Influence of accelerator pedal force feedback on truck drivers' speed control

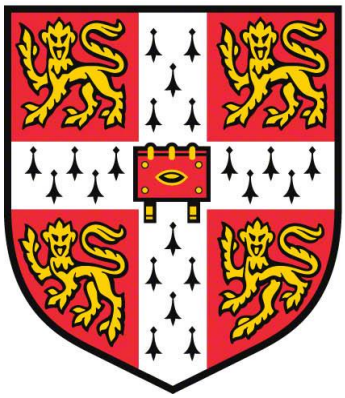

Christopher William Rosier

Department of Engineering

University of Cambridge

This dissertation is submitted for the degree of

Doctor of Philosophy

Churchill College

August 2018

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Christopher William Rosier

August 2018

Acknowledgements

The research described in this dissertation was part of a larger programme of work undertaken by the Centre for Sustainable Road Freight, a collaboration between the Universities of Cambridge and Heriot-Watt, EPSRC and an industry consortium. The industry consortium provided funding for a full studentship (fees, maintenance and expenses). In particular, I would like to thank Turners of Soham Ltd for providing access to their fleet of vehicles and their data. Terry Skelton and Stephen Blackmore (both from Turners of Soham Ltd) have been very generous with their time and advice throughout the project.

I would like to thank my supervisor, Dr David Cole, for his continued help and guidance throughout. Our regular meetings have made the project flow. I would also like to thank Prof. David Cebon for his input as my advisor. Dr Xiaoxiang Na and Dr Justin Bishop both deserve a mention for their support with the SRF Logger and SRF Mapper respectively.

Finally, I would like to thank my family and friends (especially Debbie) for supporting me through the highs and lows of completing this degree.

Influence of accelerator pedal force feedback on truck drivers' speed control

Christopher William Rosier

The UK government has set clear targets for 80% reductions (compared with 1990 levels) in greenhouse gas emissions by 2050 and pressure is increasing on the road transport industry to reduce the fuel consumption and harmful exhaust emissions of Heavy Goods Vehicles (HGVs). Vehicle manufacturers and operators alike are having to investigate and find new ways of making reductions. It is thought that improving driver behaviour offers significant potential for these reductions in fuel consumption and emissions. This thesis considers the use of Active Accelerator Pedals (AAPs) and the potential for improved driver performance that they may offer by providing pedal force feedback to the driver.

In order to develop understanding of the interactions between the human driver and accelerator pedal, two near identical tractor units, operated by Turners of Soham Ltd, were fitted with the a data logger. Data was collected and stored over a period of four months as they operated on the road. This data provided the basis for a vehicle model to be developed using real-world conditions, rather than strictly controlled test track conditions. Analysis of the behaviour of the two drivers also identified differences is styles, and explained 7% fuel consumption differences between the two drivers when negotiating roundabouts.

A new mathematical model of the human driver's longitudinal control was also developed to include the driver's cognitive control of the accelerator pedal. Model Predictive Control theory, commonly used for modelling the driver's steering control, was used and different driving styles were replicated by varying the weightings in a cost function, and a series of driving simulator experiments were performed to validate the model. Nine human drivers, two of which were professionals, performed two driving scenarios (drive cycle and car-following). The driver model was fitted to each driver individually to mathematically express the differences in their styles. The simulated RMS pedal forces from the fitted driver models lay within 20% of the measured simulator values.

The driver model was also extended to include the interactions between a human driver and an AAP using mathematical game theory. Three frameworks were proposed: decentralised, cooperative and one-sided cooperative, but, as the cooperative framework would have been very

difficult to implement experimentally, it was only considered theoretically. The same nine human drivers were presented with drive cycle and car-following scenarios whilst being assisted by pedal feedback to validate the model. Both decentralised and one-sided cooperative frameworks were applied to the fitting and compared. In the drive cycle scenario, the one-sided cooperative framework output an identical controller to the decentralised framework. In the car-following scenario, the one-sided cooperative framework produced the best fit, suggesting that the human drivers adapted their strategy to reflect the guidance from the AAP. It was noted in both scenarios that the peak pedal displacement decreased by approximately 20% with the presence of pedal force feedback.

Further work is suggested to improve the mass and road gradient data obtained from the data loggers in vehicles in order to reduce the uncertainty in the traction force and fuel rate maps. With a model for the interactions of a human driver with an AAP now in place, the pedal feedback strategy can now be optimised to improve the performance of the human driver.

Table of contents

List of figures i

List of tables vii

Nomenclature ix

Chapter 1 - Introduction and Literature Review 1

1.1 Background and Research Aim 1

1.2 Literature Review: Pedal Force Feedback 3

1.2.1 Car-following Support 3

1.2.2 Speed Limiting 7

1.2.3 Summary 12

1.3 Literature Review: Human Driver Longitudinal Control 14

1.4 Literature Review: Autonomous Speed Control 18

1.5 Research Objectives 22

1.6 Thesis Structure 22

Chapter 2 - Analysis of on-road data 25

2.1. Data Collection 25

2.1.1. Data Logging Apparatus 26

2.1.2. Vehicles 28

2.2. Identifying the vehicle model 31

2.2.1. Data extraction, resampling and filtering 31

2.2.2. Speed ratios 32

2.2.3. Road gradient 32

2.2.4. Drag properties 34

2.2.5. Mass estimation 34
2.2.6. Traction forces 37
2.2.7. Fuel consumption 44
2.2.8. Modelling Fuel Consumption 47
2.3. Identification of driver behaviour 50
2.3.1. Classification of driving activities 50
2.3.2. Junction detection 51
2.3.3. Roundabout Analysis 53
2.3.4. Statistical analysis over multiple journeys 61
2.4. Conclusions 69

Chapter 3 - Driver Model 71
3.1 Vehicle and Pedal Model 72
3.1.1 Linear Vehicle Model 72
3.1.2 Accelerator Pedal Dynamics 76
3.2 Linear Model Predictive Control Theory 88
3.3 Control Tasks 91
3.3.1 Drive-Cycle Following 91
3.3.2 Car-Following 94
3.4 Driving Experiments and Data Analysis 95
3.4.1 Experiment Design 95
3.4.2 Participants 98
3.4.3 Procedure 99
3.4.4 Drive Cycle Results 99
3.4.5 Car-Following Results 103
3.5 Model Identification 108
3.5.1 Driver model drive-cycle task 108
3.5.2 Driver model car-following task 116
3.6 Conclusions 126

Chapter 4 - Pedal Feedback 129

4.1 Two Player Decentralised Control Theory 129

4.2 Two Player Cooperative Feedback Theory 131

4.3 One-sided Cooperative Feedback 136

4.4 Driving Experiments and Data Analysis 138

4.4.1 Experiment Design 138

4.4.2 Participants 139

4.4.3 Procedure 139

4.4.4 Drive Cycle Results 139

4.4.5 Car-Following Results 144

4.5 Driver model identification 150

4.5.1 Driver model drive-cycle task with feedback 150

4.5.2 Driver model car-following task with feedback 156

4.6 Conclusions 164

Chapter 5 – Conclusions and Future Work 167

5.1 Conclusions 167

5.1.1 Chapter 2 167

5.1.2 Chapter 3 168

5.1.3 Chapter 4 169

5.2Future Work 169

5.2.1 Mass Estimation 169

5.2.2 Road Gradient 169

5.2.3 Non-linear vehicle and pedal dynamics 170

5.2.4 Limb dynamics 170

5.2.5 Pedal Feedback Strategy Optimisation 170

References 171

Appendix 1 – Active Accelerator Pedal Mechanism Analysis 181

Appendix 2 – Non-Linear Model Predictive Control 189

List of figures

Figure 1-1 – Force against pedal position for force feedback and stiffness feedback 4
Figure 1-2 – Mean time headway (THW) for different feedback types when car following 4
Figure 1-3 – Frequency Domain driver behaviour model 6
Figure 1-4 – Renault Diesel Engine Optimal acceleration against vehicle speed 9
Figure 1-5 – Difference between stiffness feedback and adaptive feedback 11
Figure 1-6 – Switched control mode driver model 16
Figure 1-7 – Errorable car following model 16
Figure 1-8 – Supervised Adaptive Dynamic Programming Control System 20
Figure 1-9 – MPC car following system 21

Figure 2-1 – SRF Logger block diagram 28
Figure 2-2 - 2015 DAF CF85, fitted with SRF Logger during test phase 29
Figure 2-3 - 2016 DAF CF85 with bulk powder trailer 29
Figure 2-4 – Map illustrating routes used by both phase 2 trucks 30
Figure 2-5 – Raw and filtered speed data for a sample drive 31
Figure 2-6 – Gear ratios 32
Figure 2-7 – Comparing GPS elevation to Google Maps elevation 33
Figure 2-8 –Estimated maximum traction forces 35
Figure 2-9 – Estimated vehicle mass distribution 37
Figure 2-10 – Force diagram for truck driving up an incline θ 38
Figure 2-11 – Estimated traction force components against time 39
Figure 2-12 – Estimated traction forces against pedal position and vehicle speed 40
Figure 2-13 - Least squares polynomial fit for the dynamic traction force map 41
Figure 2-14 – The bin standard deviations of the steady state traction force map 42
Figure 2-15 – Simple traction force map 43
Figure 2-16 – Recorded fuel use and fuel rate 44
Figure 2-17 - 3D scatter plotter of fuel rate against pedal position and vehicle speed 45
Figure 2-18 - Least squares polynomial fit for the steady state fuel map 46
Figure 2-19 – The bin standard deviations of the steady state fuel map 47

Figure 2-20 – Measured and estimated fuel rate 48
Figure 2-21 - Measured and estimated fuel rate during Cruise Control 49
Figure 2-22 – Results of the applying the Douglas-Peucker algorithm to vehicle speed 51
Figure 2-23 – Map of junction detections on Thurrock-Weymouth route 52
Figure 2-24 – Map of selected roundabout 53
Figure 2-25 - Driver 1 States 55
Figure 2-26 – Driver 2 States 56
Figure 2-27 – Ground elevation over the slip road 58
Figure 2-28 – Estimated traction forces against vehicle speed. 59
Figure 2-29 – Measured fuel use of slip road 60
Figure 2-30 – PDF histogram of gross truck masses 62
Figure 2-31 – Group A: Speed-Acceleration states over multiple journeys 63
Figure 2-32 – Group B: Driver Control states over multiple journeys 63
Figure 2-33 – PDF histogram of vehicle speeds, given non-zero vehicle speed 65
Figure 2-34 – PDF histogram of vehicle acceleration 65
Figure 2-35 - PDF histogram of pedal position, given non-zero pedal position 66
Figure 2-36 - Vehicle speed against pedal position bivariate histograms 67
Figure 2-37 – PDF histogram of Cruise Control with vehicle speed 69
Figure 2-38 – PDF histogram of Cruise Control with vehicle acceleration 69

Figure 3-1 - Driver Model Block Diagram 72
Figure 3-2 - Vehicle Model Block Diagram 72
Figure 3-3 – Vehicle Model Forces 73
Figure 3-4 – Quadratic and Linearized aerodynamic drag 74
Figure 3-5 – Free body diagram of the pedal arm 76
Figure 3-6 – Volvo accelerator pedal measurement rig 77
Figure 3-7 – Torque-rotation characteristics of the Volvo accelerator pedal 78
Figure 3-8 – Pedal assembly including motor and linkages 79
Figure 3-9 – Pedal mechanism diagram 79
Figure 3-10 – Free body diagram of linkage CD 80
Figure 3-11 – Free body diagram of linkage BC 80
Figure 3-12 – Free body diagram of linkage AB 81
Figure 3-13 – Pedal and feedback force transducer calibration 84
Figure 3-14 –Recorded and modelled force displacement relationship for accelerator pedal 86
Figure 3-15 – Closed loop vehicle control block diagram. 87
Figure 3-16 – Millbrook Suburban Drive Cycle for HGVs 92
Figure 3-17 – Modified Millbrook Suburban Drive Cycle for HGVs 93

Figure 3-18 – Driving simulator plan view 96
Figure 3-19 – Driving simulator set up for a car following scenario 97
Figure 3-20 – Audio generation in the driving simulator 98
Figure 3-21 – Mean vehicle speed error against time for the drive cycle following task 100
Figure 3-22 – Mean pedal position against time for the drive cycle following task 101
Figure 3-23 - Mean pedal force against time for the drive cycle following task 101
Figure 3-24 – Cost diagram for drive cycle scenario 102
Figure 3-25 – Driver mean vehicle speeds 103
Figure 3-26 - Mean vehicle speed against time for the car-following task 104
Figure 3-27 - Mean pedal position against time for the car-following task 104
Figure 3-28 - Mean pedal force against time for the car-following task 105
Figure 3-29 - Mean following distance against time for the car-following task 106
Figure 3-30 - Mean Time Headway (THW) against time for the car-following task 106
Figure 3-31 – Cost diagram for car following scenario. 107
Figure 3-32 - Driver mean following distances 108
Figure 3-33 – Fitting function sensitivity study for the drive cycle task 111
Figure 3-34 –Summed driver cost function sensitivity study for drive cycle task 111
Figure 3-35 – Fitted cost function weightings for the Drive Cycle scenario 112
Figure 3-36 – Total driver-model driver squared errors for Drive Cycle scenario 113
Figure 3-37 – Measured human driver and modelled driver speed profile for Driver 9 114
Figure 3-38 - Measured human driver and modelled driver pedal position for Driver 9 114
Figure 3-39 - Measured human driver and modelled driver pedal force for Driver 9 115
Figure 3-40 – Trade off graph for drive-cycle scenario 116
Figure 3-41 – Mean THW by driver including error bars of ±1 standard deviation 117
Figure 3-42 – Fitting function sensitivity study for the drive cycle task 119
Figure 3-43 – Summed driver cost function sensitivity study for car-following task 120
Figure 3-44 - Fitted cost function weightings for the car-following scenario 121
Figure 3-45 – Total driver-model driver squared errors for car-following scenario 121
Figure 3-46 – Measured human driver and modelled driver vehicle speed for Driver 1 123
Figure 3-47 - Measured human driver and modelled driver following distance for Driver 1 123
Figure 3-48 – Measured human driver and modelled driver pedal force for Driver 1 124
Figure 3-49 – Measured human driver and modelled driver pedal displacement for Driver 1 124
Figure 3-50 – Trade off graph for the car following scenario 126

Figure 4-1 – Decentralised Control block diagram 130
Figure 4-2 – Cooperative Control with independent reference signals block diagram 132
Figure 4-3 – One-sided cooperative control block diagram 136
Figure 4-4 – Mean speed error against time for the drive cycle following task 140
Figure 4-5 – Mean pedal position against time for the drive cycle following task 141
Figure 4-6 – Mean driver pedal force against time for the drive cycle following task 142
Figure 4-7 – Mean feedback force against time for the drive cycle following task 142
Figure 4-8 – Cost diagram for the drive cycle scenario with feedback 143
Figure 4-9 - Driver mean vehicle speeds 144
Figure 4-10 – Mean vehicle speed against time for the car-following task 145
Figure 4-11 – Mean pedal position against time for the car-following task 145
Figure 4-12 – Mean driver pedal force against time for the car-following task 146
Figure 4-13 – Mean feedback force against time for the car-following task 147
Figure 4-14 – Mean following distance against time for the car-following task 147
Figure 4-15 – Mean THW against time for the car-following task 148
Figure 4-16 – Cost diagram for the car-following scenario with feedback 149
Figure 4-17 - Driver mean following distances 150
Figure 4-18 – Fitting function sensitivity for the one sided cooperative cases 151
Figure 4-19 – Fitted cost function weightings for the Drive Cycle task 152
Figure 4-20 – Total driver-model driver squared errors for the Drive Cycle task 152
Figure 4-21 - Measured human driver and modelled vehicle speed for Driver 9 153
Figure 4-22 – Measured human driver and modelled pedal force for Driver 9 154
Figure 4-23 – Measured human driver and modelled feedback force for Driver 9 154
Figure 4-24 – Measured human driver and modelled pedal displacement for Driver 9 155
Figure 4-25 –Trade off graph for the drive cycle with feedback scenario 156
Figure 4-26 – Fitted cost function weightings for the car-following task 157
Figure 4-27 – Fitted ρ1 values for the one sided cooperative car-following task 158
Figure 4-28 – Total driver-model driver squared errors for car-following scenario 158
Figure 4-29 – Measure human and modelled vehicle speed for Driver 1 160
Figure 4-30 – Measured human and modelled following distance for Driver 1 160
Figure 4-31 – Measured human and modelled driver pedal force for Driver 1 161
Figure 4-32 – Measured human and modelled feedback force for Driver 1 161
Figure 4-33 – Measured human and modelled pedal displacement for Driver 1 162
Figure 4-34 –Trade off graph for the drive cycle with feedback scenario 164

Figure A1-1 – Diagram representation of the AAP mechanism 181

Figure A2-1 – Linearised aerodynamic drag 190
Figure A2-2 – Force optimisation points with a sample set of forces 192
Figure A2-3 – Vehicle speed when set a 0 to 10ms-1 ramp target over 20s. 193
Figure A2-4 – Driver pedal forces for a ramp speed demand 193

List of tables

Table 1-1 – Pedal Force Feedback Summary 12
Table 1-2 – Driver Models Summary 18
Table 1-3 – Speed Control Algorithms Summary 21

Table 2-1 – SRF Logger Data 27
Table 2-2 – Engine Specification 35
Table 2-3 – Traction force map polynomial coefficients for the dynamic surface 41
Table 2-4 – Steady state fuel map polynomial coefficients 46
Table 2-5 – Group A: Speed-Acceleration States 50
Table 2-6 – Group B: Driver Control States 50
Table 2-7 – Roundabout Data Summary 61
Table 2-8 – Gross vehicle mass statistics 62

Table 3-1 – Vehicle Model Parameters 75
Table 3-2 – Pedal Assembly Parameters 83
Table 3-3 – Effective Pedal Assembly Parameters 85
Table 3-4 – Average fitting function values for different driver cost functions 109
Table 3-5 – Average fitting function values for different driver cost functions 118

Nomenclature

$\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$	Vehicle model state matrices (discrete time)
$\mathbf{A_C}, \mathbf{B_C}, \mathbf{C_C}, \mathbf{D_C}$	Vehicle model state matrices (continuous time)
A_{noise}	Vehicle noises amplitude
a	Vehicle acceleration
a_{BC}	Pedal linkage acceleration
C	Pedal linearised damping coefficient
$C_d A$	Drag area
C_{rr}	Coefficient of rolling resistance
d	Following distance
h	Road elevation
F_{Aero}	Air resistance
$F_{Rolling}$	Rolling resistance
F_{Trac}	Traction force
$F_{TracMap}$	Traction force map
f_1, f_2	Driver/feedback force
f_B, f_C, f_D	Pedal joint forces
f_{eng}	Engine noise centre frequency
f_{Pedal}	Pedal force
f_R	Pedal pivot reaction force
f_{tyre}	Tyre noise centre frequency
G_{FD}	Final drive ratio
G_{Gear}	Gear ratio
g	Acceleration due to gravity
I_{AB}, I_{BC}, I_{CD}	Pedal linkage moments of inertia
$\mathbf{K_1}, \mathbf{K_2}$	Driver/pedal controller
k	Linearised aerodynamic drag constant of proportionality
l_{BC}, l_{CD}, l_1	Pedal characteristic lengths

l_{tread}	Characteristic tyre tread length
l_{Pedal}	Pedal arm length
M	Vehicle gross mass
M_{est}	Estimated vehicle gross mass
m_{AB}, m_{BC}, m_{CD}	Pedal linkage masses
N_P	Prediction horizon (steps)
P_{eng}	Engine Power
$P_{EngineMax}$	Maximum engine power
p_{ij}	Polynomial coefficients
$\mathbf{Q_1}, \mathbf{Q_2}, \mathbb{Q}_\mathbf{1}, \mathbb{Q}_\mathbf{2}$	Driver/feedback cost function matrices
q_{1d}, q_{2d}	Driver/feedback cost function weighting on following distance error
q_{1s}, q_{2s}	Driver/feedback cost function weighting on speed error
$q_{1\phi}, q_{2\phi}$	Driver/feedback cost function weighting on pedal displacement
$q_{1\dot{\phi}}, q_{2\dot{\phi}}$	Driver/feedback cost function weighting on pedal velocity
$\mathbf{R_1}, \mathbf{R_2}, \mathbb{R}_\mathbf{1}, \mathbb{R}_\mathbf{2}$	Driver/feedback pedal force cost function matrix
$\mathbf{r_1}, \mathbf{r_2}$	Speed demand vectors
r_1, r_2	Driver/feedback cost function weighting on pedal force
r_{Wheel}	Wheel radius
$\mathbf{S_{Q_1}}, \mathbf{S_{Q_2}}$	Driver/feedback cost function matrix
$\mathbf{S_{R_1}}, \mathbf{S_{R_2}}$	Driver/feedback pedal force cost function matrix
$\mathbf{T_1}, \mathbf{T_2}$	Driver/sustainable target
T_{Engine}	Engine torque
$T_{EngineMax}$	Maximum engine torque
T_{Pedal}	Pedal torque
T_{Wheel}	Wheel torque
t_s	Time step
t_{THW}	Time headway
$\mathbf{U_1}, \mathbf{U_2}$	Driver/feedback force vector
V_1, V_2	Driver/sustainable cost function
v	Vehicle speed
v_{wind}	Wind speed
$\boldsymbol{x}$	Vehicle model states
y	Vehicle displacement
$\mathbf{z}$	Vehicle output states
α	Road gradient

Γ_{DC}, Γ_{CF}	Drive cycle/car following fitting functions
$\boldsymbol{\varepsilon_1}, \boldsymbol{\varepsilon_2}$	Predicted errors
η	Power train efficiency
$\boldsymbol{\Theta}$	Control matrix
θ_{road}	Road gradient
Λ	Pedal displacement error function
λ_1, λ_2	Pedal displacement error function weightings
λ_{Map}	Fuel rate map
ρ_1, ρ_2	Driver/feedback dominance
ρ_{air}	Air density
σ_{FTrac}	Traction force standard deviation
ϕ_1, ϕ_3	Pedal assembly reference angles
ϕ_{Max}	Maximum pedal displacement
ϕ_m	Pedal motor angle
ϕ_{Pedal}	Pedal displacement
$\boldsymbol{\Psi}$	Control matrix
$\boldsymbol{\Omega}$	Control matrix
$\omega_{FDin}, \omega_{FDout}$	Final drive input/output angular velocity
$\omega_{Gin}, \omega_{Gout}$	Gearbox input/output angular velocity
ω_{P1}, ω_{P2}	Maximum power engine regime angular velocity bounds
ω_{T1}, ω_{T2}	Maximum torque engine regime angular velocity bounds

Chapter 1 - Introduction and Literature Review

1.1 Background and Research Aim

Heavy Goods Vehicles (HGVs) contribute 21% of surface transport CO_2 emissions, 28% of NO_x emissions and 16% of particulate matter emissions (Low Emission HGV Task Force, 2014). Despite this, they only make up 1.5% of road vehicles. Pressure is now increasing on the transport industry to reduce fuel consumption and harmful exhaust emissions. HGVs need to contribute to the government's 80% reduction in greenhouse gas emissions by 2050 (when compared with 1990 levels). There are also tightening EU regulations, specifically targeting the emissions and fuel consumption of HGVs.

It is not only the vehicle manufacturers that tackle these reductions, but also the end users: the drivers the companies that run the vehicles. To achieve these necessary reductions, several approaches can be taken. The engine and vehicle can be redesigned to be more efficient – for example, the engine could be tuned to produce the same power with lower emissions and fuel consumption, or the cab reshaped to reduce aerodynamic drag forces. These modifications would be undertaken by the vehicle manufacturer. The vehicle manager could adopt a more fuel efficient logistical plan, optimising delivery planning to reduce road miles for instance. The driver could also adopt a more sustainable driving style.

Driving style differences can make differences of tens of percent to emissions and fuel consumption (Ericsson, 2000) (DeVlieger, 1997). There are four main approaches to altering driver behaviour: driver training, visual feedback, auditory feedback and haptic feedback

There are many ways of encouraging drivers to behave more efficiently and sustainably. Several companies exist to train drivers in specialist driving styles designed to reduce fuel consumption. However, these are sometimes sold with grand claims of large fuel consumption reductions that can be difficult to reach in reality (Wahlberg, 2002). A training course can be good at motivating a driver to behave well, but the trainer cannot remain in the cab with a truck driver all the time. Many drivers will have conflicting objectives; productivity and efficiency do not always go hand in hand, and the weighting of these two in the driver's mind may change over time. Training

is shown to have strong benefits in the short-run (Huang and Ford, 2012), but diminishing effects over time as a driver's motivation decreases (Johansson et al., 2003).

Support can be offered to drivers in the form of post-driving feedback (Gonder, 2011). Drivers could be scored or ranked against others in a fleet to add a competitive element to motivation. Financial rewards could also be offered to drivers with good performance to increase motivation (Lansdown and Saunders, 2012). Constructive feedback informing the drivers of what they did well and what they can improve on is likely to achieve better results than a simple score on its own. For this to work though, driving data would be needed.

In-vehicle feedback can continuously support and motivate the driver. There are different channels within which feedback can be communicated to the driver. Visual feedback is the most used feedback channel. A feedback device could be used to indicate to the driver when high demands are on the engine (Vagg et al., 2013), or when the vehicle is speeding (Varhelyi et al., 2014). The effect of visual feedback, as with many feedback channels, can be dependent on the type of driver (Stillwater et al., 2012). A driver that wants to improve will achieve better results than one that does not. However, improvements in one area of driver performance may result in detrimental effects in other areas due to driver distraction (Recarte and Nunes, 2003).

Alternatively, the audio channel may be used to communicate feedback to the driver. In order to prevent speeding, a simple beeper could be enough to inform the driver (Adell et al., 2008). The noise of the engine itself is a good source of auditory feedback. When under high demand, the engine becomes much louder than when idling.

The final practical channel of communication to a driver is through haptics. Haptic devices have been used in steering wheels, foot pedals and even seat belts to communicate with the driver. Foot pedals potentially offer a significant improvement in fuel consumption and emissions as they are directly linked to the longitudinal control of the vehicle.

The aim of the research described in this dissertation is to extend knowledge and understanding of how force feedback at the accelerator pedal can reduce HGV energy use. This research is part of the programme of the Centre for Sustainable Road Freight - a collaboration between the University of Cambridge, Herriot Watt University, and a consortium of road haulage operators. The centre's aim is the research of low carbon vehicle technologies and logistics in the road haulage sector and runs a series of projects to achieve this aim.

The next section reviews the published literature in the field of pedal force feedback and section 1.4 reviews the literature in the field of human driver longitudinal control. Published work on autonomous vehicle speed control is reviewed in section 1.5.

1.2 Literature Review: Pedal Force Feedback

This first section of the review examines the current state of research on pedal force feedback and how it can be used to affect driver behaviour. Focussing on three major areas: car following; speeding; and fuel consumption and emissions, research has so far neglected the HGV industry but is showing promise for the car and light goods vehicle industries.

An accelerator pedal fitted with a feedback device is known as either a force feedback foot pedal, or an Active Accelerator Pedal (AAP) (Varhelyi et al., 2004). Gonder, (2011) identified the potential benefits of pedal feedback but acknowledged the complexity involved in equipping a vehicle with the technology. It was estimated that by improving driver techniques, a fuel saving of the order of 20% could be achieved for the most aggressively driven trips. Generally, in-vehicle feedback systems can be one of two categories: a system fitted to a vehicle when built or a system fitted to a vehicle already in service. An AAP is most likely to fall into the first category.

1.2.1 Car-following Support

For car-following support, the aim is to use pedal feedback to assist the driver in keeping a safe distance from the car in front. Whilst considering a car-following support system, Mulder et al., (2011) identify the three main design issues as: the measuring of parameters; choosing the type of haptic feedback; and deciding on a relationship between the feedback and the measured parameters. For car following, relevant parameters might be inter-vehicle separation and relative vehicle speeds or time-to-collision (Mulder, 2007). There are many possible parameters and combinations of parameters, leading to a range of car-following performance. Each combination comes with both advantages and disadvantages relating to measurement methods and how well the driver interacts with the controller.

In pedal feedback, the resistive force that the driver experiences when pushing on the pedal can be modified either by varying the stiffness, or by overlaying an additional force. These two methods are known as stiffness feedback and force feedback respectively, see Figure 1-1. Mulder et al., (2011) uses a stiffness feedback approach with parameters of time-to-collision and time-headway in order to assist in the control of vehicle acceleration. The apparent effectiveness of an AAP may depend on the distance to the car in front. The driver's visual perception of closing speed is more accurate at closer distances, so their need for assistance from the pedal will be less than at greater distances (Mulder et al., 2008).

Mulder et al., (2008) examined feedback of time headway (THW) for the car-following situation. THW is length of time required for the following vehicle to travel the distance between it and the car in front. It can be considered as the maximum length of time a driver has to react to a decrease in speed of the vehicle in front. A force or stiffness, proportional to the THW error (between current THW and a target of 1.5 s) was added or subtracted from the pedal. It was found

that the introduction of feedback increased the mean time headway, away from the target, and that stiffness feedback, over force feedback, offered the greater increase, see Figure 1-2. The standard deviation of THW was significantly lower for the cases with feedback.

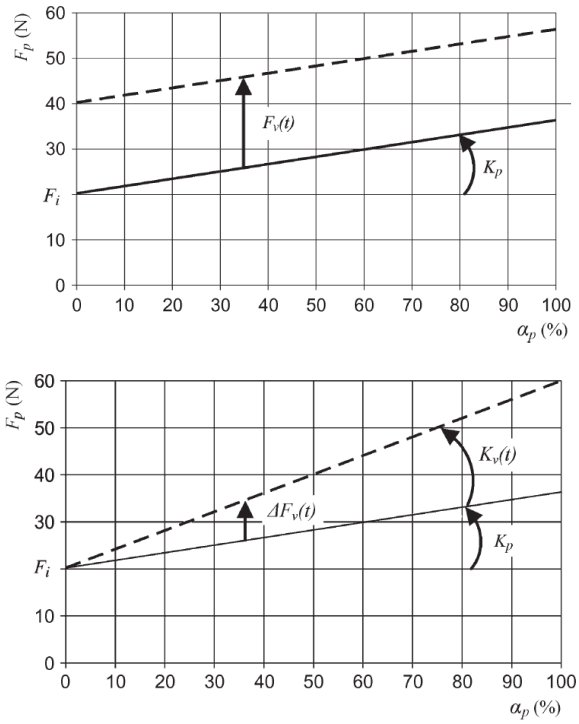

Figure 1-1 – Force against pedal position for force feedback (top) where a force offset $F_v(t)$ is applied to a constant stiffness of K_p and stiffness feedback (bottom) where a stiffness offset $K_v(t)$ is applied to a constant stiffness K_p (Mulder 2011)

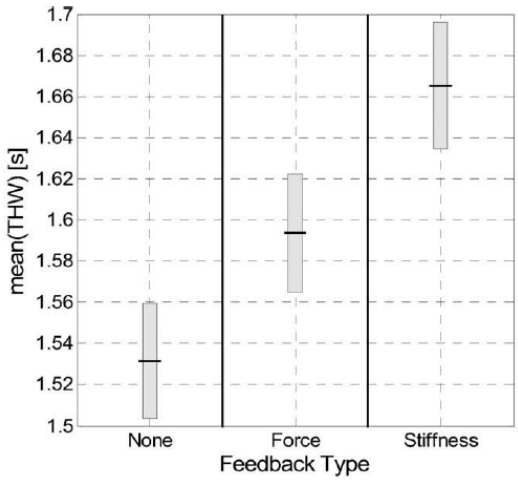

Figure 1-2 – Mean time headway (THW) for different feedback types when car following (Mulder, 2008)

Mulder et al., (2004) examined the influence of secondary tasks on car-following when supported by an AAP. Given a target THW, Drivers were also asked to press a button near the steering wheel when a shape appeared on the driving simulator screen. Drivers supported by stiffness feedback had the fastest reaction times, whilst drivers supported by force feedback had

slower reactions than the base line, unsupported case. Car-following performance was compared with and without the secondary task. It was found that the secondary task had little effect on car-following performance, but the driver's control activity increased in the supported cases, but the mean THW increased with the secondary task in the base-line test.

It is possible to extend the acceleration strategy to include deceleration support as well (Mulder et al., 2010). To do this, an algorithm was designed to increase the braking force when the driver lifts off the accelerator pedal. A relatively small experiment, with twenty participants, found that this adaptation reduced the driver's maximum braking effort during a drive cycle task in a driving simulator, whilst maintaining the frequency of braking, when compared with the standard acceleration pedal feedback device. It was also found that this deceleration support had a bigger effect when cars were closer together. In situations where cars are close together, for example in traffic jams, car-following support from an AAP can help smooth the flow of traffic and reduce the level of congestion (van Driel and van Arem, 2007). This system could also offer increased safety by encouraging the driver to approach a traffic jam more slowly.

The work by Mulder, who focuses on the effects of different feedback approaches, is quite closely supported by Abbink, who models the driver-pedal interaction. Abbink et al., (2008) models the driver and vehicle as a closed loop system and observes that continuous haptic feedback offers more support to the driver than a binary warning system. Continuous feedback gives the driver a better situation awareness and provides enough information to the driver to compensate for temporary loss of a visual feedback loop due to driver distraction.

A more detailed driver model (Abbink *et al.*, 2006), shown in Figure 3, models the driver and vehicle as five subsystems: central nervous system (CNS), spinal system (together these two form the driver), pedal, driver support system (DSS) and vehicle dynamics. A relative velocity error signal is input into the CNS (target relative velocity is zero for car following) and converted to a supra-spinal command by what is known as the visual controller. This is modelled as a simple proportional-derivative (PD) controller. The spinal system converts this to a torque on the pedal by taking into account the limb inertia, muscle activation dynamics, neural time delays, intrinsic viscoelasticity due to muscle co-contraction and other neuromuscular properties. A second pedal torque is superposed on top of the driver torque from a driver support system (DSS). The pedal dynamics, including inertia, variable stiffness (from the DSS) and damping, results in a pedal position that is the input to the vehicle dynamics, which in turn determines the speed of the vehicle.

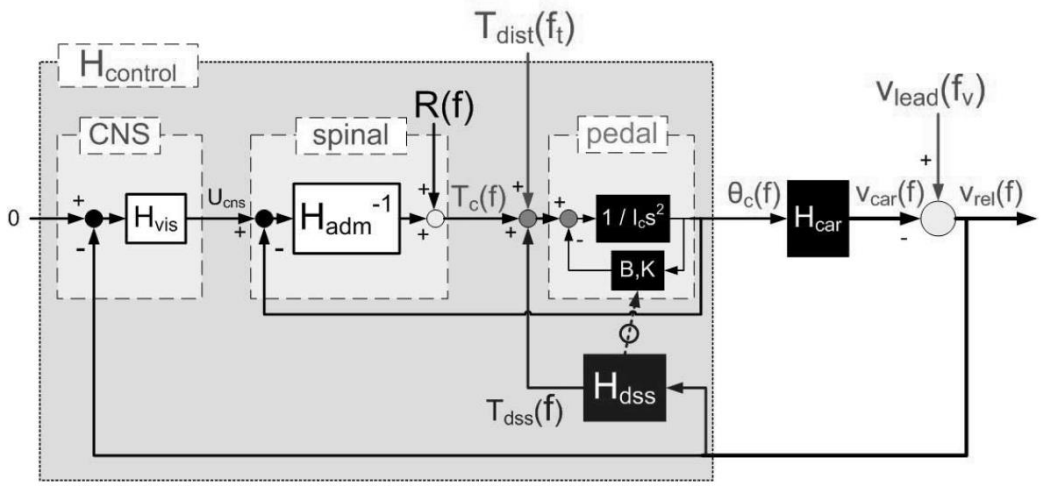

Figure 1-3 – Frequency Domain driver behaviour model (Abbink 2006)

With torque feedback in the pedal enabled, three main effects were identified by Abbink *et al.* (2006). Higher admittance (velocity response of a system to force input) of the ankle-foot complex was observed as the driver tried to maintain a constant pedal torque, meaning that the foot moved readily in response to torques from the driver support system. The driver achieved this by decreasing muscle co contraction, indicated in the experiment by reduced electromyography (EMG) activity for all muscles. This behaviour is categorised as a 'force control' task, where the driver attempts to maintain constant force at the foot, rather than a position control task (driver attempts to maintain constant foot position) or relaxation (driver tries to relax all muscles). Secondly, better coupling between perception and action was observed, because the addition of the DSS caused fewer nonlinear pedal displacements compared to the driver operating with visual feedback alone. Thirdly, performance metrics such as THW were only marginally improved, which suggested the DSS was used by the driver to reduce workload rather than improve car-following performance. Although Abbink's work represents the effects of pedal force feedback on the neuromuscular dynamics, it does not account for the driver's cognitive response to the force feedback (there is no force feedback to the CNS in Figure 1-3).

de Rosario et al., (2010) took a slightly different approach and used a vibrating accelerator pedal as a warning to the driver that the distance to the vehicle in front was below a threshold. Numerous vibration characteristics were assessed. Driver response was judged on reaction times and how quickly the speed of the car in front was matched. This highlighted a trade-off between effectiveness and driver satisfaction. As warnings became stronger and more annoying, responses were quicker, but drivers became annoyed at the system and experienced discomfort. A successful design, therefore, must find a balance. The author also suggests that staged warnings may offer better results because the problem of familiarity is less likely to arise for second or third stage warnings.

