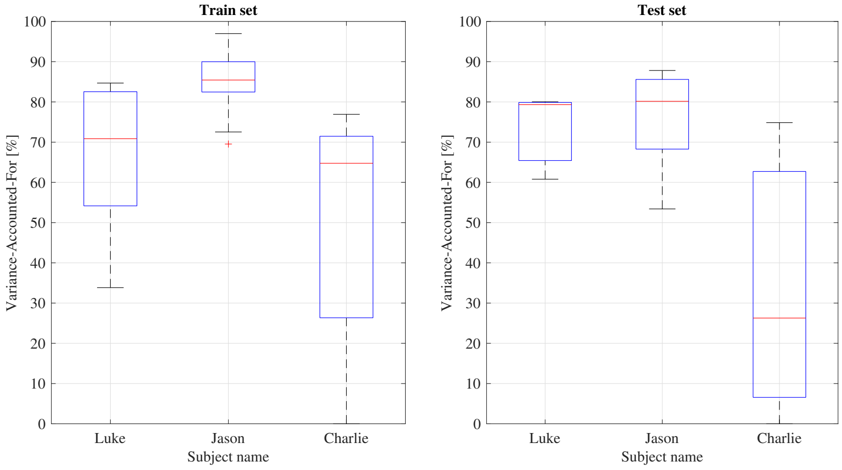

Figure 5.9: The train and test performance in terms of VAF visualized in boxplots for the pavilion data.

Table 5.9: The standard deviation σ and the mean μ VAF values for the treadmill data.

5.9.1: Steer angle VAF for both the train and test data sets.

	Train δ VAF		Test δ VAF	
	$\mu[\%]$	$\sigma[\%]$	$\mu[\%]$	$\sigma[\%]$
All	73.7	22.0	65.6	25.8
Luke	66.2	20.1	73.4	10.9
Charlie	49.3	31.8	33.7	38.0
Jason	85.0	7.8	75.9	12.4

5.9.2: Roll angle VAF for the test data set.

	Test ϕ VAF	
	$\mu[\%]$	$\sigma[\%]$
All	68.5	19.0
Luke	71.4	16.9
Charlie	67.8	5.6
Jason	67.5	24.8

3) Distribution of the fitted parameters

The distribution of the fitted free parameters are visualized in figure 5.10 and 5.11. Table 5.10 contains the values of the median controllers of the different subjects and the test set.

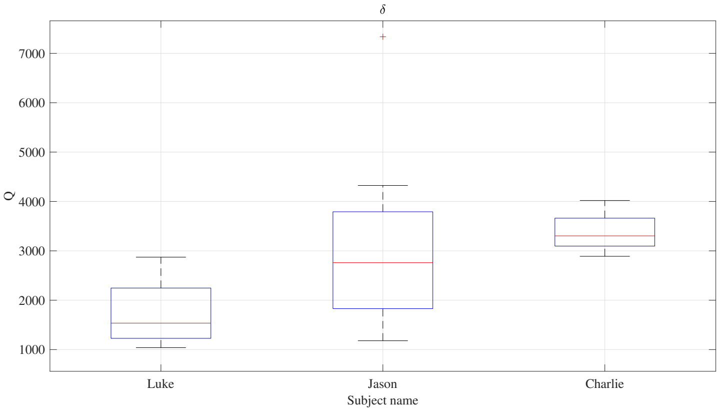

Figure 5.10: The distribution of the fitted free parameters visualized in boxplots.

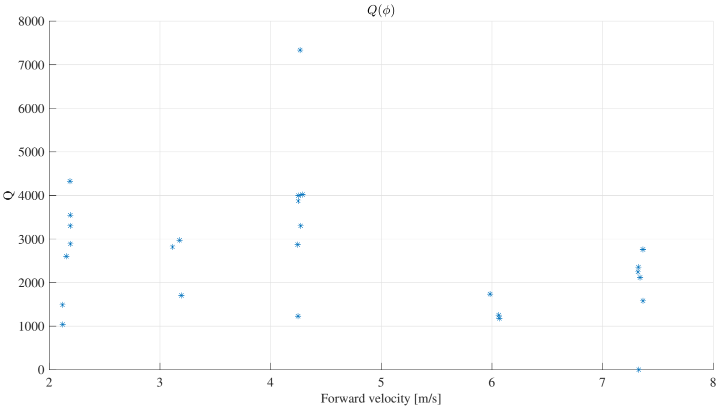

Figure 5.11: The distribution of the fitted roll angle LQR weights are plotted as a function of the forward velocity.

Table 5.10: The median LQR weights for the test set and the individual participants.

	$Q(\phi)$
Test set - median controller	795
Luke data - median controller	878
Charlie data - median controller	476
Jason data - median controller	3090

4) Velocity dependent feedback gains

From the analysis of the results in the previous two points, the fitted parameters appear to not be forward velocity dependent. However, the feedback gains that are calculated with the LQR algorithm are forward velocity dependent. In figure 5.12 these feedback gains are plotted together with the results from the individual model fits. It can be seen that the median controllers (both from the test-set and the entire data set) approximate the trends visible in the data.

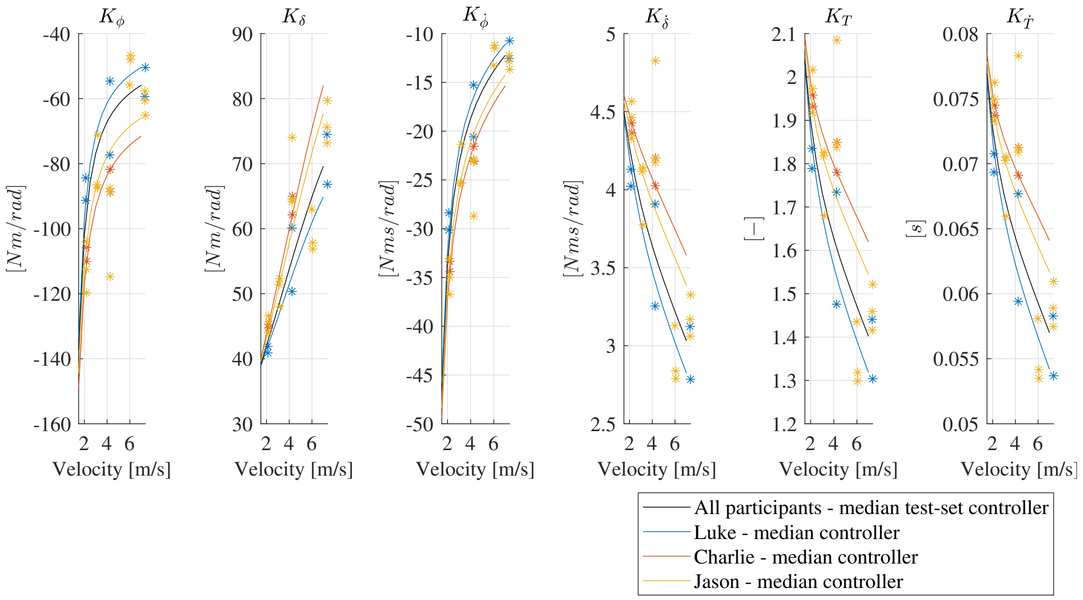

Figure 5.12: The median LQR controllers fitted to the test set and to all data are displayed here in terms of the resultant velocity dependent feedback gains. The results from the training procedure are displayed with dots of the appropriate color.

5.4.2. UC Davis experiment - Pavilion

The pavilion data-set consists of 68 runs conducted at forward velocities between 1.8 and 5.7 m/s. Runs of all three participants are included. The same steps are that are used for the treadmill data, are used here.

1) FIR model distribution and outliers

Reviewing the FIR models in 5.13, it can be seen that the optimization procedure failed to find fits with the experimental data for several runs (marked with blue crosses). Three of these are high speed (5.6 m/s) runs done by Jason. None of the outlier detection algorithms marked these runs, so the failure to fit to these runs is likely a failure of the model structure choice.

Just as for the treadmill data, there are no obvious outliers present in the data, but runs were nevertheless marked by the MAD and IQR outlier detection algorithms. The naive VAF threshold outlier detection algorithm marked runs below a forward velocity of approximately 3.5 m/s. Again, not of much use. It just tells us that the control behaviour for balancing a bicycle is speed dependent.

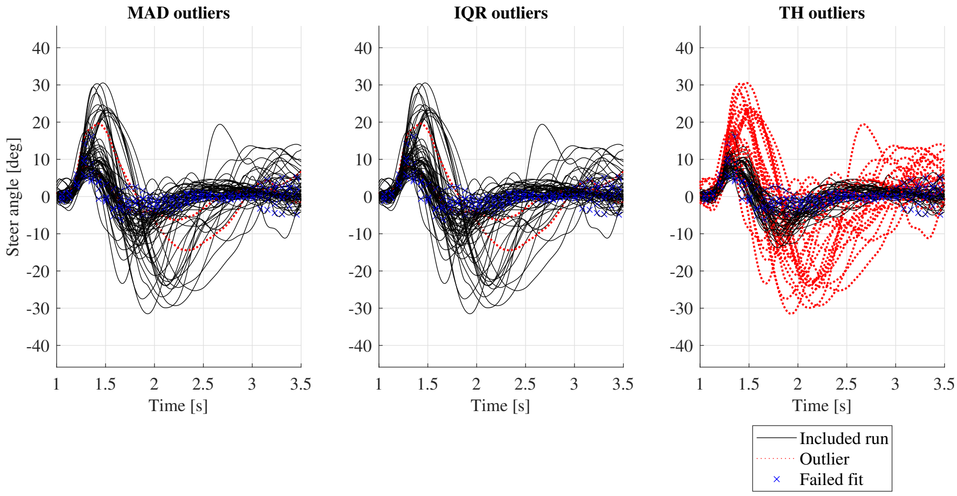

Figure 5.13: A comparison of the FIR models fitted to all raw pavilion data sets. The outliers for the different outlier detection algorithms are marked in red. The runs for which the rider-model failed to find a good fit are marked with blue crosses.

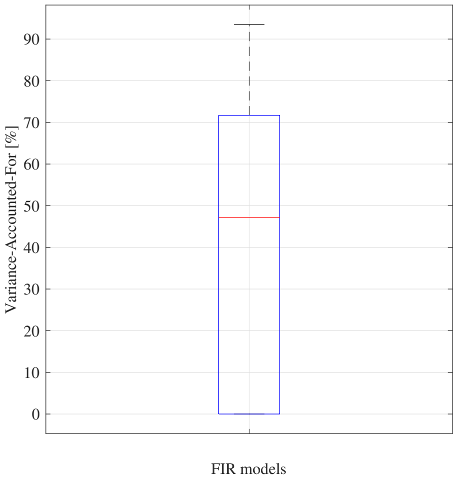

Figure 5.14: This plot shows the VAF distribution of the fitted FIR models to the median FIR model. The variance of this data set is $\sigma^2 = 34.1\%^2$

.

2) Train and test results

Reviewing both figure 5.15 and 5.16, the following remarks can be made:

- There is no clear forward velocity - LQR weight relationship.
- The train and test performance in terms of VAF is about 10 % worse when compared to the treadmill data.
- The train and test performance is quite decent for Luke and Jason, but not for Charlie. This is similar to the findings with the treadmill data.

The standard deviation σ and the mean μ VAF values are presented in table 5.11. Just as with the treadmill results, the pavilion test results are similar for the roll and steer angle.

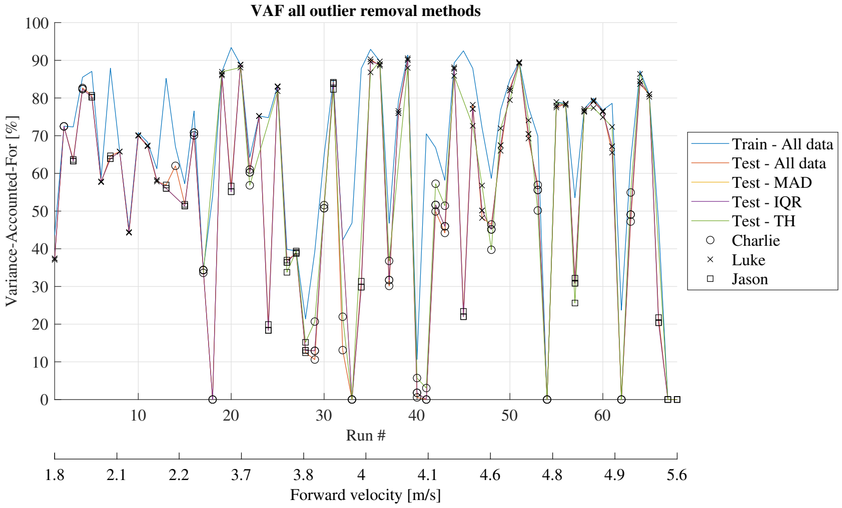

Figure 5.15: The performance of the train and test procedures in terms of the VAF . On the second x-axis, the velocity distribution associated with the run numbers is shown. For each test result, the participant that generated the data is marked.

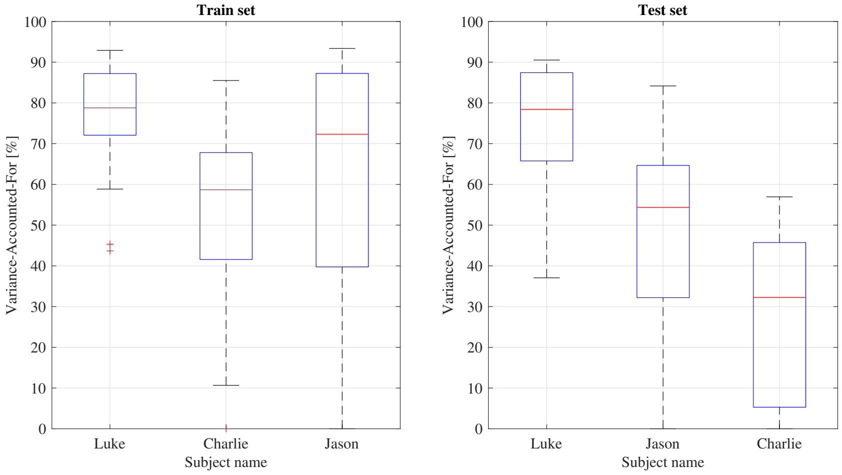

Figure 5.16: The train and test performance in terms of VAF visualized in boxplots for the pavilion data.

Table 5.11: The standard deviation σ and the mean μ VAF values for the pavilion data.