Work in this area, especially by Abbink and Mulder at TU Delft, has studied the interaction between the driver and pedal, however, little has said on the cooperative interaction of the two. Na and Cole, (2013) have applied game theory to modelling the cooperative (and non-cooperative) interaction between the driver and active front steering. Mosbach introduces the concept of a cooperative interaction between an AAP and the human driver's longitudinal control of the vehicle, but expresses the difficulty in tuning the fully cooperative AAP to each individual driver (Mosbach et al., 2017).

1.2.2 Speed Limiting

Another application of the Active Accelerator Pedal is to prevent a vehicle from exceeding the speed limit (Schumann et al., 1992). By using the haptic channel, rather than visual channel, the visual demand on the driver is reduced and less time is spent looking at the speedometer. When the driver's gaze is on the speedometer their attention on other driving tasks, such as staying in lane, reduces (Summala et al., 1996).

A three state system on the pedal (off, soon and imminent) allows the driver to drive independently most of the time, the pedal only becoming active when a potential danger is detected (Adell et al., 2011). By using two categories of warning – soon and imminent – the pedal can suggest how urgently the driver needs to react. In the EU-financed SASPENCE (SAfe SPEed safe distaNCE) project, two-stage pedal feedback was integrated into a warning system with other feedback channels, including a visual display and auditory warning, to assist drivers behave safely. The system aimed to assist in maintaining safe speeds and distances from the vehicle in front, but other safety factors were also considered. The need for a warning was assessed by using an array of sensors to detect surroundings and then understand the driver's objective (Bertolazzi et al., 2010). By comparing with a library of safe manoeuvres that would achieve the objective, unsafe manoeuvres can be flagged and an alarm sounded. The foot pedal exerted a 30N counterforce for a first warning and vibrated to warn that a dangerous event was imminent. In a test with twenty drivers, this system was observed to increase driving safety in some aspects but worsen it in others. For example, drivers were perceived to have improved interactions with vulnerable road users but the number of centre line crossings increased and harder braking occurred at traffic lights.

A study in the city of Lund, Sweden, looked into the long term effects of use of an active accelerator pedal (Varhelyi et al., 2004). 284 vehicles were fitted with an active accelerator pedal for a year. The pedal exerted a counterforce if the vehicle was speeding within the geographical boundaries of Lund. A GPS system tracked which road the vehicle was on and the corresponding speed limit. The magnitude of the pedal feedback force was such that it could be overridden by the driver if needed in emergency situations. It was found that the AAP reduced speeding and also the speed variation. Perhaps surprisingly, the AAP had a negligible effect on journey time, but beneficial effects on emissions were detected. The effect the pedal had over time decreased slightly

as drivers became used to the device. When leaving the city, the system became inactive, but data was still recorded. Many drivers slipped back to their old driving styles and the number of speeding occurrences increased. When interviewed on their experiences with the system, the AAP was judged by drivers to be effective but not attractive (Adell and Varhelyi, 2008). On a sliding scale, drivers considered the system to be more 'effective' than 'ineffective', 'clear' than 'unclear' and 'informing' than 'confusing' Young male company car drivers were identified as being the least enthusiastic about the AAP. Most of the time, this was a preconceived bias against the system, from before first use, that was maintained throughout the trial.

In the same Lund study, Hjalmdahl and Varhelyi, (2004) also observed that drivers' behaviour adapted over time. They noticed that in the long term, drivers were safer, but when the AAP was not operating, many drivers failed to adapt their speed to the surroundings. In some instances, drivers were recorded as exceeding the speed limit. This suggested that the drivers became reliant on the system to maintain safe speeds in traffic.

A study in Belgium in 2002 (Vlassenroot et al., 2007) aimed to investigate the effects of an AAP on speed change, traffic safety and drivers' attitudes. Force resistance was activated in the pedal when the driver tried to exceed the speed limit. Some drivers who frequently exceeded the limit before installation of the system would still speed as they would override the pedal. Less frequent speeders would accelerate up to the speed limit and then drive exactly on the limit rather than just below. This resulted in the unintended consequence of increasing the average vehicle speed. Even though the average speed increased, the number of speeding detections decreased.

The level of driver acceptance of AAP speed control devices is hard to measure. There are many variations of the definition adding to the difficulties. One definition of acceptance is 'the degree to which an individual intends to use a system and, where available, incorporates the system into their driving' (Adell, 2010). Many drivers don't like the system and even though they find the system useful, they would 'rather not use it' (Hof et al., 2013).

Adell et al., (2008) investigated the effects that an AAP system and a similar auditory system had in two countries: Spain and Hungary. Drivers in Hungary perceived the workload to be higher in than drivers in Spain but performance results showed little difference between the two countries. Of note is that more drivers wanted to keep the auditory system rather than the AAP even though drivers judged the AAP to be more effective.

2.1.3 Emissions and Fuel Consumption

As mentioned in the preceding section on speed control, the use of an AAP has resulted in reduced emissions whilst aiming for safer driving (Varhelyi et al., 2004). Birrell et al., (2013) used an AAP to provide a counterforce when the driver exceeded 50% throttle with the aim of reducing emissions. No emissions data were recorded, but the mean and maximum accelerations were reduced when compared with drivers being asked to drive 'economically' in a simulator.

Azzi et al., (2011) devised a feedback strategy around a mathematical model of a Renault diesel engine. This allowed feedback to be based on comparing vehicle acceleration with an optimal acceleration (calculated from the engine model) for current road speed. The optimal acceleration was calculated from a proprietary fuel consumption model of the Renault diesel engine in which vehicle speed, but not engine speed, was taken into account (Figure 1-4). If the current acceleration exceeded the optimal, then a counter force proportional to the position of the pedal was applied. On a simulated test track, with no other cars, 28 drivers were encouraged to change gear below 2000rpm and were supported with the acceleration feedback from the AAP and a screen on the console. When utilising the combined feedback, the engine model calculated a 7% reduction in polluting emissions. It was also concluded that, on the simulated test track, there was no detectable difference between emissions when only one out of two feedback systems were used (visual or a haptic). A visual device may have had a comparative advantage on the test track rather than public roads as there was no visual stimulus from other vehicles to distract the driver. Further studies were suggested on driving in traffic and reactions to critical conditions.

Larsson and Ericsson, (2009) fitted AAPs in four postal vehicles that drove a variety of routes from urban to rural. Pedal feedback was proportional to the speed of the accelerator pedal. Emissions were calculated using the Veto micro mechanical emission model and in contrast to Azzi et al., (2011), they found that an acceleration advisory AAP had no significant reduction in fuel consumption and only slight improvements in emissions for two out of three routes. The urban routes were found to offer a slight reduction in polluting emissions, whereas the rural routes did not. The difference in findings between Larsson and Eriscsson and Azzi et al. could be down to the difference in feedback, emissions model or road conditions.

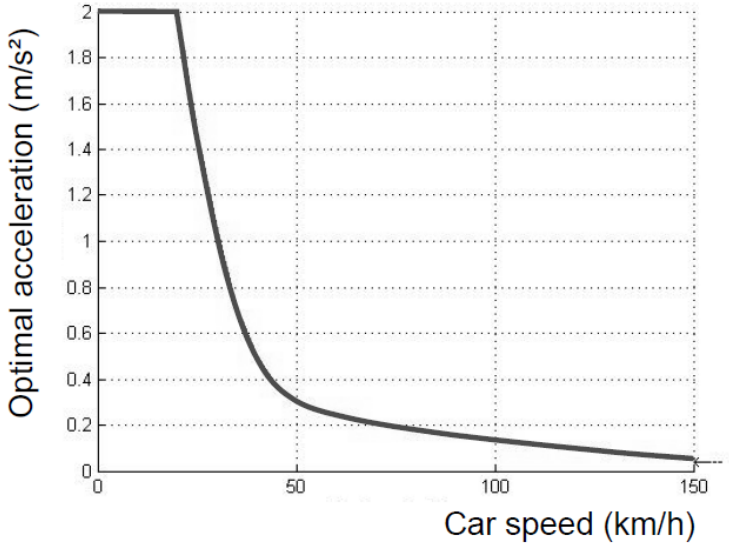

Figure 1-4 – Renault Diesel Engine Optimal acceleration against vehicle speed (Azzi, 2011)

Jamson et al., (2013) examined three different feedback options for the accelerator pedal as part of the ecoDriver Project. The aim was to discover which feedback offered most assistance to drivers aiming for maximum fuel economy. Force feedback, stiffness feedback and an 'adaptive' stiffness feedback design were compared using a paired comparison approach. In force feedback, step changes in force were applied to the pedal. In stiffness feedback, changes in pedal firmness were applied, and in adaptive stiffness, the spring stiffness of the pedal was removed so that a constant force was applied until a maximum throttle point was reached, and then additional stiffness was implemented (Figure 1-5). Twenty drivers compared the three designs when being given a target accelerator pedal position (7% for cruise and 23% for acceleration phase). The designs were rated and results showed that the drivers preferred force feedback over stiffness feedback. This seems to contradict an earlier study by Mulder, (2007) which found that stiffness feedback was preferred. The difference could arise from the slightly different driving objectives that were being used. In Jamson's study, the aim was to reach a target throttle position, whilst in Mulder's, a vehicle separation was targeted. Jamson acknowledged that the study was not conclusive and that further investigation was required on the longer term effects of force feedback versus stiffness feedback. Driver fatigue may cause a change in conditions that would result in higher physical workload (Abbink, 2006).

Also part of the ecoDriver project, Toffetti et al., (2014) investigated a visual feedback device. Comparisons were made between the visual feedback system alone and the system supported by an AAP. He concluded that when supported by the AAP as well as the visual feedback system, drivers reacted more promptly to feedback. It was also noted that there was little preference from the drivers for one system or the other.

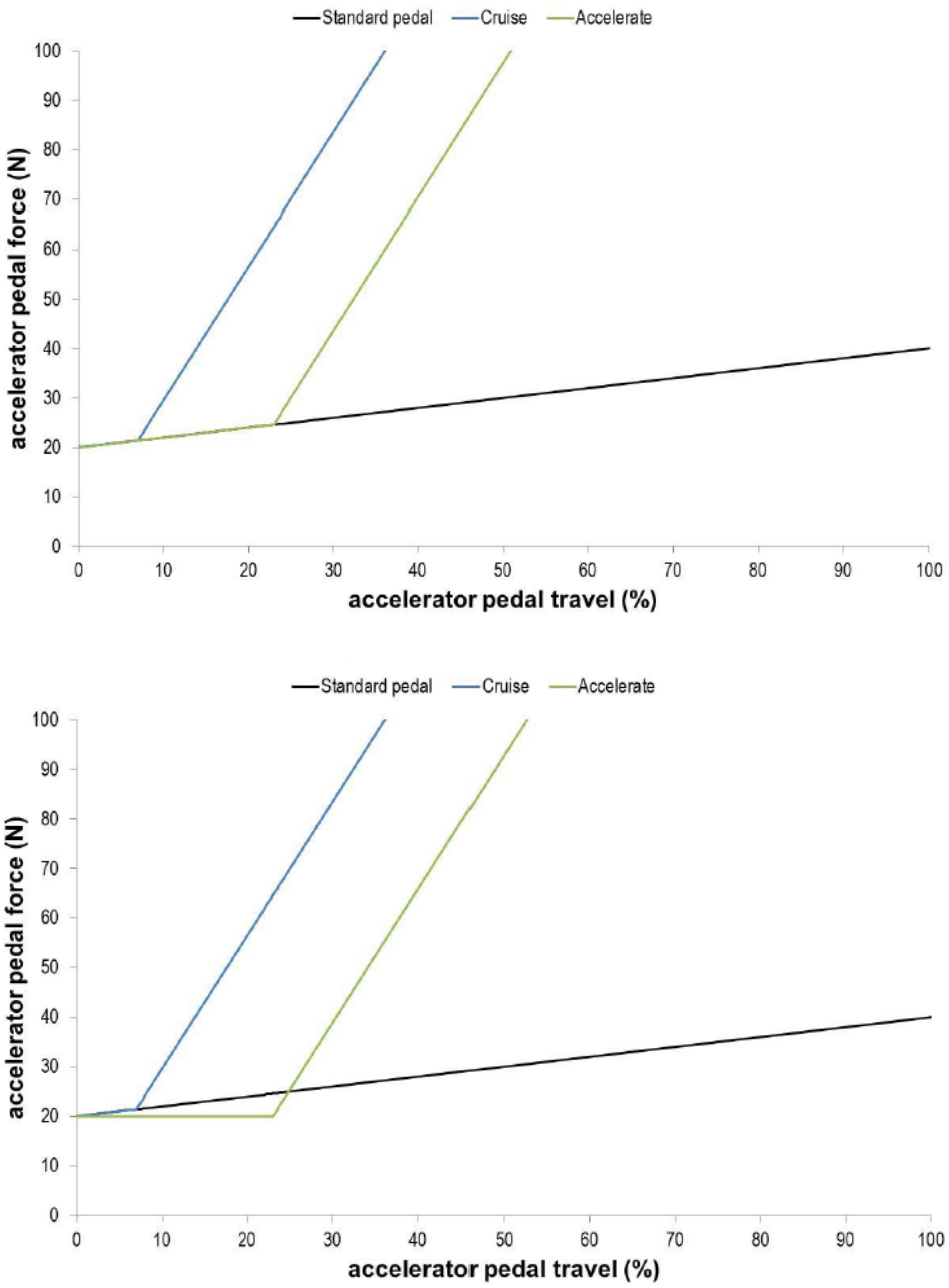

Figure 1-5 – Difference between stiffness feedback (top) and adaptive stiffness feedback (bottom) (Jamson, 2013)

An alternative to looking at the acceleration phase of the drive cycle for fuel savings is to look at the deceleration phase. A warning vibration from the pedal could be used to prompt the driver to lift off the accelerator pedal (Hajek et al., 2011). By anticipating when the vehicle will have to stop, a coast to stop could prevent energy being wasted through braking. This study claimed to find fuel consumption reductions of up to 7.5%. The test participants were interviewed and 70% viewed the device positively. Rather than viewing the warning as a signal to coast to a halt, some drivers simply saw it as a warning that deceleration would be required in the future, and still resorted to braking. The device was more warmly welcomed in rural and highway environments, rather than urban environments, where start-stop driving is prevalent. Extending this, McIlroy (2017) also used

a warning vibration to encourage coasting, but found in some situations, where the warning came too close to event occurring, performance decreased, and the driver's acceptance of the device was low. With warnings of 8-12s, driver acceptance was higher, and fuel savings of up to 11% observed (McIlroy, 2017).

1.2.3 Summary

Many of the papers reviewed offer similar opinions on the AAP. It is frequently deemed effective, but not attractive to the driver. This is summarised in Table 1. There has also been very little said on use of such a system in HGVs, where the vehicle is often very limited in available power, so driving strategy may be different.

Key: ✓ = agrees, × = disagrees, - = not mentioned

Table 1 – Pedal Force Feedback Summary

	Paper	Summary	Effective	Attractive
Fuel Consumption and Emissions	Azzi (2011)	-Investigated haptic pedal feedback and visual feedback in cars -Haptic feedback judged more effective and stabilised foot pedal	✓	-
	Birrel (2013)	-Pedal vibrated when throttle exceeded 50% -Reduced mean acceleration and maximum throttle	✓	-
	Hajek (2011)	-Anticipated when to lift off accelerator pedal to coast to stop -Calculated 7.5% fuel reduction from simulations	✓	✓
	Hof (2013)	-Review of various driver support systems for reducing emissions -No haptic devices for trucks	✓	-
	Jamson (2013)	-Compared force, stiffness and adaptive stiffness feedback -Drivers preferred high levels of force feedback	✓	-
	Larsson (2009)	-Acceleration advisory AAP fitted to post vans -No significant drop in fuel consumption across a range of routes	×	-
	Toffetti (2014)	-Investigated haptic pedal feedback and visual feedback in cars -No preference from the driver for either system	✓	✓
	Varhelyi (2004)	-Speed limiting AAP long term test -11% reduction in CO, 7% reduction in NOx	✓	✓
	McIlroy (2017)	-Vibrating pedal warnings given to indicate coasting opportunities -Drive acceptance was low; fuel savings up to 11%	✓	×
Car Following	Abbink (2006)	-Neuromuscular analysis of haptic pedal for car following -Developed driver-haptic feedback driver model	✓	-
	Abbink (2008)	-Models driver as a closed loop system -Continuous feedback offers better support than binary for car following	✓	-

	Abbink (2011)	-Neuromuscular analysis of haptic pedal for car following -High admittances of ankle foot complex	✓	-
	Adell (2011)	-Haptic, visual and audio warning system for speed and distance control -Potential for safety enhancement	✓	✓
	Bertolazzi (2010)	-AAP and audio warning system for unsafe manoeuvres -Judged as effective and well accepted	✓	-
	Mulder (2004)	-Secondary task test whilst car following with AAP -Little impact on car following performance but increased control activity	✓	-
	Mulder (2007)	-Use of AAP for car-following -Driver adopts a 'force-task'	✓	-
	Mulder (2008)	-Investigated car following behaviour changes with AAP -Performance improved and control activity decreased	✓	-
	Mulder (2010)	-Extends AAP strategy to include extra braking control -Reduced maximum breaking and increased safety	✓	-
	Mulder (2011)	-Identifies three design issues for AAP: -Type of feedback, relationship between feedback and distance and quantification	✓	-
	Mosbach (2017)	-Proposes cooperative AAP system -Considers difficulties of fully cooperative system	✓	-
Speed	Adell (2008)	-Comparison of AAP with audio warning for speeding -Drivers preferred the audio warning even though AAP was judged to be more effective	✓	×
	Adell (2008)	-Acceptance analysis of AAP -Judged as effective but not attractive	✓	×
	Hjalmdahl (2004)	-Long term effects of driving with AAP -Safer when using AAP, but frequent speeding when turned off	✓	-
	Schumann (1992)	-AAP can help reduce speeding -Feedback in pedal prevents driver from looking at the speedometer for too long	✓	-
	Varhelyi (2004)	-Effects of AAP decreased over time -Negligible effect on journey time	✓	✓
	Varhelyi (2014)	-Visual and Haptic warnings for blind spot, collision, speeding -Confusion about which haptic warning was issued	×	×
	Vlassenroot (2007)	-AAP to prevent speeding saw large variation in driver performance -Average speed increased due to quick acceleration to speed limit	✓	✓
Congestion	van Driel (2007)	-AAP and automated 'stop-and-go' reduced congestion -Safety increased as cars approached traffic jams slower	✓	-

Collision Avoidance	de Rosario (2010)	-Vibration warnings for separation distance -Unsafe distances produced annoying warnings	✓	-

1.3 Literature Review: Human Driver Longitudinal Control

Drivers are inherently difficult to model accurately due to the complexity of the human mind. An understanding of how the decision making process works and how drivers choose and process their objectives is important if feedback is to be provided effectively. This review examines the current state of research on human driver models, with a focus on those for longitudinal vehicle control.

There are many driver models already in use. These driver models take a variety of inputs and produce an equally varied array of outputs, depending on the objective of the model. Most models include the dynamics of the vehicle in the model, so will output a vehicle acceleration or speed.

Driver model objectives may be to mimic human drivers whilst following a car in front, or may be to produce an ideal and perfectly safe driver. Many of these driver models have a very specific application, due to the complex nature of modelling human behaviour.

In a thorough review of many driver and vehicle models, (Mastinu and Ploechl, 2014), it is pointed out that the vehicle is much better understood than the driver and that the weak point in any driver/vehicle simulations is the driver model. Driver models have very specific applications, so the best model for the task is very dependent on the individual applications. In some cases, the driver model is identified as the ‘test driver’. In some others, however, such as in development of new vehicle components, the role demanded of the model is closer to that of a ‘driver robot’, needing to maintain a target speed. For that, a different type of model may be more appropriate.

An early driver model, that is now used as a benchmark by many others, is Chandler’s (Chandler et al., 1958). In this paper, the author uses a simple controller to determine the desired acceleration of the vehicle by applying a gain to the relative speed of the two vehicles in order to study the stability of traffic flow. The response is not instantaneous, however, and a 1.5 s delay is implemented between the time the relative speed is assessed by the driver and the time the vehicle accelerates to respond. This is due to driver processing time and the dynamics of the engine and vehicle.

As a slight extension to the proportional controller, limits can be added to parameters. For example, a driver may choose to operate at a maximum acceleration until reaching a maximum speed (Gipps, 1981). This could resemble any acceleration to cruise scenario. The selection of these limits is dependent on the surroundings and the driver's attitudes.

The gains used in driver models can be varied to simulate the attributes of different drivers (Tang et al., 2014). Common classifications of drivers are aggressive, normal and conservative, but real drivers sit across a spectrum. These classifications were used to study the effects of different drivers on individual vehicles and hence traffic flow as these vehicles interacted with each other. It was found that there was little difference between the cases with homogenous neutral drivers and mixed drivers.

Newell, (2002) proposed a simple driver and vehicle model of cars in traffic. As one car follows another, its own speed profile is simply a translation in time and space of the vehicle in front. This allows for simplified analysis of traffic dynamics.

Alternatively, optimal control can be used to model the human driver. This is where a cost function is minimised or maximised by the controller. Burnham et al., (1974) used a minimum energy and relative distance function to model car following behaviour. Using data collected from Interstate 71, it was identified that driver behaviour differed between acceleration and cruise phases of a drive cycle as different cost functions resulted in better fits to recorded data.

Hunt et al., (2011) produced a driver model with two control modes: acceleration control and speed control – see Figure 1-6. In both modes of control, the driver's output is throttle position. In acceleration control, the driver operates an integrator controller on acceleration error. The reference acceleration is set at a maximum allowable acceleration for driver comfort. The driver model interacts with a vehicle model to convert the throttle input into vehicle acceleration. Once the target speed has been reached by the vehicle, the driver model switches from acceleration control mode to speed control mode. In speed control, the driver operates a PI controller on speed error. The model also includes a gear selector, using an array of maximum allowable engine speeds to prompt a shift.

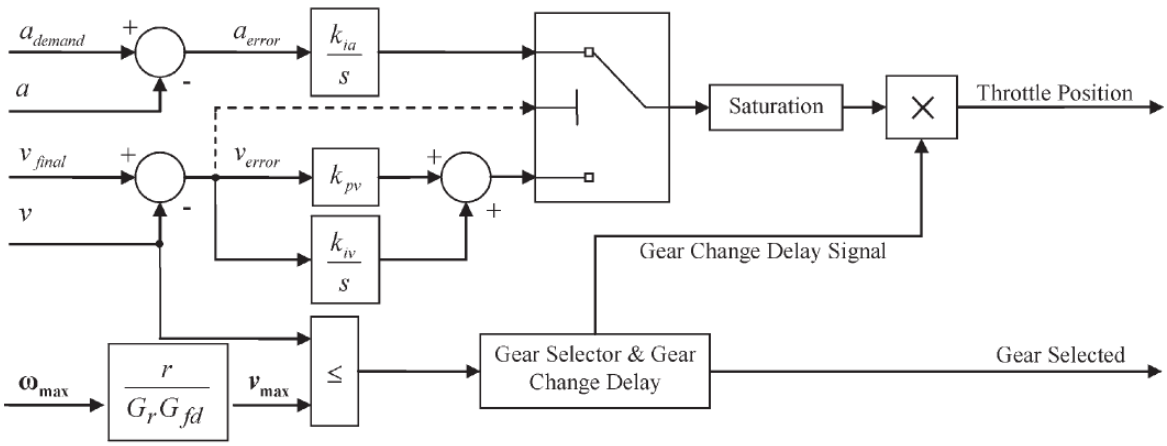

Figure 1-6 – Hunt et al., (2011) switched control mode driver model

Unlike the previous models, the Hunt system models a driver's response to a drive cycle, rather than a car following task. The objective of the driver model is to predict the fuel consumption of an HGV during set simple drive cycles (Odhams et al., 2008). Of note is that, unlike some models that incorporate the driver's behaviour, there is a clear output of the driver model (the throttle position), rather than just a combined driver-vehicle system output.

Drivers are not perfect, and make mistakes. To cover this, an 'errorable' car following model was developed by Yang and Peng, (2010) to take into account three potential driver errors: perceptual limitations, distraction and driver time delays – see Figure 1-7.

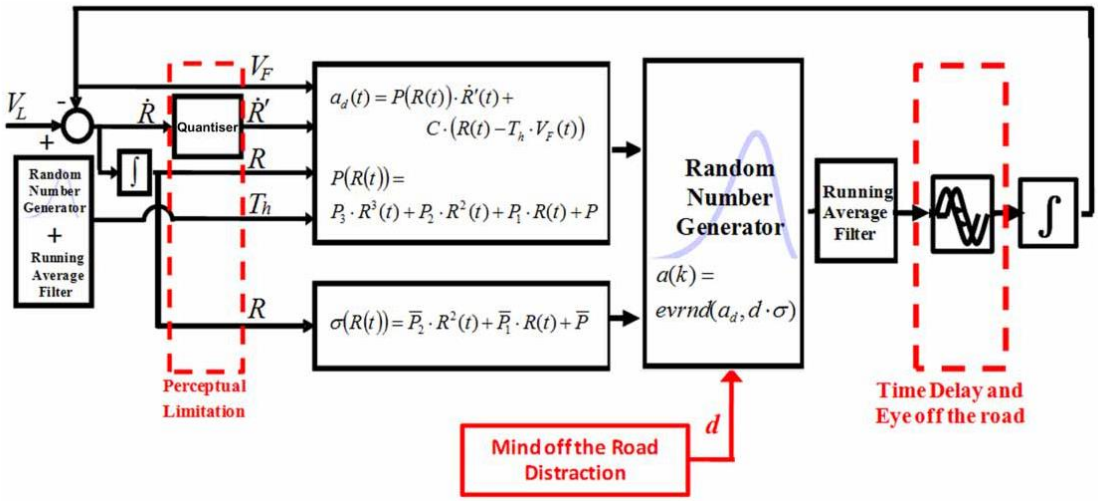

Figure 1-7 – Errorable car following model (Yang 2010)

Perceptual limitations were modelled as a quantisation on the inputs to the driver and driver distraction was modelled by randomising the acceleration of the vehicle. An extreme value distribution was applied to the vehicle acceleration with a mean determined as a function or range, range rate and THW. The standard deviation of the distribution was defined as a second order polynomial of range. Large discrepancies from the mean implied distraction.

The final type of error is defined by a time delay on the vehicle acceleration. This time delay is due to internal delays within the driver in the neuromuscular system for example, as well as eyes off the road distraction. If the driver is not paying attention to the road, then the feedback loop is opened and no new data is received by the driver. A probability distribution was also used to recreate the variable length time delay. Others have also looked at using an errorable model for similar purposes: (Przybyla et al., 2012) and (Nishiwaki et al., 2007).

Rather than introducing random errors into the driver model, McGordon et al., (2011), suggests a model based on a series of times to model the slow response of a driver. A driver speed check time models the 'recognition of cues' by setting a frequency at which the driver samples the vehicle speed. A transport delay allows for thinking time and comparison between real vehicle speed and a calculated target vehicle speed. Finally, a throttle correction time allows for the delay between comparing speeds and correcting the throttle position. A PID controller manages the magnitude of the response from the driver.

Model predictive control (MPC) is currently mostly used to model driver steering behaviour (Cole et al., 2006) (Qu et al., 2015) and (Flad et al., 2013). This is another optimal control method, meaning it is based on the minimisation/ maximisation of a cost function. In this case, the costs are associated with a predicted desired path in the future compared with an expected vehicle path. With a little modification, this could be used to model longitudinal control situations where the driver can preview a short way into the future.

Model predictive control has been used to propose an ideal response to cut in and cut out during car following (Okuyama and Murakami, 2012). The MPC minimised a weighted cost function of jerk, velocity error and distance error to simulate the response of a driver to a vehicle cutting in in front of them and providing a step change in following distance. The output of this system was a torque on the accelerator pedal. This ideal solution was used a reference for force feedback on the accelerator pedal - torque proportional to the difference between the driver's actual pedal torque, and the torque proposed by the ideal response was applied as the feedback torque on the pedal.

Sharp, (2007) used a similar optimal control algorithm, Linear Quadratic Regression (LQR), for speed tracking. By applying different cost function weightings (a cost function of performance error and control input), different performance characteristics could be achieved. Control became either 'tight' when error was made significant, or 'loose' when control input minimised.

As an alternative to predefined parametric model structures, it is possible to use non-parametric dynamic systems where the structure is determined by data instead. In a comparison of four parametric models (Constant Speed, Constant Acceleration, SUMO Simulator model and Intelligent Driver Model) with two non-parametric ones (Gaussian Mixture Regression and Artificial Neural Network), Lefevre et al., (2014) observed that for short prediction horizons, simple

parametric models were sufficient, but as prediction horizons increased, more complex parametric models, or non-parametric models were needed. The non-parametric models were noticed to perform better over all prediction horizons tested – 1-10s.

The Gaussian Mixture Model is an example of a non-parametric model. It has been used to predict pedal operation in Angkititrakul et al., (2009). Whilst close matches were achieved in the controlled simulation environment, results were less impressive when real world driving data was compared with the predictions. It was thought that this was due to the diverse uncontrolled factors present in the real world driving environment.

This review has highlighted the fact the complexity of the human driver by identifying many different approaches, each with its own specific application. The reviewed literature is summarised in Table 2.

Table 2 – Driver Models Summary

	Model Type			
Model Objective	**P, PI, PID**	**Linear Quadratic/MPC**	**Random Process**	**Sequential**
Car-following/ Speed	Chandler (1958) Tang (2014) Gipps (1981)	Burnham (1974) Sharp (2007)	Yang (2010) Przybyla (2012) Nishiwaki (2007)	Newell (2002)
Emissions	Hunt (2011) McGordon (2011) Rafael (2006)			

1.4 Literature Review: Autonomous Speed Control

Subtly different to the Human Driver models, speed control algorithms have been developed to achieve high driving performance. Several algorithms have also been produced in an attempt to create a perfect driver that delivers high quality driving performance, rather than to recreate the human driver. An application of such an algorithm would be an adaptive cruise controller for example. This is where the vehicle wishes to maintain a safe distance from the vehicle in front whilst also maintaining passenger comfort by limiting acceleration for example. These algorithms fall into similar categories as the human driver models.

Adaptive Cruise Controllers (ACC) operate as an extension of standard cruise controllers. ACCs adapt to the environment around them in order to maintain safe conditions. To do this, many ACC equipped vehicles utilise a radar or laser system to detect any vehicles around them. However, with many vehicles following each other stability issues can occur (Rajamani, 2011). It is possible for an individual vehicle to be stable, but for a string of vehicles to be unstable. This is due to the delay in response of each vehicle, and the lack of information available on what vehicles further

down the chain are doing. An analysis of this phenomenon results in conditions on the properties of the ACC control system (Swaroop, 1997).

A non-cooperative car following control law was proposed by Chien and Ioannou, (1992) that was free of oscillations and the 'slinky effect'. This used relative distance, relative velocity and relative acceleration between the host vehicle and the vehicle immediately in front. The key to eliminating oscillations in this case was the introduction of a safe following distance as a function of speed.

Sivaji and Sailaja, (2013) proposed an adaptive cruise controller that utilised a PID controller in a similar method to the benchmark Chandler model mentioned earlier (Chandler et al., 1958). However, in this model, the 'driver' receives inputs from the speed of the host vehicle, distance between the two vehicles and a target time headway (THW).

It is possible to customise a PID controller by introducing gain scheduling, (Shakouri et al., 2011). This allows different gains to be used at different speeds, meaning the controller can perform better over a range of speeds.

Yanakiev and Kanellakopoulos, (1996) suggest a cruise controller for use specifically in HGVs. A specialist controller is needed as an HGV is often very power restricted, and so a drive profile is determined by amount of available power, rather than just the driver's desired speed. A linearised HGV model is used and a couple of control techniques explored: standard and adaptive PID controllers. It was found that for good control characteristics, especially when in truck platoons, aggressive control actions were needed. This was due to the much heavier mass of HGVs compared with cars. However, the heavier mass reduces the discomfort of passengers by restricting acceleration and jerk of the vehicle. A non-linear adaptive PID controller was found to have the best characteristics of those tested, but did require the most tuning to the individual vehicle, even though steps were taken to extend the operating range of the controller.

In a different approach at ACC design, the standard drawbacks of a Dynamic Programming controller (such as slow convergence) were reduced by the implementation of what was called a Supervised Adaptive Dynamic Programming (Zhao and Hu, 2011) – see Figure 1-8. Machine learning, like this, can offer the potential for excellent control in a wide range of conditions, but the large training times required can be a major disadvantage. In this example, the control was split into two controllers: an upper and a lower controller. The upper controller studied the environment and used a Reinforced Learning (RL) approach to decide a desired acceleration. The lower controller utilised a fuzzy logic control algorithm to control the pedal and brake inputs to the host vehicle. In comparisons with a standard PI controller, it was noted that the SADP controller performed well, especially in an emergency stop situation, where it outperformed the standard ACC.

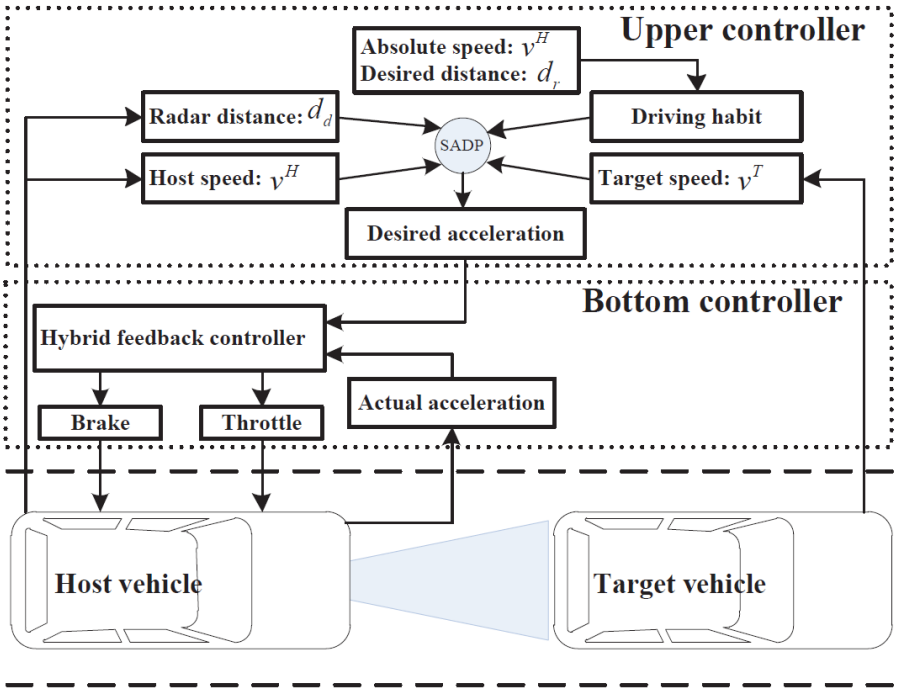

Figure 1-8 – Supervised Adaptive Dynamic Programming Control System (Zhao 2011)

As well as being used to model a human driver, Model Predictive Control has also been used in speed control algorithms for ACC's (Corona et al., 2006). A benefit of MPC is that it can utilise a cost function covering many variables and can satisfy constraints. This means that it can span a range of objectives. ACCs may wish to achieve good tracking performance, whilst reducing fuel consumption and keeping the driver comfortable.

Simulations of a heavy truck using such a Model Predictive Controller (Li et al., 2011) showed that it was possible to limit longitudinal acceleration to keep the driver comfortable, and by constraining the vehicle-following distance, fuel reductions were achieved as it could discourage other drivers to cut in and cause braking in the host vehicle - Figure 1-9. The car following model was based around a Generalised Vehicle Longitudinal Dynamics Model (following vehicle) and the interactions with a preceding vehicle. Desired acceleration $a_{fdes}$and inter-vehicle distances $d_{des}$are determined by an upper level controller and the corresponding error signals are considered as variables in the cost functions to be optimised

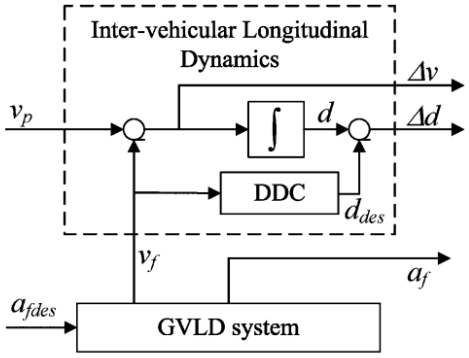

Figure 1-9 – MPC car following system (Li 2011)

Most of the reviewed literature fits into one of several categories of algorithm, as summarised by Table 3.

Table 3 – Speed Control Algorithms Summary

	Algorithm Type			
Objective	**P, PI, PID**	**Linear Quadratic**	**Dynamic Programming**	**Model Predictive**
Car following - Machine	Sivaji (2013) Shakouri (2011) Rajamani (2011)	Shakouri (2011)	Zhao (2011)	Li (2011) Corona (2006)

Several approaches have already been investigated for determining a suitable and safe speed for a vehicle. Each has its own strengths and weaknesses. A feedback controller may use one of these algorithms for deciding a reference speed for feedback.

1.5 Research Objectives

The review of published literature revealed that pedal force feedback has been investigated experimentally, and, although it is found to be effective, it is often disliked by the driver. Many mathematical models of driver speed control have been developed, but only a few include pedal force feedback. Existing models of pedal force feedback do not explicitly account for cooperation between the driver and the vehicle speed controller. The objectives of the research described in this dissertation therefore centre on applying model predictive control and cooperative control theory to understanding a driver's cognitive response to pedal feedback:

- Obtain long term measurements of the vehicle and driver states of a heavy goods vehicle and driver in real-world conditions over the period of at least one month, for at least two drivers, to better understand real-world vehicle and driver behaviour.
- Develop a mathematical model relating the longitudinal control input to the vehicle (accelerator pedal position) to the vehicle states and fuel consumption. The model should be validated using the measured data and will then be used to simulate the response of the vehicle to control inputs from drivers with and without pedal force feedback. The model should be able to estimate cumulative fuel used to within 20%.
- Investigate and quantify the differences between different human driver behaviours on the road to understand a range of driving styles.
- Develop a model of the driver's longitudinal control of heavy goods vehicles to simulate changes in driving style and behaviour. The driver model should be validated using driving simulator experiments.
- Extend the driver model to incorporate mathematical-game-theoretical interactions with an AAP. The extended driver model should be validated using driving simulator experiments and the driver model will used to simulate the driver's cognitive control response to the pedal force feedback.

1.6 Thesis Structure

Chapter 1 of this thesis provided the background and motivation to the project. A thorough literature review also reviewed previous research in the fields of pedal force feedback and vehicle speed control models.

In Chapter 2, operational heavy goods vehicles are fitted with data loggers to monitor the location and performance of the vehicle and driver. Statistical analysis is completed to quantify the differences between human drivers. The data is also used to generate a vehicle model, specific to the HGVs on the road.

A driver model is proposed in Chapter 3, utilising Model Predictive Control Theory. This driver model is designed to replicate driver behaviour in two driving tasks: drive-cycle following, and car-following. A set of driving simulator experiments were carried out to validate the model.

In Chapter 4, mathematical game theory is used to model the effect of pedal feedback on driver behaviour. Different structures are proposed for the interaction, and further driving simulator experiments are used for validation purposes.

Chapter 5 concludes the thesis with concluding remarks and further work.

Chapter 2 - Analysis of on-road data

In order to fully explore the potential of pedal force feedback in the road freight sector, first a thorough understanding of the human driver is needed to provide the baseline. In later chapters, a theoretical model is proposed and validated with simulator experiments, but to start with, in this chapter, on-road driving behaviour is observed by collecting on-road data.

The purpose of this data collection is two-fold: firstly, as already mentioned, the data will allow observation of on-road driving styles rather than the more controlled, but less realistic strategies observed in driving simulator experiments. Secondly, the data collected will help identify a suitable model of the engine and automatic gearbox to assist in accurate theoretical simulations later on.

Data loggers were placed in several vehicles, operated by Turners of Soham Ltd, over a period of time to collect and store data for a range of drivers and driving styles.

2.1. Data Collection

When in complete control of the vehicle, the driver has two outputs with which to control the longitudinal dynamics of the vehicle (assuming an automatic gearbox), the accelerator pedal and the brake pedal. The driver may choose to remove this immediate control by switching on the vehicle's cruise control, at which point they can remove their feet from the pedals. The vehicle speed and acceleration are the consequences of the driver's control. These need to be logged in order to develop a mathematical model of the vehicle. In order to assess the effectiveness of pedal feedback in reducing the fuel consumption of a truck, a fuel consumption model will also need to be developed. This means that the vehicle's fuel consumption also needs to be recorded. The environment through which the vehicle is driving can offer many disturbances to the vehicle. For example, the road gradient will affect the relationship between pedal position, vehicle speed and acceleration. For this reason, a GPS position is added to the list of required data signals.

2.1.1. Data Logging Apparatus

The SRF Logger is a tool developed by the Centre for Sustainable Road Freight to log data whilst a vehicle is driving and to transmit data back to a central database in real time. The purpose of this logger is to easily enable high frequency data logging from real-world driving conditions, rather than a simulated environment.

The logger uses a mobile phone running the Android operating system to connect to the truck's Fleet Management System (FMS) port via a Bluetooth connection. This provides access to a wide variety of signals from the truck's own management systems. The data available through this port is dependent on vehicle manufacturer and FMS standards, but data summaries for a FMS v3 and SRF Logger equipped vehicle are available in Table 2-1. Note that some signals are sourced from the phone itself, rather than via the FMS port.

Many modern 'smart' mobile phones are fitted with GPS receivers and accelerometers. An application on the Android platform uses these sensors, as well as the data received via the Bluetooth connection to collect data. This data is then transmitted over mobile data networks to a central database in the Department of Engineering at the University of Cambridge.

Table 2-1 – SRF Logger Data

Source	Description	Units
FMS	Engine fuel use low resolution	L
FMS	Engine speed	RPM
FMS	Vehicle distance high resolution	m
FMS	Engine coolant temperature	℃
FMS	Wheel-based Vehicle Speed	km/h
FMS	Clutch pedal switch	-
FMS	Brake pedal switch	-
FMS	Cruise control switch	-
FMS	Accelerator pedal position	-
FMS	Axle weight for 1st axle	kg
FMS	Axle weight for 2nd axle	kg
FMS	Axle weight for 3rd axle	kg
FMS	Ambient air temperature	℃
FMS	Fuel rate	L/hour
FMS	Instantaneous fuel economy	km/L
FMS	Engine fuel use high resolution	L
FMS	Engine percent load	%
FMS	Engine percent torque	%
FMS	Air pressure in the service brake circuit or reservoir #1	kPa
FMS	Air pressure in the service brake circuit or reservoir #2	kPa
FMS	Diesel exhaust fluid tank level	%
FMS	Combination vehicle weight	kg
FMS	Retarder percent torque	%
Phone	Latitude	deg
Phone	Longitude	deg
Phone	Altitude	m
Phone	Bearing	deg
Phone	Speed	m/s
Phone	Longitudinal acceleration	m/s^2
Phone	Lateral acceleration	m/s^2
Phone	Vertical acceleration	m/s^2

A diagrammatic representation of the logger and sensors is included in Figure 2-1 and installation of the logger in this setup is straight-forward. The phone is fixed in available space underneath the dashboard on the passenger side of the vehicle. This is to prevent any tampering from the driver. An adapter box is connected to the FMS port, and a FMS splitter is used if the truck is already fitted with telemetry equipment. This box houses the Bluetooth dongle and is also fitted with USB sockets to keep the phone permanently on charge. When installed correctly, the equipment is invisible to the driver. The logger is calibrated for vehicle orientation by experiencing a longitudinal acceleration over five seconds after installation is completed.

The logger is triggered by the vehicle ignition signal and data is transmitted back to the central server over the 4G network in data packets. Data is organised by truck, date and trip number and stored securely.

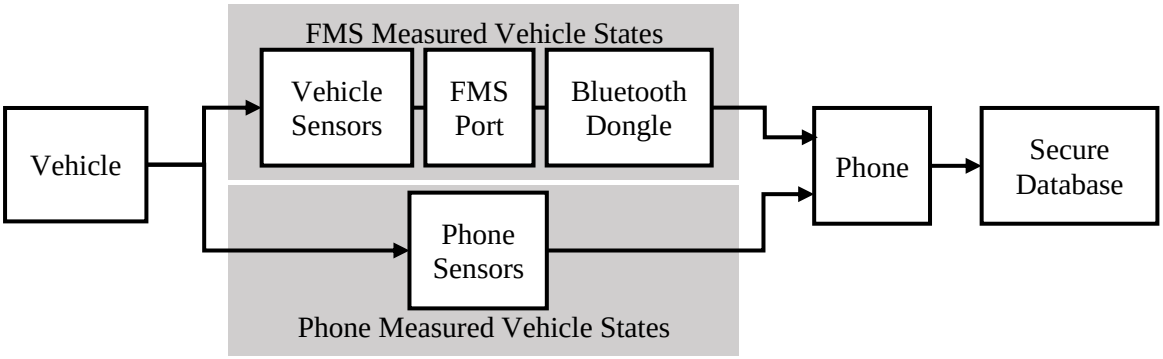

Figure 2-1 – SRF Logger block diagram

2.1.2. Vehicles

Two phases of data collection were carried out. The initial test phase was completed on a 2015 Euro 6, FMS v3, DAF CF85 - Figure 2-2. Due to the nature of this tractor's duties, the tractor did not operate with a single trailer or driver. This meant that driving style and ability were variable, along with the physical properties of the vehicle (such as drag coefficients).

In the second, and more significant phase, two trucks were fitted with the SRF Logger. These trucks, with consecutive registration numbers, were identical models and were both 2016 Euro 6, FMS v3, DAF CF85s - Figure 2-3. These trucks are operated as bulk powder transporters and were consistently paired with a single tractor and trailer and had one driver allocated to each, meaning that drag properties and driving styles would be constant for each vehicle.

Figure 2-2 2015 DAF CF85, fitted with SRF Logger during test phase

Figure 2-3 2016 DAF CF85 with bulk powder trailer

The operator's telematics system identifies Driver 1 as one of the drivers with highest fuel consumption at the depot, and Driver 2 as one of the drivers with low fuel consumption. As bulk powder trucks are weighed on leaving the depot on a weigh bridge, accurate mass information is available for these trucks on departure from their home depot. However, the exact return weight is not known. It is assumed that the return weight of the vehicle will be its empty weight i.e. no powder is returned to depot.

114 days of driving data were collected for Driver 1, and due to logger reliability issues, 44 days of driving data were collected for Driver 2. For both drivers, this is sufficient time to see the variety of routes used and to observe the driver's style and behaviour. These routes are displayed in Figure 2-4.

Where applicable, in this chapter, journey data is labelled with the following trip ID notation:

YYYY.MM.DD-t

where YYYY is the year, MM the month, DD the day, and t the trip number.

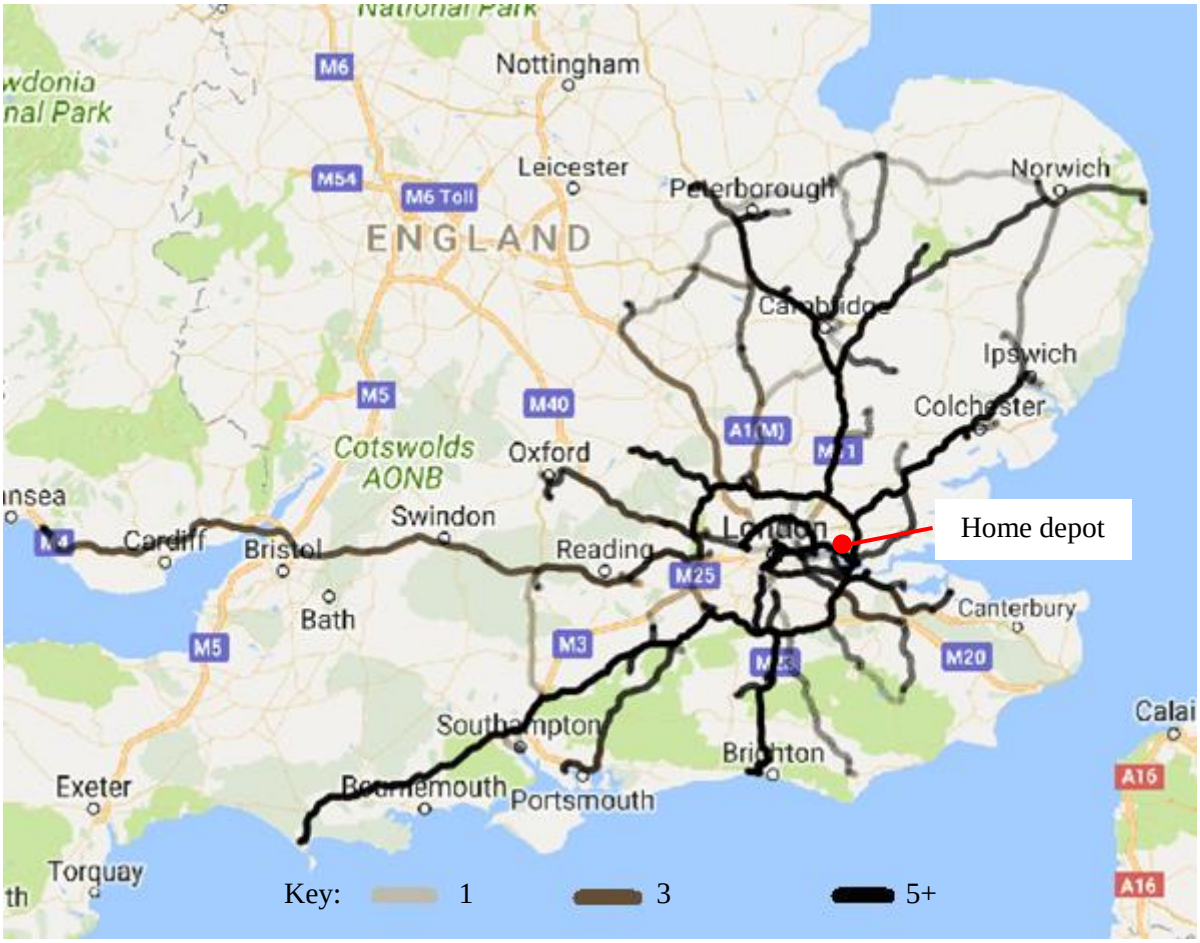

Figure 2-4 – Map illustrating routes used by both phase 2 trucks. Darkness of line indicates the frequency of use (the darker the line, the more frequent its use)

2.2. Identifying the vehicle model

In this section, the processing of data collected by the logger is explained. The data used here as an example comes from the 2015 DAF CF85.

2.2.1. Data extraction, resampling and filtering

The data is extracted from the database and resampled. The logger records the data at which ever frequency the FMS port or on-board sensors update so it is necessary to resample all data to a consistent frequency. In this case 20 Hz was selected, in order to allow suitable analysis of the driver's control of the accelerator pedal. The raw data includes many signals and these have a range of sampling frequencies from 100 Hz to 1 Hz. Where necessary, the data is interpolated to produce a consistent sample time across all data. The data is filtered to remove noise using a fourth-order Butterworth low-pass filter with cut-off frequency of 1 Hz (Figure 2-5). Due to excessive levels of noise in the acceleration data recorded from the phone's internal accelerometers, longitudinal vehicle acceleration is calculated by differentiating the filtered vehicle speed, rather than using the acceleration data from the phone.

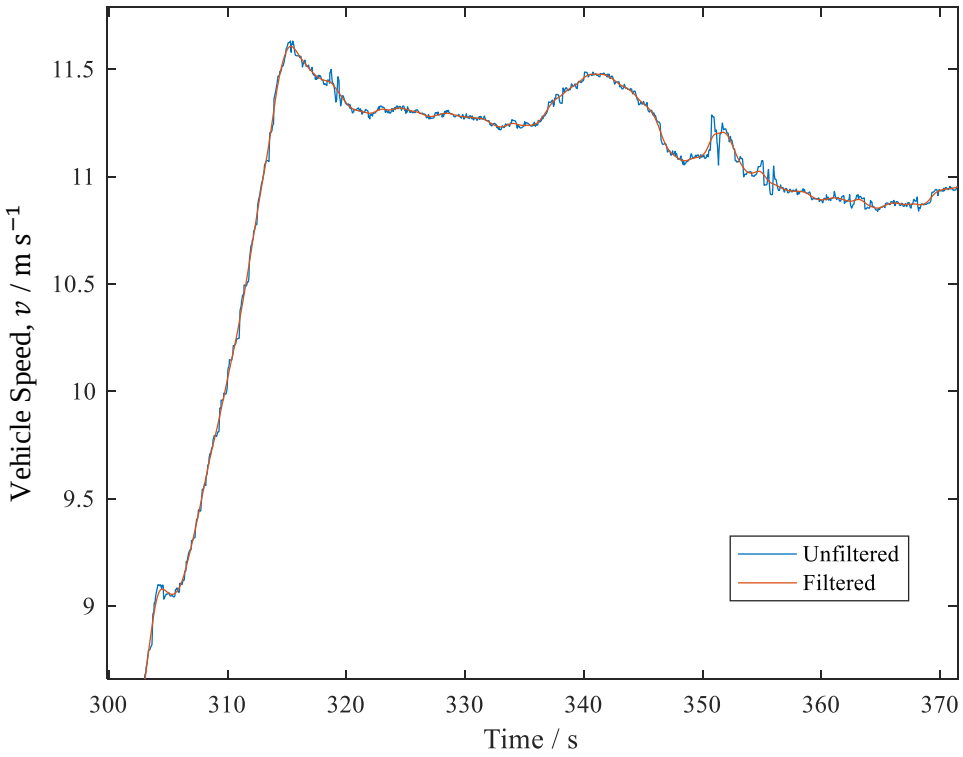

Figure 2-5 – Raw and filtered speed data for a sample drive

2.2.2. Speed ratios

At this point, the engine speed is divided by the vehicle speed to calculate speed ratios. Gear ratios from the vehicle specification are compared with the calculated values to ensure correct spacing between the vehicle speed-engine speed ratios - Figure 2-6. All data points are then allotted to the nearest ratio. Gear change events are identified in the time series and corresponding locations recorded. It can be seen from Figure 2-6 that most data points lie in good agreement with the specified gear ratios.

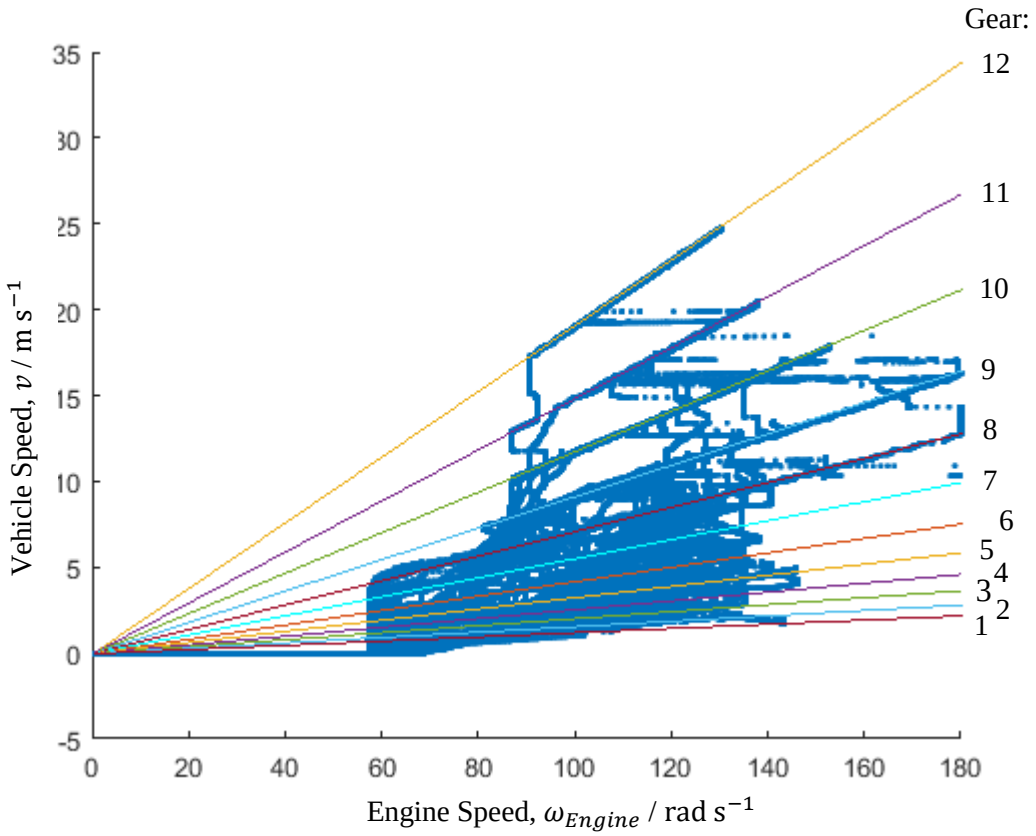

Figure 2-6 – Gear ratios are identified by comparing the easily identifiable top gear with the manufacturer's specifications. Note reverse gear is not included. Data is from the 2015 DAF CF85 Trip ID 2016.09.09-1

2.2.3. Road gradient

With a heavy truck, significant forces can be experienced when on an incline. It is therefore important to take the road gradient into account (Sentoff et al., 2015). As no direct measure of truck inclination is measured using the SRF logger, it is derived from the GPS data.

The land elevation is queried in Google Maps by inputting the GPS longitude and latitude position. The road gradient is estimated by dividing the change in elevation between two time steps by the product of vehicle speed and time step length:

$$\theta_{road} = \frac{\Delta h}{v t_s} \tag{2.1}$$

Where θ_{road} is the road gradient in radians, Δh is the change in road elevation in metres, v is the vehicle speed is m s^{-1} and t_s is the time step in seconds.

Google Maps does not offer elevation data at sufficiently high resolution in latitude and longitude to include the effects of cuttings and embankments (typical resolutions are within 20m latitude and longitude). This means that in some instances, the prevailing land gradient, and the road gradient may be quite different. The use of the GPS altitude for this calculation was ruled out due to the poor accuracy of GPS receivers in the vertical direction. A comparison between the Google Maps data and the data collected by the GPS is illustrated over a drive (trip ID 2016.09.09-1 for this vehicle) in Figure 2-7.

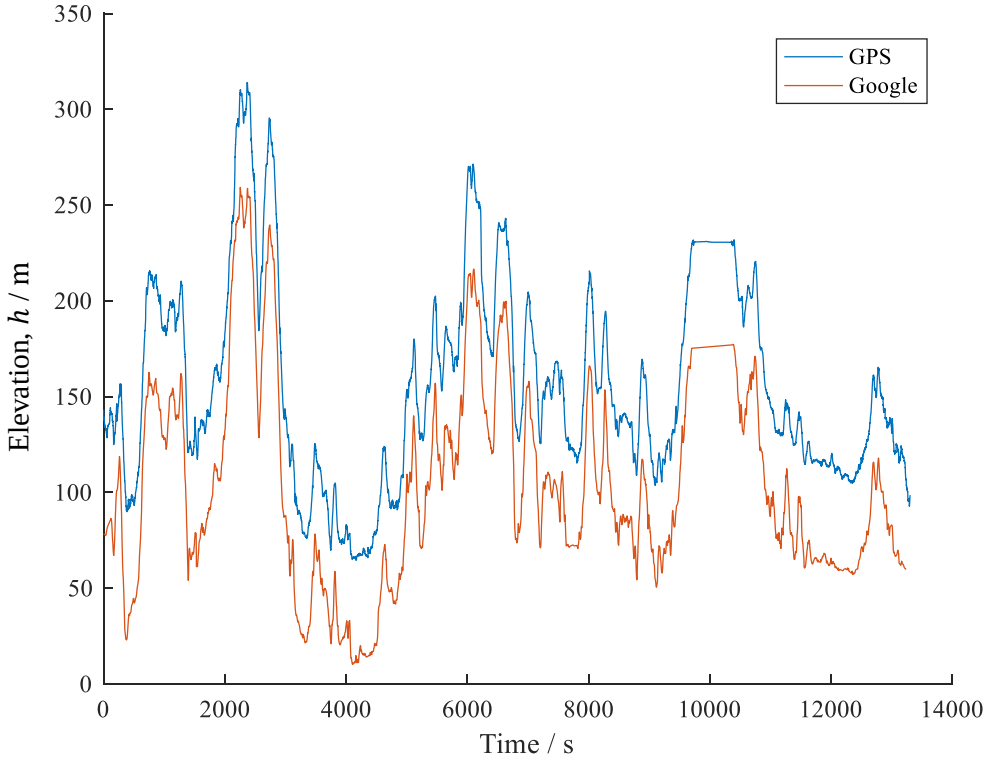

Figure 2-7 – Comparing GPS elevation to Google Maps elevation. Trip ID 2016.09.09-1

It can be seen that the elevation data from Google is consistently lower than the GPS data. This could be due to the difference in definitions of reference elevations between the GPS system and Google, but the offset is not consistent.

There are many mapping data providers offering similar datasets. Google was selected as it is openly available and widely accessible. Higher resolution datasets are available from elsewhere, but tend to be offered at high cost, and in less practical formats.

2.2.4. Drag properties

Estimates of the drag properties, the drag area, $C_d A$ and the coefficient of rolling resistance, C_{rr}, of the vehicle are taken from the output of the SRF Mapper tool. These estimates are consisted with Hunt et al. (2011). This uses a general approach to estimate the engine maps for a vehicle using only the data available from the SRF Logger (Bishop et al. 2016). The other outputs of the Mapper tool are not utilised outright here as better estimates of the vehicle characteristics are achievable in this case. This is because there is more information available on the specific trucks used here than the Mapper tool would utilise, as it takes a generalised approach – see 2.2.5.

2.2.5. Mass estimation

The mass is a very significant factor in the forces applied on the truck. For this reason, it is beneficial to have knowledge of the mass at all times. Although mass data is available for the 2016 DAF trucks in Phase 2 on departure, the mass may change during the drive after a drop has been made.

In this section, an attempt is made at estimating the vehicle mass using the data recorded using the SRF Logger. Ultimately, this approach was unsuccessful in estimating the vehicle mass, but it is recorded here for completeness, and also to outline how changes in mass are detected. Weighbridge masses are henceforth used for calculations, and any data from after a detected mass change is discarded.

In some specifications of the FMS standard, vehicle mass is reported. It is commonly derived from the air suspension pressures. However, examination of this data shows that the reported value can vary significantly over a short time. The SRF Mapper analyses this data and attempts to identify the vehicle mass by searching for data where the vehicle is stationary. When stationary, the effects of vehicle motion or road gradients are not expected to be seen, or they are at least minimised.

With more information about the specific trucks used, a further iteration on vehicle mass is made. The maximum power and maximum torque output of the engine, along with the corresponding speed ranges were sourced from the manufacturer - Table 2-2. It is assumed that the engine will operate at this point when the engine speed is within the specified range, accelerator

pedal position has been at greater than 95% for at least two seconds, the brake and cruise control are not being used, and no gear change event occurs within two seconds. Power losses through the transmission are approximated to 8% of the engine power (Hunt et al. 2010).

Table 2-2 – Engine Specification

Parameter Name	Symbol	Numerical Value	Units
Maximum Torque	$T_{EngineMax}$	2100	Nm
-low speed cut off	ω_{T1}	104	rad/s
-high speed cut off	ω_{T2}	151	rad/s
Maximum Power	$P_{EngineMax}$	320	kW
-low speed cut off	ω_{P1}	151	rad/s
-high speed cut off	ω_{P2}	178	rad/s

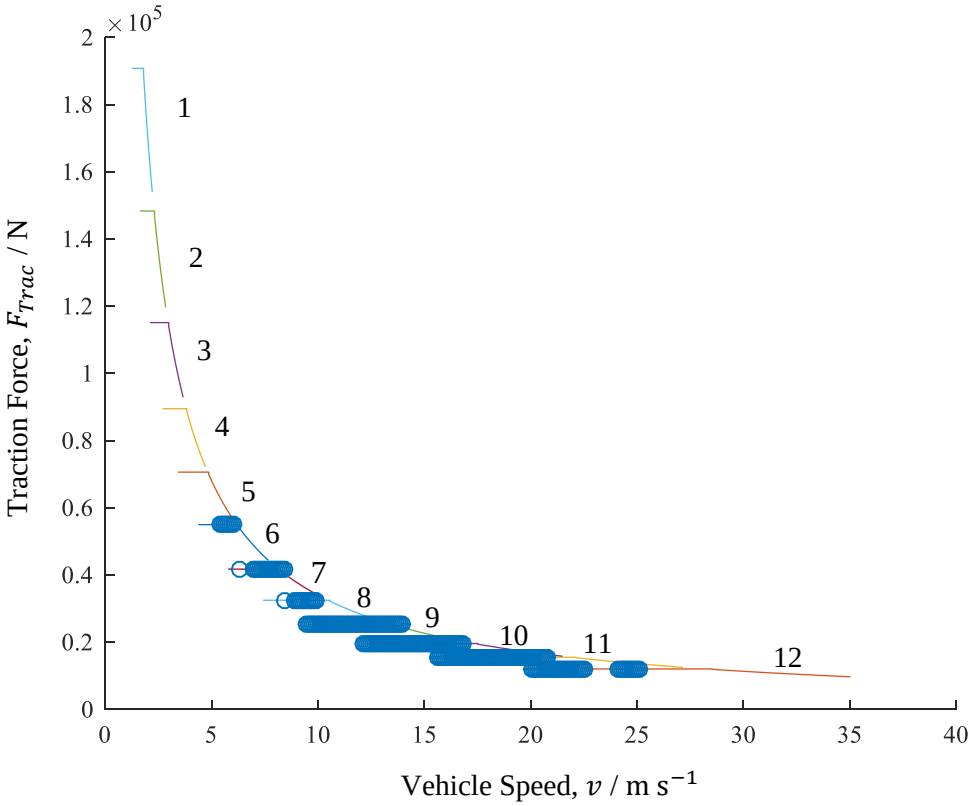

Figure 2-8 –Estimated maximum traction forces were limited by the torque regime more than the power regime on this drive. 2015 DAF trip ID 2016.09.09-1

Figure 2-8 illustrates how observed data from a trip in the 2015 DAF (trip ID 2016.09.09-1 for this vehicle) is fitted to the maximum engine curves. For each gear a maximum traction force

(horizontal line) and maximum power (curved contour) are plotted. Where suitable data exists, as determined by the criteria above, the speed and gear are assessed and the points are plotted on the corresponding gear contour. It can be seen that the vast majority of this data, from this one drive, fits more closely to the maximum torque specification of the engine, rather than its maximum power. This is a characteristic of the driver/vehicle combination.

Using a force balance on the vehicle (Figure 2-10) at every time step meeting the above criteria, an estimate of the mass is made using the equation:

$$M_{est} = \frac{\eta P_{eng} - \frac{1}{2}\rho C_d A(v + v_{wind})^2}{a + g(\alpha + C_{rr})} \quad (2.2)$$

Where η is the power train efficiency, P_{eng} is the engine power rating corresponding to the engine speed, ρ is the density of air, $C_d A$ is the drag area, v is the vehicle speed, v_{wind} is the average effective wind speed, a is the vehicle acceleration, α is the road gradient and C_{rr} is the coefficient of rolling resistance. As the drag area and coefficient of rolling resistance are constant in this calculation, and these properties are estimated in the Mapper tool through use of the mass reported over the FMS port, the calculation in (2.2) is in effect a second iteration of the vehicle mass.

Quantifying the effective wind speed is difficult, as no accurate data exists for the winds experienced by the truck. This is very environmentally dependent, and depends on other road users and driving environment. An approximation of the wind speed is made by manually searching weather logs for a prevailing wind speed and direction on the day of travel. The average direction of travel of the truck is considered and the longitudinal component of the wind relative to the truck direction is incorporated.

A log-normal curve is fitted to the distribution of mass estimates to take into account its positive skew. Figure 2-9 illustrates the mass estimate distribution for a sample journey.

The modal mass estimate is 29.7 tonnes, the median mass is 40.8 tonnes and the mean mass is 46.9 tonnes. In comparing these masses with the recorded masses from weighbridges, it soon becomes apparent that the modal and mean mass estimates are unreliable, with errors up to 25%, or 10 tonnes (the example in Figure 2-9 has a weighbridge gross mass of 39.1 tonnes), whereas the median mass estimate has an error of 1.5 tonnes, or 4%, which is quite reasonable.

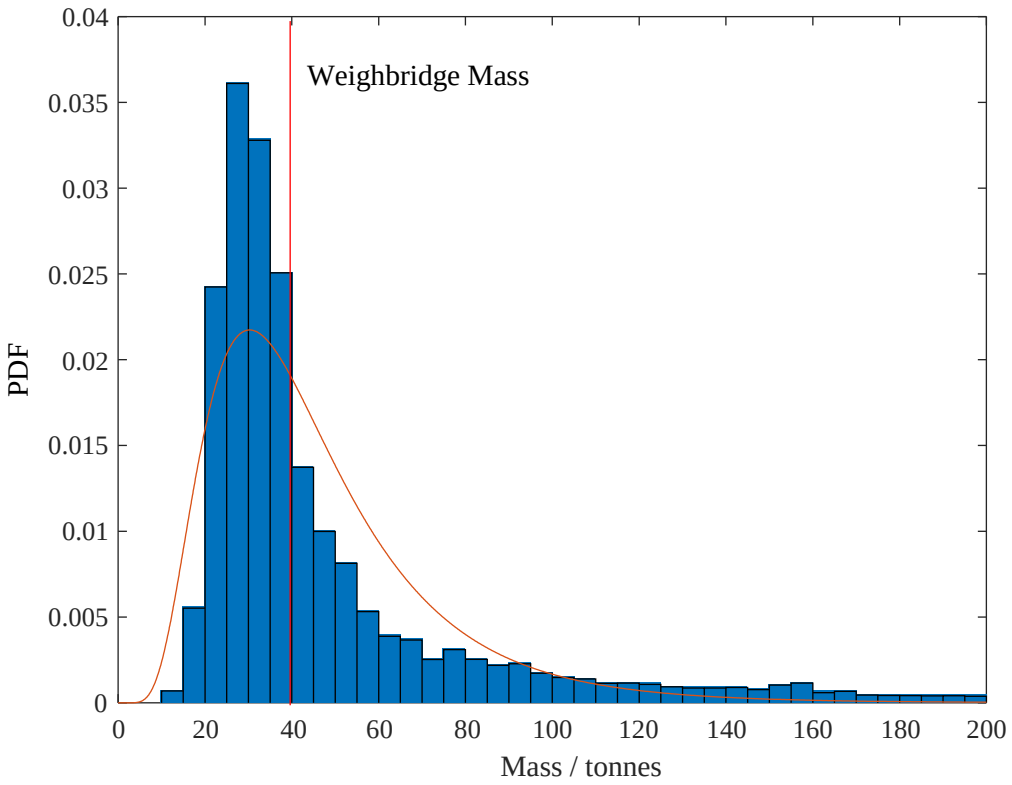

Figure 2-9 – Estimated vehicle mass distribution

The mass reported over the FMS port is useful in detecting a change in mass of the vehicle. A moving average reported mass is monitored and should a change of over 10% of the vehicle mass be detected a possible mass change is flagged. As has been shown, the mass estimation strategy explained here has not proved successful, so when a mass change is detected, the remaining journey data is discarded from analysis.

2.2.6. Traction forces

With an estimate of the mass and the other information gathered, it is possible to estimate the traction force at every time step. Figure 2-10 shows the forces on the truck. A force balance is used to calculate the traction forces.