5.11.1: Steer angle VAF for both the train and test data sets.

	Train δ VAF		Test δ VAF	
	$\mu[\%]$	$\sigma[\%]$	$\mu[\%]$	$\sigma[\%]$
All	65.6	23.7	55.6	27.7
Luke	77.3	12.5	73.1	17.1
Charlie	53.2	21.9	27.9	22.0
Jason	60.2	31.4	49.6	26.3

5.11.2: Roll angle VAF for the test data set.

	Test ϕ VAF	
	$\mu[\%]$	$\sigma^2[\%^2]$
All	58.5	24.3
Luke	72.7	16.4
Charlie	32.5	26.4
Jason	56.6	14.7

3) Distribution of the fitted parameters

The distribution of the fitted free parameters are visualized in figure 5.17 and 5.18. Figure 5.18 confirms that there is no apparent forward velocity dependency.

Table 5.12 contains the values of the median controllers of the different subjects and the test set.

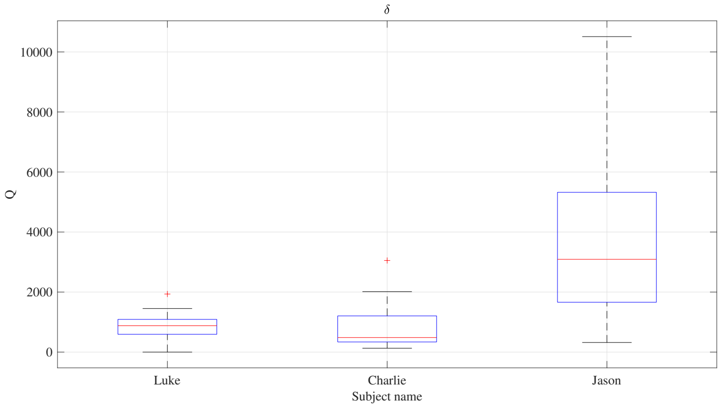

Figure 5.17: The distribution of the fitted free parameters visualized in boxplots. The first four boxplots show the data distribution of the free LQR weight factors, the last boxplot show the distribution of the fitted sensory delay.

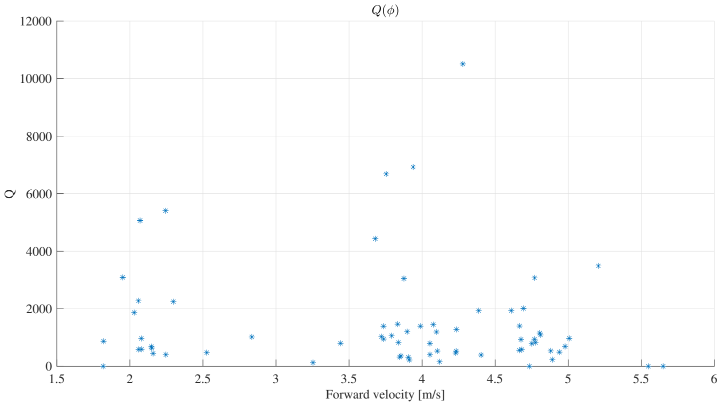

Figure 5.18: The distribution of the fitted roll angle LQR weights are plotted as a function of the forward velocity.

Table 5.12: The median LQR weights for the test set and the individual participants.

	$Q(\phi)$
Test set - median controller	1925
Luke data - median controller	1536
Charlie data - median controller	3303
Jason data - median controller	2760

4) Velocity dependent feedback gains

Plotting the median controller of the test-set and the median controllers of the different participants in figure 5.19 confirms the remarks at the previous points. The gains determined with the training data are variable, but still are grouped together for each subject. Which results in three median controllers that are different from each other (also visible in table 5.12). The median controller from the data from Luke lies closest to the median test-set controller.

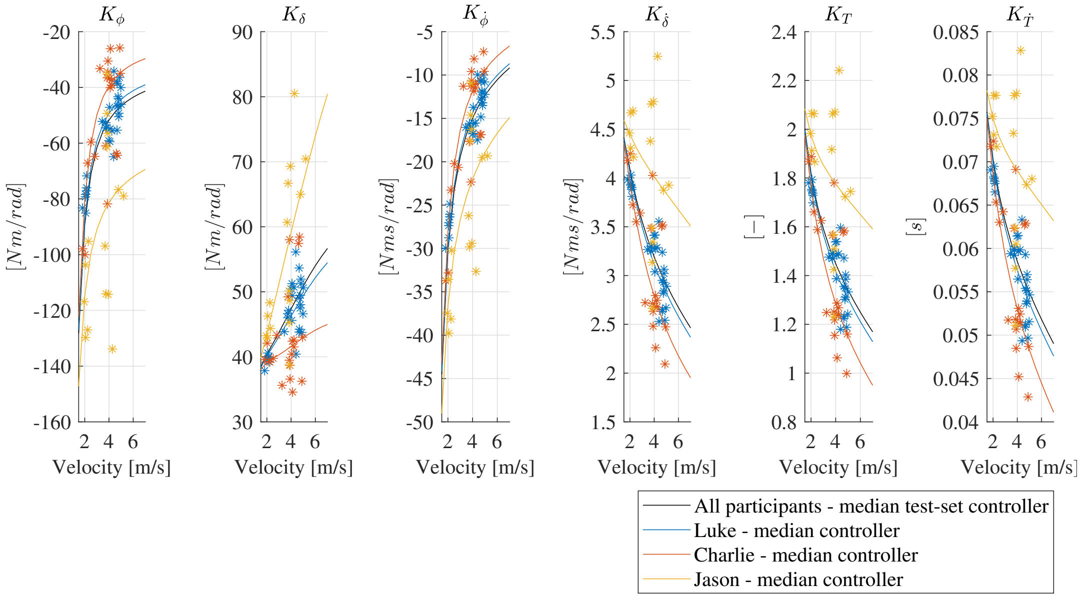

Figure 5.19: The median LQR controllers fitted to the test set and to all data are displayed here in terms of the resultant feedback gains. The results from the training procedure are displayed with dots of the appropriate color.

5.4.3. TU Delft experiment

The model structure is now applied to a part of the TU Delft experimental data, but now with the roll and roll rate LQR weights left as free parameters. The data of three participants is selected, who all did a single run at each of the four different forward velocities (2.5, 3.7, 4.5 and 5.5 m/s). One participant did another run at 5.5 m/s.

1) FIR model distribution and outliers

In figure 5.20 all the fitted FIR models are plotted. There are no clear outliers present in the selected data. The naive VAF threshold algorithm simply marks the three low forward velocity runs of all participants. Other than for UC Davis data, the fitting procedure never failed to find a fit to any of the runs.

In figure 5.21 the distribution of the fitted FIR models to the median FIR model is shown.

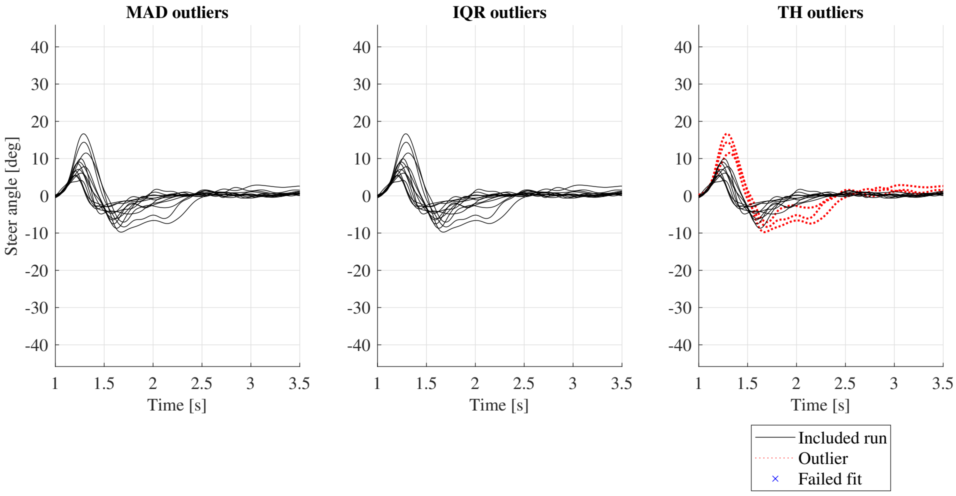

Figure 5.20: A comparison of the FIR models fitted to all raw pavilion data sets. The outliers for the different outlier detection algorithms are marked in red. The runs for which the rider-model failed to find a good fit are marked with blue crosses.

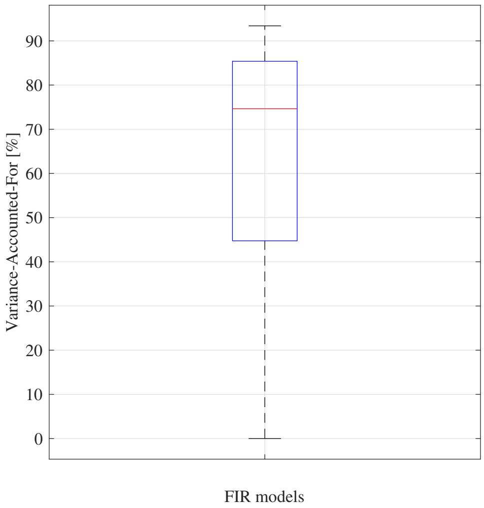

Figure 5.21: This plot shows the VAF distribution of the fitted FIR models to the median FIR model. The variance of this data set is $\sigma^2 = 35.8\%^2$

.

2) Train and test results

Reviewing both figure 5.22 and 5.23, the following remarks can be made:

- There does seem to be a forward-velocity dependency. The threshold based outlier detection algorithm, were low forward velocity runs are excluded, performs better at higher forward velocities.
- The train and test performance in terms of steer angle VAF is about 10 % worse when compared to the treadmill data.

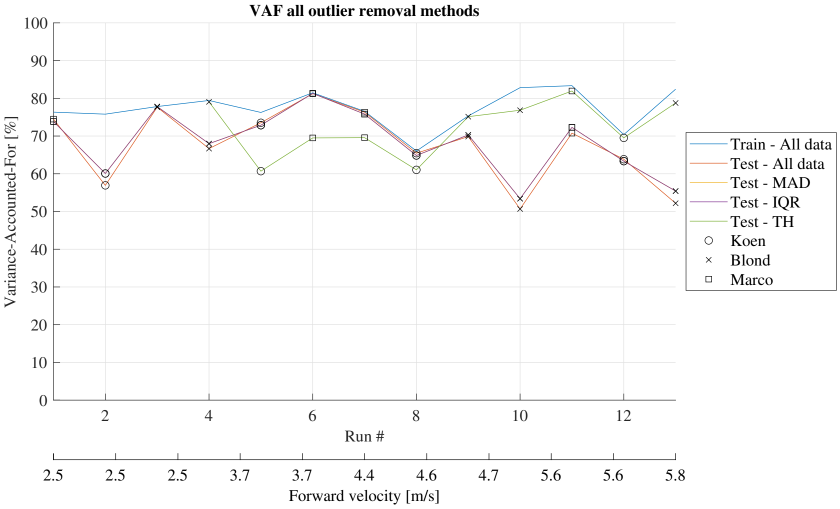

Figure 5.22: The performance of the train and test procedures in terms of the VAF . On the second x-axis, the velocity distribution associated with the run numbers is shown. For each test result, the participant that generated the data is marked.

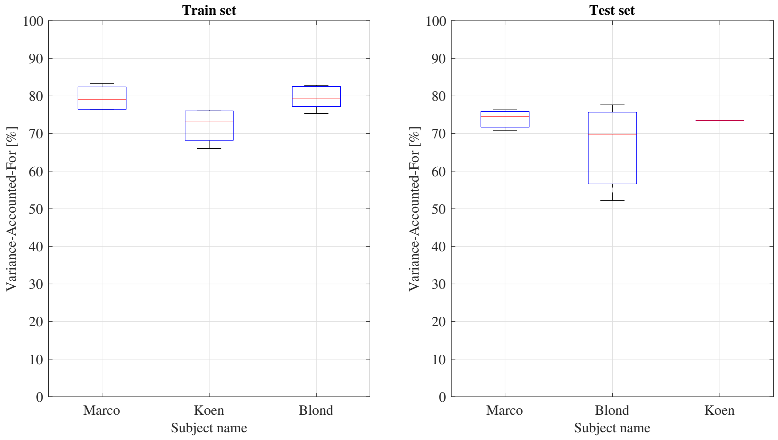

Figure 5.23: The train and test performance in terms of VAF visualized in boxplots for the cycling path data.

The standard deviation σ and the mean μ VAF values are presented in table 5.13. Comparing the test steer and roll angle VAF values confirms that the unaltered Whipple bicycle model is not valid for this experiment.

Table 5.13: The standard deviation σ and the mean μ VAF values for the cycling path data.

5.13.1: Steer angle VAF for both the train and test data sets.

	Train δ VAF		Test δ VAF	
	μ[%]	σ[%]	μ[%]	σ[%]
All	77.2	5.0	70.7	8.6
Blond	79.5	3.1	66.6	13.0
Koen	72.1	21.9	73.4	0
Marco	79.4	3.5	73.9	2.8

5.13.2: Roll angle VAF for the test data set.

	Test ϕ VAF	
	μ[%]	σ[%]
All	46.8	14.8
Blond	58.2	12.4
Koen	37.2	0
Marco	38.5	12.5

3) Distribution of the fitted parameters

The distribution of the fitted free parameters are visualized in figure 5.24 and 5.25. Figure 5.25 confirms that there is a forward velocity dependency.

Table 5.14 contains the values of the median controllers of the different subjects and the test set.

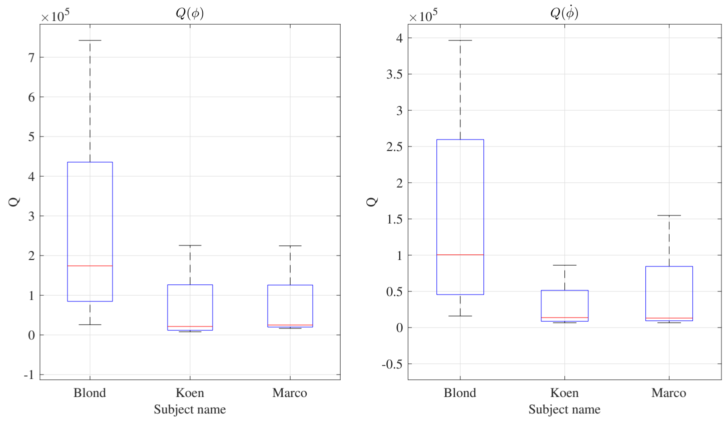

Figure 5.24: The distribution of the fitted free parameters visualized in boxplots.