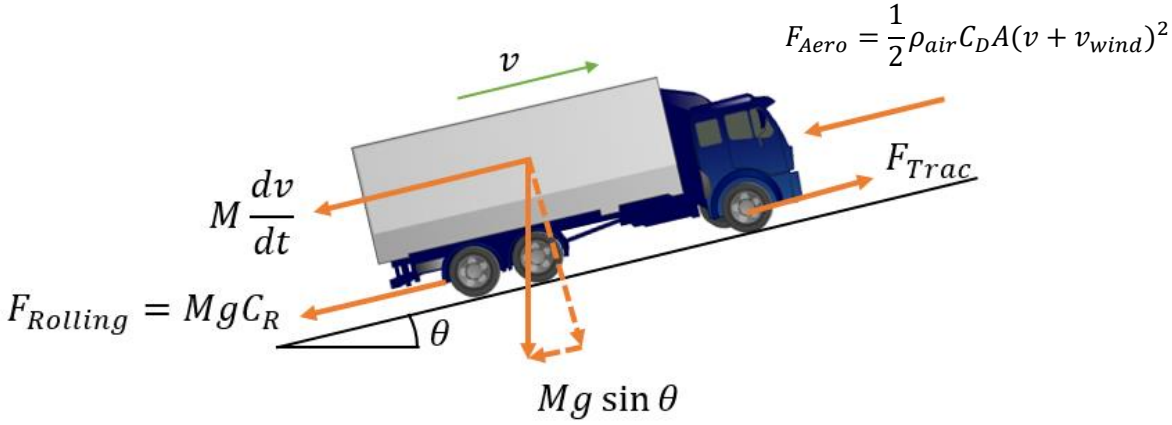

Figure 2-10 – Force diagram for truck driving up an incline θ

Examining data from a typical real road drive, the contributions to the total traction force of the four constituents (inertia, road gradient, air resistance and rolling resistance) are estimated and displayed in Figure 2-11. It can be seen that the major contributor in this case is the road gradient. This factor is very uncertain due to the limited resolution from the ground elevation data. The steep gradients here may be because the ground points that are interpolated between may be very different if they fall on the road, or immediately off the road, up or down a steep bank. It is unlikely that the real vehicle experienced such great gradients. The second most significant factor in the traction force is the vehicle mass. For a 40 tonne vehicle, these two terms dominate, and the aerodynamic drag and rolling resistance are minimal in comparison.

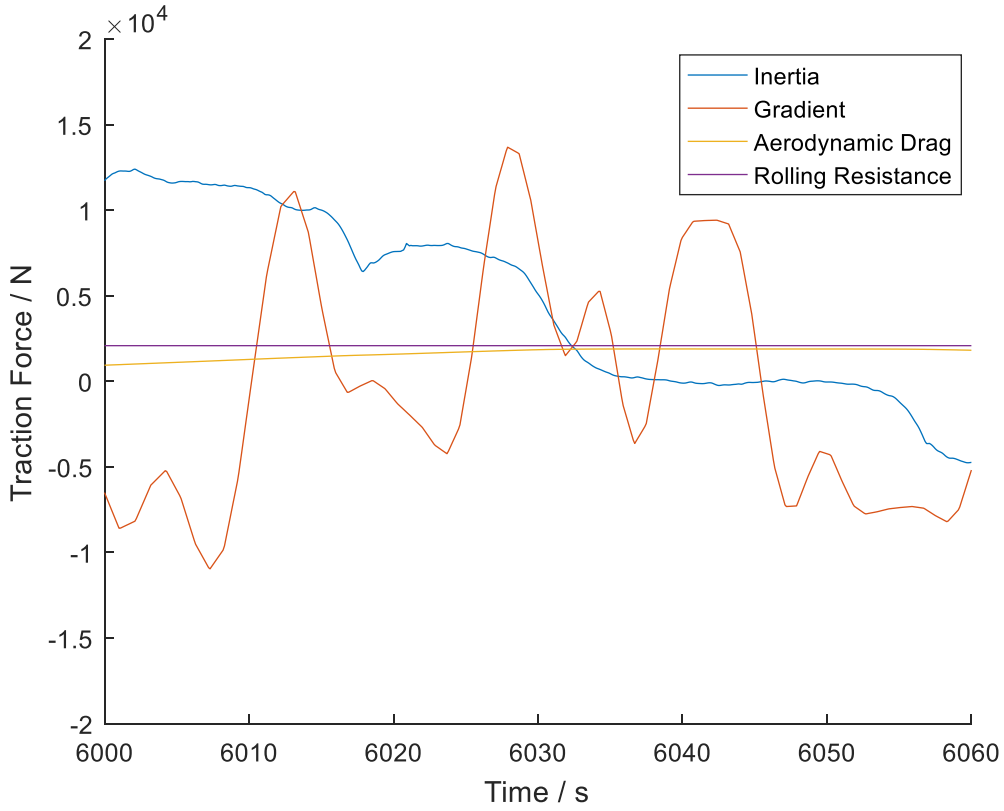

Figure 2-11 – Estimated traction force components against time. 2015 DAF trip ID 2016.09.09-1

In order to simulate the driver's control of the vehicle, a relationship between the traction force produced, vehicle speed and the accelerator pedal position is sought. Known as a traction force map, the relationship needs to encompass the behaviour of both the engine and the truck's automatic gearbox – a complicated system with its own sophisticated control algorithms. For this reason, the more common form of engine torque maps will not be used, as a gear change strategy would then need to be modelled. The traction force map removes the need for accurate gear change modelling.

Figure 2-12 illustrates in a scatter form how the estimated traction forces relate to the driver's pedal position and vehicle speed. Data from both 2016 vehicles is combined as they are identical vehicles and the different driving styles of each driver populate different areas of the map. As the relationship that incorporates the pedal position is sought, it has been necessary to eliminate some data. The included data must meet the following criteria:

- Cruise control must not be engaged
- The brake pedal must not be depressed
- There must be no gear change within two seconds

The three events listed are cases where the accelerator pedal is not directly controlling the vehicle or the vehicle is not in steady state.

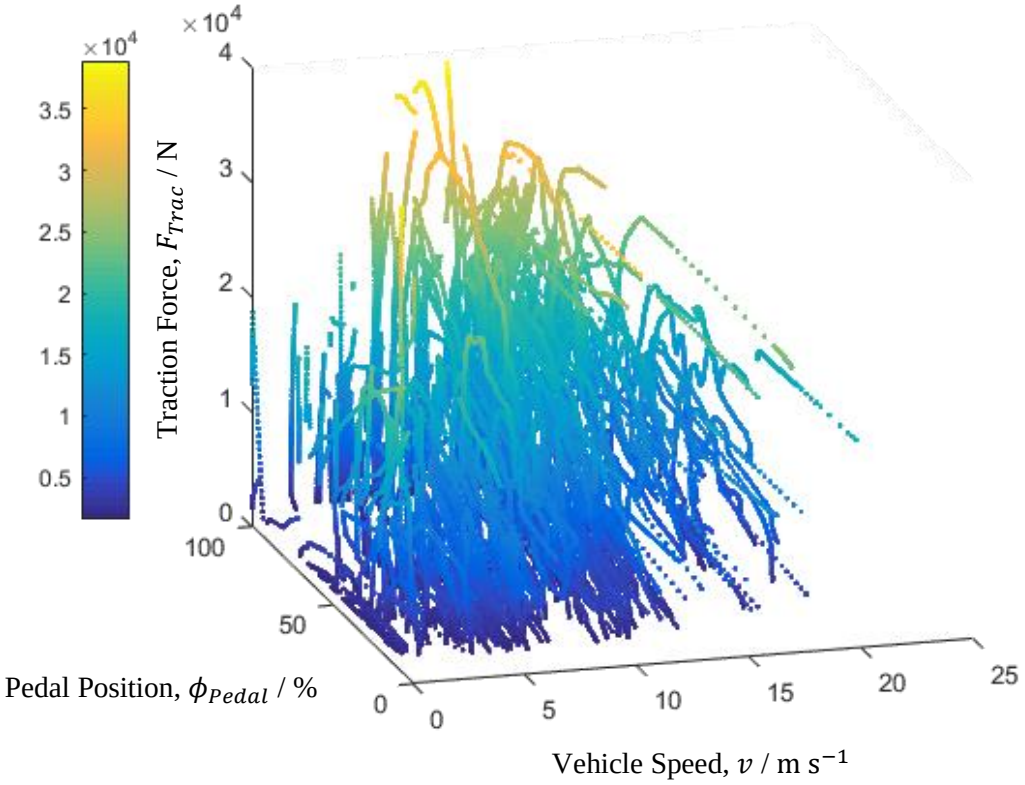

Figure 2-12 – 3D scatter plotter of estimated traction forces against pedal position and vehicle speed

The suitable data is broken down into bins and the mean and standard deviation in each bin is calculated. A two-two order polynomial, of the form below, is sought:

$$F_{TracMap} = p_{00} + p_{10}v + p_{01}\phi + p_{20}v^2 + p_{11}v\phi + p_{02}\phi^2 \quad (2.3)$$

This order of polynomial was set by a trade-off of modelling complexity against accuracy. A least squares method is then applied to the mean bin values, located at the bin centres, to identify the polynomial relationship between pedal position, vehicle speed and traction force.

The vehicle is slower to respond to changes in pedal position than the pedal position's response to changes in pedal force. This could be because it takes time for the engine controller to interpret demands and to act accordingly. It is therefore necessary to model the transient response of the vehicle to pedal demands. To model this, a first order lag filter (selected to minimise the number of parameters to fit) is applied to the pedal position before feeding it into the traction force map.

The pedal lag time constant is determined using the Matlab fmincon function. The function varies the lag frequency, and assesses the corresponding r^2 value of the second order polynomial traction force map. The calculated lag frequency is 2.84 Hz and the r^2 is 0.667, meaning a good fit to the data. The corresponding polynomial coefficients are in Table 2-3 and the surface is displayed in Figure 2-13, alongside the mean bin values (circles) and the published maximum engine power (displayed in red on the maximum throttle position plane). Assuming a perfectly efficient engine and power train (for illustrative purposes only), the engine power is equal to the product of traction force and vehicle speed, which results in the contour depicted in the figure.

Table 2-3 – Traction force map polynomial coefficients for the dynamic surface

Coefficient	Value	Units
p_{00}	1090	kg m s^{-2}
p_{10}	703	kg s^{-1}
p_{01}	106	kg m s^{-2}
p_{20}	-23.2	kg m^{-1}
p_{11}	-6.6	kg s^{-1}
p_{02}	1.1	kg m s^{-2}

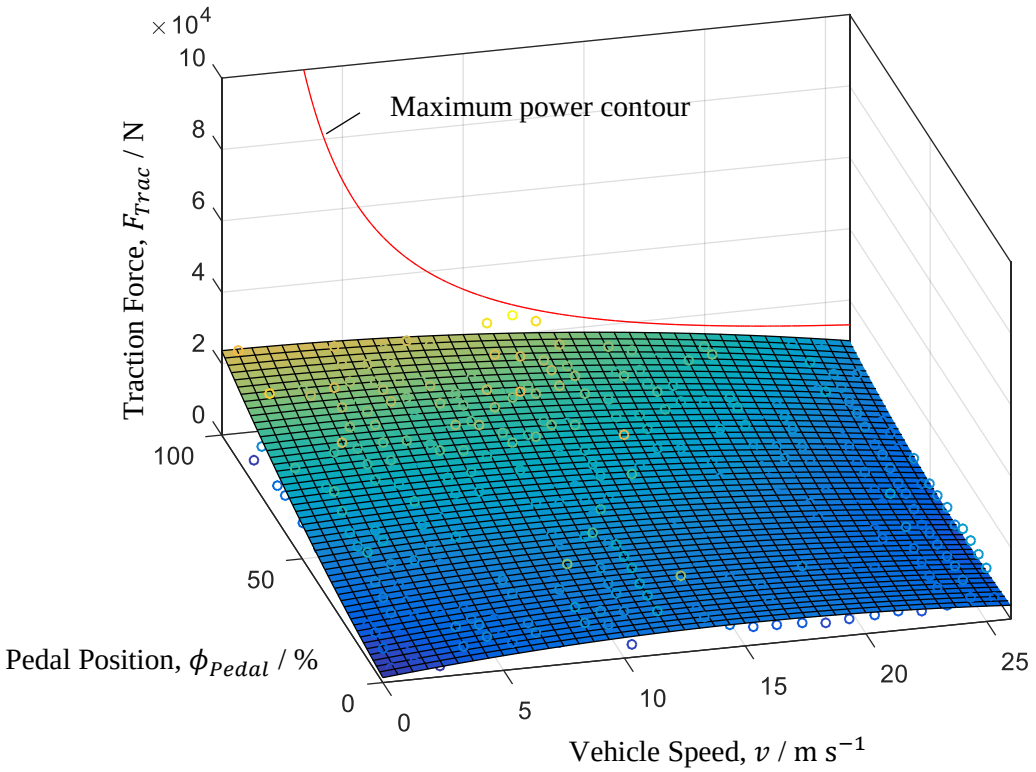

Figure 2-13 - Least squares polynomial fit for the dynamic traction force map

It can be seen that the polynomial surface falls below the maximum power contour. Powertrain inefficiencies are taken into account in this contour, but as different gears are used, the engine will not always be operating at its maximum traction force, even when pedal position is maximum. This means that the mean traction force is reduced, explaining the difference noted. At low speeds, the difference become very significant. This is thought to be because the vehicle is unlikely to be able to apply maximum power at low speeds without slipping wheels or clutch. The result is effectively reducing the maximum power.

By examining the bin standard deviations (Figure 2-14) it is noted that the larger inconsistencies occur at high pedal positions at mid-range speeds. More generally, the standard deviations are approximately 20% of the mean value. This is a significant uncertainty and reflects the uncertainty in the estimation of traction forces, primarily due to road gradient errors, and truck mass errors.

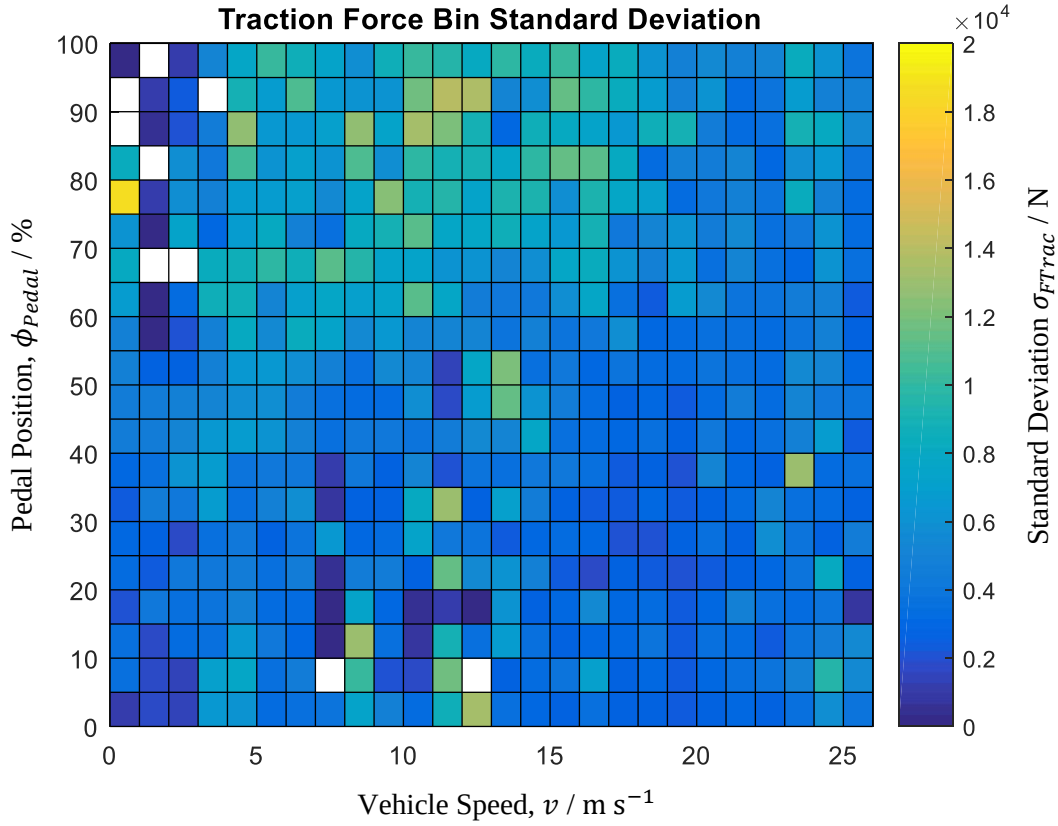

Figure 2-14 – The bin standard deviations of the steady state traction force map

In some modelling situations, it is beneficial to have a very simple traction force map to minimise the complexity of the vehicle model, and hence minimise computing expense, as is the case in Chapters 3 and 4. Here one such model is proposed, and then assessed, using the traction force estimates already discussed.

The polynomial traction force maps already produced demonstrate a definite trend of increasing traction force with pedal position. The trend is less strong in the vehicle speed direction.

The proposed simple model exploits this by assuming the traction force is independent of vehicle speed, and that the traction force is proportional to the pedal position. The purpose of this model is simplicity, in computing terms, and in generation, so the proportionality constant is set to mean that the maximum engine torque (as stated in vehicle documentation) is reached at 100% pedal position, assuming the vehicle is in a mid-range gear. This model can be expressed as:

$$F_{Trac} = \frac{\phi_{Pedal}}{\phi_{max}} T_{EngineMax} \quad (2.4)$$

A plot of the simple traction force map for the trucks fitted with the logger, in middle gear, is included in Figure 2-15. It is noted that this model predicts higher traction forces than the second order polynomial already examined, but with $r^2 = 0.38$, does correlate with the bin means reasonably well, and therefore is fit for purpose.

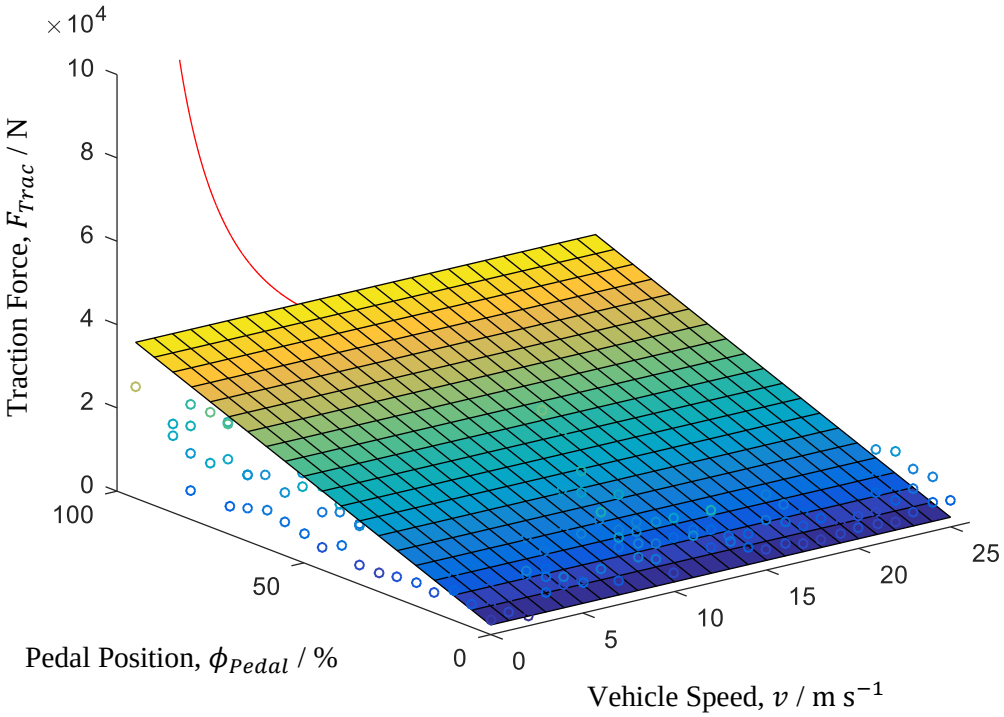

Figure 2-15 – Simple traction force map

2.2.7. Fuel consumption

The vehicle's fuel use is reported in two ways over the FMS port. Firstly, a cumulative fuel use is recorded, and secondly a high precision fuel rate is also recorded - Figure 2-16 demonstrates these for a segment of a real drive. The limited precision of the cumulative fuel used is clearly visible when compared to the high precision fuel rate. The cumulative fuel use appears to have a discretisation step of 0.5 L/min.

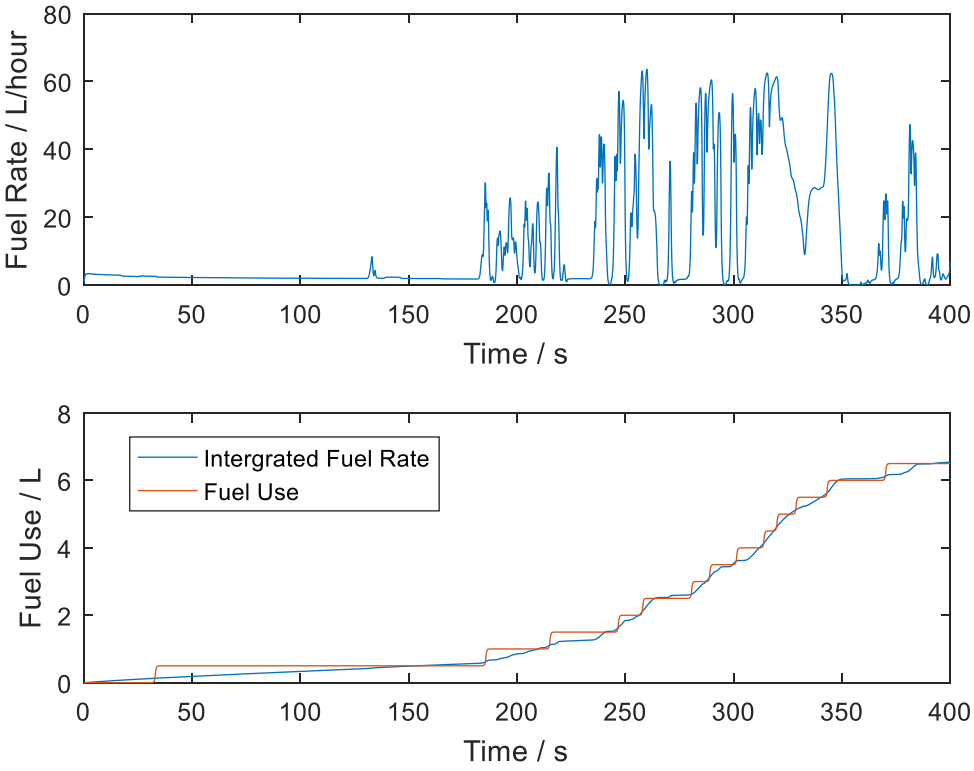

Figure 2-16 – Recorded fuel use and fuel rate

For predicting fuel use of the model vehicle, a relationship between fuel rate, vehicle speed and pedal position is sought. A steady state map is fitted to data deemed to be steady state i.e. low rate of change of pedal position and away from any gear changes. Cruise control and braking also had to be disengaged. These are the same requirements as for the traction force plot earlier. As the fuel rate is affected by vehicle mass, the map uses data from multiple journeys from both trucks, but with a very narrow gross mass band of 41.8T to 42.0T, as measured at the weighbridge. As described in 2.2.5, if a vehicle mass change occurs during the journey, data after the change is discarded. Figure 2-17 illustrates the suitable data for these requirements.

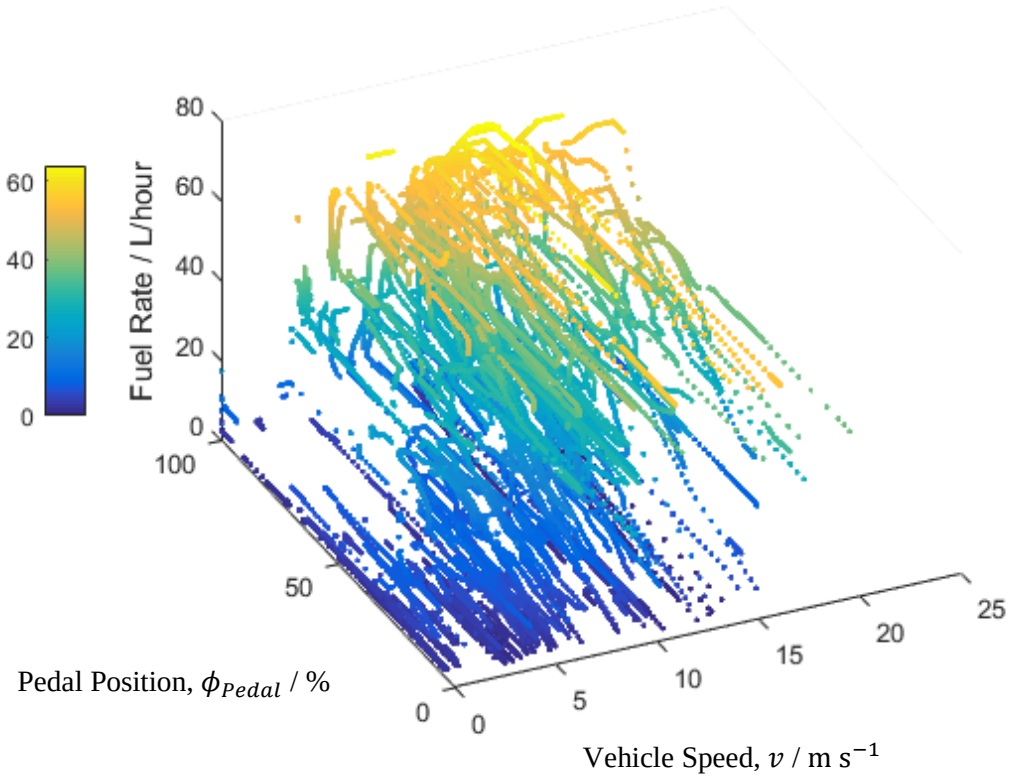

Figure 2-17 - 3D scatter plotter of fuel rate against pedal position and vehicle speed

The same method of polynomial fitting is used as in the traction force case. A first order lag was applied again, but, in this case, the polynomial is required to take on a higher order in order to better reflect the shape of the data:

$$\lambda_{Map} = p_{00} + p_{10}v + p_{01}\phi + p_{20}v^2 + p_{11}v\phi + p_{02}\phi^2 + p_{30}v^3 + p_{21}v^2\phi + p_{12}v\phi^2 \quad (2.5)$$

A dynamic fuel rate model is also sought, so the same approach is taken with the fuel rate as was taken with the traction force. The calculate fuel rate lag frequency is 5.0 Hz, with $r^2 = 0.89$, and the corresponding coefficients are in Table 2-4, and the corresponding surface for the steady state fuel map is illustrated in Figure 2-18. The value of r^2 is 0.87, indicating good correlation between surface and data. Notice here that the general trend is for increasing fuel rate with speed and pedal position, but at high speeds, the fuel rate surface gradients get very steep.

Table 2-4 – Steady state fuel map polynomial coefficients

Coefficient	Value	Units
p_{00}	-1.15	L/hour
p_{10}	1.97	L/hour $m^{-1}s$
p_{01}	0.052	L/hour
p_{20}	-0.175	L/hour $\mathrm{m}^{-2}\mathrm{s}^{2}$
p_{11}	0.015	L/hour $\mathrm{m}^{-1}\mathrm{s}$
p_{02}	0.0013	L/hour
p_{30}	0.0049	L/hour $\mathrm{m}^{-3}\mathrm{s}^{3}$
p_{21}	-0.00095	L/hour $\mathrm{m}^{-2}\mathrm{s}^{2}$
p_{12}	0.00017	L/hour $\mathrm{m}^{-1}\mathrm{s}$

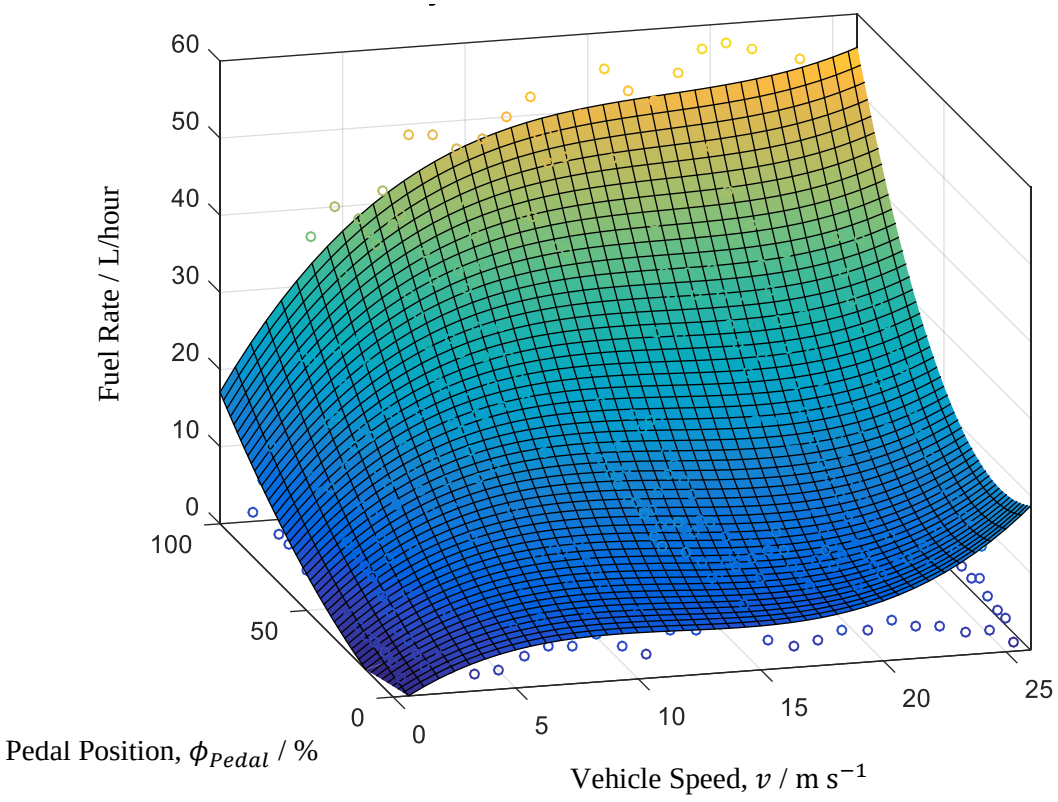

Figure 2-18 - Least squares polynomial fit for the steady state fuel map

The standard deviation of the individual bins in the fuel rate map are depicted in Figure 2-19. With the exception of some extreme variation in the top right corner, the standard deviations are generally low, between 10% and 20% of the mean values. There is much better agreement in these points than in the traction force estimates due to the fact that this data comes directly from measurements, whereas the traction force maps are based on estimates of the traction force. The estimation process introduces a degree of uncertainty.

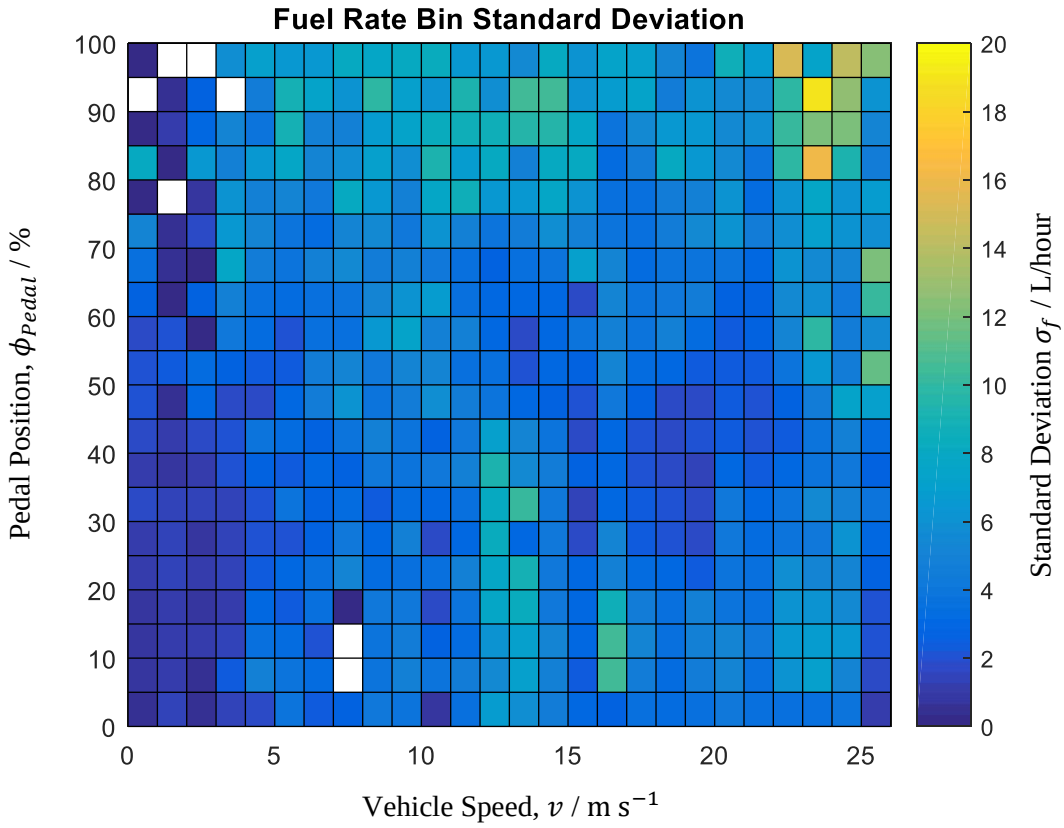

Figure 2-19 – The bin standard deviations of the steady state fuel map

2.2.8. Modelling Fuel Consumption

In 2.2.6 and 2.2.7, maps were generated to relate the traction forces and fuel rates to the pedal position, ϕ_{Pedal}, and the vehicle speed, v. In this section, these maps applied to recorded pedal position and vehicle speed data, and predicted fuel use is compared to the measured fuel used.

Firstly, the pedal map is applied to pedal position and vehicle speed data without Cruise Control. For illustration, a 500s extract Driver 1 trip 2017.03.21-1 is seen in Figure 2-20. It can be seen that there is good agreement between the two, except for a small region between 1280s and 1320s where the measured data exceeds the estimated data. In this region of the map (corresponding to vehicle speeds between 15 ms^{-1} and 20 ms^{-1}, and pedal positions over 80%), the gradient of the map does not reach the very high mean fuel rates observed in Figure 2-18. The map does,

however, reach the highest fuel rates when the pedal position is greater than 90% and vehicle speed is above 22 ms^{-1}. This explains why the simulated fuel rate is higher between 1400 s and 1450 s than between 1280 s and 1320 s. By integrating the fuel use over the segment, it is possible to quantify the performance of the fuel rate map – the estimated total fuel use is 2.19 L and the total measured fuel use is 2.63 L, an error of 16%. This is considered a reasonable performance for the fuel rate map. Similar results are achieved for Driver 2's vehicle (figures not included). The performance may be improved through the use of a higher order map, or through a surface of different formulation. The suggested surfaces may better reflect the nature of the relationship over a broader range of pedal positions and vehicle speeds.

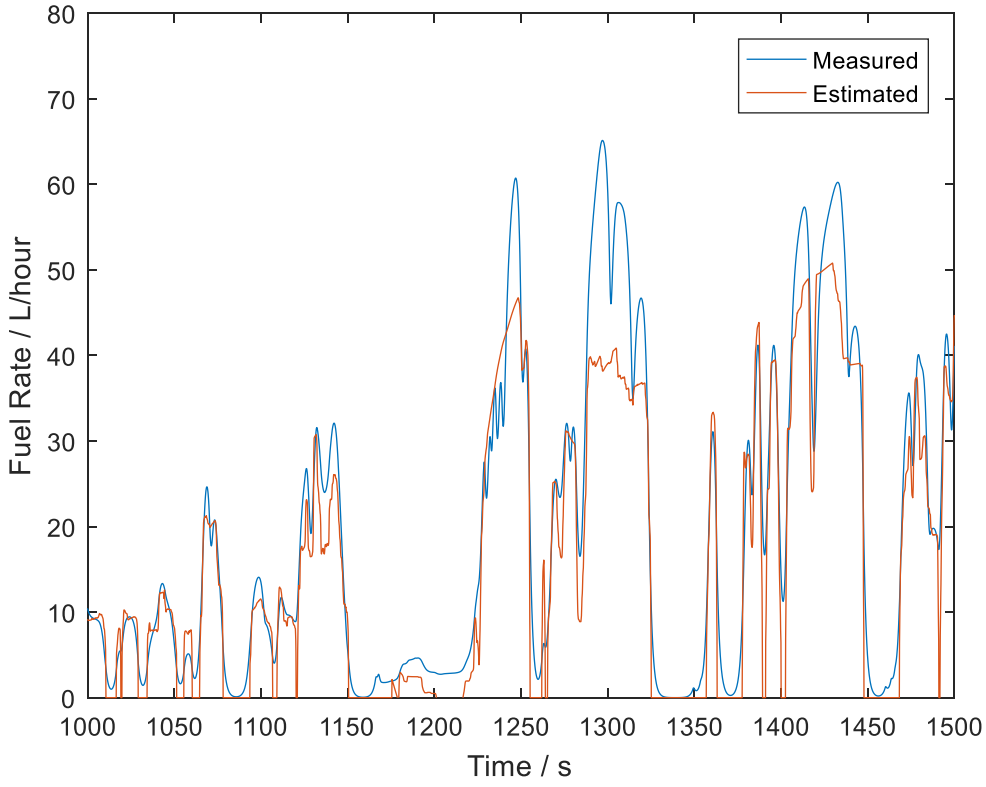

Figure 2-20 – Measured and estimated fuel rate

However, the use of Cruise Control complicates the relationship, as the accelerator pedal is no longer linked to the control of the vehicle. To overcome this, the 'effective accelerator pedal position' is calculated. Using vehicle mass, speed, acceleration and road gradient data for when cruise control is engaged, an estimate of the traction force at any instant is made - Figure 2-10. The vehicle speed and estimated traction force are input to the traction force map (Figure 2-13) to back-calculate the effective accelerator position. As the traction force map is quadratic in pedal position (2.3), there are two possible solutions for the pedal position. The solution of interest is the one solution in the range 0-100% pedal position.

This estimated pedal position is then fed into the fuel rate map (2.4) and the estimated fuel rate is calculated - Figure 2-21. Where cruise control is not engaged (white background in the figure), the model is able to utilise the measured pedal position data. Where cruise control is engaged (grey background in the figure), the estimated pedal position is used. It can be clearly seen that the estimated cruise control fuel rate does not fit the measured data so well as the non-cruise control case. This is because in this combined traction force map and fuel rate map approach, the uncertainties combine to produce very significant uncertainties. In the cruise control region of Figure 2-21, the total measured fuel used is 0.87L, and the estimated fuel used is 0.43L, an error of 51%. In comparison, in the non-cruise region, the measured fuel use is 0.31L and the estimated fuel use is 0.33L, an overestimate of 4%. Simulations over entire journeys produce similar results with very large errors in fuel rate.

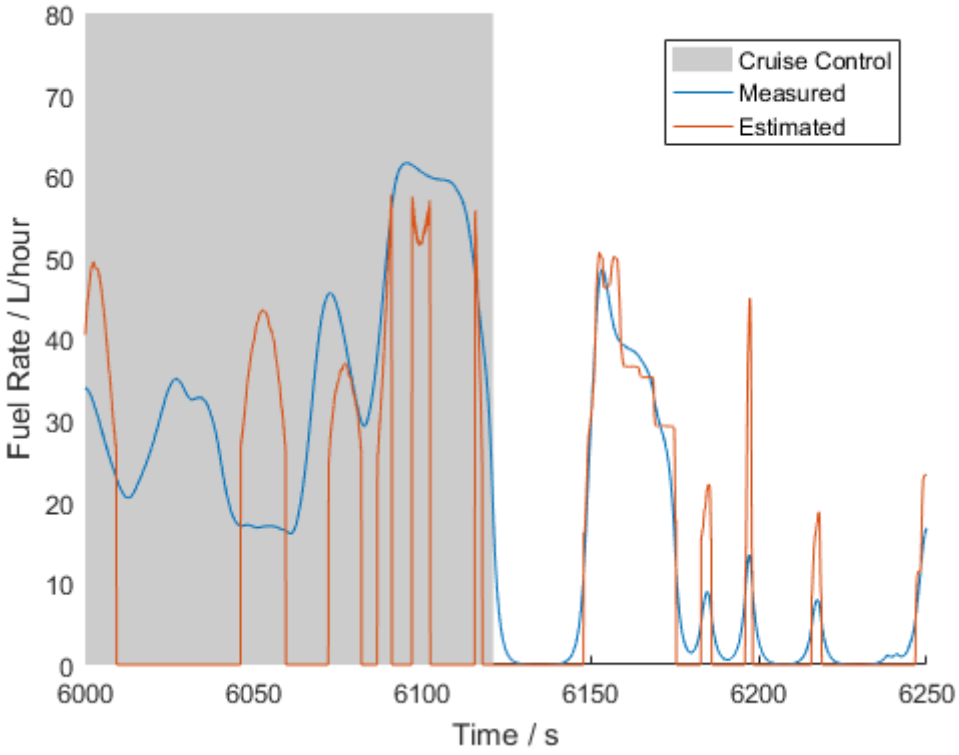

Figure 2-21 - Measured and estimated fuel rate during Cruise Control

The very high errors of the cruise control fuel estimation mean that this technique is not appropriate for use in examining the effects of changing a driver's behaviours as a whole. The results for fuel consumption with the accelerator pedal are better, and some analysis could be completed with these maps. Further work is required to improve the performance of the fuel rate and tractions force maps.

2.3. Identification of driver behaviour

In this section, the data collected from the logger is analysed with the intent of better understanding the behaviour of the two drivers of the phase 2 vehicles, Driver 1 and Driver 2. To start, the journey data is broken down into a series of driving states in 2.3.1 before examining the roundabout manoeuvre in 2.3.2 and 2.3.3.

2.3.1. Classification of driving activities

The data is broken down at this point into two groups of driving states as defined in Table 2-5 and Table 2-6. Group A states are purely based on the speed of the vehicle whereas Group B states reflect the driver's control input to the vehicle.

Table 2-5 – Group A: Speed-Acceleration States

State	Definition
Accelerating	$a > 0.2\ \mathrm{ms}^{-2}$
Constant Speed	$\lvert a \rvert < 0.2\ \mathrm{ms}^{-2}$
Decelerating	$a < -0.2\ \mathrm{ms}^{-2}$
Stationary	$\lvert v \rvert < 0.5\ \mathrm{ms}^{-1}$

Table 2-6 – Group B: Driver Control States

State	Definition
Braking	Brake pedal engaged
Cruise Control	Cruise Control engaged; no braking
Accelerator	$\phi_{Pedal} > 5\%$; No braking or Cruise Control
Coasting	$\phi_{Pedal} < 5\%$; No braking or Cruise Control

Firstly, the data is simplified to reduce the number of data points being dealt with. The Douglas-Peucker algorithm (Douglas & Peucker, 1973) is used to break down the vehicle speed data into a series of variable length segments.

The algorithm recursively breaks the data down into smaller segments by splitting the data at the point furthest from the line segment formed by connecting the start and end points of the existing segment. This process is repeated on all newly formed sub-segments until all data lies

within a threshold distance of the new line segments. The results of this applying this algorithm to the vehicle speed data are displayed in Figure 2-22. The acceleration over these new segment is calculated and each segment is classified as either acceleration, constant speed or deceleration.

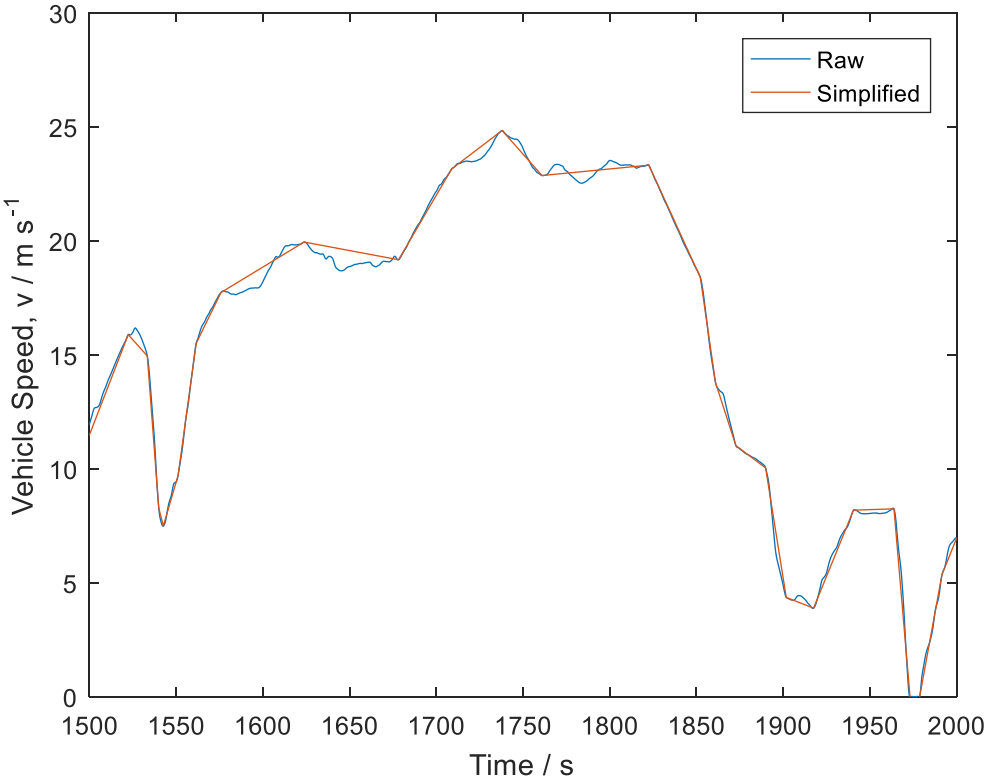

Figure 2-22 – Results of the applying the Douglas-Peucker algorithm to the vehicle speed data

A Boolean coasting signal is defined by identifying when pedal position drops below 5%. This Boolean function, along with the cruise control and braking Boolean signals are overlaid. The journey segments are then classified into the two groups of states above.

2.3.2. Junction detection

With a bank of data built up over time, it is possible to analyse the behaviour of different drivers. As it is the drivers' interaction with the accelerator pedal that is of interest, it is important to identify when the drivers commonly use it.

The situations in which the driver uses the accelerator pedal can be broken down into two groups: lateral manoeuvres and longitudinal manoeuvres. In lateral manoeuvres, such as crossing a roundabout and other junctions, the driver is strongly guided by the fixed environment and is also affected by the dynamic environment (other road users). The driver's speed choice is therefore influenced by predictable events such as bends and roundabouts. In longitudinal manoeuvres, such as passing traffic lights, or responding to speed changes in the traffic ahead, the driver is more

significantly affected by other road users, and this makes the scenarios less repeatable. The driver's speed choice is hence influenced by unpredictable evets. It is also difficult to distinguish between traffic lights and other road user dependent scenarios with the data available. The lateral manoeuvres are therefore considered more predictable and repeatable, so are explored in more depth here.

Lateral manoeuvres are detectable in two ways: by searching for high lateral accelerations, or searching for low radii of curvature of the vehicle path. The first is simply a case of searching the recorded lateral accelerations from the accelerometers in the phone for any peak accelerations beyond a threshold. The second approach uses the relation $a = v^2 / r$ rearranged to give the curvature. With these calculations completed, the threshold for each method is manually adjusted to reach a compromise between maximising the number of true detections and minimising the number of false detections (cases where the algorithm outputs a location where there is no roundabout or junction in reality). The results of both methods are combined to provide the greatest chance of identifying junctions. Figure 2-23 depicts the junctions detected using the method outlined above for a drive from Thurrock to Weymouth. All junctions are successfully identified, but several false detections occur when the vehicle is travelling along the M25. There were no other false detections in this instance.

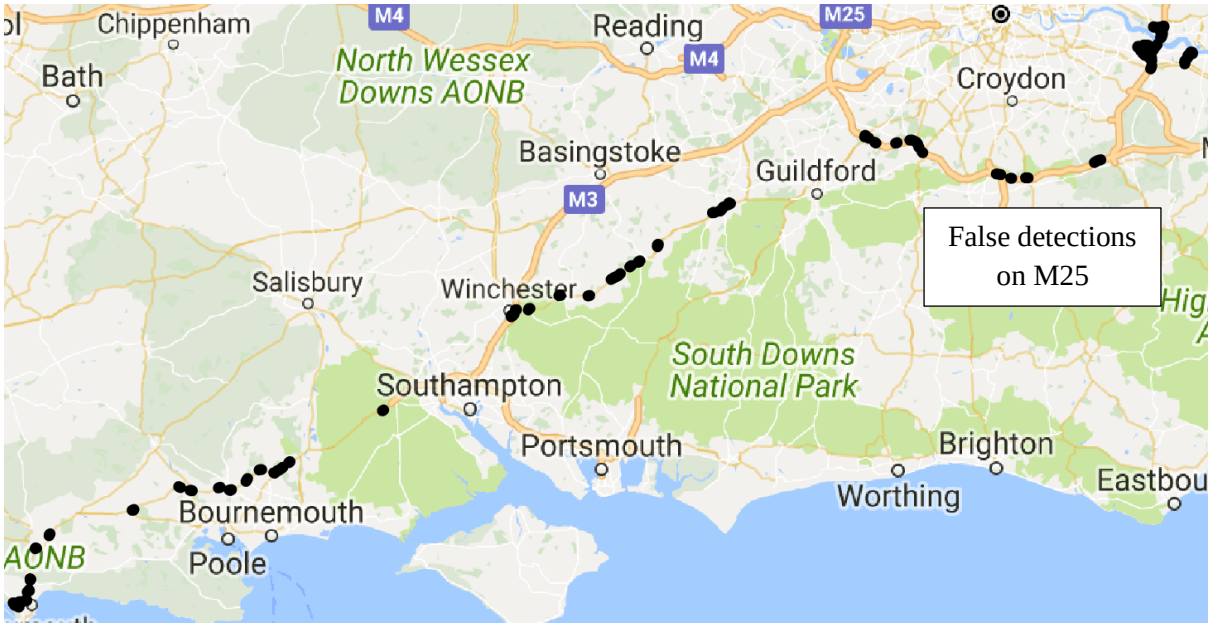

Figure 2-23 – Map of junction detections on Thurrock-Weymouth route

Each approach for detection of junctions has its own strengths and weaknesses. The first does not cope well with low speed manoeuvres, as the accelerations could be low. The radius approach better accounts for this because both the speed and lateral acceleration are lower for slow speed turning.

2.3.3. Roundabout Analysis

When roundabouts or junctions are successfully identified, they can be analysed to see how different drivers choose to travel them. As an example, a roundabout near the trucks' home depot is selected as both trucks frequently pass over it.

A border is manually defined around the junction or roundabout to encompass all entrances and exits - Figure 2-24. The latitude and longitude of the entrances and exits are noted and stored. The GPS history of all trips is then explored for any trace that falls within the boundary. The time and date of any detected crossing is recorded in a log. The distance between the first GPS position logged inside the border and each of the entry points is measured, and the smallest distance corresponds to the entry point used. Similarly, the distance of the last GPS position inside the border before leaving the roundabout and all exit points are also calculated. The smallest distance is corresponds to the exit used. The entry and exit points are added to the log.

Figure 2-24 maps the selected roundabout with its entry/exit points. The border is placed non-centrally around the roundabout in order to incorporate as much of the approach/exit roads as possible without encroaching on neighbouring junctions. The A282 is a major three lane dual carriageway in this location, so the slip roads south of the roundabout are included up into the point of merging. There is another major roundabout north of this one, before northbound traffic can merge with the A282, so the slip roads are fully included in this direction. Note that one extra exit point is marked that does not encompass the roundabout. This route from entry point 1 to exit point 2 is only included as one of the trucks travels down it once.

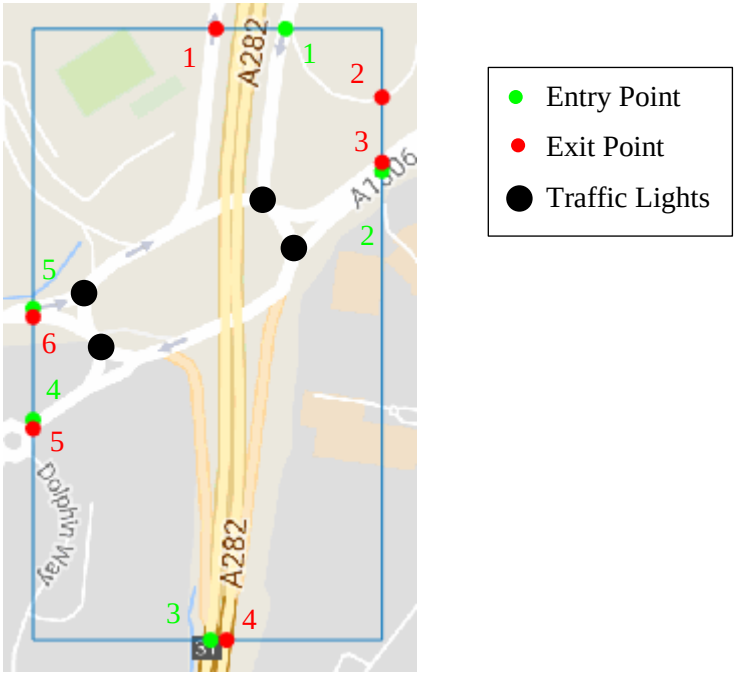

Figure 2-24 – Map of selected roundabout. Entrances onto the roundabout are marked with a green dot, exits off the roundabout with a red dot

With all roundabout crossings logged, an insight into the drivers' behaviour is sought. Obviously, the route the driver takes across the roundabout is going to have a strong effect on the observed behaviour. For a direct comparison in this analysis, the route from the home depot (entry 4) to the A282 southbound (exit 4) is selected due to a high frequency of journeys.

Approaching the roundabout from entry point 4, the speed limit is set to the national limit, but the driver does not have far to accelerate from the previous roundabout and the driver may have to stop at traffic lights as they hit the roundabout itself. There is a second set of traffic lights just before the road curves sharply to the right to allow traffic from entry point 5 to join the roundabout. A third set of traffic lights controls the traffic from entry point 1 joining and a fourth set controls the traffic from entry point 2 joining. On leaving the roundabout and entering the slip road, the speed limit becomes 50mph for all vehicles.

The driver's journey across the roundabout is therefore very dependent on how heavy the traffic is and the timing of the traffic lights. Although it is considered that lateral manoeuvres such as roundabouts are more predictable than longitudinal ones, it is still quite possible that the driver will have to stop at every set of lights, or may be able to travel straight through and onto the slip road without stopping. This means that simply taking the mean of the measured variables over multiple passes may not produce a representative example of any real drive as stopping at different locations, or not stopping at all will have a severe impact on the averaged accelerations and hence fuel use. An alternative approach is therefore proposed to account for this.

Cross correlation is a method used in signal processing to measure how similar two signals are. By using the cross correlation function between every set of pairings of journeys for each driver individually, a 'most typical' journey across the roundabout can be identified by summing the peak cross correlation coefficients for each journey. The journey with the highest sum of cross correlation coefficients is marked as the most typical for each driver.

By using one set of data in its entirety, rather than averaging data between sets, it is ensured that the accelerations and hence fuel uses examined are representative of real behaviour from the driver. Figure 2-25 and Figure 2-26 demonstrate the 'most typical' journey for each driver by depicting the speed of the vehicle and also the driving modes.

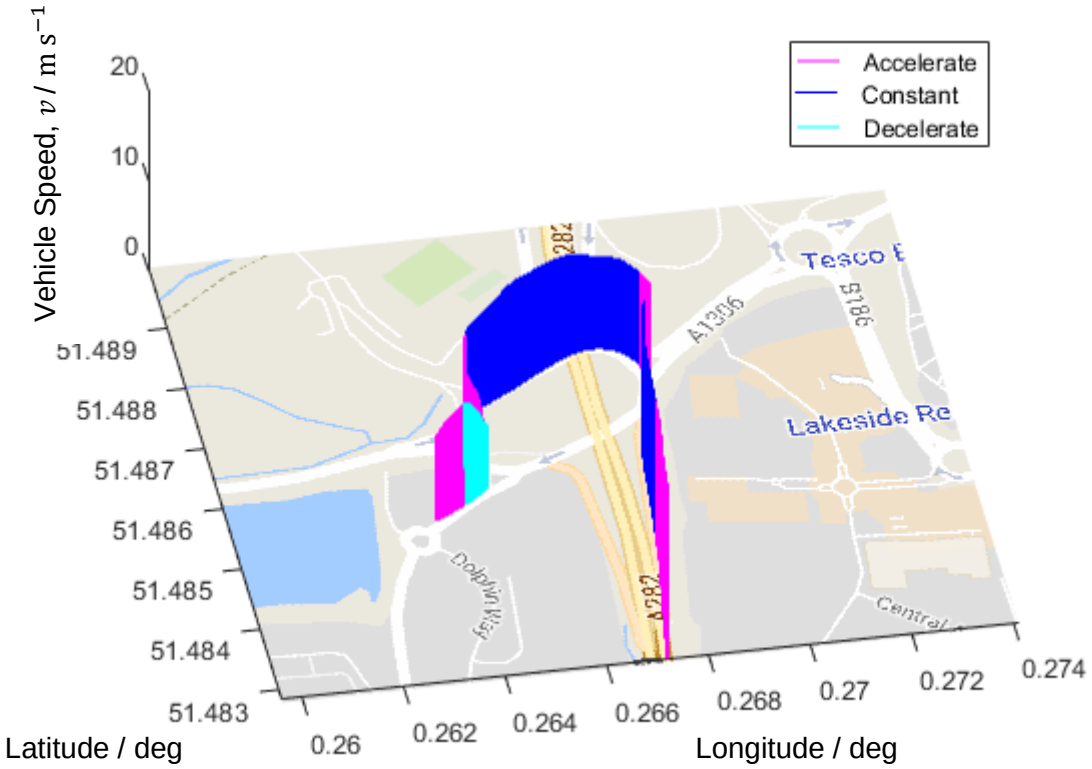

Figure 2-25a - Driver 1 Group A states

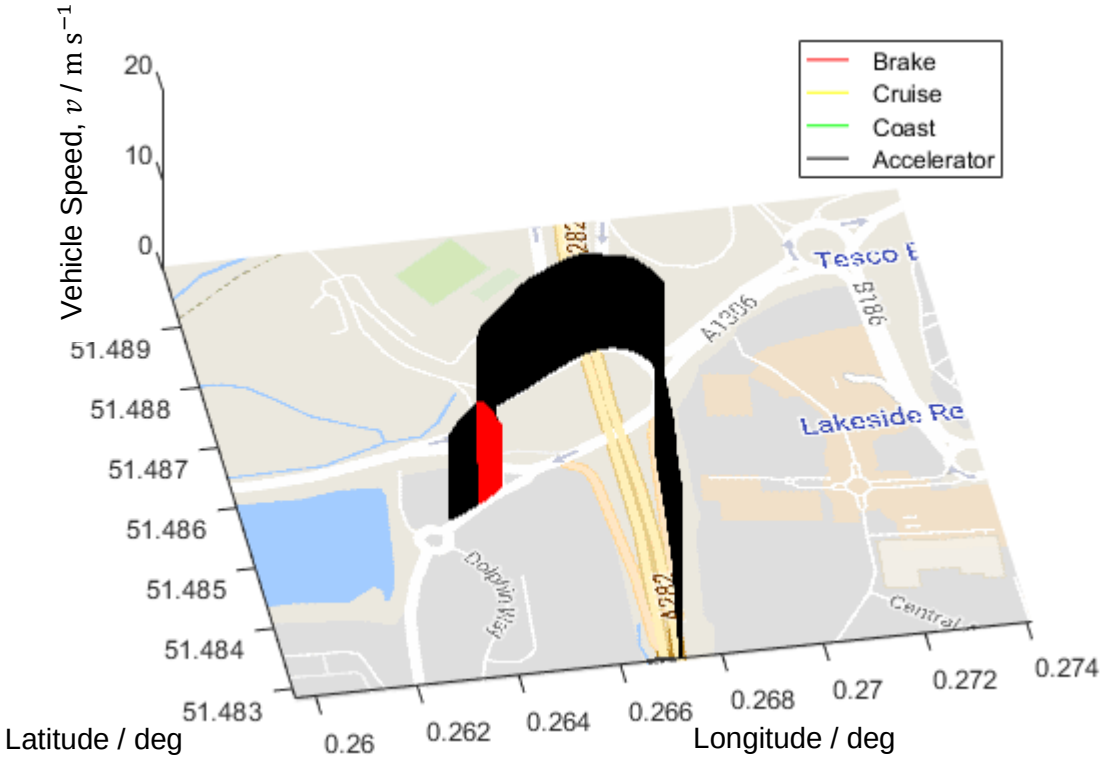

Figure 2-25b – Driver 1 Group B States

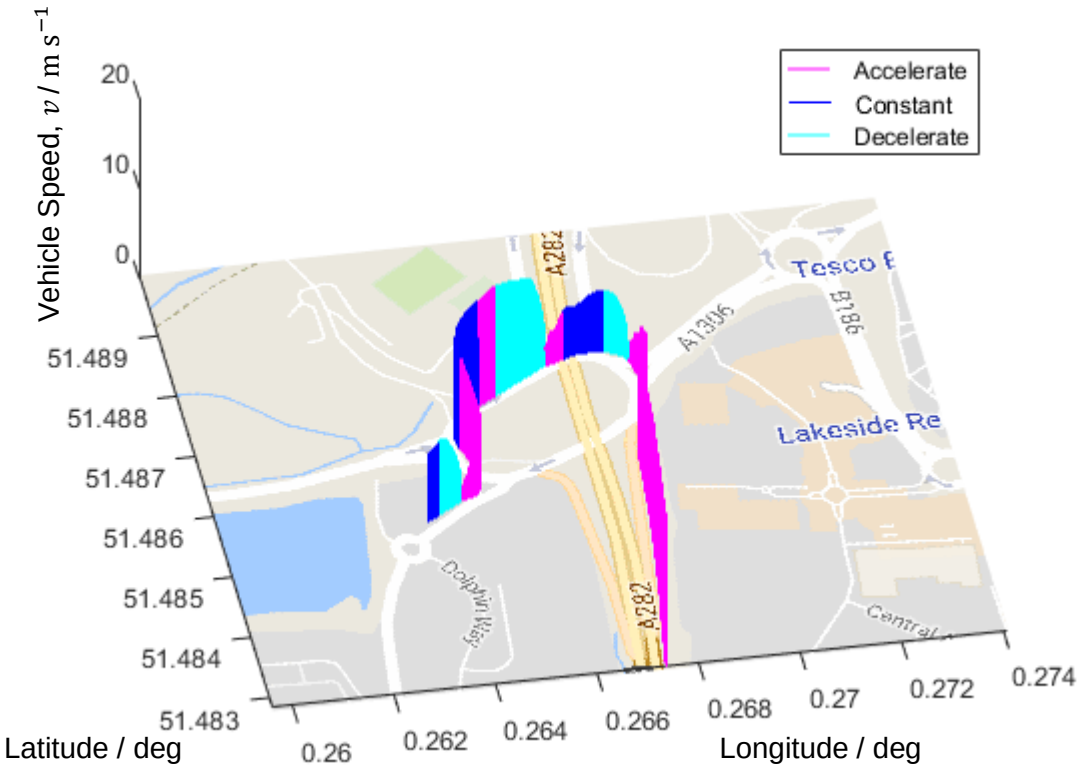

Figure 2-26a – Driver 2 Group A States

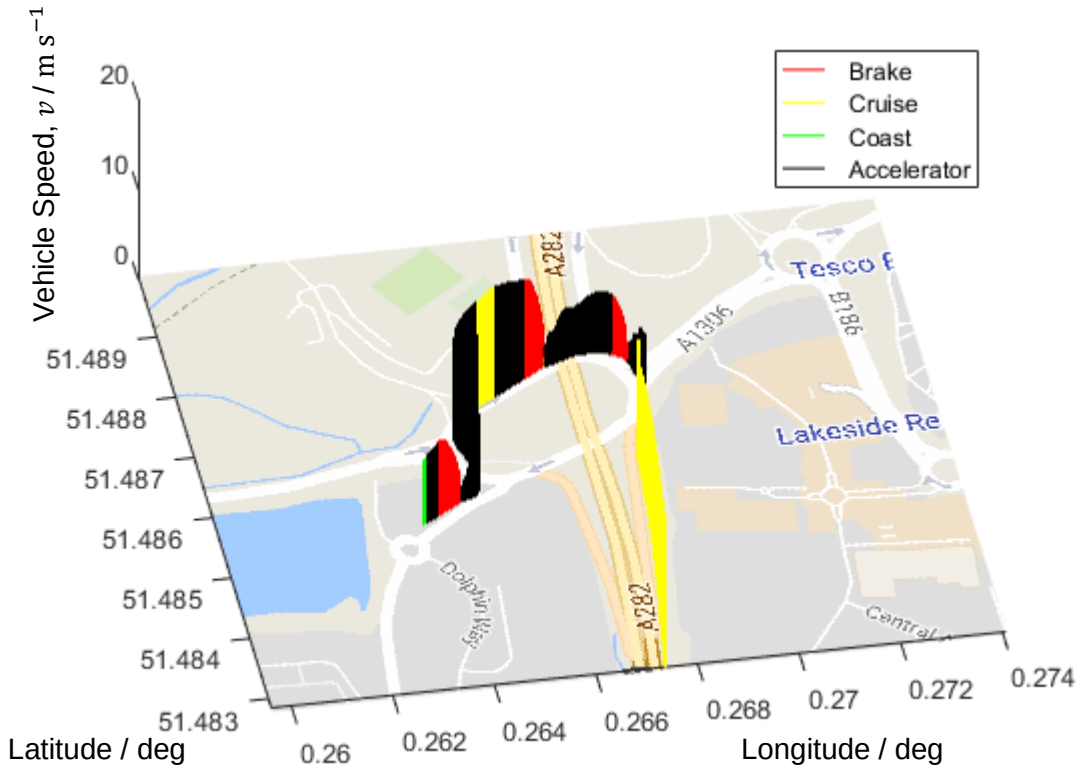

Figure 2-26b – Driver 2 Group B States

Firstly, Figure 2-25a displays the Group A states for Driver 1. The driver enters the area accelerating from the previous roundabout, but soon has to decelerate to stationary at the traffic lights to enter the roundabout. When the lights change, the vehicle accelerates up to $12\ \mathrm{ms}^{-1}$ and manages to maintain that speed through to the final lights. On passing the final lights at green, the driver accelerates the vehicle up to $15\ \mathrm{ms}^{-1}$ and exits the roundabout onto the slip road. Once out of the bend of the slip road, the driver starts accelerating again to merge with the traffic on the A-road. Figure 2-25b displays the control states, or Group B states, and it can be seen that the driver is exclusively on the accelerator pedal, apart from when braking for the red light at the entrance to the roundabout.

Figure 2-26a depicts the Group A states for Driver 2 on the roundabout. The driver here has adopted a lower target speed on approach to the roundabout, so is maintaining a constant speed on entry into the area, rather than still accelerating as the case was with Driver 1. Once again though, the driver has to stop for traffic lights on entry to the roundabout, although this time, due to traffic conditions, the vehicle stops further away from the roundabout. When clear, the vehicle accelerates up to $12\ \mathrm{ms}^{-1}$ before having to decelerate to another stop for the traffic lights managing traffic joining from entry point 1.The driver is also stopped at the traffic lights for entry point two, before then being able to accelerate onto the slip road and then onto the dual carriageway. Figure 2-26b illustrates the Group B modes for Driver 2. It is clear here that the two drivers have very different styles. When Driver 2 has accelerated out of the first set of traffic lights, they apply cruise control briefly before having to disengage it to brake to stop for traffic lights. Once Driver 1 has navigated the last of the traffic lights, he engages cruise control again to accelerate onto and through the slip road.

The analysis explained so far has shown that Driver 2 adopts a lower speed throughout the manoeuvre and engages cruise control to manage some of the vehicle acceleration, but the variations in traffic conditions and traffic lights mean that the two ‘most typical’ journeys of each driver across this roundabout are not directly comparable. For further analysis, the focus is now switched to the slip road, where both drivers are assumed to have sufficient space on the road for them to manage vehicle acceleration without any interference from other road users.

Whereas the roundabout itself was assumed flat and gradients could be ignored, the slip road has gradients that need to be taken into account - Figure 2-27 illustrates the road elevation. Using the change in elevation from this figure, the recorded vehicle speed and the vehicle model in Figure 2-10, an estimate is made of the traction forces. By plotting this information on axes of vehicle speed and traction force, the demanded tractions forces are now compared with the maximum output power and torque of the engine - Figure 2-28. This demonstrates that Driver 1 is consistently demanding higher forces at higher speeds from the engine as Driver 1 desired greater speeds and accelerations.

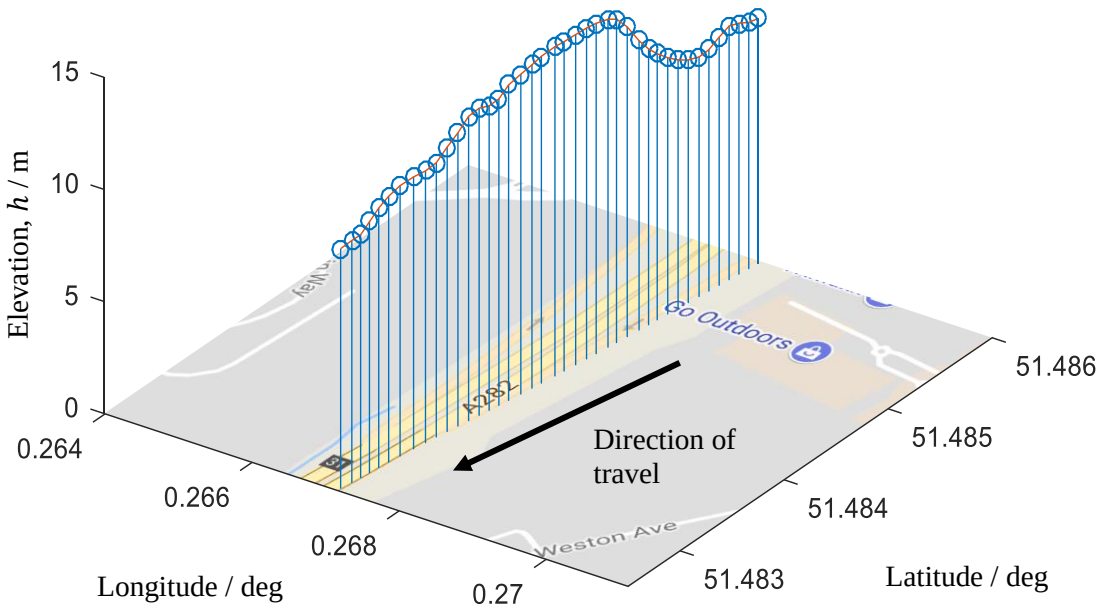

Figure 2-27 – Ground elevation over the slip road

The cyan, magenta and blue lines of Figure 2-28 are the maximum power and torque contours of the engine, as already seen in Figure 2-8. The stars (Driver 2) and circles (Driver 1) on the driver traction force traces indicate which gear the vehicle is in at that moment. Figure 2-28 clearly shows the difference in speed of the two drivers. The vehicles are accelerating up the slip road. Both traces are a similar shape, due to the dominance of the road gradient term of the traction force, but it is noted that Driver 1 is speed shifted by approximately 3ms^{-1} in the positive direction throughout. It is also noted that under Driver 1, the vehicle is changing gear at higher speeds – this could be ab indication of the high demands on the engine meaning that the gearbox chooses to shift up later.

Also of note is that the demanded traction force for both drivers goes above the maximum power of the vehicle. This particularly unlikely, especially as it falls within a gear change period, also for both drivers. On investigation, this peak in traction force is put down to unrealistically steep gradients calculated from the road elevation data in Figure 2-27. Observing the lay of the land around the slip road, there are steep embankments up to the dual carriageway on the west of the slip road, and steep embankments down towards former chalk quarry pits on the east side. It is quite probable that the quoted land elevations may drift off either side of the road and individual points may in fact be land heights of either embankment.