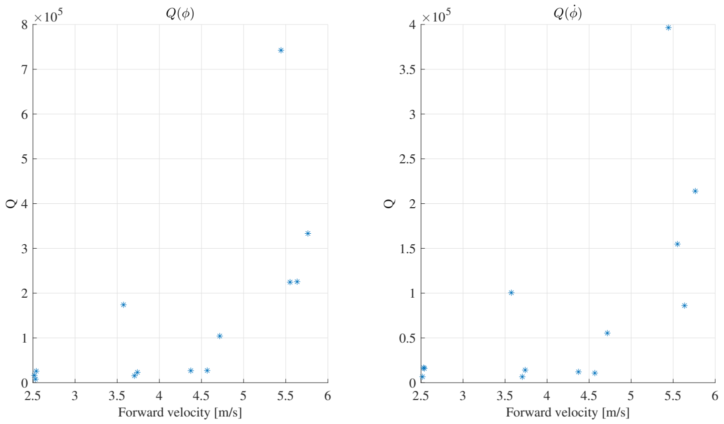

Figure 5.25: The distribution of the fitted LQR weights (roll angle on the left, roll rate on the right) are plotted as a function of the forward velocity.

Table 5.14: The median LQR weights for the test set and the individual participants.

	$Q(\phi)$	$Q(\dot{\phi})$
Test set - median controller	2706	1442
Blond data - median controller	17407	10051
Koen data - median controller	2136	1375
Marco data - median controller	2501	1311

4) Velocity dependent feedback gains

In figure 5.26 these feedback gains are plotted together with the results from the individual model fits. It is clear (especially if you look at the data from Koen and Marco) that the data collected at 5.7 m/s are distinct from the gains calculated with the median controllers.

The median controller of Blond is very different from the median controllers from Koen and Marco, indicating a different control strategy.

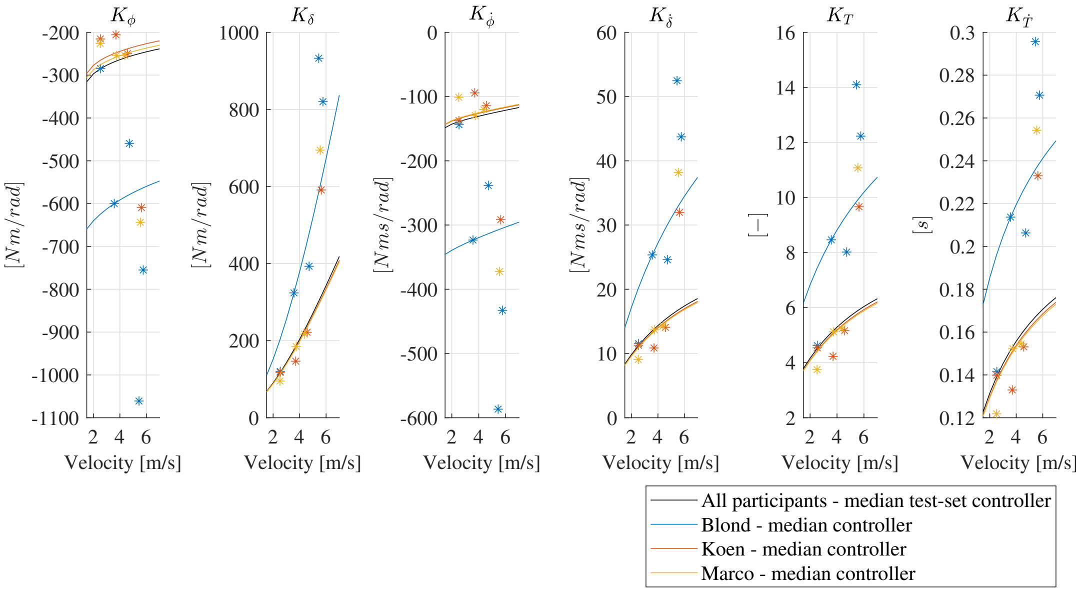

Figure 5.26: The median LQR controllers fitted to the test set and to all data are displayed here in terms of the resultant feedback gains. The results from the training procedure are displayed with dots of the appropriate color.

5.4.4. Initial results passive rider

In this section a model structure with (simple) passive rider dynamics is applied to the runs also used in [79], which are UC Davis runs done by Jason on the treadmill.

The specific settings are:

- Added ”arm” inertia: 1 $kg \cdot m^2$.
- Added ”arm” stiffness: 200 Nm/rad.
- Free parameters: LQR weights $Q(\phi)$, $Q(\dot{\delta})$ and neural sensory delay τ_s.

Train and test results

The boxplots of the steer angle VAF with the non-parametric model is in figure 5.27 and the standard deviation σ and the mean μ VAF values are presented in table 5.15.

The training performance is better when compared to the model structure without passive rider dynamics (91.4 % against 75.9 % VAF). However, the testing performance has degraded more (74.0 % against 75.9 % VAF). This can be attributed to the increase in free parameters (3 versus 1).

The roll angle VAF also is about 5 % worse when compared to the earlier results.

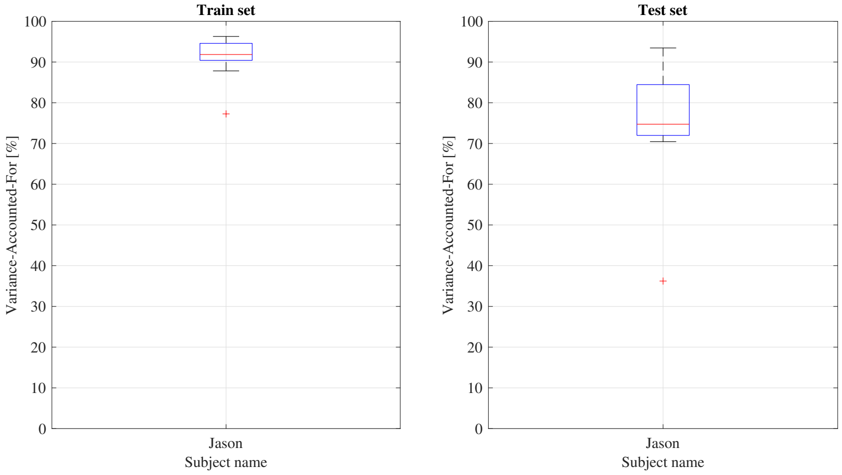

Figure 5.27: The train and test performance in terms of VAF visualized in boxplots for the treadmill data.

Table 5.15: The standard deviation σ and the mean μ VAF values for Jason on the treadmill.

5.15.1: Steer angle VAF for both the train and test data sets.

	Train δ VAF		Test δ VAF	
	μ[%]	σ[%]	μ[%]	σ[%]
Jason	91.4	4.7	74.0	17.3

5.15.2: Roll angle VAF for the test data set.

	Test ϕ VAF	
	μ[%]	σ^2[%^{2}]
Jason	62.8	33.5

Distribution of the fitted parameters

The distribution of the fitted free parameters are visualized in figure 5.28. In figure 5.29, the fitted parameters are plotted as a function of the forward velocity they are fitted at. The LQR weight on the steer rate is clearly forward velocity dependent. The LQR weight on the roll angle and the fitted neural delay also seem the have a dependency on the forward velocity, albeit less strong.

Table 5.16 contains the values of the median controllers of the train set and the test set.

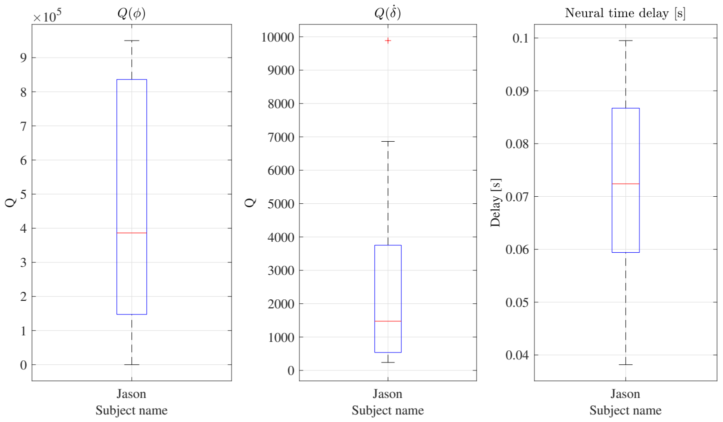

Figure 5.28: The distribution of the fitted free parameters visualized in boxplots.

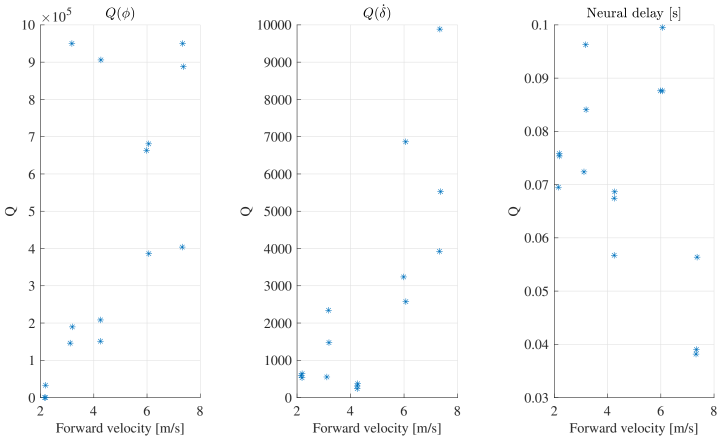

Figure 5.29: The distribution of the fitted LQR weights and neural delay (left to right: roll angle, steer rate, neural delay) are plotted as a function of the forward velocity.

Table 5.16: The median LQR weights for the test set and the individual participants.

	$Q(\phi)$	$Q(\dot{\phi})$	Neural delay [s]
Test set - median controller	208278	1475	0.076
Jason data - median controller	386168	1475	0.072

Velocity dependent feedback gains

In figure 5.30 these feedback gains are plotted together with the results from the individual model fits. The fits are approximated reasonably well by the the median test-set controller, but especially the gains on the steer rate, torque and torque rate could be captured better by the median controllers.

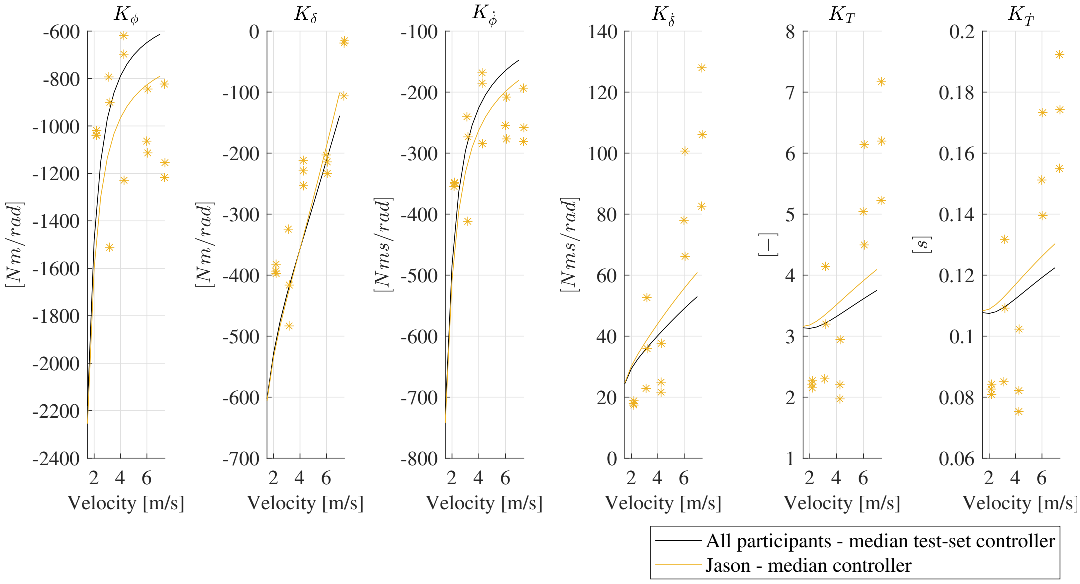

Figure 5.30: The median LQR controllers fitted to the test set and to all data are displayed here in terms of the resultant feedback gains. The results from the training procedure are displayed with dots of the appropriate color.

6

Discussion

In the previous chapter the findings that can be attributed to single model components or experimental settings were presented and interpreted. In this chapter the different experiments are compared and discussed.

6.1. Model structure selection

Control and estimation algorithms

It is interesting that a combination of optimal control (the LQR) and optimal estimation (the Kalman filter and the TDL (Tapped-Delay-Line) predictor) yields good results. While I've seen multiple papers hypothesizing or proving this ([93] being the most similar, but for human stance control), I have not seen this model structure implemented and validated for human control of an element anywhere in literature.

In this model structure, the Kalman filter noise co-variance matrices are the main thing that may need to be specified further. Based on the already good results they are left as identity matrices, but further investigating their influence may yield more realistic results. It could be possible that there is an interaction between the selected sensory channels and the co-variances of the noise associated with them.

They could also be used to investigate sensory degradation, as a higher noise co-variance corresponds to a more unreliable sensor.

Internal model

Results show that the internal model has to incorporate all dynamics present in the model for the best fitting performance. For both the bicycle dynamics and the neuromuscular dynamics this does make some sense: we need to know the relationship between the bicycle states and steer input. Also, knowing the our own control bandwidth limitations seems logical.

Sensory channels

The number of sensory channels is limited to two: the steer angle and the roll rate. Increasing the number of channels did not necessarily improve the results or even degraded performance. Apparently, some degree of inaccuracy is beneficial for the results. The exact reason for this is unclear. It could be that with other Kalman filter noise co-variances matrices this effect does not occur. It could also be that it can be influenced by the addition of a passive rider model, as the over estimation of the roll angle occurs only in the short moments right after the perturbation has been applied.

The results all had one thing in common: the roll rate channel is crucial in order to reach realistic results. Note that stabilizing a bicycle is perfectly possible without a roll rate measurement, it just doesn't yield results comparable to the experimental data.

6.2. Identified rider models

Individual subjects

Even though only one free parameter is fitted, which should limit the risk of over-fitting, from the results it is clear that there are differences between the subjects. These differences are both present for the

UC Davis and TU Delft data.

For the UC Davis data, the model structure does well at capturing the control behavior of Luke and Jason for both environments, besting the variability and performance of the median Finite-Impulse-Response (FIR) model. But, the results of Charlie are very variable. This is to be expected, considering there also is a greater variability in the raw data from Charlie (figure 4.6 in section 4.1.2). The strict and LTI framework of this model structure and data-prepossessing procedure is not able to grasp the control behaviour of Charlie.

The TU Delft data required the number of free LQR weight parameters to be increased to two, because no successful fits where found with only one free LQR weight. The training performance is stable, with Variance-Accounted-For (VAF) values around 75-80 %. This is quite decent, given the dependence of the LQR and Kalman filter algorithms on a correct plant model. This indicates that the control behaviour is different, but that the model structure is able to account for this.

The roll angle VAF values are significantly lower, confirming that the for experiment conducted at the TU Delft, the Whipple bicycle model should be extended to account for the free upper body movement.

The fitted parameters where found to be velocity dependent, with both parameters fitted to the high speed 5.7 m/s runs to be distinctly higher than those fitted to the lower speed runs. Therefore, this jump was not linear.

UC Davis treadmill and pavilion comparison

The results across both environments can be compared because the runs are conducted by the same subjects. This is done in figure 6.1, were the forward velocity dependent feedback gains are plotted for the median controllers of all subjects for both environments. This shows:

1. There are clear differences in control strategy between the different subjects.
2. The median controllers of Luke and Jason are similar between both environments. This indicates that the control strategy changes little.

We can further confirm this by applying the median controllers of the treadmill data to the pavilion data and visa-versa. The results of this procedure are shown in table 6.1. It is clear that that Charlie switches control strategies for the different environments. However, due to the variability of his results this probably is not a significant result.

For Luke, and especially for Jason, the controllers can be switched without a large penalty in VAF with the experimental data.

Table 6.1: In this table the original test results and the switched median controller test results are shown for all runs and for the runs separated according to the participants.

Test	Treadmill	Pavilion
Original mean test VAF	69.9 %	53.2 %
Switched controller mean test VAF	53.0 %	46.7 %
Luke original mean test VAF	64.1 %	74.4 %
Switched Luke controller mean test VAF	57.3 %	67.4 %
Charlie original mean test VAF	49.1 %	40.8 %
Switched Charlie controller mean test VAF	20.5 %	18.9 %
Jason original mean test VAF	79.8 %	50.8 %
Switched Jason controller mean test VAF	78.9 %	50.7 %

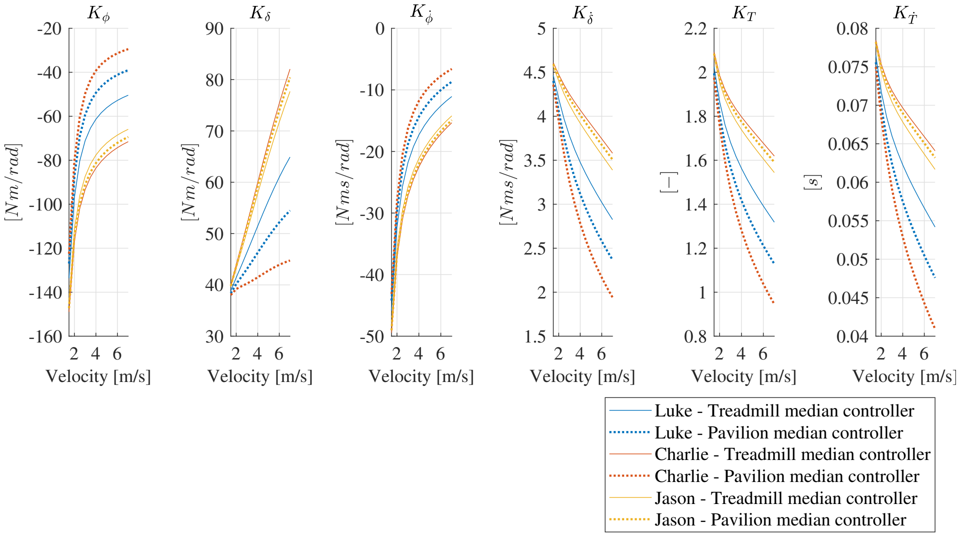

Figure 6.1: Forward velocity dependent feedback gains on the bicycle states. In this plot, the median test-set controllers and the median controllers from both the treadmill and pavilion data are included.

Comparison of all identified controllers

By calculating the forward velocity dependent feedback gains that result from the fitted LQR weights, we can compare the controllers in all three environments (treadmill, pavilion and cycling path)[1]. This is plotted in figure 6.2. It becomes very clear from this figure that the identified controllers of the TU Delft data are very different than those identified by the UC Davis data.

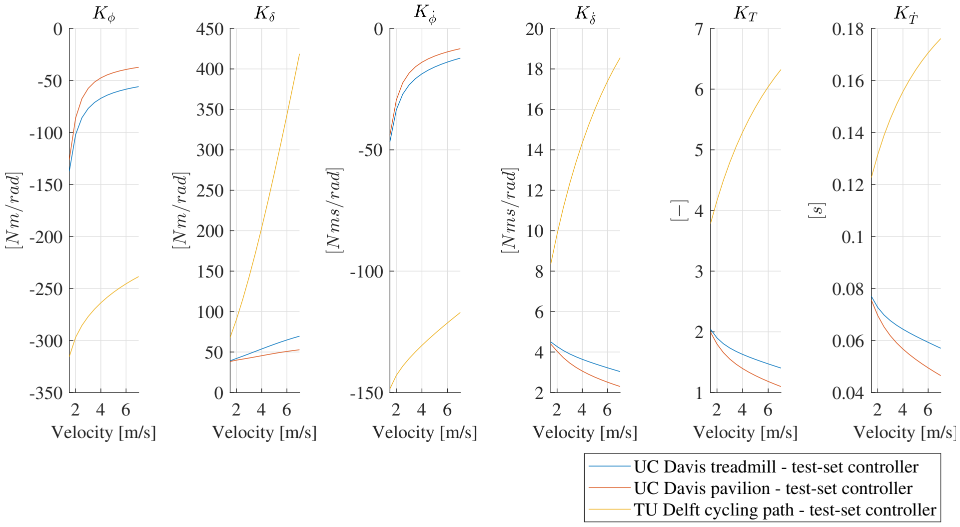

Figure 6.2: The forward velocity dependent feedback gains calculated with the median controllers from the UC Davis and TU Delft experimental data.

Comparison with the rider control identification paper by [79]

The results from the treadmill data of Jason can be directly compared with the results from [79]. The average training performance is about 5 to 10 % worse for the model structure that is presented in this thesis. However, one must consider that this model structure is only fitted with one free parameter for the entire forward velocity range and is also optimized for testing performance. The model in [79] uses four free parameters which are optimized for each individual run.

---

[1] The passive rider results are excluded because the bicycle model parameters are very different for those results.

The addition of a passive rider model

The results indicate that including a passive rider model increases the fidelity of the identified models: a sensory delay can be left as a free parameter and yields a fitted delay that is within human control capabilities. This is done with a very simplistic passive rider model.

There fitted parameters show signs of a forward velocity dependency. For the neural delay (which decreases with increasing forward velocity) this could point to tightening control behaviour of the rider.

Better parameters with the current passive rider model, or a more advanced (e.g., steer angle dependent and/or forward velocity dependent) passive rider model are likely to improve the testing performance.

6.3. Discussion on the experimental methods of the UC Davis and TU Delft experiments

It is difficult to come to a final judgment on quality of the experiments. Both have their positives and negatives, which will be discussed below.

Remarks on the UC Davis pavilion experiment

The results show that the variability between the FIR models fitted to the treadmill data is similar to the variability of those fitted to the pavilion data (treadmill $\sigma^2 = 34.7\%^2$, $N = 35$ and pavilion $\sigma^2 = 35.8\%^2$, $N = 68$). This while the evaluated forward velocity range is smaller (treadmill: 2.2-7.4 m/s, pavilion: 1.8-5.6 m/s) and the data set is bigger (treadmill: 35 runs, pavilion: 68 runs). Therefore, this is an indication that the pavilion data is more variable.

This could be caused by the experimental method. The runs in the pavilion were of a rather short duration during which usually only two or three perturbations are applied. Also, the participants immediately had to make a 90 degree turn at the end of the their run on the slippery pavilion surface.

Remarks on the TU Delft pavilion experiment

I think the experimental procedure and setup of the TU Delft experiment was not strict enough to be really useful for this research. The participants where not given time to train before the experimental data was acquired. This increases the intrasubject variability. Only pull-type perturbations were given and the person giving the perturbations was visible to the participant. This increases the risk that the rider trains a response specific to the experiment. Lastly, the upper body movement was left free, but not measured. This causes the identified controllers to be dissimilar but also makes it impossible to determine the control strategy guiding the upper-body movement.

I do like the environment of the TU Delft experiment: a long and undisturbed cycling path. Also, given the difference in identified controllers, upper body movement has an influence on control behaviour: gains of a greater magnitude are required to match the found steer angles. It is interesting to investigate this effect on stability.

Discussion on the type perturbation used during the experiments

Designing a proper perturbation to enable the unique identification of a human controller can be quite difficult. Usually, for the identification of a (assumed to be linear) human controller for pure tracking or disturbance compensation tasks, a sum-of-sinusoids perturbation signal is applied. If done properly, this method allows to filter out noise introduced by the test subject and results in models valid around the frequency ranges of the applied sinusoids [96]. It is assumed that linearity can be approximated when the subjects are well trained and the perturbation signal is random appearing and devoid of any large spikes.

The impulse-like perturbations applied during the UC Davis and TU Delft experiments do not meet those conditions. The risk of this is, is that a bicycle rider model is identified which is only valid for this particular perturbation. Due to the objective of this research, which is to identify a rider model that can be used for fall prevention, this does not need to be a big problem. At least, if the perturbation (and other experiment conditions) match the scenario we are interested in. Well, falling over essentially reduces to an increasing roll angle from which the rider fails to recover. The perturbation used initiates this process, and would result in a fall in the case of an uncontrolled bicycle[2]. So at least, the perturbation triggers behaviour that ought to be a bit realistic. However, in practice, the fall is usually not initiated

---

[2] The perturbation is large enough that even a bicycle in it's self-stable velocity range will fall over

by a roll perturbation, but by a a disturbance of the steer assembly. Furthermore, a lot of the falls occur during mounting and dismounting of the bicycle. Both are scenarios which might not be properly excited with a roll perturbation.

A more valid perturbation would be a steer perturbation. However, the effect of such a perturbation greatly depends on the posture of the bicycle-rider. A ”stiff” posture, for example if the rider leans on the handlebars or has co-contracted his arm muscles, will cause a way lower steer angle change (assuming a steer torque perturbation) when compared to a rider with a ”soft” posture.

Also, this research assumes that the human response scales linearly with the perturbation. Because all of the perturbations where roughly of the same magnitude, it is not possible to verify this.

6.4. Other findings

In this section I briefly touch upon some findings during my search towards to best bicycle rider control model.

FIR Model

Initially the same two data preprocessing steps were taken as in [22], these steps are:

1. Fit a non-parametric FIR model to the experimental data. Note that, when applicable, the experimental data is already detrended and calibrated.

2. Smooth the resulting data with a 8th order low-pass Butterworth filter with a cut-off frequency of 20 Hz.

The smoothing step incurs a group delay between 0.08 and 0.15 seconds (at a 200 Hz sampling frequency), which is significant for human control and skewed a lot of early results.

The fitting performance of the FIR model is also a bit haphazard[3]. The UC Davis Pavilion and TU Delft cycling path experimental data is approximated by the FIR model with VAF values between 70 and 100 %. This indicates that little potentially useful data is discarded.

However, the FIR model fitting performance is worse for the UC Davis treadmill data with VAF values ranging between 35 and to 80 %.

It is difficult to say something smart about this. The good performance of the FIR model with the UC Davis pavilion data can in part be attributed to the short pavilion run durations, during which only few (usually less than four) perturbations are applied. This increases the risk of the FIR model also fitting to variable human control behaviour, which would be discarded in runs of a longer duration.

This could explain the worse performance of the FIR model with the UC Davis treadmill runs, which are of a longer duration. But the same could be said for the TU Delft cycling path runs, which are also of a longer duration, but are approximated well with the FIR model.

An explanation could be, that the treadmill experiment did cause some non-linear rider control behaviour, which can’t be captured by the linear FIR model. This important to consider when attributing a certain confidence to the findings.

The fitting performance of the FIR model shows the same trend of variability that also is visible in the fitting performance of the final model structure. So, the FIR models fitted the Luke and Jason runs are less variable than those fitted to Charlie runs.

As an alternative to the non-parametric FIR model, ARX and ARMAX type models were tried as an alternative. However, these models performed worse than the FIR model, unless very high degree (N = 300+) models are used, which increases the computational time beyond a practical limit.

Adaptive Kalman filter

Because it generally is assumed that humans fuse the measured states optimally *and* do some sort of on-line signal reliability estimation, the use of an adaptive Kalman filter is a logical avenue to pursue.

Such an adaptive Kalman filter algorithm [1] was implemented with some initial success. But for runs done at velocities lower than 6 m/s, it struggled to achieve stable results. The reason seemed to be that

---

[3] The detailed results can be looked up in section 4.1.2

the noise co-variance matrix (which now follows from the algorithm in stead of being provided by the user) converges to a form during the period before the perturbation that results in wrong state estimates once the perturbation is applied to the system. Hence, the adaptive Kalman filter was dropped.

Multiple delay Kalman filter

With a normal Kalman filter, each sensory pathway is combined as they arrive irrespective of the sensory delay. But, because there are multiple different delays associated with the different sensory pathways, it could make sense to account for this. To do this, the sensory channels should be combined in such a way that measurements done at the same time are combined instead of combining the measurements as they arrive in the ”brain”. Unfortunately, this is easier said than done, because as a more delayed measurement is introduced to the Kalman filter, the noise co-variance matrices have to be updated to the current time recursively. The method described by [2] is used to do this in a way that is somewhat easy on memory and computational requirements. The algorithm worked fine for a minimum working example, but applying it to the more complex model proved to be difficult. After some tinkering and given that the impact of the method on the results was doubtful, it was scratched. With time, I'm confident that the algorithm can be implemented successfully, but I think just assuming a single effective delay is a more scientifically founded approach.