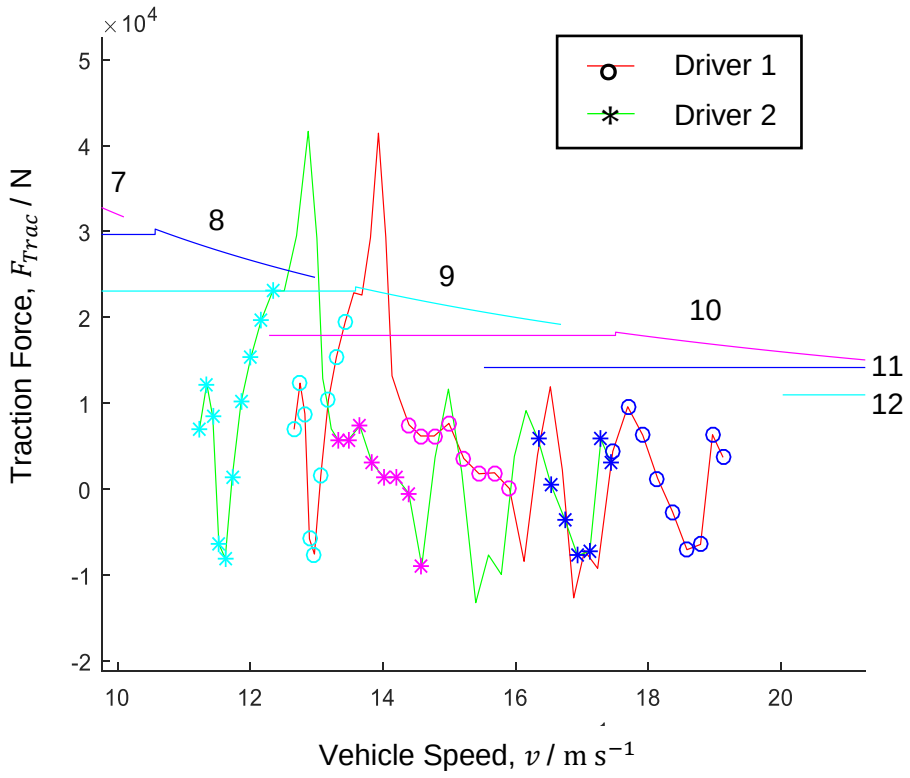

Figure 2-28 – Estimated traction forces against vehicle speed. The force profile for Driver 1 is plotted in red with ∗, and for Driver 2 in green with ∘. Maximum power/torque contours are plotted in blue, cyan and magenta for gears 7 to 12 and asterisks and circles are coloured for the corresponding gear.

The final characteristic to be examined is the fuel use. Figure 2-29 illustrates the cumulative fuel use across the slip road. By the end of the slip road, Driver 2 has made a 7% fuel saving compared with Driver 1, but this is at the cost of 5s journey time, or an increase in 15% of the time taken to cover the slip road.

In summary, the analysis of this roundabout and slip road has shown that by a combination of slightly lower speeds and much greater use of cruise control, Driver 2 has been able to achieve significant fuel use savings of 7%.

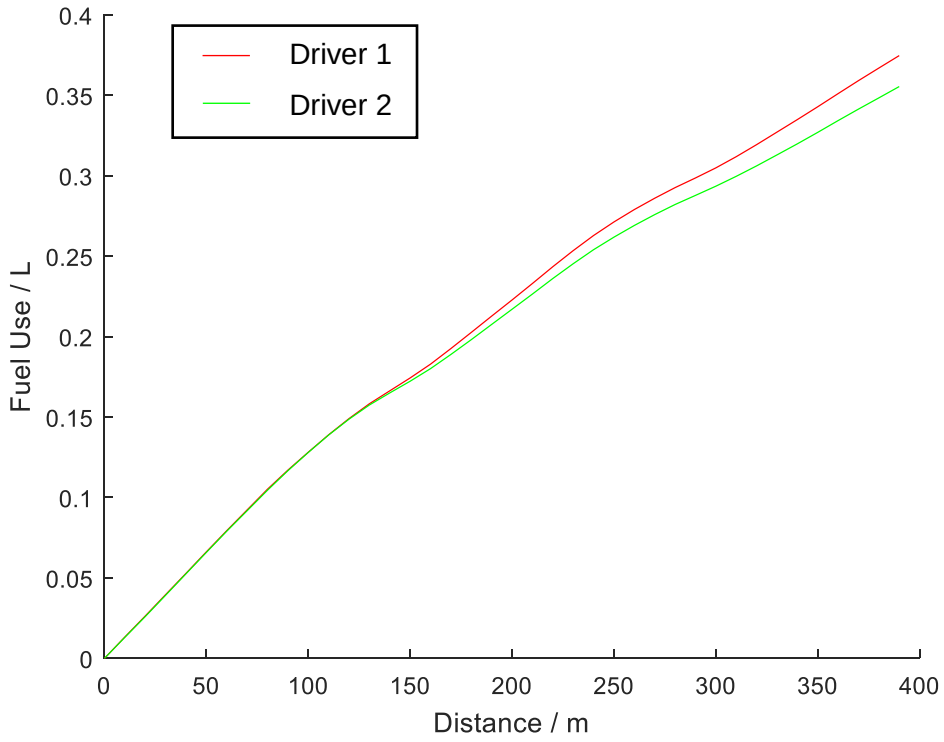

Figure 2-29 – Measured fuel use of slip road

The roundabout analysis already covered in this section is extended to cover 9 other roundabouts. The route over each roundabout is determined by identifying the most frequent route across both drivers, and then the 'most typical' pass for each driver is identified in the manner already described. Table 2-7 summarises the different cruise control use for each driver and the fuel used in the most typical journey for each driver. The difference in fuel is calculated as a percentage of Driver 1's fuel use. Note that in not one of the roundabouts examined, does Driver 1 engage Cruise Control, whereas Driver 2 is frequently using it for approximately 40% of the distance covered over the roundabout. In terms of fuel use, Driver 2 is consistently using up to 7% less fuel than Driver 1. Detailed analysis of the roundabouts, in the same manner as Figure 2-25 and Figure 2-26 (not included for the extra roundabouts for the sake of conciseness), reveals that Driver 2 adopts a lower speed through the manoeuvre and makes better utilisation of cruise control.

Table 2-7 – Roundabout Data Summary

Latitude / deg	Longitude / deg	Driver 1 Percentage distance in Cruise Control	Driver 2 Percentage distance in Cruise Control	Driver 1 Fuel Used / L	Driver 2 Fuel Used / L	Fuel Used Percentage Difference
51.487	0.266	0	42	0.97	0.90	-7%
51.525	0.073	0	36	1.04	0.97	-7%
50.867	-0.155	0	47	0.61	0.58	-5%
51.462	-0.185	0	31	1.10	1.06	-4%
51.531	-0.293	0	40	0.75	0.71	-5%
51.258	-0.198	0	44	0.92	0.86	-7%
51.259	-0.065	0	39	1.23	1.17	-5%
51.727	-1.260	0	43	0.91	0.85	-7%
51.219	-0.778	0	46	0.71	0.76	-6%
51.426	0.238	0	42	1.05	1.00	-5%

2.3.4. Statistical analysis over multiple journeys

Journey data has been amassed over a significant period of time for each driver. It is therefore possible to analyse statistics spanning multiple journeys as any differences in routes or duties should average out over time, leaving the underlying driver behaviour in the statistics.

Firstly, as a test of this, Figure 2-30 illustrates a Probability Density Function (PDF) histogram of the gross vehicle masses and Table 2-8 presents gross vehicle mass statistics. It can easily be seen that Driver 2 (lower fuel consumption) has a higher average vehicle gross mass. The mean, median and mode of both drivers are well within one standard deviation of each other, suggesting the two data sets could have insignificant difference. A heteroscedastic Student T-Test, however, produces a t-value of 0.0027, meaning that the two means are statistically significantly different. With these mass distributions, the driver behaviours should still be observable through the statistics.

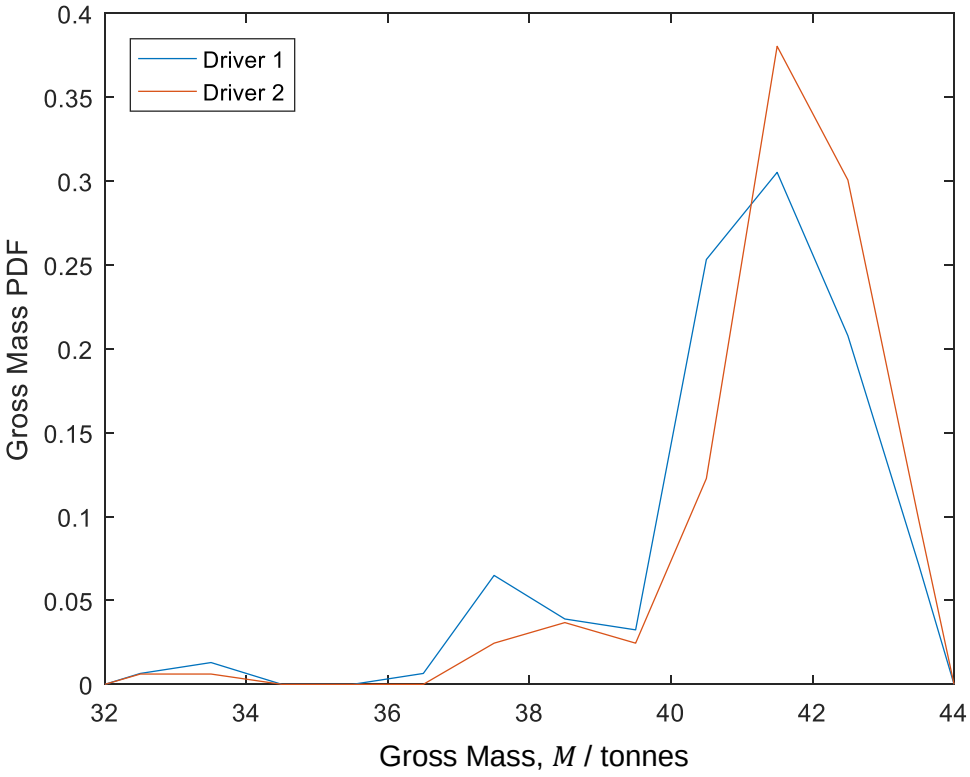

Figure 2-30 – PDF histogram of gross truck masses

Table 2-8 – Gross vehicle mass statistics (all units in tonnes)

	Driver 1	**Driver 2**
Mean	40.96	41.53
Median	41.30	41.78
Mode	40.56	41.22
σ	1.88	1.61

The driving data is broken down into the two groups as first described in 2.3.1. Figure 2-31 and Figure 2-32 illustrate the proportion of time each driver spent in each state for Group A and Group B. In terms of Group A states, both drivers spend the greatest amount of time with the vehicle stationary. Whilst some stationary time is inevitable (at traffic lights for example), there is a six percent gap between Driver 1 and Driver 2, suggesting that there may be scope for Driver 1 to reduce stationary time and hence engine idling.

Both drivers spend an identical proportion of the time maintaining a constant speed. Driver 2, however, spends longer accelerating. On first inspection, this may seem a negative trait, as

acceleration is well known as a high-fuel consumption task, but the inference could also be that the driver is accelerating slower than Driver 1, but is still aiming for the same target speed. In this case, the higher proportion of time accelerating is a positive, as acceleration rates are lower.

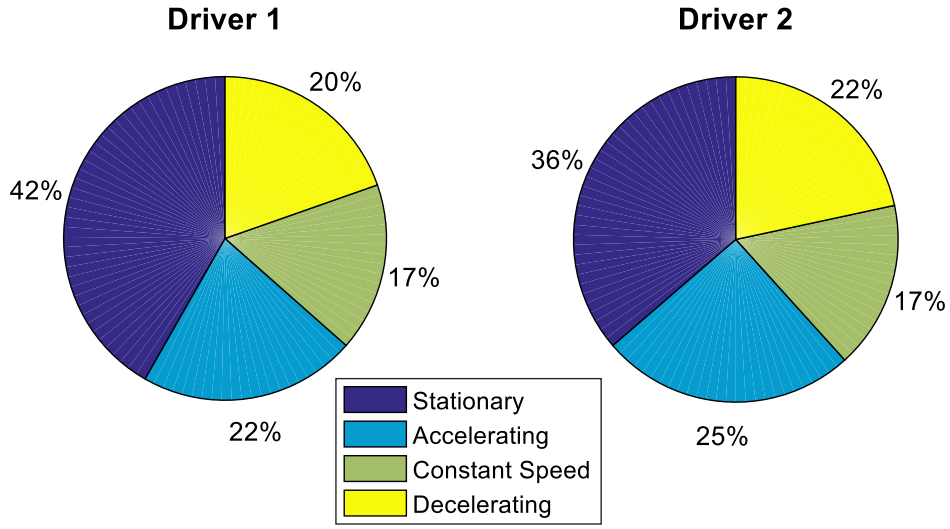

Figure 2-31 – Group A: Speed-Acceleration states over multiple journeys. Percentages are of engine-on time

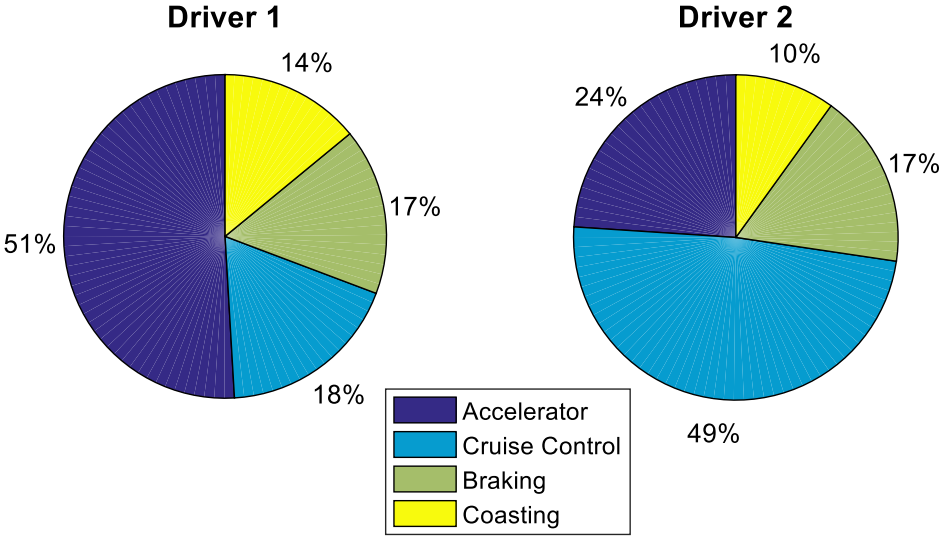

Figure 2-32 – Group B: Driver Control states over multiple journeys. Percentages are of non-stationary time

Now observing the Group B tasks, it can be seen that both drivers spend a similar proportion of the time braking and coasting, but there is a significant difference in the accelerator pedal and cruise control use. Driver 1 demonstrates a clear preference for controlling the vehicle with the accelerator pedal, whereas Driver 2 uses the Cruise Control. Although these pie charts are useful in giving an overview of the drivers' preferences, they do not give a thorough analysis of their styles.

The first characteristics to be examined in more detail are the vehicle speed and acceleration logs. Figure 2-33 illustrates the PDF histogram of vehicle speeds, with probabilities given the vehicle is non-stationary. At low speeds, the two histograms are similar, with Driver 2's PDF's slightly higher than Driver 1. This changes significantly, however, at high speeds. For vehicle speeds great than $24\ \mathrm{ms}^{-1}$, Driver 1 has much high probability density than Driver 2. This implies that Driver 2 is choosing slightly slower speeds, probably with the aim of better fuel consumption, where Driver 1 is determined to minimise journey time.

In Figure 2-34, the vehicle acceleration is examined. Of interest is the apparently higher accelerations of Driver 2 than Driver 1. However, what this figure does not illustrate well is the fact that Driver 1 does have high acceleration rates, but as their target speeds are therefore reached quickly, the probabilities (which are related to time), are seemingly lower. At the tails of the distribution, Driver 1 has higher acceleration probabilities.

The next section of analysis focusses on the control states of the vehicle. A PDF histogram of pedal position (Figure 2-35) shows clear differences in driver preferences between the two drivers. The probabilities are calculated given a non-zero pedal position, so this histogram does not include the use of cruise control, or coasting (these are split into separate statistics). It is very apparent, however, that Driver 2 spends a much higher proportion of the time when they are using the pedal at low pedal positions, whereas Driver 1 is prepared to push the vehicle with high pedal demands. This figure, however, does not show the full picture, as the driver's pedal control is also very dependent on vehicle speed. The bivariate histograms in Figure 2-36 explore the relationship in more detail.

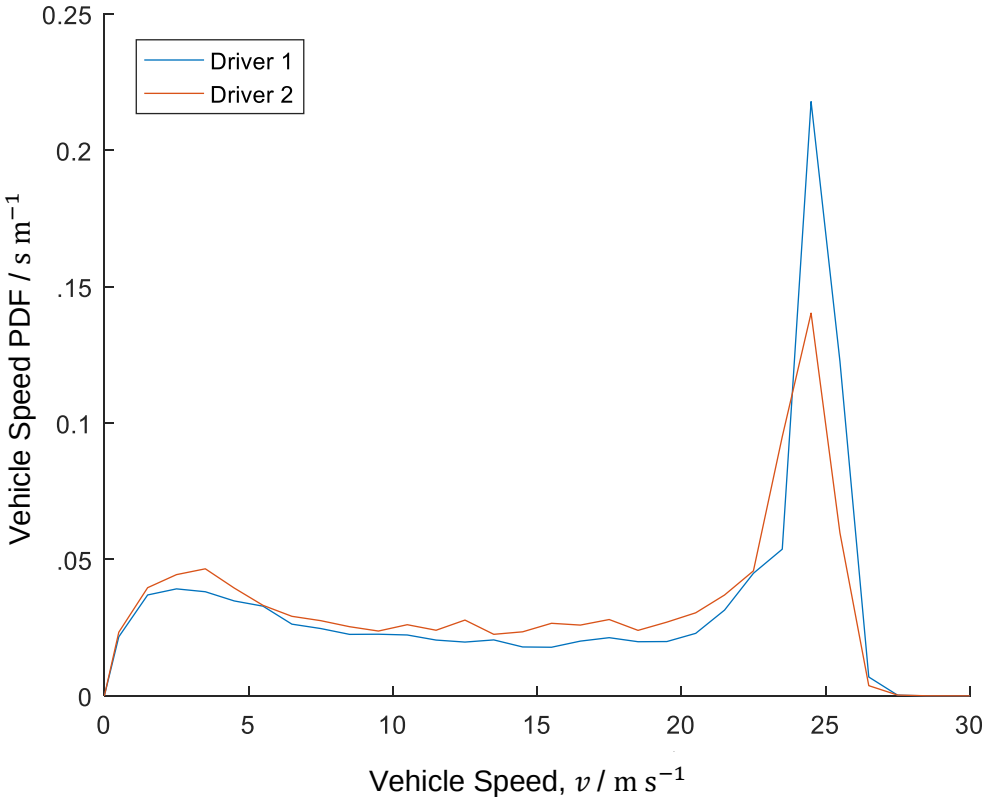

Figure 2-33 – PDF histogram of vehicle speeds, given non-zero vehicle speed

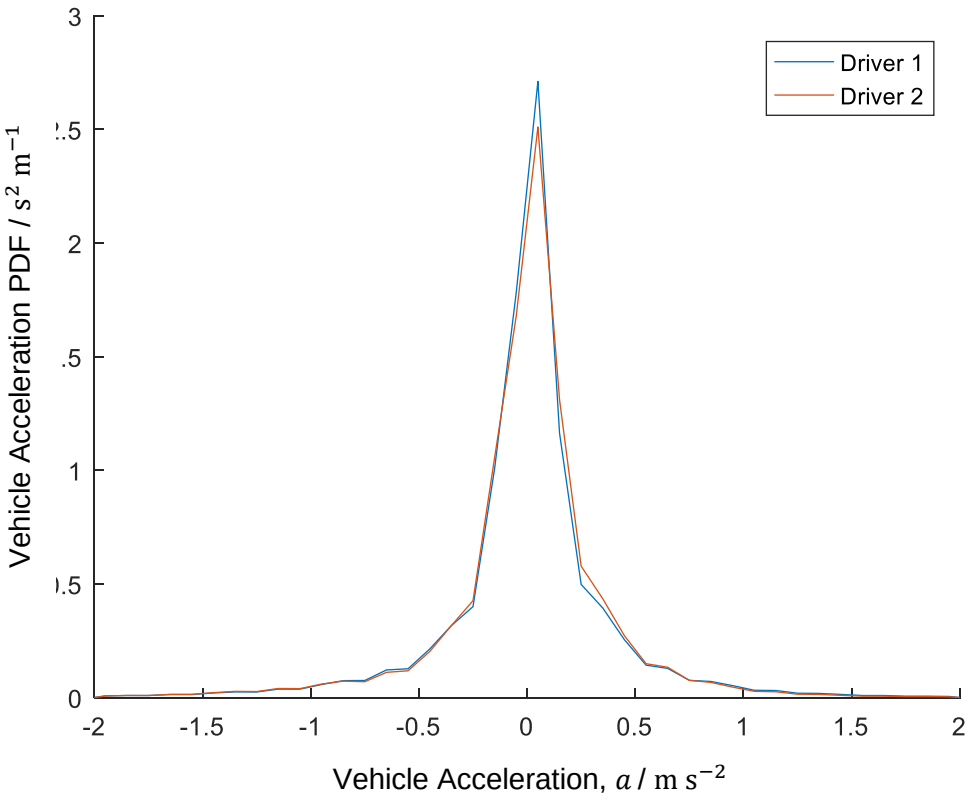

Figure 2-34 – PDF histogram of vehicle acceleration

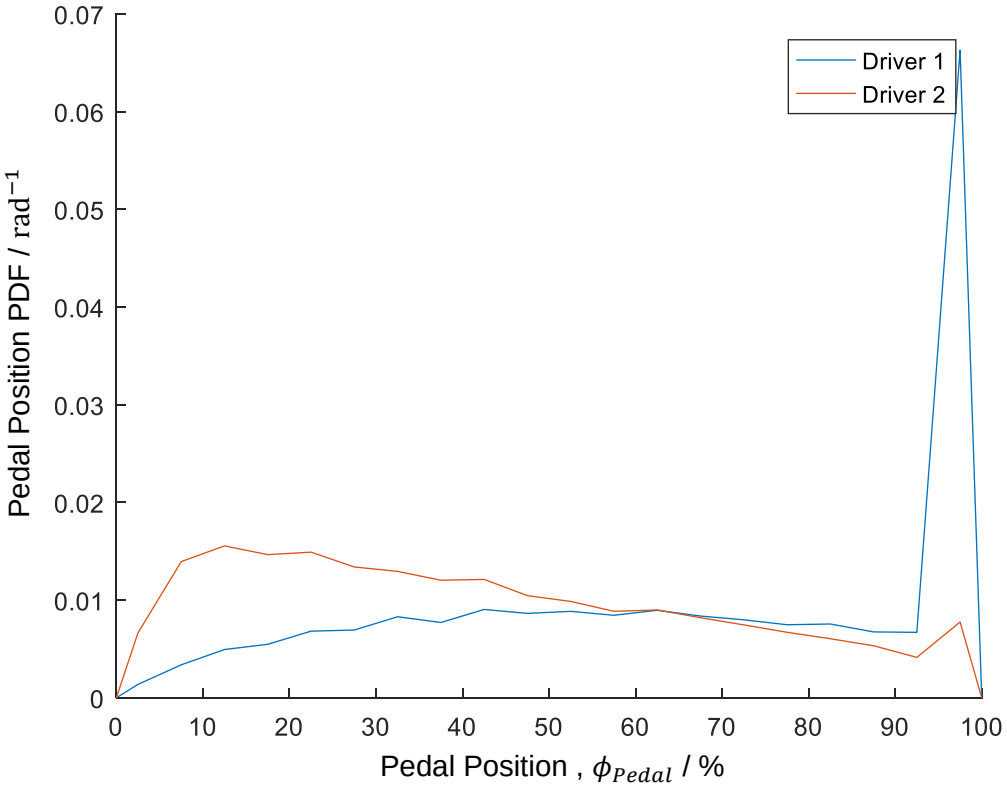

Figure 2-35 - PDF histogram of pedal position, given non-zero pedal position

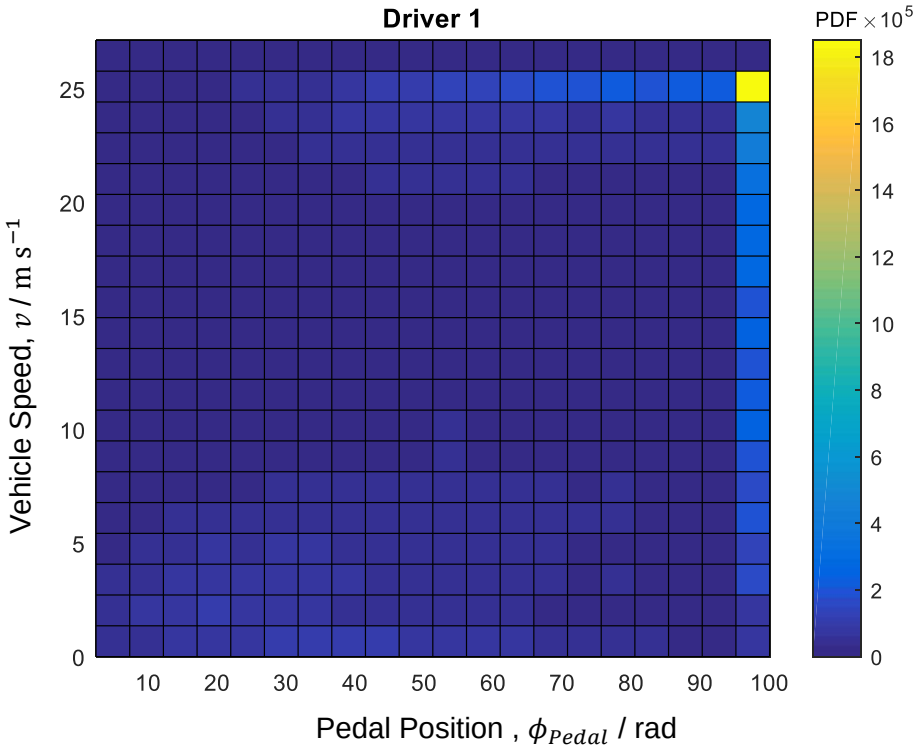

Figure 2-36a– Vehicle speed against pedal position bivariate histograms given pedal position is non-zero for Driver 1

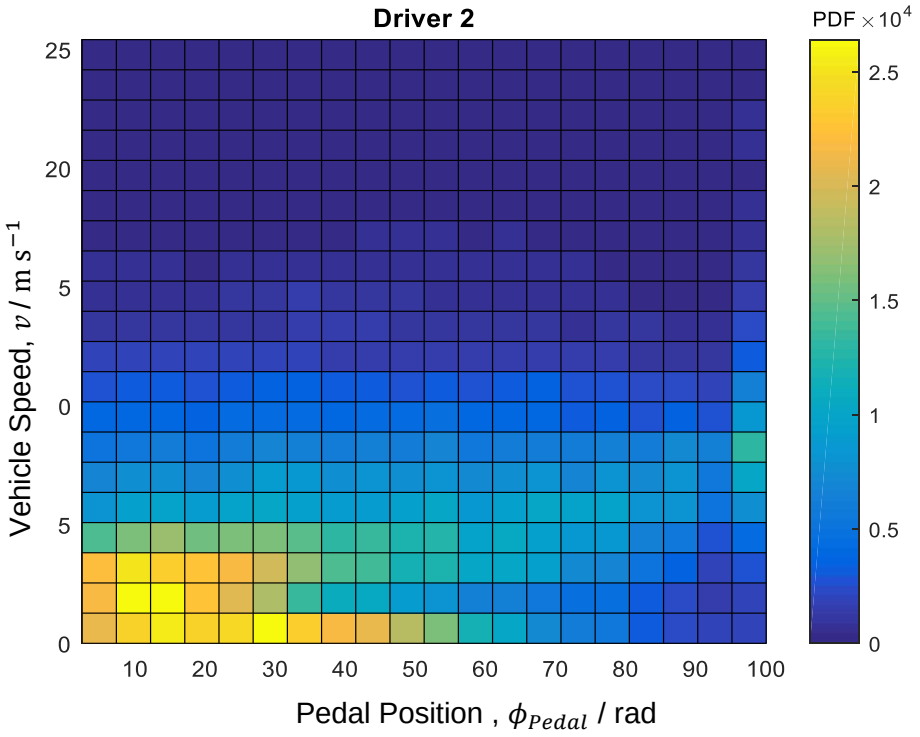

Figure 2-36b – Vehicle speed against pedal position bivariate histograms given pedal position is non-zero for Driver 2

The bivariate PDF histograms in Figure 2-36 clearly demonstrate the differences in the drivers' control of the accelerator pedal with speed. Driver 1's histogram is completely dominated by very high pedal positions at very high speeds. Slightly less apparent is the habit of lifting off the pedal at high speeds when target speed is reached as well as the depression of the pedal to near full depression and the acceleration up from low speeds to maximum speed. Driver 2, on the other hand, chooses to keep pedal position low over a range of speeds, but there is still a trend apparent to depress the pedal and then wait for acceleration through to higher speeds. Driver 2 doesn't show high use of large pedal positions because they use Cruise Control a lot more than Driver 1.

Finally, the drivers' use of Cruise Control is examined in more detail. Figure 2-37 illustrates the probability of Cruise Control usage at any speed given the vehicle is non-stationary, and Figure 2-38 illustrates the accelerations experienced under Cruise Control. The main difference noted from the speed figure, is that Driver 2 is using Cruise Control over a much greater range of speeds – from $8\ \mathrm{ms}^{-1}$ upwards. Driver 1 on the other hand only really starts using cruise control at $20\ \mathrm{ms}^{-1}$. It can be seen from Figure 2-38 that Driver 1 mostly uses Cruise Control to maintain a constant speed, whereas driver 2 is using it to accelerate as well as maintain constant speed. Note that for both Figure 2-37 and Figure 2-38, the areas under the curve do not sum to one for each driver, as the probabilities plotted to not include the condition of cruise control disengaged – ie the greater area implies that Driver 2 uses Cruise Control significantly more often than Driver 1.

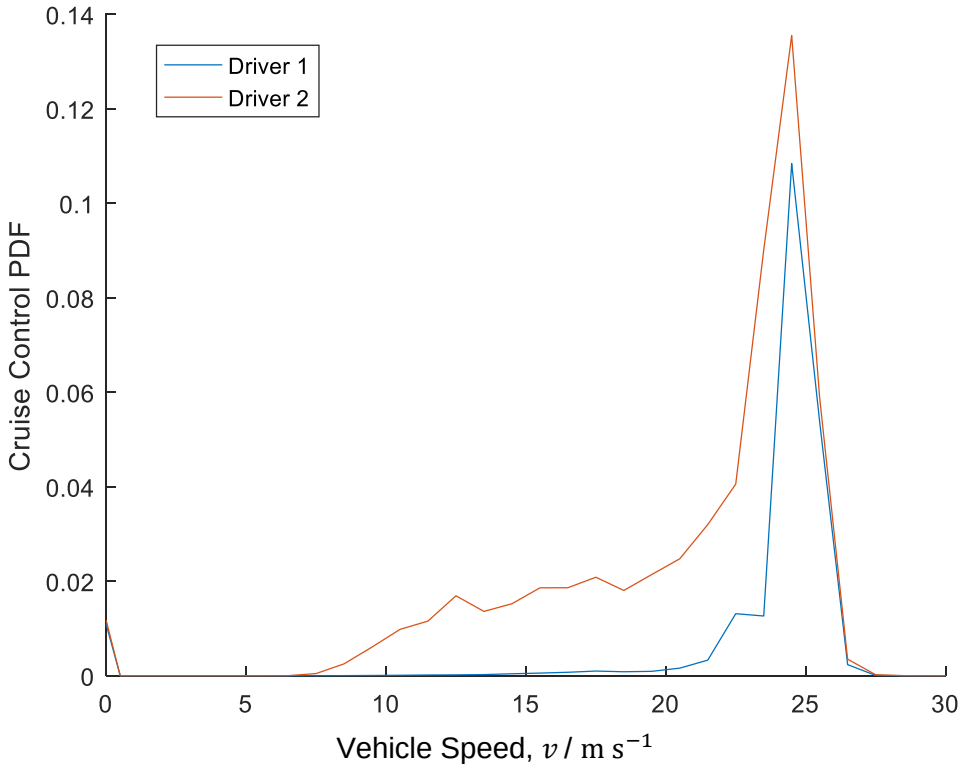

Figure 2-37 – PDF histogram of Cruise Control with vehicle speed, given non-zero vehicle speed

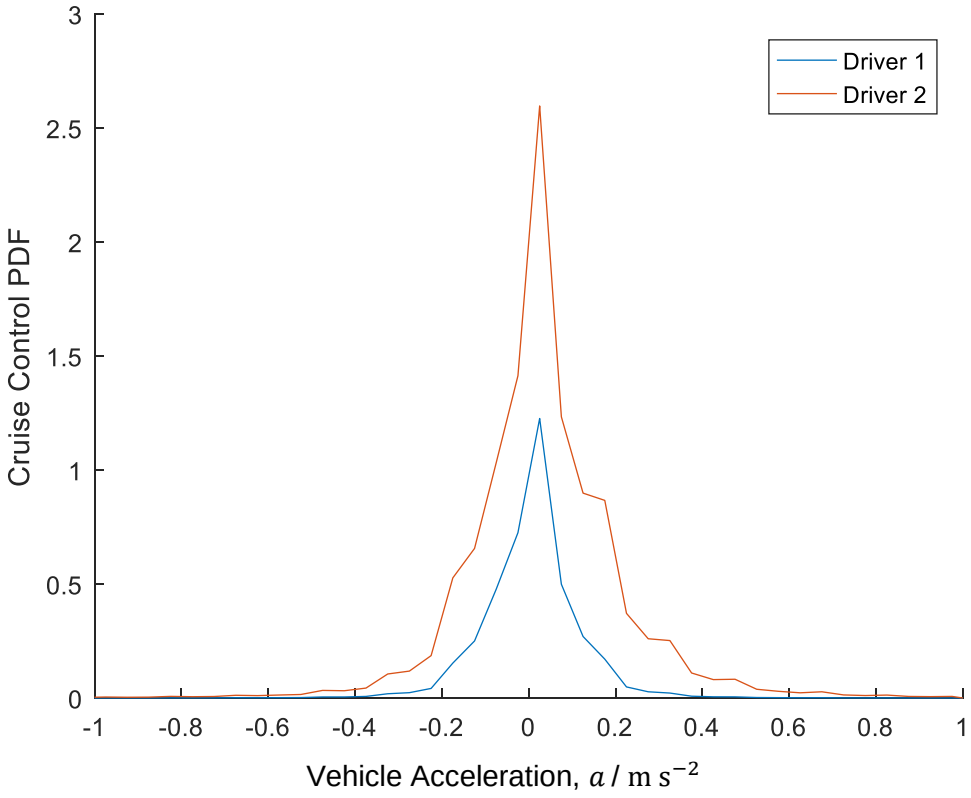

Figure 2-38 – PDF histogram of Cruise Control with vehicle acceleration, given non-zero vehicle speed

2.4. Conclusions

In this chapter, two heavy goods vehicles were fitted with the SRF Logger and on-road data collected over a period of four months. A large variety of routes and journey types have enabled identification of a vehicle model representative of the vehicle during its working life, rather than on a test track in very controlled conditions – an extension of Hunt's work (2011). Improvements in road gradient estimation were made by using an alternative source of ground elevation data; rather than the inbuilt SRF Logger sensors, Google Maps data was used, although there is still further room for improvement. Traction force and fuel maps were generated by fitting multi-dimensional polynomials to recorded data. However, when applying these maps to recorded vehicle speed and pedal position data, significant errors were identified when simulating Cruise Control.

The bank of data has also been used for new analysis of truck driver behaviour. Firstly, two groups of driving tasks (Group A and B) were defined to specify the driving tasks. Through use of the Douglas Peucker algorithm, the number of data points is reduced. This analysis demonstrated

that a large proportion of the distance covered by the two drivers is done so under cruise control on dual carriageway, so there is little effect from the driver's behaviour. The drivers do, however, use the accelerator pedal extensively throughout junctions and roundabouts, so a technique was developed to identify the crossings of the junctions and roundabouts over a route. This proved very successful, with a 100% of all junctions suitably identified but also a 7% false detection rate over a sample route from Thurrock to Weymouth.

With junctions successfully detected, analysis of individual roundabouts showed that through use of slightly lower speeds and increased cruise control engagement, Driver 2 was able to achieve fuel savings of 7% over the roundabout.

Statistical analysis over the entire bank of data was also completed to explore the differences in driver behaviour. Both vehicles were found to be operating at a similar average gross mass, and the difference in driver preferences was once again made clear by the differences in accelerator pedal against cruise control usage. This analysis has identified the potential to bring Driver 1 closer to Driver 2's style and corresponding fuel savings by modifying Driver 1's behaviour. Increasing use of their cruise control, reducing the extreme high pedal positions, reduction harsh accelerations and adopting a smoother speed profile are all identified as positive potential actions for Driver 1.