Steer torque feedback

During the research, there was a lot of discussion on the importance of torque feedback. Personally, I think that it is unlikely that torque feedback provides direct information on the complete state of the bicycle. If we review the Whipple bicycle equations, it is impossible to reconstruct an unique set of bicycle states only knowing the steer torque. The same is true of one wants to reconstruct the magnitude of the perturbation, again only knowing the steer torque. The feedback off[4] TU Delft experiment reinforces my belief that this is true.

However, what I can imagine to be true, is that torque is a powerful mechanism to provide information on the angle of the handlebars, as it is on the ”level” of the acceleration of the handlebars. In the case of this model, this information is essentially supplied in the form of the control input.

6.5. Relevance

The results presented in this thesis do not yield results that are relevant to society immediately. However, a high fidelity model describing human bicycle balancing can be used to investigate the stability margins before an unrecoverable fall occurs for various situations. For example, to investigate the effect of sensory degradation, or the effect of bicycle design on the stability margins. The results of these experiments can be used for tailored measures to reduce the injury rate among cyclists.

The results also indicate that human cycling experiments done on a treadmill are a viable alternative to experiments done in environments with a less strict restriction on lateral movement. This is of course useful if you want to do further bicycle research in a controlled and convenient environment.

Lastly, having a good model structure that can be fitted to experimental results is important when developing steer assist bicycles as it allows the investigation of behavioral changes and the simulation of combined effects.

---

[4] In this experiment only the combined inertia of the handlebars and arms of the participants was left as a ”source” of feedback torque

7

Conclusion and further research

7.1. Conclusion

Based on the results and discussion, a few things become clear:

- Using LQR and estimation (Kalman filter and TDL predictor) algorithms in combination with modelled neuromuscular dynamics and the linear Whipple bicycle model yields a model structure that is able to describe human bicycling for all evaluated forward velocities.

  For the UC Davis data, one rider dependent parameter is sufficient: a LQR weight placed on the roll angle. Successfully fitting a model structure to the TU Delft data requires two rider dependent parameters: LQR weights placed on the roll angle and the roll rate.

- Leaving the neural delay is a free parameter in the model structure without passive rider dynamics does not yield a realistic delay. For model structure with passive rider dynamics, the estimated delay ranges between 0.04 and 0.1 seconds. This is in agreement in earlier studies into human bicycle balancing [87][27].

- Roll rate and the handlebar angle of the bicycle provide enough information to the rider to control the bicycle in a realistic manner. With the roll rate being the primary measurement as realistic results can't be achieved with only handlebar information. Using other states either reduces the performance or does not increase performance.

  The handlebar angle might originate both from the muscle spindles as from the Golgi tendon sensor organs.

  The roll rate can both be estimated by the vision and vestibular organs. It is likely that it is a combination of both, but given the fast control behaviour and circumstantial evidence, the vestibular organ is likely to be the most important.

- Preliminary results indicate that there is not a lot of difference between the UC Davis treadmill and pavilion experiments.

- Even though the above points are true for all subjects, the participants tend to have distinct control behaviour.

7.2. Further research and recommendations

1. The most interesting avenue for further improvement of this model are passive rider models. I've had some success with a simple model, but a more elaborate model or better parameters of the simple model are likely to improve the results.

2. The influence of the spectral density matrices of the Kalman filter can be investigated. This in order to increase the fidelity of the model, but also investigate the impact of sensory degradation. An adaptive Kalman filter might be worthwhile to try here, but the problem touched upon in the discussion must be addressed.

3. From the TU Delft experiment, it also became clear that unrestricted upper body movement caused a significant change in control behaviour. This conflicts with earlier research which was done without perturbations. It is definitely interesting to research the impact of upper body movement on perturbed cycling: is it merely a passive change in the dynamics of the bicycle or is it active control behaviour?

4. There might be better alternatives to the FIR model used for data pre-processing. The biggest problem is that the bicycle-rider system is open-loop unstable, making it difficult to select a valid black-box model structure.

5. Knowing that a large part of the bicycle accidents are single vehicle accidents that either happen during mounting and dismounting of the bicycle or are caused by a short perturbation to a part of the steering assembly, conducting an experiment tailored to such conditions might provide new insights into the actual cause of a lot of cycling accidents.

6. At forward velocities that are lower than approximately 3.5 m/s, the steer angles start to exceed 15 degrees. This causes the Whipple bicycle model to be increasingly non-correct due to non-linear effects. If more research is done in this low forward velocity range, it makes sense to test the impact of using a non-linear Whipple bicycle model.

   As a next step, it is then interesting to test this non-linear plant model with a linear and a non-linear internal model. If the non-linear internal model performs a lot better, this indicates that our internal model is non-linear both in forward velocity *and* in the roll and steer angle. If not, then it is only non-linear in forward velocity.

7. Because the conclusion of this thesis indicates that the vestibular organ is the primary source of information for the rider, it is interesting to conduct an experiment with blindfolded cyclists.

8. It is interesting to do a more extensive comparison between the treadmill and pavilion experiments. If they truly are similar, further research into bicycle control becomes a lot easier. Although I wouldn't want to ride on a treadmill, blindfolded.

9. It is interesting to conduct the same experiments again, but with perturbations of varying magnitudes. This should yield information on how the human response scales. When further research is conducted in accident prevention, this information is important.

A

Appendix

A.1. Derivation of system state-space matrices

A.1.1. Plant model

Whipple bicycle model

The canonical linearised equations which describe the lateral dynamics for the Whipple bicycle model with coefficients as derived by Meijaard et al. [59] can be written as:

$$\mathbf{M}\ddot{\mathbf{q}} + v\mathbf{C}_1\dot{\mathbf{q}} + [g\mathbf{K}_0 + v^2\mathbf{K}_2]\mathbf{q} = \mathbf{T} \tag{A.1}$$

Where $\mathbf{q} = [\phi, \delta]^\top$ and $\mathbf{T} = [T_\phi, T_\delta]^\top$. This translates in the continuous time following state-space model:

$$\begin{bmatrix} \dot{\mathbf{q}}(t) \\ \ddot{\mathbf{q}}(t) \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{I} \\ -\mathbf{M}^{-1}(g\mathbf{K}_0 + v^2\mathbf{K}_2) & -\mathbf{M}^{-1}v\mathbf{C}_1 \end{bmatrix} \begin{bmatrix} \mathbf{q}(t) \\ \dot{\mathbf{q}}(t) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{M}^{-1} \end{bmatrix} \mathbf{f} \tag{A.2}$$

In short and making a distinction between perturbation torque $\mathbf{T}^e$ and rider input torque $\mathbf{T}^c$:

$$\begin{aligned} \dot{\mathbf{x}}_b &= \mathbf{A}_b\mathbf{x}_b + \mathbf{B}_b(\mathbf{T}^e + \mathbf{T}^c) \\ \mathbf{y}_b &= \mathbf{C}_b\mathbf{x}_b \end{aligned} \tag{A.3}$$

Additional bicycle model states

The heading angle and lateral displacement, velocity and acceleration of the bicycle can be interesting information for the rider to use for the control task. These states can be expressed as a function of the lateral bicycle model states (A.2):

$$\begin{aligned} \dot{\psi}(t) &= \frac{v\cos(\lambda)}{w}\delta(t) + \frac{c}{w}\dot{\delta}(t) \\ \dot{y}_p(t) &= v\psi(t) \\ \ddot{y}_p(t) &= v\dot{\psi}(t) \end{aligned}$$

The heading angle and the global coordinates can now be found found by integrating the heading rate $\dot{\psi}$ of the bicycle and the lateral velocity $\dot{y}_p$ of the rear wheel contact point over time. A.4 is the state-space representation with can be appended to the state equations of the linear bicycle model (A.2).

$$\begin{bmatrix} \dot{\psi}(t) \\ \dot{y}_p(t) \\ \ddot{y}_p(t) \end{bmatrix} = \begin{bmatrix} 0 & \frac{v\cos(\lambda)}{w} & 0 & \frac{c\cos(\lambda)}{w} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & \frac{v^2\cos(\lambda)}{w} & 0 & \frac{vc\cos(\lambda)}{w} & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \phi(t) \\ \delta(t) \\ \dot{\phi}(t) \\ \dot{\delta}(t) \\ \psi(t) \\ y_p(t) \\ \dot{y}_p(t) \end{bmatrix} \tag{A.4}$$

Written in a compact form, with the new states denoted as $\mathbf{x}_{hy}$:

$$\begin{bmatrix} \dot{\mathbf{x}}_b \\ \dot{\mathbf{x}}_{hy} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_b & \mathbf{0} \\ \mathbf{A}_{hy,b} & \mathbf{A}_{hy,hy} \end{bmatrix} \begin{bmatrix} \mathbf{x}_b \\ \mathbf{x}_{hy} \end{bmatrix} \tag{A.5}$$

Neuromuscular dynamics

The relation between the steer torque desired by the rider ($\hat{T}_\delta^c$) and the realized steer torque (T_δ^c) can described by a neuromuscular model which lumps the neuromuscular and control inceptor systems together [21]. In the Laplace domain, the neuromuscular dynamics are described by the following transfer function:

$$H_{nm}(s) = \frac{T_\delta^c(s)}{\hat{T}_\delta^c(s)} = \frac{\omega_c^2}{s^2 + 2\zeta\omega_c s + \omega_c^2} \tag{A.6}$$

For this research shoulder muscle parameters from from Happee et al. [38] are used: the cut-off frequency $\omega_c = 2.17 \cdot 2\pi$ rad/s and the damping coefficient $\zeta = \sqrt{0.5}$. This transfer function can be written in a convenient (as it leaves an intact physical interpretation of the states) continuous state-space phase variable form:

$$\begin{bmatrix} \dot{T}_\delta^c(t) \\ \ddot{T}_\delta^c(t) \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\omega_c^2 & -2\zeta\omega_c \end{bmatrix} \begin{bmatrix} T_\delta^c(t) \\ \dot{T}_\delta^c(t) \end{bmatrix} + \begin{bmatrix} 0 \\ \omega_c^2 \end{bmatrix} \hat{T}_\delta^c(t) \tag{A.7}$$

$$T_\delta^c(t) = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} T_\delta^c(t) \\ \dot{T}_\delta^c(t) \end{bmatrix} \tag{A.8}$$

Written in a compact form, with y_{nm} as the realized steer torque.

$$\begin{aligned} \dot{\mathbf{x}}_{nm} &= \mathbf{A}_{nm}\mathbf{x}_{nm} + \mathbf{B}_{nm}\hat{T}_\delta^c \\ y_{nm} &= \mathbf{C}_{nm}\mathbf{x}_{nm} \end{aligned} \tag{A.9}$$

A.1.2. Sensory model

Semi-circular-canals

Fernandez and Goldberg [31] formulates a Laplace domain transfer function which describes the relationship between stimulus angular acceleration and the afferent neural response:

$$H_{SCC}(s) = \frac{\tau_a s}{(1 + \tau_a s)} \cdot \frac{K(1 + \tau_L s)}{(1 + \tau_1 s)(1 + \tau_2 s)} \tag{A.10}$$

With parameters:

- Adaptation time constant τ_a
- Neural lead term τ_L
- Time constant τ_1
- Time constant τ_2
- Gain K

This response of this transfer function correlates with rotational velocity denoted as $\tilde{\omega}$. The transfer function can be rewritten in the following form:

$$H_{SCC}(s) = \frac{\tilde{\omega}(s)}{\alpha(s)} = \frac{T_0 s^2 + T_1 s}{s^3 + T_2 s^2 + T_3 s + T_4} \tag{A.11}$$

With:

- $T_0 = \frac{K\tau_L}{\tau_1\tau_2}$
- $T_1 = \frac{K}{\tau_1\tau_2}$

- $T_2 = \frac{\tau_a \tau_1 + \tau_a \tau_2 + \tau_1 \tau_2}{\tau_a \tau_1 \tau_2}$
- $T_3 = \frac{\tau_a + \tau_1 + \tau_2}{\tau_a \tau_1 \tau_2}$
- $T_4 = \frac{1}{\tau_a \tau_1 \tau_2}$

This transfer function can be written in observable canonical form which also outputs a derivative of the afferent neural response. This derivative correlates with rotational acceleration denoted as $\tilde{\alpha}$. The output is delayed with a neural delay τ_{SCC}.

$$\begin{bmatrix} \tilde{\alpha}(t) \\ \dot{x}_{SCC,2}(t) \\ \dot{x}_{SCC,3}(t) \end{bmatrix} = \begin{bmatrix} -T_2 & 1 & 0 \\ -T_3 & 0 & 1 \\ -T_4 & 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{\omega}(t) \\ x_{SCC,2}(t) \\ x_{SCC,3}(t) \end{bmatrix} + \begin{bmatrix} T_0 \\ T_1 \\ 0 \end{bmatrix} \alpha(t)$$
$$\begin{bmatrix} \tilde{\omega}(t + \tau_{SCC}) \\ \tilde{\alpha}(t + \tau_{SCC}) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ -T_2 & 1 & 0 \end{bmatrix} \begin{bmatrix} \tilde{\omega}(t) \\ x_{SCC,2}(t) \\ x_{SCC,3}(t) \end{bmatrix} + \begin{bmatrix} 0 \\ T_0 \end{bmatrix} \alpha(t) \tag{A.12}$$

Short:

$$\begin{aligned} \dot{\mathbf{x}}_{SCC} &= \mathbf{A}_{SCC}\mathbf{x}_{SCC} + \mathbf{B}_{SCC}\alpha \\ \mathbf{y}_{SCC} &= \mathbf{C}_{SCC}\mathbf{x}_{SCC} + \mathbf{D}_{SCC}\alpha \end{aligned} \tag{A.13}$$

Otoliths

Where the SCC's respond to rotational accelerations, the otoliths respond to translational accelerations. Fernandez and Goldberg [31] proposed the following transfer function which describes the relation between translational acceleration and the afferent neural response:

$$H_{oto}(s) = K \frac{1 + K_a \tau_a s}{1 + \tau_a s} \cdot \frac{1 + K_v (\tau_v s)^{K_v}}{1 + \tau_m} \tag{A.14}$$

The (fractional) exponent in the numerator makes implementation difficult. Therefore, in a lot of research on human modelling a simplified transfer function based on the physical nature of the otoliths is used A.15. This transfer function is also used for this research. It functions as an accelerometer, with the translational acceleration denoted by a and the measured translational acceleration denoted by $\tilde{a}$.

$$H_{oto}(s) = \frac{\tilde{a}(t)}{a(t)} = K \frac{1 + \tau_n s}{(1 + \tau_1 s)(1 + \tau_2 s)} \tag{A.15}$$

This transfer function can be converted to a state-space representation (OCF):

$$\begin{bmatrix} \dot{\tilde{a}}(t) \\ \dot{x}_{oto}(t) \end{bmatrix} = \begin{bmatrix} -T_2 & 1 \\ -T_3 & 0 \end{bmatrix} \begin{bmatrix} \tilde{a}(t) \\ x_{oto,2}(t) \end{bmatrix} + \begin{bmatrix} T_0 \\ T_1 \end{bmatrix} a(t)$$
$$\tilde{a}(t + \tau_{oto}) = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \tilde{a}(t) \\ x_{oto,2}(t) \end{bmatrix} \tag{A.16}$$

With:

- $T_0 = \frac{K}{\tau_1 \tau_2}$
- $T_1 = \frac{K \tau_n}{\tau_1 \tau_2}$
- $T_2 = \frac{\tau_1 + \tau_2}{\tau_1 \tau_2}$
- $T_3 = \frac{1}{\tau_1 \tau_2}$

Short:

$$\begin{aligned} \dot{\mathbf{x}}_{oto} &= \mathbf{A}_{oto}\mathbf{x}_{oto} + \mathbf{B}_{oto}a \\ y_{oto} &= \mathbf{C}_{oto}\mathbf{x}_{oto} \end{aligned} \tag{A.17}$$

Visual system

The visual system estimating position and velocity can be represented by a low-pass filter with cut-off frequency $\omega_c = 0.810$ rad/s [65]:

$$H_{vis}(s) = \frac{\omega_c}{s + \omega_c} \tag{A.18}$$

In state-space representation:

$$\begin{bmatrix} \dot{\tilde{\theta}}(t) \\ \dot{\tilde{\omega}}(t) \\ \dot{\tilde{\psi}}(t) \\ \dot{\tilde{y}}_{rw}(t) \end{bmatrix} = \begin{bmatrix} -\omega_c & 0 & 0 & 0 \\ 0 & -\omega_c & 0 & 0 \\ 0 & 0 & -\omega_c & 0 \\ 0 & 0 & 0 & -\omega_c \end{bmatrix} \begin{bmatrix} \tilde{\theta}(t) \\ \tilde{\omega}(t) \\ \tilde{\psi}(t) \\ \tilde{y}_{rw}(t) \end{bmatrix} + \begin{bmatrix} \omega_c & 0 & 0 & 0 \\ 0 & \omega_c & 0 & 0 \\ 0 & 0 & \omega_c & 0 \\ 0 & 0 & 0 & \omega_c \end{bmatrix} \begin{bmatrix} \theta(t) \\ \omega(t) \\ \psi(t) \\ y_{rw}(t) \end{bmatrix} \tag{A.19}$$

$$\begin{bmatrix} \tilde{\theta}(t+\tau_{vis}) \\ \tilde{\omega}(t+\tau_{vis}) \\ \tilde{\psi}(t+\tau_{vis}) \\ \tilde{y}_{rw}(t+\tau_{vis}) \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \tilde{\theta}(t) \\ \tilde{\omega}(t) \\ \tilde{\psi}(t) \\ \tilde{y}_{rw}(t) \end{bmatrix} \tag{A.20}$$

Short:

$$\begin{aligned} \dot{\mathbf{x}}_{vis} &= \mathbf{A}_{vis}\mathbf{x}_{vis} + \mathbf{B}_{oto} \begin{bmatrix} \mathbf{x}_b \\ \mathbf{x}_{hy} \end{bmatrix} \\ \mathbf{y}_{vis} &= \mathbf{C}_{vis}\mathbf{x}_{vis} \end{aligned} \tag{A.21}$$

A.1.3. Combined description

The combined plant takes the form of:

$$\begin{bmatrix} \dot{x}_b \\ \dot{x}_{nm} \\ \dot{x}_s \end{bmatrix} = \mathbf{A} \begin{bmatrix} x_b \\ x_{nm} \\ x_s \end{bmatrix} + \mathbf{B}_1 \hat{T}_\delta^c + \mathbf{B}_2 \begin{bmatrix} T_\phi^e \\ T_\delta^e \end{bmatrix} \tag{A.22}$$

$$\begin{bmatrix} y_b \\ y_s \end{bmatrix} = \mathbf{C} \begin{bmatrix} x_b \\ x_{nm} \\ x_s \end{bmatrix} + \mathbf{D} \begin{bmatrix} T_\phi^e \\ T_\delta^e \end{bmatrix} \tag{A.23}$$

Where input matrices $\mathbf{B}_1$ and $\mathbf{B}_2$ describe the relation between a state change and input steer torque by the rider and the perturbation torques respectively. The states of the sensory dynamics (x_{SCC}, x_{vis} and x_{oto}) are denoted by x_s.

The sensory dynamics can be incorporated by noting that the derivative of the sensory state vector $\mathbf{x}_s$ can be written as:

$$\begin{aligned} \dot{\mathbf{x}}_{SCC} &= \mathbf{A}_{SCC}\mathbf{x}_{SCC} + \mathbf{B}_{SCC}\alpha \\ \alpha = \ddot{\phi} &= \mathbf{A}_b(3,:)\mathbf{x}_b + \mathbf{B}_b(3,:)\left(\mathbf{T}^c + \mathbf{T}^e\right) \end{aligned} \tag{A.24}$$

$$\begin{aligned} \dot{\mathbf{x}}_{vis} &= \mathbf{A}_{vis}\mathbf{x}_{vis} + \mathbf{B}_{vis} \begin{bmatrix} \theta & \omega & \psi & y_{rw} \end{bmatrix}^\mathsf{T} \\ \begin{bmatrix} \theta & \omega & \psi & y_{rw} \end{bmatrix}^\mathsf{T} &= \begin{bmatrix} x_{b,1} & x_{b,2} & x_{hy,1} & x_{hy,2} \end{bmatrix}^\mathsf{T} \end{aligned} \tag{A.25}$$

$$\begin{aligned} \dot{\mathbf{x}}_{oto} &= \mathbf{A}_{oto}\mathbf{x}_{oto} + \mathbf{B}_{oto}a \\ a = \ddot{y}_{rw} &= \mathbf{A}_{hy,b}(3,:)\mathbf{x}_b \end{aligned} \tag{A.26}$$

The neuromuscular dynamics can be incorporated by noting that:

$$\begin{aligned} \dot{\mathbf{x}}_b &= \mathbf{A}_b\mathbf{x}_b + \mathbf{B}_b(:,2)T_\delta^c + \mathbf{B}_b\mathbf{T}^e \\ \dot{\mathbf{x}}_{nm} &= \mathbf{A}_{nm}\mathbf{x}_{nm} + \mathbf{B}_{nm}\hat{T}_\delta^c \end{aligned} \tag{A.27}$$

Resulting in a final continuous time state-space representation of the plant:

$$\begin{bmatrix} \dot{\mathbf{x}}_b \\ \dot{\mathbf{x}}_{hy} \\ \dot{\mathbf{x}}_{nm} \\ \dot{\mathbf{x}}_{SCC} \\ \dot{\mathbf{x}}_{oto} \\ \dot{\mathbf{x}}_{vis} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_b & \mathbf{0} & \begin{bmatrix} \mathbf{B}_b(:,2) & \mathbf{0} \end{bmatrix} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{A}_{hy,b} & \mathbf{A}_{hy,hy} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_{nm} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{B}_{SCC}\mathbf{A}_b(3,:) & \mathbf{0} & \begin{bmatrix} \mathbf{B}_{SCC}\mathbf{B}_b(3,2) & \mathbf{0} \end{bmatrix} & \mathbf{A}_{SCC} & \mathbf{0} & \mathbf{0} \\ \mathbf{B}_{oto}\mathbf{A}_{hy,b}(3,:) & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{A}_{oto} & \mathbf{0} \\ \begin{bmatrix} \mathbf{B}_{vis}(1:2,1:2) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} & \begin{bmatrix} \mathbf{B}_{vis}(3:4,3:4) & 0 \\ \mathbf{0} & 0 \end{bmatrix} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{A}_{vis} \end{bmatrix} \begin{bmatrix} \mathbf{x}_b \\ \mathbf{x}_{hy} \\ \mathbf{x}_{nm} \\ \mathbf{x}_{SCC} \\ \mathbf{x}_{oto} \\ \mathbf{x}_{vis} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{B}_{nm} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \hat{T}_\delta + \begin{bmatrix} \mathbf{B}_b \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{B}_{SCC}\mathbf{B}_b(3,:) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} T_\phi^e \\ T_\delta^e \end{bmatrix} \tag{A.28}$$

$$\begin{bmatrix} \mathbf{y}_{SCC} \\ \mathbf{y}_{oto} \\ \mathbf{y}_{vis} \end{bmatrix} = \begin{bmatrix} \mathbf{D}_{SCC}\mathbf{A}_b(3,:) & \mathbf{0} & \begin{bmatrix} \mathbf{D}_{SCC}\mathbf{B}_b(3,2) & \mathbf{0} \end{bmatrix} & \mathbf{C}_{SCC} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{C}_{oto} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{C}_{vis} \end{bmatrix} \begin{bmatrix} \mathbf{x}_b \\ \mathbf{x}_{hy} \\ \mathbf{x}_{nm} \\ \mathbf{x}_{SCC} \\ \mathbf{x}_{oto} \\ \mathbf{x}_{vis} \end{bmatrix} + \begin{bmatrix} \mathbf{D}_{SCC}\mathbf{B}_b(3,:) \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} T_\phi^e \\ T_\delta^e \end{bmatrix} \tag{A.29}$$

A.2. FIR model

With a FIR model, output data $y(t)$ can be estimated using the measured input $w(t)$ and output $y(t)$ data. The general model is described by:

$$\hat{y}(t) = \sum_{k=1}^{m} \hat{g}(k) \cdot w(t-k) + v(t) \tag{A.30}$$

With $\hat{y}$ as the (optimaly) estimated output data, the (optimally selected) coefficients $\hat{g}$, the measured inputs w, the noise v and the number of coefficients m. Equation A.30 can be written out in a matrix form: $\hat{\mathbf{y}} = \mathbf{F}\mathbf{x} + \mathbf{v}$, with:

$$\begin{aligned} \hat{\mathbf{y}} &= \begin{bmatrix} \hat{y}(1) & \dots & \hat{y}(n) \end{bmatrix}^T \\ \mathbf{x} &= \begin{bmatrix} \hat{g}(1) & \dots & \hat{g}(m) \end{bmatrix}^T \\ \mathbf{y} &= \begin{bmatrix} w(1-1) & \dots & w(1-m) \\ \vdots & \ddots & \vdots \\ w(n-1) & \dots & w(n-m) \end{bmatrix} \\ \mathbf{v} &= \begin{bmatrix} v(1) & \dots & v(n) \end{bmatrix}^T \end{aligned}$$

We want choose the parameters in **x** such a way, that the estimated output $\hat{y}$ matches the measured output $y(t)$ as close as possible. A common method for this is to pose this problem as the following linear least squares problem, were the objective function S (which has to be minimized) is given by:

$$S(\mathbf{x}) = \|\mathbf{y} - \mathbf{F}\mathbf{x}\|^2 \tag{A.31}$$

The optimal solution, in a linear least squares sense, then is:

$$\hat{\mathbf{x}} = \left(\mathbf{F}^T\mathbf{F})^{-1}\mathbf{F}^T\mathbf{y}\right) \tag{A.32}$$

A.3. Optimization algorithms

Pattern-search optimization

Pattern-search optimization methods do not require a cost-function that is differentiable. Instead it evaluates the cost function around the provided initial condition (i.e., the initial point). Which points it exactly evaluates depends on the mesh size. It does this until the cost of a certain point is lower than the initial cost (so not all points around the initial condition might be evaluated). The process then repeats for the new point, but with a mesh size that is larger by some expansion factor. In the case that at the new point, no lower cost around that point is found, the mesh is sized down with some contraction factor. This process repeats until a stopping criteria is reached. The stopping criteria usually are a limit on the amount of function evaluations, the minimum mesh size, or the decrease in cost.

Genetic-Algorithm optimization

The GA is an iterative optimization method that is based on natural selection. It generates candidate solutions and, at each iteration, selects the most fit candidates to generate new solutions. Because a certain amount of randomization is allowed by also allowing some less fit candidates to survive and through candidate mutations, a GA can find the solution that corresponds to the global optimum.

A.4. Graphs

SCC - Roll rate comparison

In figures A.1 and A.2 the results of two separate optimization procedures are included. The VAF of the non-SCC optimization is 94.8 % with the SCC optimization. The VAF with the experimental data dropped with 0.2 %.

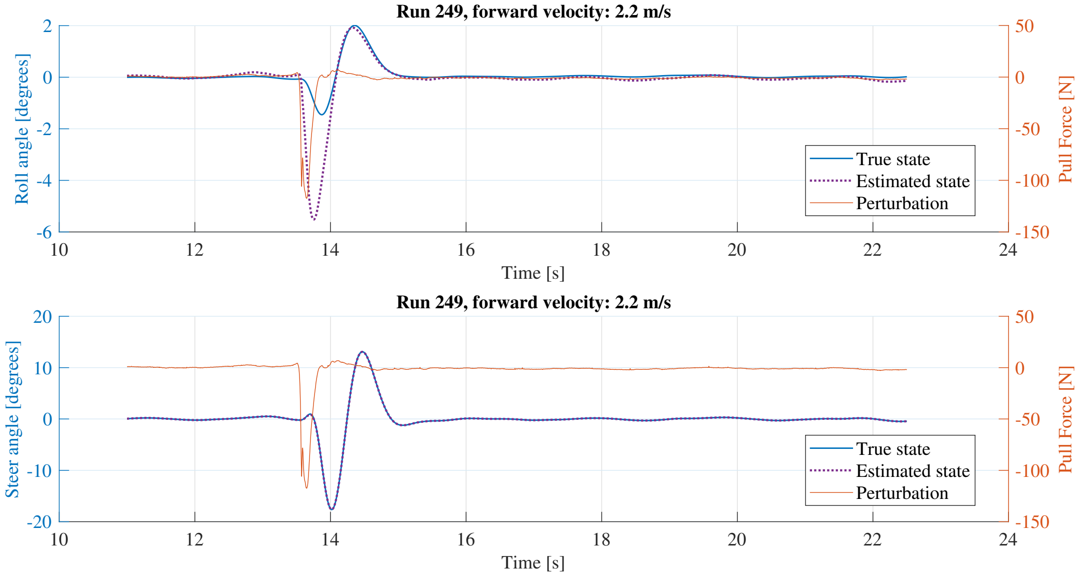

Figure A.1: Roll and steer angle as a function of time with no sensory dynamics.