Chapter 3 - Driver Model

As highlighted in the literature review, human drivers are difficult to model due to their complexity. Most existing models, as reviewed in the literature review in Chapter 1, break the driver's control of the vehicle down into several tasks and will only simulate certain aspects of it, such as the lateral or longitudinal control of the vehicle. In this project, the driver's interaction with the accelerator pedal is of interest so a longitudinal control model is required. Many of the existing models incorporate the vehicle as well and directly relate the driver's target to the states of the vehicle (eg. vehicle speed). As already stated, it is the driver's interaction with the accelerator pedal that is of interest, so these models are of little use in their published form. A Model Predictive Control (MPC) control framework was selected as it can account for the preview the driver has of the upcoming speed demands, but can also take into account constraints, such as actuator limitations and can easily reflect different driving styles (Sharp, 2007). However, Sharp's model did not directly incorporate the driver's control of the accelerator pedal.

For the interaction of the driver with the accelerator pedal to be examined a vehicle model is also required. The structure of the combined driver and vehicle model is displayed in Figure 3-1. The driver model, represented by the controller $\mathbf{K_1}$ (here the subscript one indicates the human driver, in later chapters, subscript 2 is used to indicate the pedal feedback controller), compares the current vehicle speed with the driver's target vehicle speed and decides to exert a force on the pedal in order to accelerate the vehicle if necessary. The pedal converts this force into a demand signal (pedal displacement) and feeds this into the engine and vehicle model and the vehicle responds and accelerates.

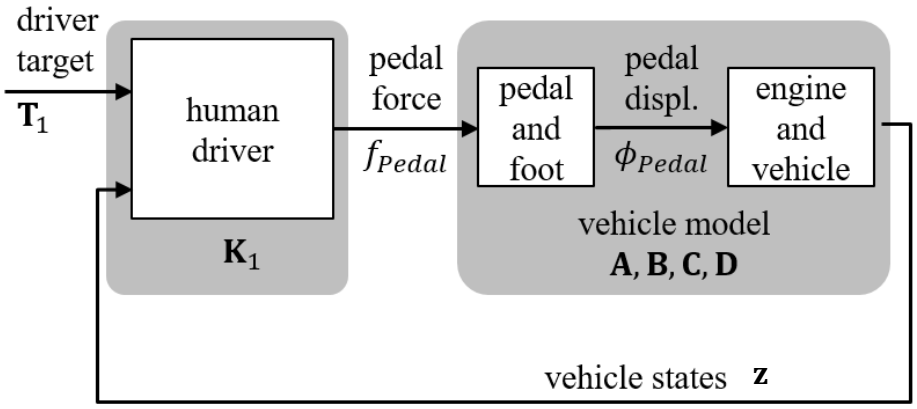

Figure 3-1: Driver Model Block Diagram

Model Predictive Control relies on a plant model to calculate control actions. Whether this model is linear, or non-linear, has significant effects on the method of control, and complexity of control. This chapter will explore linear optimal control (notes on non-linear control are included as an appendix). Firstly, the linear vehicle and pedal models are defined in 3.1. Linear Model Predictive Control theory is documented in 3.2. In 3.3, the control tasks are defined as a drive-cycle scenario and a car-following scenario, before a parameter study explores the behaviour of the driver model. In 3.4, driving simulator experiments are described and the results analysed before model identification in 3.5. Conclusions are drawn in 3.6

3.1 Vehicle and Pedal Model

The vehicle and pedal models are key components of the driver-vehicle model. They are described here in full.

3.1.1 Linear Vehicle Model

The vehicle dynamics model is composed of engine and vehicle dynamics, and a pedal and foot dynamics model. Firstly, the engine and vehicle are modelled: Figure 3-2.

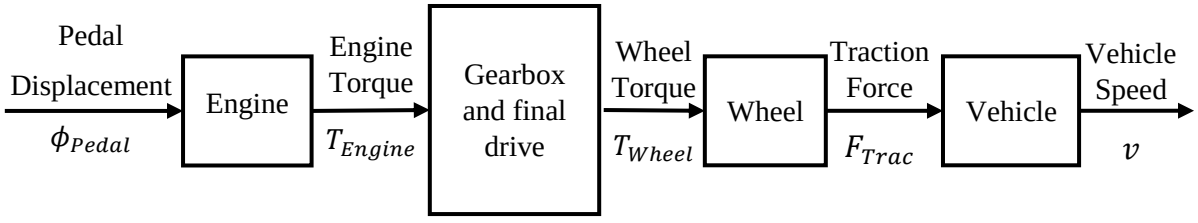

Figure 3-2 - Vehicle Model Block Diagram

For the engine, the engine torque was set as proportional to the throttle displacement – see 2.2.6. Zero torque would be output at zero throttle displacement and maximum torque at maximum throttle displacement. The maximum torque was set to 2000Nm – the maximum engine torque of the DAF CF85 trucks logged in Chapter 2, as specified by the manufacturer. The Engine block can be expressed as:

$$T_{Engine} = \frac{\phi_{Pedal}}{\phi_{max}} T_{EngineMax} \quad (3.1)$$

where T_{Engine}, is the engine torque, ϕ_{Pedal} is the pedal displacement, ϕ_{max} is the maximum pedal displacement, and $T_{EngineMax}$ is the maximum engine torque. Assuming no losses, the engine torque to total wheel torque conversion is a multiplication by gear ratio, $G_{Gear} = \omega_{Gin}/\omega_{Gout}$, and the final drive ratio, $G_{FD} = \omega_{FDin}/\omega_{FDout}$, where $\omega_{Gout} = \omega_{FDin}$:

$$T_{Wheel} = G_{Gear} G_{FD} T_{Engine} \quad (3.2)$$

To convert from wheel torque, to traction force, the wheel torque is divided by the wheel radius:

$$F_{Trac} = \frac{T_{Wheel}}{r_{Wheel}} \quad (3.3)$$

Four opposing forces were applied to the vehicle: an inertial force, a gravitational component from road inclination, a constant rolling resistance and a linearized aerodynamic drag. The forces were defined as illustrated in Figure 3-3 (very similar to Figure 2-10, with the exception of aerodynamic drag).

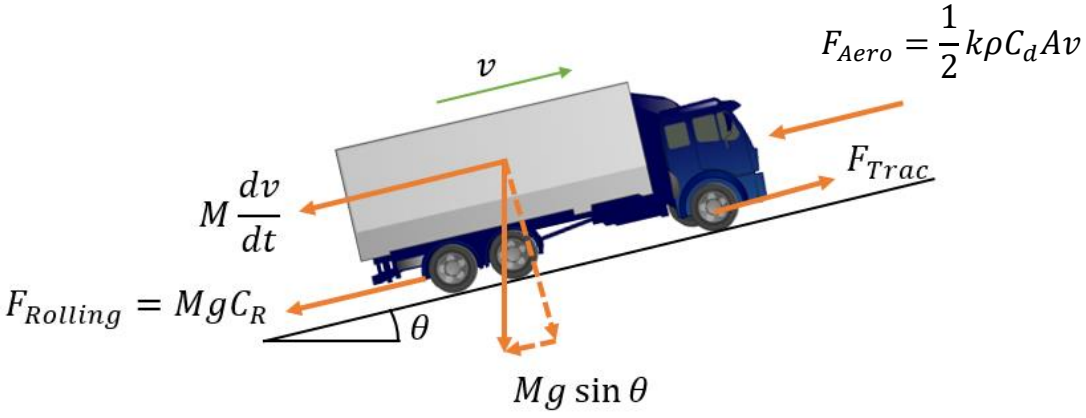

Figure 3-3 – Vehicle Model Forces

Applying Newton's Second Law to the vehicle gives:

$$M\frac{dv}{dt} = F_{Trac} - F_{Aero} - F_{Rolling} - Mg\sin\theta \qquad (3.4)$$

The conventional aerodynamic drag equation is

$$F_{Aero} = \frac{1}{2}\rho v^2 C_d A \qquad (3.5)$$

As a quadratic, this doesn't meet the requirement of a linear model, so a first order polynomial, $\frac{1}{2}k\rho C_d A v$, constrained to pass through the origin was sought. This constraint ensures that aerodynamic drag always resists motion (and will not accelerate the vehicle backwards) and by linearising about a fixed operating point, the model does not need to be linearised at every time step, reducing computational cost. A least squares regression of the errors between the quadratic drag term, and a linear drag term over the vehicle speed range of interest (0 to 30ms^{-1}) calculated the desired constant, k, as $22.5\ \text{ms}^{-1}$ (Figure 3-4).

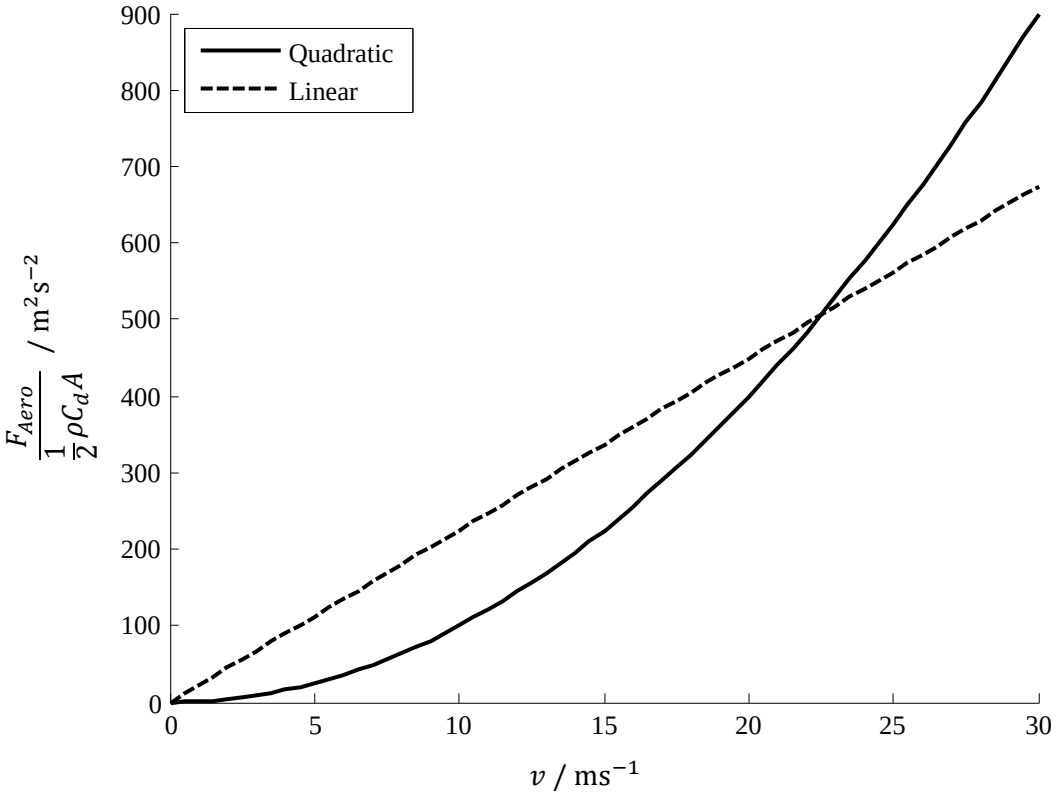

Figure 3-4 – Quadratic and Linearized aerodynamic drag

Table 3-1 summarises all vehicle parameters used in the vehicle model to simulate a generic tractor. With the exception of maximum engine torque and the linearised aerodynamic drag constant of proportionality as already discussed, all remaining vehicle parameters were set to those stated by Hunt et al., (2011), or taken from the output of the SRF Mapper tool discussed in Chapter 2. Hunt et al. measured the properties of a three axle Volvo tractor unit, similar to the DAF trucks logged in Chapter 2, albeit operating with a three axle box trailer, rather than bulk powder trailer. For this reason, the SRF Mapper tool was used to estimate the drag properties of the DAF tractor, bulk powder trailer combination from Chapter 2. The maximum engine torque was not defined in the Hunt paper, so was extracted here from the manufacturer's specifications of the Chapter 2 trucks as already described.

Table 3-1 – Vehicle Model Parameters

Parameter Name	**Symbol**	**Numerical Value**	**Units**	**Source**
Maximum engine torque	$T_{EngineMax}$	2000	Nm	Engine specification
Gear ratio	G_{Gear}	5.74	-	Hunt et al.
Final drive ratio	G_{FD}	3.44	-	Hunt et al.
Wheel radius	r_{Wheel}	0.5	m	Hunt et al.
Air density	ρ_{Air}	1.225	kg m^{-3}	Hunt et al.
Drag Area	$C_D A$	6.62	m^2	SRF Mapper
Mass	M	36600	kg	Hunt et al.
Coefficient of rolling resistance	C_r	662×10^{-5}	-	SRF Mapper
Linearised aerodynamic drag constant of proportionality	k	22.5	ms^{-1}	-

3.1.2 Accelerator Pedal Dynamics

The dynamic properties of the accelerator pedal will play a large part in any simulation of the driver's speed control of a vehicle. A Volvo truck accelerator pedal was deemed representative of a typical truck pedal and one was acquired for use in experiments. The objective of this section is to determine a linear dynamic model of the pedal for addition to the linear model of the vehicle.

Due to the curvature of the pad of the pedal (noticeable in Figure 3-6) meaning that the contact point between foot and pedal may move as the foot rotates, and the complication of the foot rotation about the contact point between heel and floor, an approximation of the pedal –foot system is required. Figure 3-5 illustrates the dimensions and layout of such a simplified pedal. The force from the driver's foot is applied normal to the pad of the pedal. It is assumed that the centre of pressure of the contact between the foot and the pedal is at the centre of the pad, a distance $l_{Pedal} = 0.16\text{m}$ from the pedal pivot point, and that this distance remains constant for all pedal positions. As this remains a constant distance, the driver's input can be considered either a force, or torque, independent of pedal displacement.

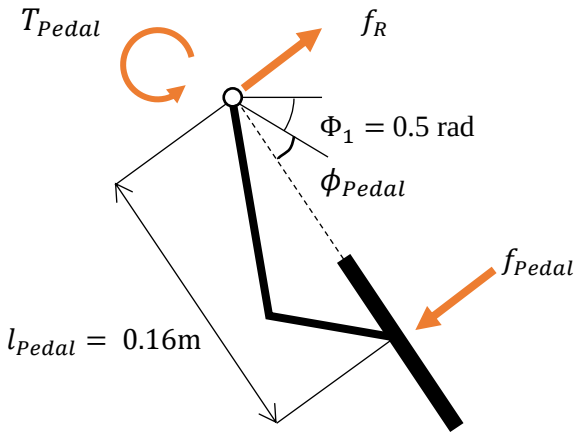

Figure 3-5 – Free body diagram of the pedal arm

The first property of the pedal to be examined was the static force-displacement relationship. The pedal was installed in a rig to angle the pedal such that the direction of travel of the centre of the pad remained as near vertical as possible throughout the full range of travel. Known masses were hung from the centre of the pedal pad and the corresponding displacements were measured using a protractor, as illustrated in Figure 3-6.

Figure 3-6 – Volvo accelerator pedal measurement rig

Figure 3-7 illustrates the recorded torque displacement relationship. There are two clear regions where the pedal is static – at zero travel and at maximum travel. The remaining data then falls neatly onto a loading and an unloading line. The different regions of the curve are manually identified and a line of best fit is then fitted in each by minimising the square of deviation, resulting in the hysteretic loop. This hysteresis pattern indicates a constant friction with preloaded, constant stiffness. Nearly 2.9 Nm of torque is needed before any motion occurs, but on return to the zero position, the measured torque is only 1.5 Nm. This means that there is an effective 0.7 Nm friction torque (corresponding to a 2.5 N friction force on the centre of the pedal) opposing any motion of the pedal. The maximum achievable pedal depression was measured as 0.36 rad. The dynamic stiffness, determined by averaging the gradients of the top (5.01 Nm/rad) and bottom 4.95 Nm/rad) of the hysteresis loop, was calculated as 4.98Nm/rad. It is worth noting here that the pedal force-displacement behaviour is clearly non-linear, but the objective is to identify equivalent linear stiffness, damping and inertia terms.

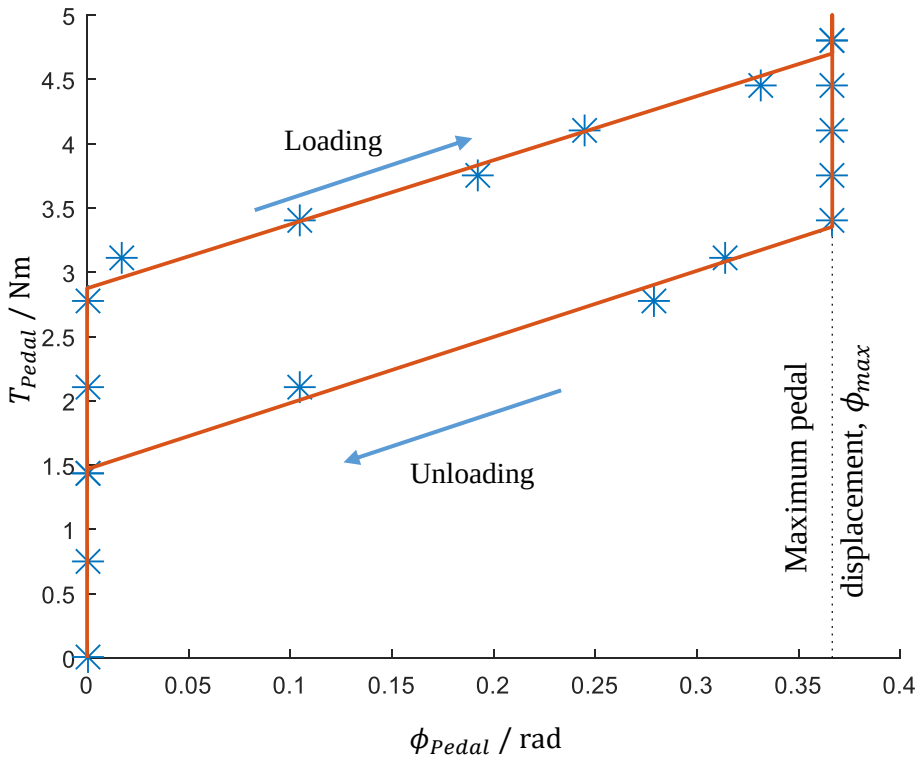

Figure 3-7 – Torque-rotation characteristics of the Volvo accelerator pedal

When installed in the driving simulator, this pedal has additional linkages attached to allow additional forces to be applied to the pedal, by a motor, for the purposes of pedal feedback (see Figure 3-8 and Chapter 4). As a direct comparison is sought between the driver-only model developed in this chapter, and the driver-and-feedback models developed in Chapter 4, the pedal dynamics are here extended to include the additional linkages and motor. The major difference this makes is through the additional inertia added to the system by the motor (later, it is shown that the motor inertia is minor compared to the additional linkages). Figure 3-9 illustrates the resulting mechanism produced when the motor is connected to the pedal. An estimate of the effective moment of inertia of the mechanism about the point A is made in the method explained below.

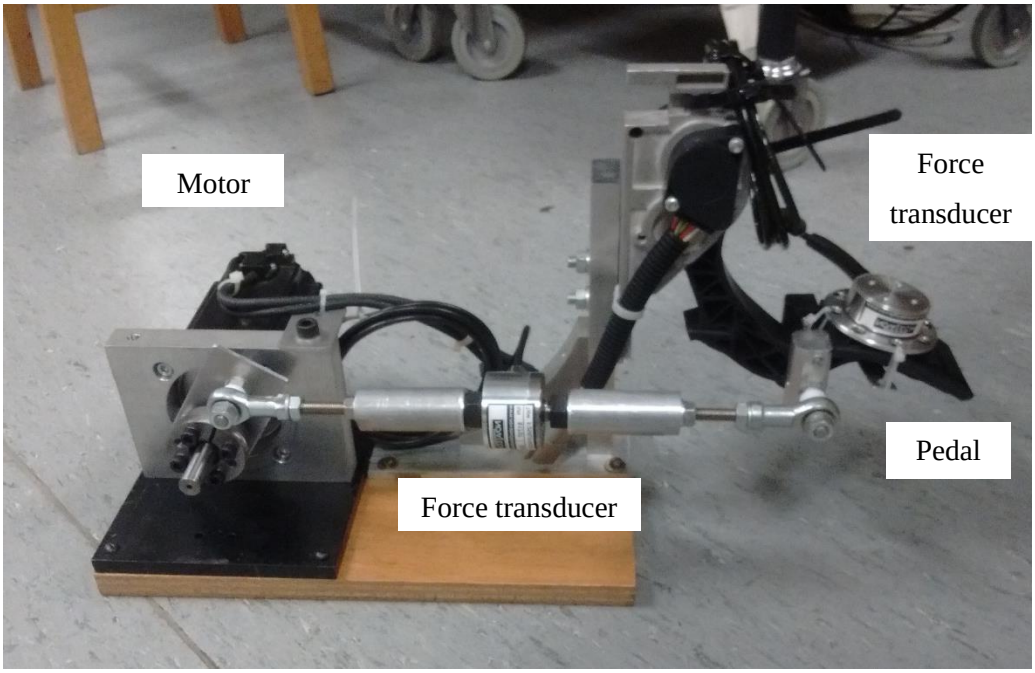

Figure 3-8 – Pedal assembly including motor and linkages

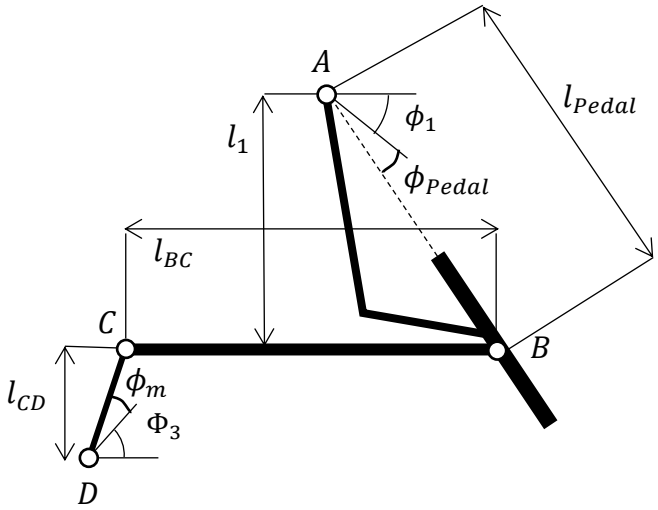

Figure 3-9 – Pedal mechanism diagram

The original pedal forms linkage *AB*, the motor lever arm forms linkage *CD* and the connecting rod between the motor and pedal forms linkage *BC*. Linkage *AB* has a linearised stiffness coefficient, κ, damping coefficient C and moment of inertia I_{AB}, all about A. C and κ are linearised parameters determined from experimental data. Linkage *CD* has the moment of inertia I_{CD} and the effects of damping and stiffness are assumed negligible. Due to the relative lengths l_{Pedal} and l_{CD} to the length l_{BC}, the linkage *BC* remains within 0.06 rad of horizontal at all times (measured experimentally). The vertical travel of the linkage (5mm) is also much smaller than the

horizontal travel (50mm). Linkage BC is therefore modelled as a mass constrained to move horizontally. As the rotation is small, the moment of inertia of this rod becomes irrelevant, as inertial forces are dominated by the rod's inertia in the horizontal directions. A D'Alembert force balance equation on the linkages can identify the effective moment of inertia of the whole system, as experienced at the pedal.

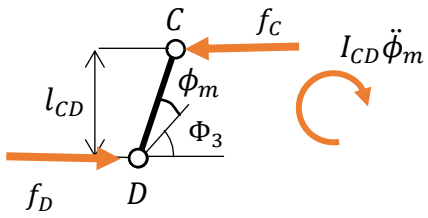

Figure 3-10 – Free body diagram of linkage CD

Summing moments about point D of Figure 3-10:

$$I_{CD}\ddot{\phi}_m = f_C l_{CD} \quad (3.6)$$

Figure 3-11 – Free body diagram of linkage BC

Examining horizontal equilibrium for Figure 3-11:

$$f_C + m_{BC}a_{BC} = f_B \quad (3.7)$$

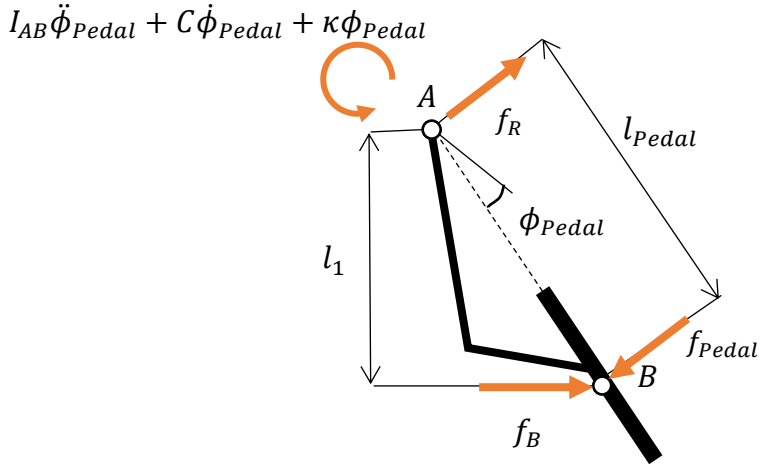

Figure 3-12 – Free body diagram of linkage *AB*

Figure 3-12 illustrates the forces acting on the core pedal. The forces from the feedback assembly are applied at *B*, and the foot forces are applied at the centre of the pad. At the other end of the pedal, the dynamic pedal torque is applied, and a reaction force at the pivot point. Summing moments about point A of Figure 3-12:

$$I_{AB}\ddot{\phi}_{Pedal} + C\dot{\phi}_{Pedal} + \kappa\phi_{Pedal} + l_1 f_B = l_{Pedal} f_{Pedal} \quad (3.8)$$

Combining (3.6) to (3.8):

$$l_1\left(\frac{I_{CD}\ddot{\phi}_m}{l_{CD}} + m_{BC}a_{BC}\right) + I_{AB}\ddot{\phi}_{Pedal} + C\dot{\phi}_{Pedal} + \kappa\phi_{Pedal} = l_{Pedal} f_{Pedal} \quad (3.9)$$

As the linkage *BC* is modelled as a rigid rod, is horizontal, and assumed to be constrained to move horizontally, the speed of point *B* is equal to the speed of point *C*:

$$l_{CD}\dot{\phi}_m = l_1\dot{\phi}_{Pedal}$$

which rearranges to:

$$\dot{\phi}_m = \frac{l_1}{l_{CD}}\dot{\phi}_{Pedal} \quad (3.10)$$

Differentiating (3.10) with respect to time produces:

$$\ddot{\phi}_m = \frac{l_1}{l_{CD}} \ddot{\phi}_{Pedal} \quad (3.11)$$

The horizontal displacement, in the direction of increasing pedal displacement, of the centre of mass of *BC* is given by:

$$x_{BC} = l_1 \phi_{Pedal} \quad (3.12)$$

Differentiating (3.12)with respect to time produces:

$$a_{BC} = l_1 \ddot{\phi}_{Pedal} \quad (3.13)$$

Now, substituting (3.11) and (3.13) into (3.9) gives:

$$l_1 \left(\frac{I_{CD} \frac{l_1}{l_{CD}} \ddot{\phi}}{l_{CD}} + m_{BC} l_1 \ddot{\phi}_{Pedal} \right) + I_{AB} \ddot{\phi}_{Pedal} + C \dot{\phi}_{Pedal} + \kappa \phi_{Pedal} = l_{Pedal} f_{Pedal} \quad (3.14)$$

which rearranges to:

$$\left(I_{AB} + \left(\frac{l_1}{l_{CD}} \right)^2 I_{CD} + l_1^2 m_{BC} \right) \ddot{\phi}_{Pedal} + C \dot{\phi}_{Pedal} + \kappa \phi_{Pedal} = l_{Pedal} f_{Pedal} \quad (3.15)$$

Hence

$$I_{Pedal} = I_{AB} + \left(\frac{l_1}{l_{CD}} \right)^2 I_{CD} + l_1^2 m_{BC} \quad (3.16)$$

It is impractical to disassemble the pedal assembly to weigh the linkage *AB*, including the pad of the pedal, so it's mass is estimated as third of the total mass of the pedal assembly (excluding any additional linkages applied). This results in a mass estimate of 0.30kg. The cross section of the

linkage *AB* changes significantly along its length. Because of this, the mass is considered equally distributed along its length when the pad of the pedal and the force transducer are taken into account. For the sake of this model, although not strictly accurate due to the changing cross sections, linkage *AB* is modelled as a rod of length 0.16m and mass 0.30kg in order to take into account the pad of the pedal. This corresponds to a moment of inertia of 0.00256kg m^2. The remaining values are summarised in the Table 3-2.

Table 3-2 – Pedal Assembly Parameters

Parameter name	Symbol	Numerical value	Units
Pedal arm length	l_{Pedal}	0.16	m
Pedal lever arm length	l_1	0.14	m
Connecting rod length	l_{BC}	0.33	m
Motor lever arm length	l_{CD}	0.04	m
Pedal arm moment of inertia	I_{AB}	2.56×10^{-3}	kg m^2
Motor moment of inertia	I_{CD}	8.53×10^{-5}	kg m^2
Connecting rod mass	m_{BC}	0.4	kg
Pedal reference angle	Φ_1	0.50	rad
Motor reference angle	Φ_2	1.05	rad

Using the values from Table 3-2 produces an effective moment of inertia of the pedal assembly, $I_{Pedal} = 0.0596$ kg m^2.

The stiffness and inertia of the pedal assembly has now been measured or estimated, but the damping is not yet quantified. To estimate the damping of the assembly, data is needed to relate the pedal force to the pedal motion.

The accelerator pedal has an inbuilt potentiometer used to provide a pedal displacement signal. The voltage varies linearly with pedal displacement. A force transducer was installed on the pad of the pedal to measure the force between the driver's foot and the pedal. As already described, the pedal is connected via a mechanical linkage to a motor for the application of pedal force-feedback at a later date. This linkage has a second force transducer built into it to measure the feedback forces in later experiments. To calibrate the force measurements, known masses were applied in compression to the pedal force transducer. Masses were applied in tension on the

feedback force transducer, and the transducer was assumed to have the same calibration factor in compression. The force calibration data is illustrated in Figure 3-13.

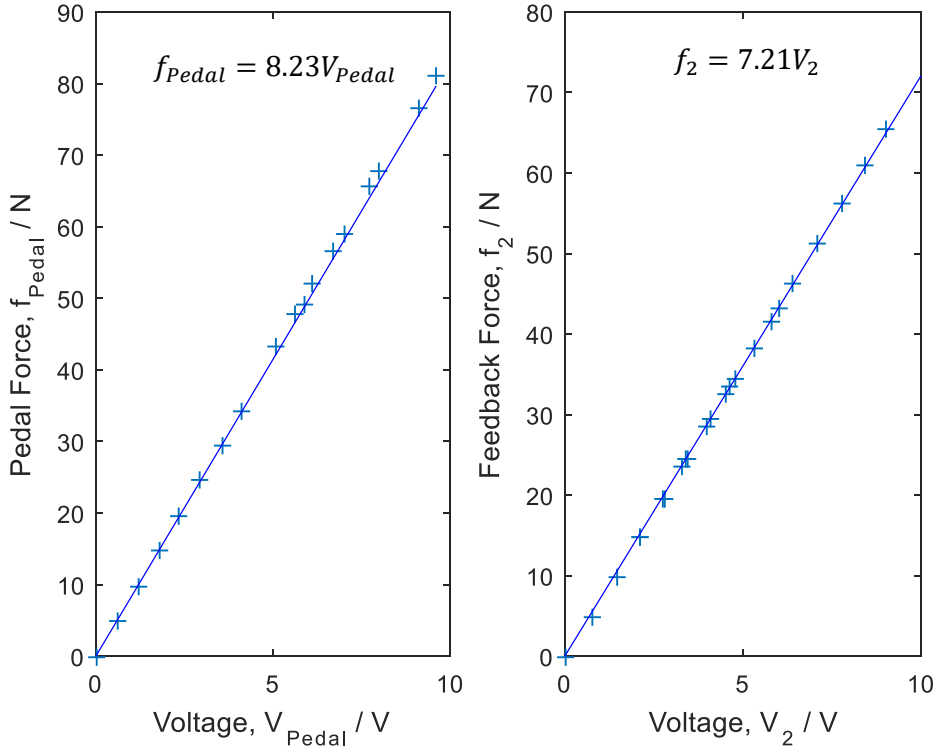

Figure 3-13 – Pedal and feedback force transducer calibration

A set of data from a drive cycle task (see Section 3.4.4) was selected to represent the motion of the pedal in normal use and the force on the pad of the pedal was recorded, along with the pedal displacement. Values of the linear inertia, stiffness and damping were optimised to minimise the error between the recorded force-displacement relationship and the modelled relationship. The recorded pedal force was applied to the model pedal and the squared errors between the modelled pedal displacement and velocity and the recorded ones were calculated. The Matlab function fmincon was used to search for the three values that produce the minimum weighted square error:

$$\Lambda = \lambda_1 \sum (\phi_{measured} - \phi_{modelled})^2 + \lambda_2 \sum (\dot{\phi}_{measured} - \dot{\phi}_{modelled})^2 \quad (3.17)$$

where the weightings λ_1 and λ_2 were defined as the reciprocal of the RMS measured pedal displacement, $\phi_{measured}$, and pedal velocity, $\dot{\phi}_{measured}$, respectively.

All three parameter values were constrained to be positive. The initial value of inertia was I_{Pedal} from earlier, the initial value of stiffness was taken from the linear gradients of Figure 3-7 and the initial value of the damping factor was set to zero, as no estimates are easily acquired.

The inertia and stiffness were allowed to vary in this optimisation, even though they were already estimated. In the case of the stiffness, this is because the recorded pedal force-displacement relationship was very non-linear and a linear approximation of the relationship was sought. Simply taking the gradient of the constant stiffness region of the static curve may not be the best approximation to the dynamic curve. The optimal values are as recorded in Table 3-3 below.

Table 3-3 – Effective Pedal Assembly Parameters

Parameter name	Symbol	Numerical value	Units
Effective pedal moment of inertia	I_{Pedal}	2.96×10^{-9}	$\mathrm{kg\,m^2}$
Effective pedal damping coefficient	C	0.219	$\mathrm{Nms\,rad^{-1}}$
Effective pedal stiffness	κ	18.9	$\mathrm{Nm\,rad^{-1}}$

The effective pedal stiffness calculated here is approximately three and a half times greater than the static value (the gradient of the hysteresis loop). This is because of the large friction forces experienced in the pedal. The forces in the real pedal are a sum of the linear stiffness, and the non-linear friction forces. This optimisation has tried to approximate this over the used region of pedal displacement and speeds.

The effective pedal moment of inertia, on the other hand, is significantly smaller than the estimated value. This is because the inertia acts on an acceleration, and acceleration is not measured directly, but differentiated twice from the recorded pedal displacement signal. The result is a very noisy acceleration signal. This noise is not replicated in the measured force signal, so the optimisation algorithm sets the inertia to a low value. This value is much lower than expected, so the initial estimated value is used instead. The same problem applies to the velocity data and damping, but to a lesser extent as the displacement data is only integrated once.

Figure 3-14 illustrates the recorded torque-displacement relationship for the static and dynamic cases, alongside the modelled torque-displacement relationship using the parameters above. It is interesting to note that dynamic cases looks slightly different to the static test case. In the dynamic test, there is a slight offset in torque, but it is also noted that there is a slight stiffening effect with displacement. This is because as the pedal is depressed, the contact patch between the

foot and the pedal shifts slightly towards the pivot point of the pedal. This effect is small, so is not incorporated into the model.

Due to the fact that the foot would lose contact with the pedal at negative forces, the pedal measured force is considered to be constrained to be positive. Constraints are not placed on the simulated pedal position, however, as the forces are expected to be dominated by the pedal stiffness, which will correspond to the constraint.