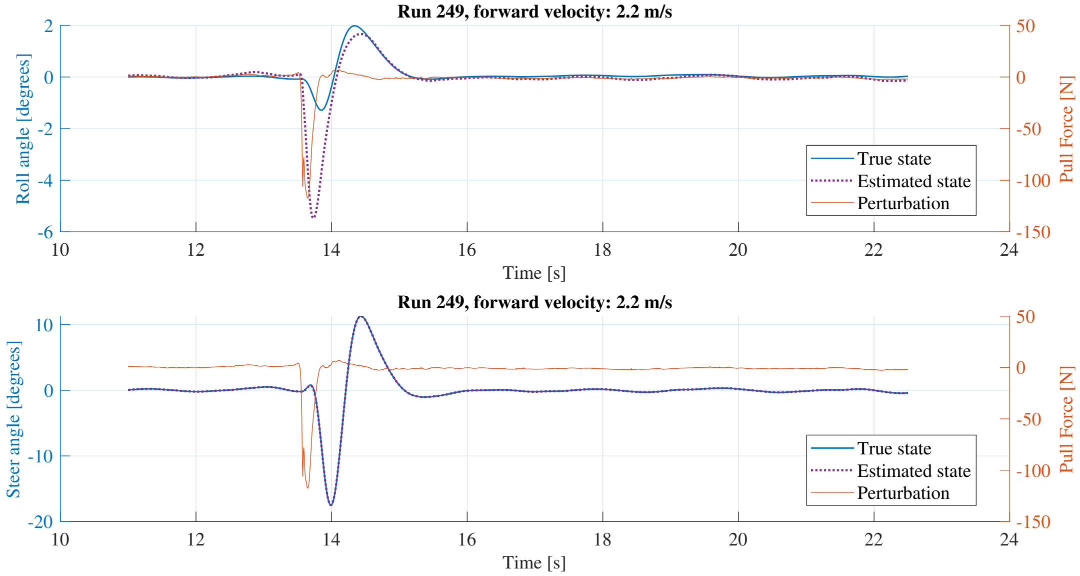

Figure A.2: Roll and steer angle as a function of time with SCC sensory dynamics.

A.5. Experimental runs used

The run ID's according to the number scheme by Moore of the evaluated runs are included here.

UC Davis treadmill

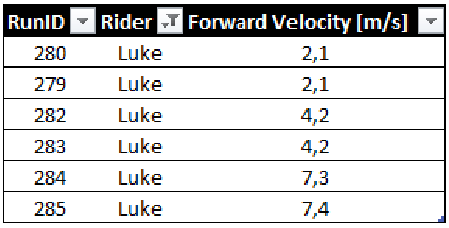

RunID	Rider	Forward Velocity [m/s]
280	Luke	2,1
279	Luke	2,1
282	Luke	4,2
283	Luke	4,2
284	Luke	7,3
285	Luke	7,4

Figure A.3: Runs from Luke used for validation.

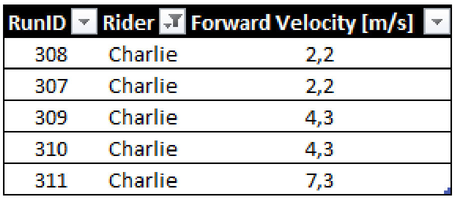

RunID	Rider	Forward Velocity [m/s]
308	Charlie	2,2
307	Charlie	2,2
309	Charlie	4,3
310	Charlie	4,3
311	Charlie	7,3

Figure A.4: Runs from Charlie used for validation.

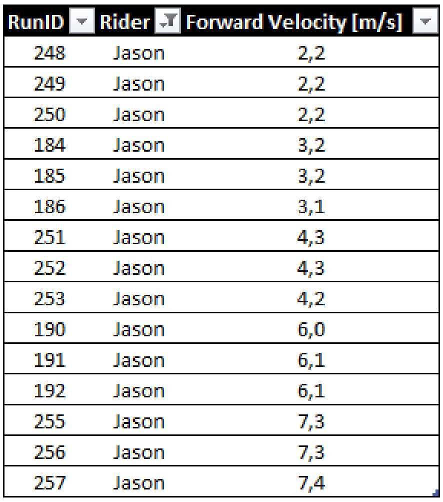

RunID	Rider	Forward Velocity [m/s]
248	Jason	2,2
249	Jason	2,2
250	Jason	2,2
184	Jason	3,2
185	Jason	3,2
186	Jason	3,1
251	Jason	4,3
252	Jason	4,3
253	Jason	4,2
190	Jason	6,0
191	Jason	6,1
192	Jason	6,1
255	Jason	7,3
256	Jason	7,3
257	Jason	7,4

Figure A.5: Runs from Jason used for validation.

UC Davis pavilion

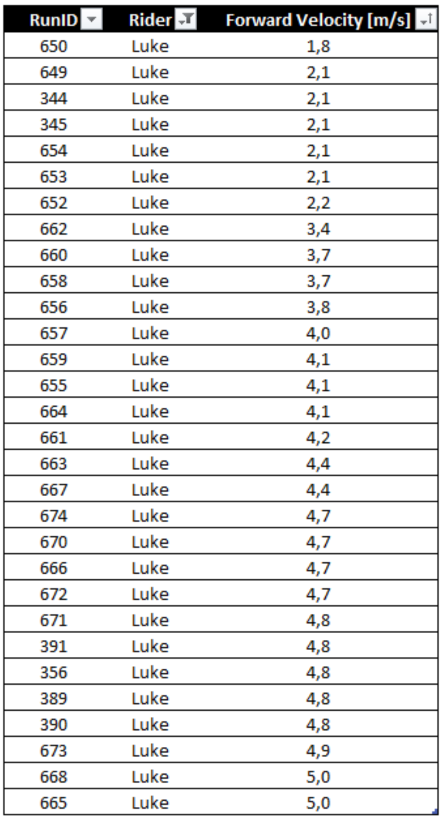

RunID	Rider	Forward Velocity [m/s]
650	Luke	1,8
649	Luke	2,1
344	Luke	2,1
345	Luke	2,1
654	Luke	2,1
653	Luke	2,1
652	Luke	2,2
662	Luke	3,4
660	Luke	3,7
658	Luke	3,7
656	Luke	3,8
657	Luke	4,0
659	Luke	4,1
655	Luke	4,1
664	Luke	4,1
661	Luke	4,2
663	Luke	4,4
667	Luke	4,4
674	Luke	4,7
670	Luke	4,7
666	Luke	4,7
672	Luke	4,7
671	Luke	4,8
391	Luke	4,8
356	Luke	4,8
389	Luke	4,8
390	Luke	4,8
673	Luke	4,9
668	Luke	5,0
665	Luke	5,0

Figure A.6: Runs from Luke used for validation.

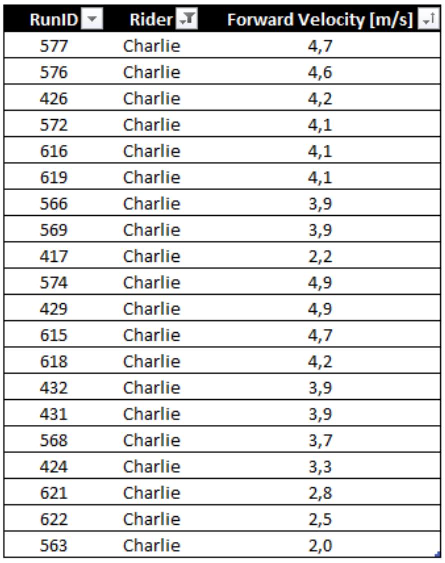

RunID	Rider	Forward Velocity [m/s]
577	Charlie	4,7
576	Charlie	4,6
426	Charlie	4,2
572	Charlie	4,1
616	Charlie	4,1
619	Charlie	4,1
566	Charlie	3,9
569	Charlie	3,9
417	Charlie	2,2
574	Charlie	4,9
429	Charlie	4,9
615	Charlie	4,7
618	Charlie	4,2
432	Charlie	3,9
431	Charlie	3,9
568	Charlie	3,7
424	Charlie	3,3
621	Charlie	2,8
622	Charlie	2,5
563	Charlie	2,0

Figure A.7: Runs from Charlie used for validation.

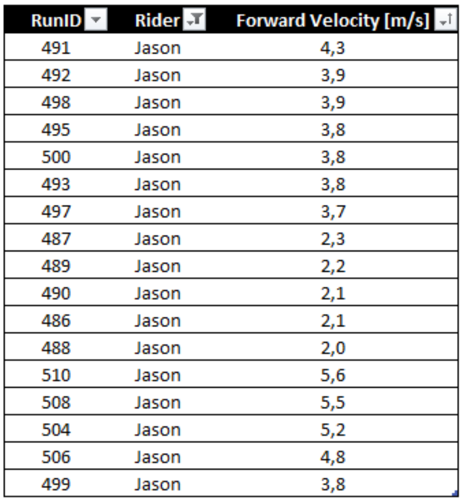

RunID	Rider	Forward Velocity [m/s]
491	Jason	4,3
492	Jason	3,9
498	Jason	3,9
495	Jason	3,8
500	Jason	3,8
493	Jason	3,8
497	Jason	3,7
487	Jason	2,3
489	Jason	2,2
490	Jason	2,1
486	Jason	2,1
488	Jason	2,0
510	Jason	5,6
508	Jason	5,5
504	Jason	5,2
506	Jason	4,8
499	Jason	3,8

Figure A.8: Runs from Jason used for validation.

Bibliography

[1] S. Akhlaghi, N. Zhou, and Z. Huang. Adaptive adjustment of noise covariance in kalman filter for dynamic state estimation. 2017.

[2] Harold L Alexander. State estimation for distributed systems with sensing delay. In *Data Structures and Target Classification*, 1991.

[3] Houshyar Asadi, Shady Mohamed, Chee Peng Lim, and Saeid Nahavandi. A review on otolith models in human perception. *Behavioural brain research*, 309:67–76, 2016.

[4] Swee T Aw, Michael J Todd, and G Michael Halmagyi. Latency and initiation of the human vestibuloocular reflex to pulsed galvanic stimulation. *Journal of neurophysiology*, 96(2):925–930, 2006.

[5] Michael Barnett-Cowan. Vestibular perception is slow: a review. *Multisensory research*, 26(4): 387–403, 2013.

[6] Pradipta Basu-Mandal, Anindya Chatterjee, and J.M. Papadopoulos. Hands-free circular motions of a benchmark bicycle. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, (463):1983–2003, 2007.

[7] R. S. Bigler. *Automobile driver sensory system modeling*. PhD thesis, Cambridge University, 2013.

[8] Carlo Bourlet. Traité des bicycles et bicyclettes suivi d'une application a la construction des vélodromes. 1894.

[9] Carlo Bourlet. Nouveau traité des bicycles et bicyclettes, tome i: équilibre et direction. 1898.

[10] Carlo Bourlet. La bicyclette, sa construction et sa forme. 1899.

[11] Carlo Bourlet. Étude théorique sur la bicyclette. *Bull. Soc. Math. France*, 27:47–67, 76–96, 1899.

[12] Jospeh-Valentin Boussinesq. Aperçu sur la théorie de la bicyclette. *J. Math. Pures Appl.*, (5): 117–135, 1899.

[13] Jospeh-Valentin Boussinesq. Complément a une étude récente concernant la théorie de la bicyclette: influence, sur l'équilibre, des mouvements latéraux spontanés du cavalier. *J. Math. Pures Appl.*, (5):217–232, 1899.

[14] Th Brandt, Jo Dichgans, and E Koenig. Differential effects of central versus peripheral vision on egocentric and exocentric motion perception. *Experimental brain research*, 16(5):476–491, 1973.

[15] Kenneth H Britten. Mechanisms of self-motion perception. *Annu. Rev. Neurosci.*, 31:389–410, 2008.

[16] Vera E Bulsink, Alberto Doria, Dorien van de Belt, and Bart Koopman. The effect of tyre and rider properties on the stability of a bicycle. *Advances in mechanical engineering*, 7(12), 2015.

[17] CBS StatLine. Mobiliteit; per persoon, persoonskenmerken, vervoerwijzen en regio's. Technical report, 2018.

[18] TD Chu and CK Chen. Modelling and model predictive control for a bicycle-rider system. *Vehicle system dynamics*, 56(1):128–149, 2018.

[19] J.D Cole and H.Y Khoo. rediction of vehicle stability using a ”back to back” tyre test method. *International Journal of Vehicle Design - INT J VEH DES*, 26, January 2001.

[20] Matthew Cook. It takes two neurons to ride a bicycle. *Demonstration at NIPS*, 4, 2004.

[21] Herman Damveld, David Abbink, Mark Mulder, Max Mulder, Marinus (René) Van Paassen, Frans Van der Helm, and R. J. Hosman. Identification of the feedback component of the neuromuscular system in a pitch control task. 2010. doi: 10.2514/6.2010-7915.

[22] P.D.L. de Lange. Rider control identification in bicycling. Master’s thesis, TU Delft, 2011.

[23] KN De Winkel, Florian Soyka, M Barnett-Cowan, HH Bülthoff, EL Groen, and PJ Werkhoven. Integration of visual and inertial cues in the perception of angular self-motion. *Experimental brain research*, 2013.

[24] Georgios Dialynas, Christos Christoforidis, Riender Happee, and AL Schwab. The effect of haptic feedback in the balance task of bicycling. In *Symposium on the Dynamics and Control of Single Track Vehicles*. Bicycle and Motorcycle Dynamics, 2019.

[25] Alberto Doria and Mauro Tognazzo. The influence of the dynamic response of the rider’s body on the open-loop stability of a bicycle. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 228(17):3116–3132, 2014.

[26] Alberto Doria, Mauro Tognazzo, Gianmaria Cusimano, Vera Bulsink, Adrian Cooke, and Bart Koopman. Identification of the mechanical properties of bicycle tyres for modelling of bicycle dynamics. *Vehicle system dynamics*, 51(3):405–420, 2013.

[27] Anthony John Redfern Doyle. *The skill of bicycle riding*. PhD thesis, University of Sheffield, 1987.

[28] Andrew Dressel and Adeeb Rahman. Measuring sideslip and camber characteristics of bicycle tyres. *Vehicle system dynamics*, 50(8):1365–1378, 2012.

[29] Marc O Ernst and Heinrich H Bülthoff. Merging the senses into a robust percept. *Trends in cognitive sciences*, 8(4):162–169, 2004.

[30] European Commission. Quality of transport. Technical report, 2014.

[31] Cesar Fernandez and Jay Goldberg. Physiology of peripheral neurons innervating semicercular canals of the squirrel monkey ii. response to sinusoidal stimulation and dynamics of peripheral vestibular system. 1971.

[32] Cesar Fernandez and Jay M Goldberg. Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. i. response to static tilts and to long-duration centrifugal force. *Journal of neurophysiology*, 39(5):970–984, 1976.

[33] Neil H Getz. Internal equilibrium control of a bicycle. In *Proceedings of 1995 34th IEEE Conference on Decision and Control*, volume 4, pages 4285–4287. IEEE, 1995.

[34] Neil H Getz. *Dynamic Inversion of Nonlinear Maps with Applications to Nonlinear Control and Robotics*. PhD thesis, University of California at Berkeley, 1995.

[35] James J Gibson. *The perception of the visual world*. Houghton Mifflin Company, 1950.

[36] Donald A Gordon. Static and dynamic visual fields in human space perception. *Josa*, 55(10): 1296–1303, 1965.

[37] Jose Luis Guzmán, Pedro Garcia, Tore Hägglund, Sebastian Dormido, Pedro Albertos, and Manuel Berenguel. Interactive tool for analysis of time-delay systems with dead-time compensators. *Control Engineering Practice*, 16(7):824–835, 2008.

[38] R. Happee, E. de Vlugt, and A. Schouten. Posture maintenance of the human upper extremity; identification of intrinsic and reflex based contributions. *SAE international journal of passenger cars: mechanical systems*, 2009.

[39] Ronald Hess, Jason Keith Moore, and Mont Hubbard. Modeling the manually controlled bicycle. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 42(3): 545–557, 2012.

[40] Ruud Hosman. *Pilot's Perception and Control of Aircraft Motions*. PhD thesis, Delft University of Technology, 1996.

[41] Ruud Hosman and Henk Stassen. Pilot's perception in the control of aircraft motions. *Control engineering practice*, 7(11):1421–1428, 1999.

[42] James Houk and William Simon. Responses of golgi tendon organs to forces applied to muscle tendon. *Journal of neurophysiology*, 30(6):1466–1481, 1967.

[43] Léna Jami. Golgi tendon organs in mammalian skeletal muscle: functional properties and central actions. *Physiological reviews*, 72(3):623–666, 1992.

[44] Marc Jeannerod. Action monitoring and forward control of movements. In Micheal A. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 83–85. The MIT Press, 2nd edition, 2003. ISBN 0-262-011097-2.

[45] Faisal Karmali, Koeun Lim, and Daniel M Merfeld. Visual and vestibular perceptual thresholds each demonstrate better precision at specific frequencies and also exhibit optimal integration. *Journal of neurophysiology*, 111(12):2393–2403, 2014.

[46] Osamu Kawakami, Yoshiki Kaneoke, Koichi Maruyama, Ryusuke Kakigi, Tomohisa Okada, Norihiro Sadato, and Yoshiharu Yonekura. Visual detection of motion speed in humans: spatiotemporal analysis by fmri and meg. *Human brain mapping*, 16(2):104–118, 2002.

[47] David L. Kleinman. Optimal control of linear systems with time-delay and observation noise. *IEEE Transactions on automatic control*, 14:524–527, October 1969. ISSN 0018-9286. doi: 10.1109/tac.1969.1099242.

[48] J. D. G. Kooijman, A. L. Schwab, and J. P. Meijaard. Experimental validation of a model of an uncontrolled bicycle. 19:115–132, 2008. ISSN 1384-5640. doi: 10.1007/s11044-007-9050-x.

[49] JDG Kooijman, AL Schwab, and Jason K Moore. Some observations on human control of a bicycle. In *ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 2021–2028. American Society of Mechanical Engineers, 2009.

[50] JDG Kooijman, AL Schwab, and Jason K Moore. Some observations on human control of a bicycle. In *ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, pages 2021–2028. American Society of Mechanical Engineers Digital Collection, 2010.

[51] Johannes Dionisius Gerardus Kooijman. *Bicycle Rider Control - Observations, Modeling & Experiments*. PhD thesis, TU Delft, 2012.

[52] David J.N. Limebeer and Robin S. Sharp. Bicycles, motorcycles and models. *IEEE Control Systems Magazine*, pages 34–61, October 2006.

[53] Lennart Ljung. *System Identification: Theory for the User*. Prentice Hall PTR, Upper Saddle River, 2nd edition edition, 1999. ISBN 978-0136566953.

[54] A. van Lunteren and H. G. Stassen. Investigations on the bicycle simulator. *Annual report 1969 of the man-machine systems group*, pages 3–55, 1970.

[55] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11):1432–1438, 2006.

[56] Matteo Massaro, Roberto Lot, and Vittore Cossalter. A virtual motorcycle driver to simulate real manoeuvres from experimental data. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of automobile engineering*, 226(9):1211–1219, 2012.

[57] P. B. C. Matthews. Evidence from the use of vibration that the human long-latency stretch reflex depends upon spindle secondary afferents. *The Journal of physiology*, 348(1):383–415, 1984.

[58] Duane T McRuer and Henry R Jex. A review of quasi-linear pilot models. *IEEE transactions on human factors in electronics*, 3:231–249, 1967.

[59] J. P. Meijaard, Jim M. Papadopoulos, Andy Ruina, and A. L. Schwab. Linearized dynamics equations for the balance and steer of a bicycle: a benchmark and review. 463:1955–1982, 2007. ISSN 1364-5021. doi: 10.1098/rspa.2007.1857.

[60] J.P. Meijaard and A.L. Schwab. Linearized equations for an extended bicycle model. pages 772–772, 2006. doi: https://doi.org/10.1007/1-4020-5370-3_772.

[61] RC Miall, D Jo Weir, Daniel M Wolpert, and JF Stein. Is the cerebellum a smith predictor? *Journal of motor behavior*, 25(3):203–216, 1993.

[62] Milana P Mileusnic and Gerald E Loeb. Mathematical models of proprioceptors. ii. structure and function of the golgi tendon organ. *Journal of Neurophysiology*, 96(4):1789–1802, 2006.

[63] Jason K. Moore. *Human Control of a Bicycle*. PhD thesis, University of California at Davis, 2012.

[64] Max Mulder, Daan M Pool, David A Abbink, Erwin R Boer, Peter MT Zaal, Frank M Drop, Kasper van der El, and Marinus M van Paassen. Manual control cybernetics: State-of-the-art and current trends. *IEEE Transactions on Human-Machine Systems*, 48(5):468–485, 2017.

[65] Christopher J. Nash, David J. Cole, and Robert S. Bigler. A review of human sensory dynamics for application to models of driver steering and speed control. 110:91–116, 2016. ISSN 0340-1200. doi: 10.1007/s00422-016-0682-x.

[66] CJ Nash and David James Cole. Modelling the influence of sensory dynamics on linear and nonlinear driver steering control. *Vehicle system dynamics*, 56(5):689–718, 2018.

[67] AC Newberry, MJ Griffin, and M Dowson. Driver perception of steering feel. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 221(4):405–415, 2007.

[68] Charles Clark Ormsby. *Model of human dynamic orientation*. PhD thesis, Massachusetts Institute of Technology, 1974.

[69] Manfred Plöchl, Johannes Edelmann, Bernhard Angrosch, and Christoph Ott. On the wobble mode of a bicycle. *Vehicle system dynamics*, 50(3):415–429, 2012.

[70] RE Poppele and RJ Bowman. Quantitative description of linear behavior of mammalian muscle spindles. *Journal of Neurophysiology*, 33(1):59–72, 1970.

[71] Fred Previc and Tom Mullen. A comparison of the latencies of visually induced postural change and self-motion perception. *Journal of Vestibular Research*, 1:317–323, 1991.

[72] Dale Purves, George J Augustine, David Fitzpatrick, William C Hall, Anthony-Samuel LaMantia, James O McNamara, and Leonard E White, editors. *Neuroscience*. Sinauer Associates, Inc, 5th edition, 2012.

[73] Florian Raudies and Heiko Neumann. A review and evaluation of methods estimating ego-motion. *Computer Vision and Image Understanding*, 116(5):606–633, 2012.

[74] Robert W Roig. A comparison between human operator and optimum linear controller rms-error performance. *IRE Transactions on Human Factors in Electronics*, (1):18–21, 1962.

[75] R.D. Jr. Roland and J.P. Lynch. *Bicycle dynamics: tire characteristics and rider modeling*. Cornell Aeronautical Laboratory, Inc, 1972.

[76] Alfred C Schouten, Winfred Mugge, and Frans CT van der Helm. Nmclab, a model to assess the contributions of muscle visco-elasticity and afferent feedback to joint dynamics. *Journal of biomechanics*, 41(8):1659–1667, 2008.

[77] A. L. Schwab, J. P. Meijaard, and J. D.G. Kooijman. Lateral dynamics of a bicycle with a passive rider model: stability and controllability. *Vehicle System Dynamics*, 50(8):1209–1224, aug 2012. doi: 10.1080/00423114.2011.610898.

[78] AL Schwab, JDG Kooijman, and Jacob Philippus Meijaard. Some recent developments in bicycle dynamics and control. *4th European Conference on Structural Control, ECSC 2008*, pages 695–702, 2008.

[79] AL Schwab, PDL de Lange, R Happee, and Jason K Moore. Rider control identification in bicycling using lateral force perturbation tests. *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics*, 227(4):390–406, 2013. doi: 10.1177/1464419313492317.

[80] Arend L Schwab and JP Meijaard. A review on bicycle dynamics and rider control. *Vehicle System Dynamics*, 51(7):1059–1090, 2013.

[81] Robin S Sharp. Optimal stabilization and path-following controls for a bicycle. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 221 (4):415–427, 2007.

[82] R.S. Sharp. The stability and control of motorcycles. *Joarnal Mechanical Engineering Science*, 13(5):316–329, 1971.

[83] RS Sharp. A review of motorcycle steering behavior and straight line stability characteristics. Technical report, SAE Technical Paper, 1978.

[84] O. J. M. Smith. Closer control of loops with dead time. *Chemical Engineering Progress*, 53: 217–219, 1957.

[85] Damoon Soudbakhsh, Yang Zhang, and Jingang Yi. Stability analysis of human rider's balance control of stationary bicycles. In *2012 American Control Conference (ACC)*, pages 2755–2760, 2012.

[86] F Soyka, M Barnett-Cowan, P Robuffo Giordano, and HH Bülthoff. Integration of translational and rotational vestibular cues for direction detection during eccentric rotations. *2th Conference of Junior Neuroscientists of Tübingen (NeNA 2011)*, 2011.

[87] H. G. Stassen and A. van Lunteren. The bicycle simulator. *Progress report january 1970 until january 1973 of the man-machine systems group*, pages 93–120, 1973.

[88] Robert J Telban and Frank M Cardullo. Motion cueing algorithm development: Human-centered linear and nonlinear approaches. *NASA/CR-2005-213747*, 2005.

[89] Stephen Timoshenko and Donovan Harold Young. *Advanced Dynamics*. McGraw-Hill, New York, 1948.

[90] Alen Turnwald and Steven Liu. A nonlinear bike model for purposes of controller and observer design. *IFAC-PapersOnLine*, 51(2):391–396, 2018.

[91] WA Van de Grind. The possible structure and role of neuronal smart mechanisms in vision. *Cognitive Systems*, 2(2):163–180, 1988.

[92] WA Van de Grind, JJ Koenderink, and AJ Van Doorn. The distribution of human motion detector properties in the monocular visual field. *Vision research*, 26(5):797–810, 1986.

[93] Herman van der Kooij, Ron Jacobs, Bart Koopman, and Henk Grootenboer. A multisensory integration model of human stance control. *Biological cybernetics*, 80(5):299–308, 1999.

[94] Herman van der Kooij, Ron Jacobs, Bart Koopman, and Frans van der Helm. An adaptive model of sensory integration in a dynamic environment. *Biological Cybernetics 103-115*, 84, 2001.

[95] Herman van der Kooij, Bart Koopman, and Frans C.T. van der Helm. Human motion control reader. TU Delft course wb2407, 2008.

[96] MM René van Paassen and Max Mulder. Identification of human operator control behaviour in multiple-loop tracking tasks. *IFAC Proceedings Volumes*, 31(26):455–460, 1998.

[97] VeiligheidNL. Seh-bezoeken door fietsongevallen. Technical report, 2018.

[98] Everett Wang, Juncheng Zou, Gengping Xue, Lin Yuan, Zhaoheng Zeng, Gengfeng Xue, Yijun Liu, Gary Zhang, and Qun Fan. Symbolic derivation of bicycle kinematics with toroidal wheels. *MATEC Web of Conferences*, 34, 2015.

[99] Pengcheng Wang and Jingang Yi. Dynamic stability of a rider-bicycle system: Analysis and experiments. In *2015 American Control Conference (ACC)*, pages 1161–1166, 2015.

[100] F. J. W. Whipple. The stability of the motion of a bicycle. *Quarterly Journal of Pure and Applied Mathematics*, (30):312–348, 1899.

[101] Laurence R Young and Jacob L Meiry. A revised dynamic otolith model. *Aerospace medicine*, 39(6):606–608, 1968.

[102] Peter MT Zaal, Daan M Pool, Jaap De Bruin, Max Mulder, and Marinus M van Paassen. Use of pitch and heave motion cues in a pitch control task. *Journal of Guidance, Control, and Dynamics*, 32(2):366–377, 2009.

[103] Peter MT Zaal, Frank M Nieuwenhuizen, Marinus M van Paassen, and Max Mulder. Modeling human control of self-motion direction with optic flow and vestibular motion. *IEEE transactions on cybernetics*, 43(2):544–556, 2013.

[104] Greg L Zacharias, Alper K Caglayan, and John B Sinacori. A model for visual flow-field cueing and self-motion estimation. *IEEE Transactions on Systems, Man, and Cybernetics*, (3):385–389, 1985.

[105] Lotfi A Zadeh. Fuzzy sets. *Information and control*, 8(3):338–353, 1965.

[106] Lotfi A Zadeh. Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on systems, Man, and Cybernetics*, (1):28–44, 1973.