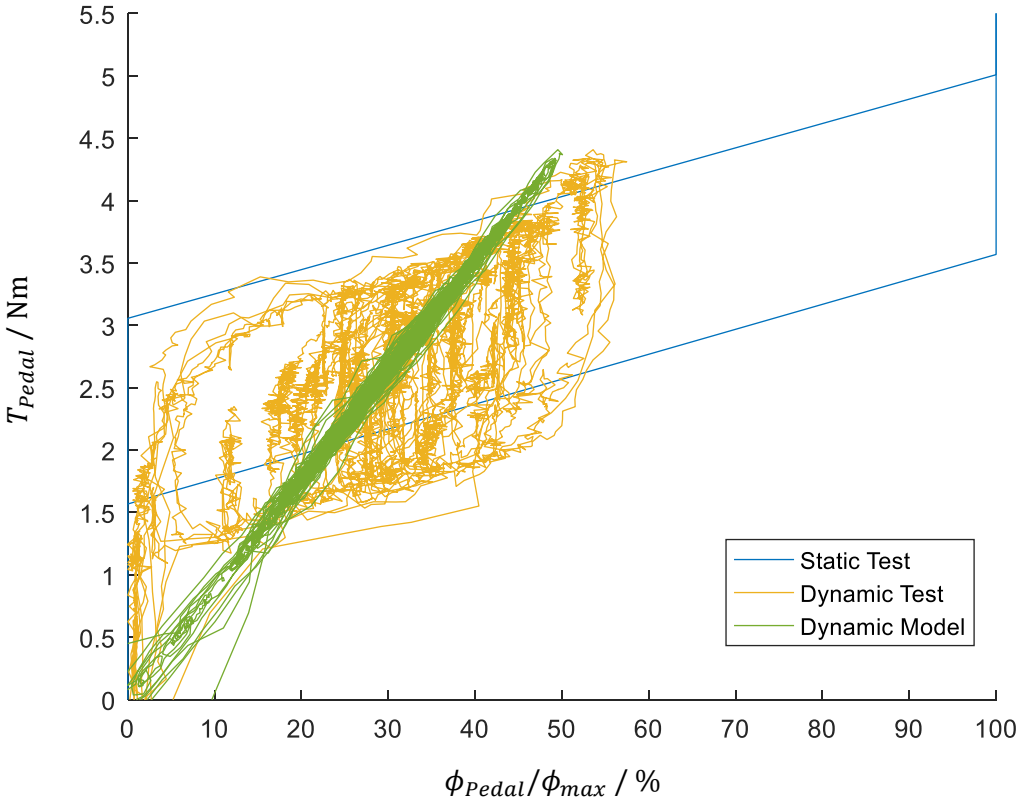

Figure 3-14 – The recorded and modelled force displacement relationship for the accelerator pedal

The model can now be expressed in equation form:

$$I_{Pedal}\ddot{\phi}_{Pedal} + C\dot{\phi}_{Pedal} + \kappa\phi_{Pedal} = T_{Pedal} \quad (3.18)$$

and solved using standard second order differential equation techniques.

The pedal torque can now be related to the force exerted on the pedal:

$$f_{Pedal} = \frac{T_{Pedal}}{l_{Pedal}} \quad (3.19)$$

It is worth noting that this is just one way of linearising the pedal dynamics. It is also possible to use a non-linear model and linearise the model at every time step – this is examined in Appendix 2.

With the addition of the pedal and driver blocks, the block diagram first illustrated in Figure 3-2 is now complete and illustrates the driver acting in closed loop - Figure 3-15.

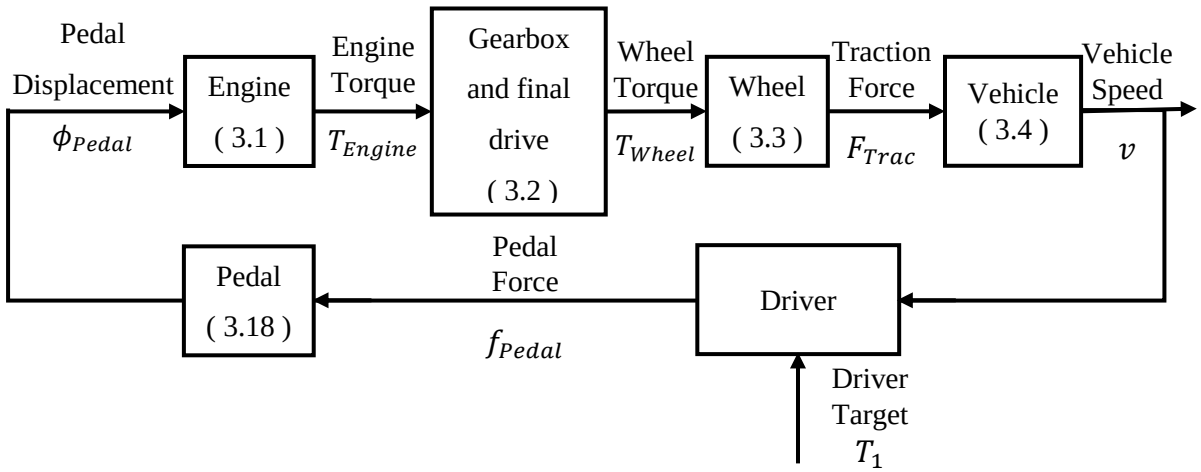

Figure 3-15 – Closed loop vehicle control block diagram. Relevant equation numbers are included in brackets

The combined linear vehicle and pedal model (equations (3.1) to (3.4) and (3.18)) can then be expressed in matrix form:

$$\dot{\boldsymbol{x}} = \mathbf{A}_C \boldsymbol{x} + \mathbf{B}_C f_{Pedal} \qquad (3.20)$$

where

$$\boldsymbol{x} = \begin{Bmatrix} \phi_{Pedal} \\ \dot{\phi}_{Pedal} \\ v \\ y \\ 1 \end{Bmatrix},$$

$$\mathbf{A}_C = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ -\frac{\kappa}{I_{Pedal}} & -\frac{C}{I_{Pedal}} & 0 & 0 & 0 \\ \frac{1}{\phi_{Max}} \times \frac{T_{EngineMax} G_{Gear} G_{FD}}{r_{Wheel} M} & 0 & -\frac{1}{2}\frac{\rho_{air} C_D A k}{M} & 0 & -g\sin\theta + C_r \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$B_C = \begin{bmatrix} 0 \\ \frac{r_1}{I_{Pedal}} \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

3.2 Linear Model Predictive Control Theory

The driver model is developed using Model Predictive Control Theory. The objective of the driver model is to mimic the driver's cognitive interpretation of current vehicle speed and target speed in order to produce a plan of pedal forces to reach the target speed, exploiting the driver's knowledge of the vehicle.

In linear Model Predictive Control, a linear internal vehicle and pedal model is used, and then, taking the current vehicle state and a previewed speed demand, a control action is calculated. The preview time for a real driver is very dependent on the road conditions, the number of other road users present for example. Conventionally, the prediction horizon is set as a time, meaning the driver is aware of the speed demand for the next t seconds. The controller calculates control actions up until the control horizon. In this case, the prediction horizon and the control horizon were set equal. The control actions are determined by minimising a cost function. For a longitudinal control scenario, the speed error, pedal displacement and pedal force are suitable variables for the cost function as they can be interpreted as a function of accuracy (speed error), energy consumption (pedal displacement is used an indicator) and driver effort (pedal force). The cost function used was defined as a weighted sum of the mean square pedal force, mean square pedal position and mean square speed error. The optimum pedal force control action up to the end of the control horizon was calculated using a least squares solution. It is the weightings of the variables in the cost function that is likely to result in the variation in driving styles observed in human drivers. Weighting values will be sought experimentally.

Due to the linearity of the vehicle model and that fact that parameters remain constant, a set of state gains for a set preview can be calculated and then applied again without needing any recalculation. Once these gains are known, they are applied to the current states and preview information, control actions decided, and then only the first time step of control is used, before the process is repeated on the next time step. The derivation of the control law is included below.

The controller works in discrete time, so the vehicle model must also be in discrete time steps. The model discussed in the previous section, (3.20), becomes:

$$\boldsymbol{x}(k+1) = \mathbf{A}\boldsymbol{x}(k) + \mathbf{B}f_1(k) \qquad (3.21)$$

where **A** and **B** are the discrete time state-space matrices, $\mathbf{x}(k)$ the state vector at time step k and f_1 is the driver pedal force. Notice that this is different to f_{Pedal} from the previous section. This subtlety of notation is defined to highlight the difference between contributions to total pedal force from the human driver and the contributions from the pedal force feedback controller in Chapter 4. The outputs of the vehicle model are given the symbol $\mathbf{z}(k)$ and are defined as

$$\mathbf{z}(k) = \mathbf{C}\mathbf{x}(k) \tag{3.22}$$

where **C** is a matrix mapping the states, $\mathbf{x}(k)$, to the desired outputs, $\mathbf{z}(k)$.

By iterating every time step, it is then possible to predict the state of the vehicle at every time step up until the prediction horizon, N_p steps ahead:

$$\begin{Bmatrix} \boldsymbol{x}(k+1) \\ \boldsymbol{x}(k+2) \\ \vdots \\ \boldsymbol{x}(k+N_p) \end{Bmatrix} = \begin{bmatrix} \mathbf{A} \\ \mathbf{A}^2 \\ \vdots \\ \mathbf{A}^{\mathrm{N_p}} \end{bmatrix} \mathbf{x}(k) + \begin{bmatrix} \mathbf{B} & 0 & \cdot & \cdot & \cdot & 0 \\ \mathbf{AB} & \mathbf{B} & 0 & \cdot & \cdot & \cdot \\ \mathbf{A^2B} & \mathbf{AB} & \mathbf{B} & 0 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \mathbf{A^{N_p-1}B} & \cdot & \cdot & \cdot & \mathbf{AB} & \mathbf{B} \end{bmatrix} \begin{Bmatrix} f_1(k) \\ f_1(k+1) \\ \vdots \\ f_1(k+N_p-1) \end{Bmatrix} \tag{3.23}$$

For the case where the prediction horizon and control horizon are equal (the controller plans control actions up until the prediction horizon), the following can be shown by following Maciejowski (2002) and Cole (2006).

$$\mathbf{Z}(k) = \mathbf{\Psi}\mathbf{x}(k) + \mathbf{\Theta}\mathbf{U_1}(k) \tag{3.24}$$

where

$$\mathbf{Z}(k) = \begin{Bmatrix} \mathbf{z}(k+1) \\ \mathbf{z}(k+2) \\ \vdots \\ \mathbf{z}(k+N_p) \end{Bmatrix} \qquad \mathbf{U_1}(k) = \begin{Bmatrix} f_1(k) \\ f_1(k+1) \\ \vdots \\ f_1(k+N_p-1) \end{Bmatrix}$$

$$\mathbf{\Psi} = \begin{bmatrix} \mathbf{CA} \\ \mathbf{CA}^2 \\ \vdots \\ \mathbf{CA}^{N_p} \end{bmatrix} \qquad \mathbf{\Theta} = \begin{bmatrix} \mathbf{CB} & 0 & \cdot & \cdot & \cdot & 0 \\ \mathbf{CAB} & \mathbf{CB} & 0 & \cdot & \cdot & \cdot \\ \mathbf{CA^2B} & \mathbf{CAB} & \mathbf{CB} & 0 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \mathbf{CA^{N_p-1}B} & \cdot & \cdot & \cdot & \mathbf{CAB} & \mathbf{CB} \end{bmatrix}$$

Carrying on with Maciejowski's approach, the next stage is to calculate a control action that optimises a chosen cost function. In this instance, a cost function, $V(k)$, of mean square pedal force and speed error was chosen.

$$V_1(k) = \sum_{i=1}^{N_p} \|\mathbf{z}(k+i) - \mathbf{r_1}(k+i)\|^2_{\mathbf{Q}(i)} + \sum_{i=0}^{N_p-1} \|f_1(k+i)\|^2_{\mathbf{R}(i)} \quad (3.25)$$

where $\|\mathbf{y}\|^2_{\mathbf{Q}(i)} = \mathbf{y}^{\mathrm{T}}\mathbf{Q}(i)\mathbf{y}$ and $\mathbf{r_1}(k+i)$ are the future speed demands for the vehicle as previewed by the driver. $\mathbf{Q}(i)$ is the cost function matrix associated with time step i ahead of the current time, and $\mathbf{R}(i)$ is the cost function matrix associated with the future pedal force values. Defining

$$\mathbb{Q}_\mathbf{1} = \begin{bmatrix} \mathbf{Q_1}(1) & & \\ & \ddots & \\ & & \mathbf{Q_1}(N_p) \end{bmatrix} \qquad \mathbb{R}_\mathbf{1} = \begin{bmatrix} \mathbf{R_1}(0) & & \\ & \ddots & \\ & & \mathbf{R_1}(N_p-1) \end{bmatrix}$$

$$\mathbf{T_1}(k) = \begin{Bmatrix} \mathbf{r_1}(k) \\ \mathbf{r_1}(k+1) \\ \vdots \\ \mathbf{r_1}(k+N_p-1) \end{Bmatrix}$$

Note the subscript 1 indicates that these apply to the driver. It is now assumed that $\mathbf{Q}(i)$ and $\mathbf{R}(i)$ are constant throughout the preview. Maciejowski states it is now possible to rewrite (3.25) without the summation:

$$V_1(k) = \|\mathbf{Z}(k) - \mathbf{T}_1(k)\|^2_{\mathbf{Q_1}} + \|\mathbf{U}_1(k)\|^2_{\mathbf{R}_1} \quad (3.26)$$

The predicted error, $\boldsymbol{\varepsilon_1}(k)$, is

$$\boldsymbol{\varepsilon_1}(\mathrm{k}) = \mathbf{T_1}(\mathrm{k}) - \boldsymbol{\Psi}\mathbf{x}\,(\mathrm{k}) \quad (3.27)$$

Substituting (3.27) into (3.26) results in

$$V_1(\mathrm{k}) = \|\boldsymbol{\Theta}\mathbf{U_1}(\mathrm{k}) - \boldsymbol{\varepsilon_1}(\mathrm{k})\|^2_{\mathbf{Q_1}} + \|\mathbf{U_1}(\mathrm{k})\|^2_{\mathbf{R_1}} \quad (3.28)$$

Rewriting:

$$V_1(k) = \left\| \begin{bmatrix} \mathbf{S_{Q_1}}\{\boldsymbol{\Theta}\mathbf{U_1}(k) - \boldsymbol{\varepsilon_1}(k)\} \\ \mathbf{S_{R_1}}\mathbf{U_1}(k) \end{bmatrix} \right\|^2 \quad (3.29)$$

where $\mathbf{S_Q^T S_Q} = \mathbf{Q}$ and $\mathbf{S_R^T S_R} = \mathbf{R}$.

In order to minimise the cost function, the least squares solution of the following is sought:

$$\begin{bmatrix} \mathbf{S_{Q_1}}\{\boldsymbol{\Theta}\mathbf{U_1}(k)_{\mathrm{opt}} - \boldsymbol{\varepsilon}_1(k)\} \\ \mathbf{S_{R_1}}\mathbf{U_1}(k)_{\mathrm{opt}} \end{bmatrix} = 0 \quad (3.30)$$

This can be solved using QR decomposition:

$$\mathbf{U}_{1\mathrm{opt}}(k) = \mathbf{K}_{1\mathrm{full}}\boldsymbol{\varepsilon_1}(k) \quad (3.31)$$

where

$$\mathbf{K}_{1\mathrm{full}} = \begin{bmatrix} \mathbf{S_{Q_1}}\boldsymbol{\Theta} \\ \mathbf{S_{R_1}} \end{bmatrix} \backslash \begin{bmatrix} \mathbf{S_{Q_1}} \\ 0 \end{bmatrix} \quad (3.32)$$

where '\' is the Matlab command for QR decomposition.

This produces a set of future pedal forces that would minimise the cost function. The first of these is taken, and then new set of optimised pedal forces, $\mathbf{U}(k+1)_{\mathrm{opt}}$ are calculated, and the process repeated. The time independence of $\mathbf{K}_{\mathrm{full}}$ means that the optimum pedal force can be determined by a linear time-invariant controller:

$$f_1(k) = \mathbf{K}_{1\mathrm{w}}\boldsymbol{\varepsilon_1}(k) \quad (3.33)$$

where $\mathbf{K}_{1\mathrm{w}}$ is the first row (corresponding to first control action) of $\mathbf{K}_{1\mathrm{full}}$.
Using equation (14), and the substitution $\mathbf{K}_{1p} = [-\mathbf{K}_{1\mathrm{w}}\boldsymbol{\Psi} \quad \mathbf{K}_{1\mathrm{w}}]$, the control law can be written as:

$$f_1(k) = \mathbf{K}_{1p}\begin{Bmatrix} \mathbf{x}(k) \\ \mathbf{T_1}(k) \end{Bmatrix} \quad (3.34)$$

3.3 Control Tasks

The real-world driving task can be complex. In one journey a single vehicle may experience empty roads with few other vehicles, or may encounter heavy traffic jams. The presence of pedestrians or other road side obstacles may also affect the driver's choice of speed in a complex manner. In this project, two simple driving scenarios are considered – drive-cycle following and car-following.

3.3.1 Drive-Cycle Following

In the first scenario, drive-cycle following, it is assumed that the driver has a target speed profile that is unaffected by external influences. Therefore, the driver is free to aim for their own target without having to respond to a changing environment.

The Millbrook Suburban Cycle for HGVs (Barlow et al., 2009) represents a typical HGV drive in a suburban area with low to medium speeds - Figure 3-16. As the purpose of this project is

to look at the use of the accelerator pedal, the drive cycle is modified slightly to eliminate the need for use of a brake pedal. This is achieved by selecting sections of the drive cycle that do not include a decelerating to stop. The identified sections are highlighted in grey. The sections are re-joined to form a new drive-cycle. The join produces a discontinuity in acceleration, but is continuous in speed demand. The acceleration discontinuity is consistent with the rest of the drive cycle. The resulting four minute drive cycle is displayed in Figure 3-17.

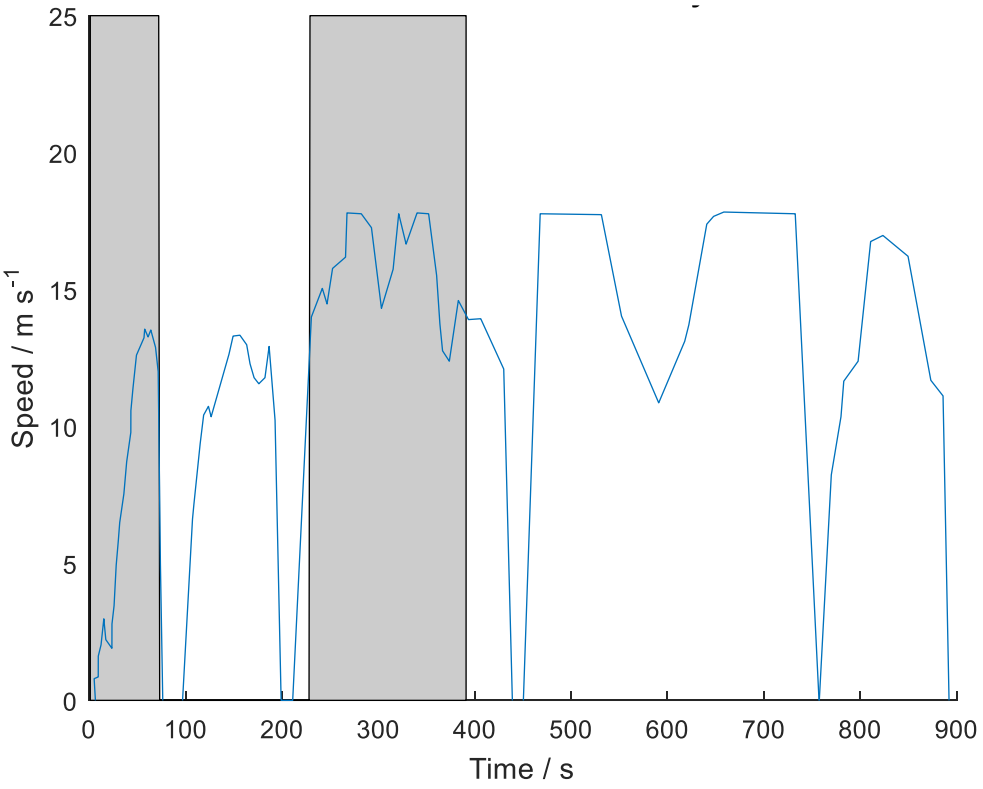

Figure 3-16 – Millbrook Suburban Drive Cycle for HGVs

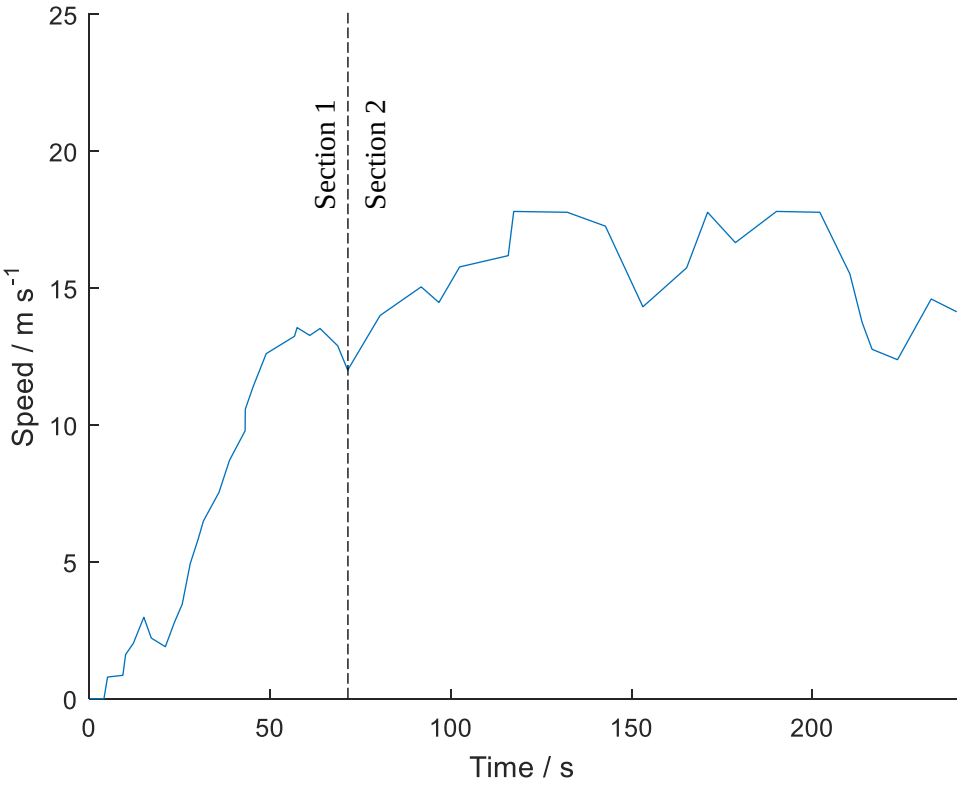

Figure 3-17 – Modified Millbrook Suburban Drive Cycle for HGVs

For a driver undergoing a drive-cycle task, their style may be determined by a cost function of performance and effort. As the target is clearly defined as a speed demand, an obvious measure of performance is the error between the vehicle's current speed and the target. There are many other measures of performance that the driver may choose to adopt, such as fuel economy. The relationship between fuel use and the driver's control of the pedal position and vehicle speed is non-linear (see Chapter 2), making it unsuitable for the current linear driver model. However, Chapter 2 also demonstrates that for a given speed, the fuel rate increases with pedal position, so pedal position is therefore used as an indicator of fuel use. The physical effort the driver has to put into the control of the vehicle is a function of the force applied on the pedal and the pedal motion. Pedal force is used as the measure of driver effort for this controller. Another potential measure of driver effort is pedal speed. The more the driver is having to move the pedal, the more effort is being required. The weightings of these variables are given the symbols q_{1s}, $q_{1\phi}$, r_1 and $q_{1\dot{\phi}}$respectively. The subscript 1 indicates these weightings are applied to the driver's cost function. As the pedal force is the output of the proposed driver model, the weighting on this is given the symbol r_1.

This chapter will examine a variety of potential combinations of these variables and will seek to identify a suitable cost function to represent driver behaviour once simulator experiments are analysed. To encompass all variable discussed, the final model predictive control matrices, $\mathbf{C}, \mathbf{Q_1}(i)$ and $\mathbf{R_1}(i)$ are hence defined as

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad (3.35)$$

$$\mathbf{Q_1}(i) = \begin{bmatrix} q_{1s} & 0 & 0 \\ 0 & q_{1\phi} & 0 \\ 0 & 0 & q_{1\dot{\phi}} \end{bmatrix} \quad (3.36)$$

$$\mathbf{R_1}(i) = r_1 \quad (3.37)$$

This scenario will allow the comparison of the driver's trade-off between a clearly defined target and the effort put in by the driver.

3.3.2 Car-Following

The second driving scenario modelled is the car-following scenario. In a lot of real-world driving, the driver's speed choice is very dependent on other road users around them. The driver's target is therefore determined by their own speed targets and the behaviour of other vehicles. The aim of analysing this scenario was to evaluate the driver's trade-off between a more complicated target and driver effort.

In this scenario, the driver may have a cost function trading off matching the speed of the vehicle in front, maintaining a safe following distance (both indicators of performance), pedal displacement (as an indicator of fuel use again), pedal force (as an indicator of driver effort) or pedal speed. The weightings of these three variables are given the symbols q_{1s}, $q_{1\phi}$, q_{1d}, r_1 and $q_{1\dot{\phi}}$ respectively. The subscript 1 indicates these weightings are applied to the driver's cost function.

The safe following distance is dependent on the speed of the driver's vehicle. A common method for determining the safe following distance is to use a constant time headway (THW) method (Mulder et al., 2011). This determines the safe following distance target to be the distance travelled by the following vehicle in a set period of time. This method breaks down at low speeds, where the following distance will fall to near zero. Therefore, it is a good idea to add a constant distance offset to the target to eliminate low speed collisions. The safe following distance target is therefore defined as:

$$x_{THW} = vt_{THW} + l_{THW} \quad (3.38)$$

For the purposes of the driver model, the controller matches the target speed to the vehicle in front, and matches the displacement of the vehicle in front to the safe following distance added

to the following vehicle's displacement. The final model predictive control matrices, $\mathbf{C}, \mathbf{Q_1}(i)$ and $\mathbf{R_1}(i)$ are therefore defined as

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & t_{THW} & 1 & l_{THW} \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad (3.39)$$

$$\mathbf{Q_1}(i) = \begin{bmatrix} q_{1s} & 0 & 0 \\ 0 & q_{1\phi} & 0 \\ 0 & 0 & q_{1d} \end{bmatrix} \quad (3.40)$$

$$\mathbf{R_1}(i) = r_1 \quad (3.41)$$

3.4 Driving Experiments and Data Analysis

The longitudinal control driver model was validated in the CUED driving simulator using both the drive-cycle following and car following scenarios.

3.4.1 Experiment Design

The CUED driving simulator has been used for a variety of experiments since its conception in 2003. At the start of this project, its software and hardware were upgraded by the author of this thesis to reflect improvements in technology.

The simulator has a fixed base and three 4K 65 inch screens providing the driver with a 120° field of view. The seat is positioned so that the vertical distance between base of the seat and the pad of the pedal is the same as that measured in a Volvo truck cab. The horizontal distance between the seat and the pedal is also set to match the Volvo cab, taking into account that both the Volvo seat, and simulator seat are adjustable in the longitudinal direction of the vehicle. No other controls are available to the driver. Figure 3-18 illustrates the layout of the driving simulator.

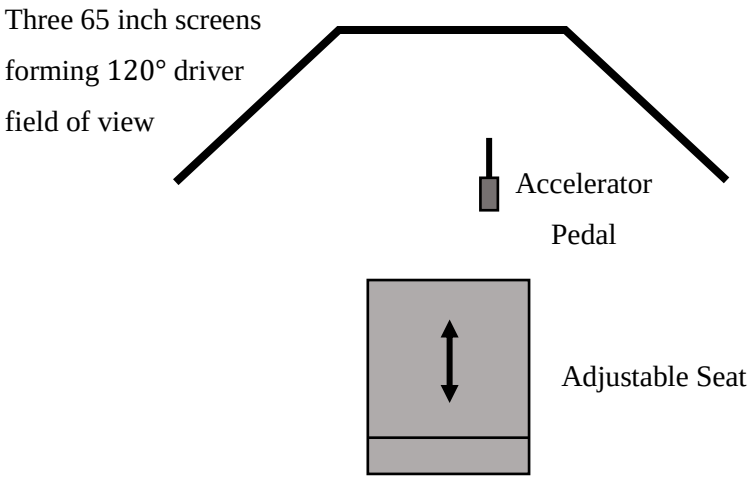

Figure 3-18 – Driving simulator plan view

The simulator operates using Simulink Real-Time over three computers – the Host PC, Target PC, and Audio PC. The Target PC operates the vehicle model in real time and interfaces with the analogue sensors; the Host PC produces the graphical interface of the simulator and the Audio PC produces an engine sounds audio track to provide extra speed feedback to the driver.

In the Drive-Cycle following task, only the centre screen is used. The driver is presented with a graph previewing the next four seconds of speed demand and illustrating the past four seconds of speed demand and achieved speed.

In the car-following scenario (Figure 3-19), all three screens are used. A virtual driving environment was created using Simulink 3D Animation. The virtual road consists of a straight road with houses lining either side. The wide field of view assists with the driver's perception of speed.

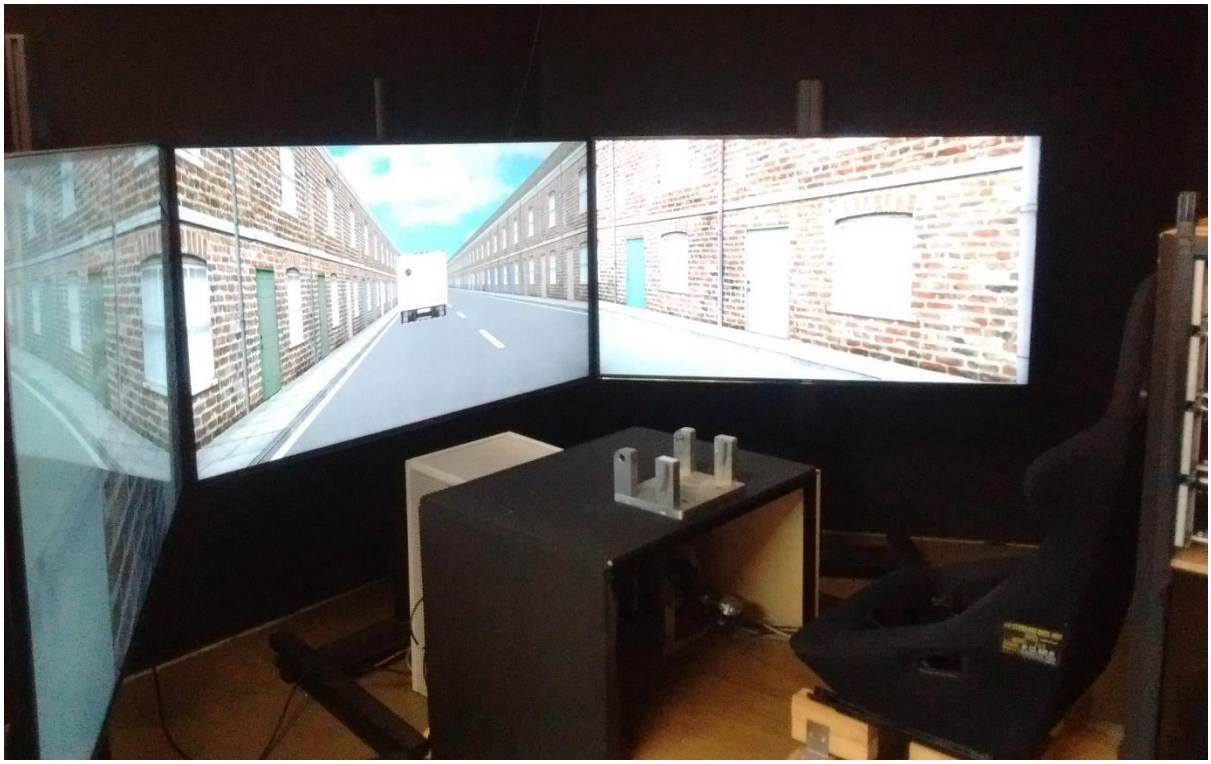

Figure 3-19 – Driving simulator set up for a car following scenario

The Target PC records the three pedal signals, pedal position, pedal force, and feedback force, at 50Hz. The simulated vehicle speed and displacement are also recorded at 50Hz. These signals are fed to the Host PC for storage.

The noise synthesizer's purpose is to produce an additional feedback loop to the driver, incorporating engine and vehicle information through the frequency spectrum and amplitude of engine and vehicle sounds. Specifically, engine noise and tyre noise and recreated, with the aim of improving the driver's perception of speed. The Host PC sends on the simulated vehicle speed and recorded pedal positon via Ethernet connection to the Audio PC.

At 0.02s intervals, the Audio PC assesses the speed and pedal position from the Host PC. In order to represent the engine and tyre noise of the vehicle, two key frequencies, the engine noise centre frequency and tyre noise centre frequency, are then calculated using equations (3.42) and (3.43) below:

$$f_{eng} = \frac{v G_{Gear} G_{FD}}{r_{Wheel}} \quad (3.42)$$

$$f_{tyre} = \frac{v}{l_{tread}}$$

(3.43)

where l_{tread} is a characteristic length of the tyre tread. For the trucks modelled here, the tyre tread length is set to 5cm.

Two tenth order band pass Butterworth filters are defined to produce engine and tyre noise. The centre frequencies from (3.42) and (3.43) are used and each filter has -3dB cut-off frequencies at 90% and 110% of the centre frequency. Each filter is then applied to a 0.2s clip of white noise, before the two clips are summed in the time domain.

A gain is applied to define the amplitude of the new signal as

$$A_{noise} = \frac{1}{2}\left(\frac{\phi_{pedal} v}{\phi_{max} v_{max}} + 1\right)$$

(3.44)

Figure 3-20 illustrates the complete sound process in block diagram form. The noise signal is output to a set of headphones to provide the driver with audio feedback.

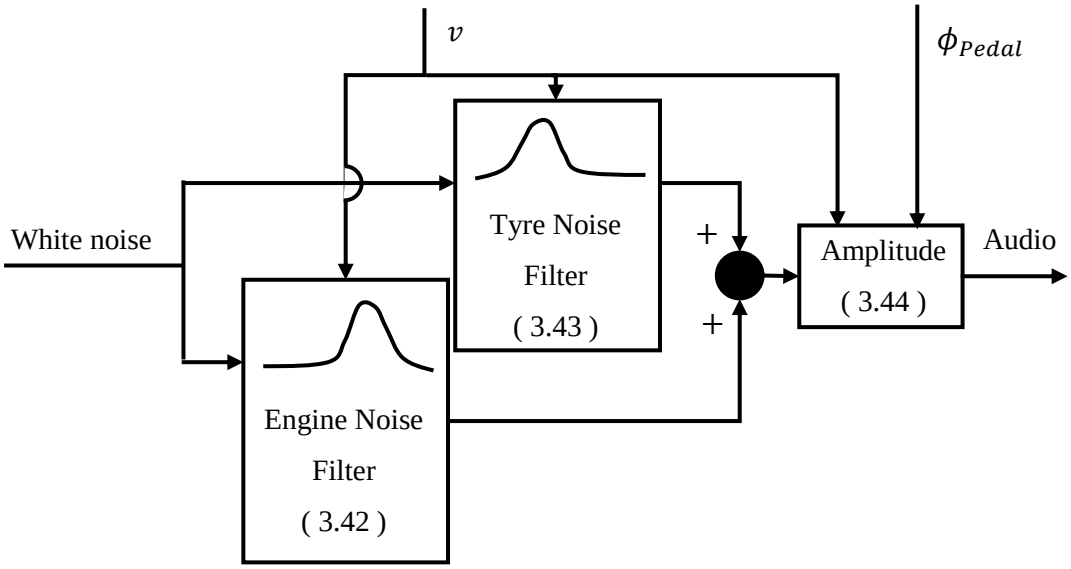

Figure 3-20 – Audio generation in the driving simulator. Equation numbers are included for the centre frequencies of the noise filters, and of the amplitude function

3.4.2 Participants

Nine human drivers were selected to participate in the validation of the driver model. The drivers were between 23 and 55 years old and all held driving licences. One out of the nine drivers were

female. Drivers 1 to 7 were non-HGV drivers, and drivers 8 and 9 were professional, and experienced, HGV drivers.

3.4.3 Procedure

The experiment was divided into two halves – the drive cycle scenario and the car-following scenario.

In the first scenario, drive cycle following, the Modified Millbrook Suburban Drive Cycle for HGVs was set as the target speed. The driver was given a practice run through to familiarise themselves with the behaviour of the vehicle and the dynamics of the pedal. Drivers were asked to follow the target speed. The test was then completed three times, to improve the reliability of the results.

In the second scenario, car following, the speed of the target vehicle was set as the Modified Millbrook Suburban Drive Cycle for HGVs. Again, the driver was allowed a familiarisation run to get used to the virtual world. Drivers were asked to follow the vehicle in front at a safe distance, as they normally would on the road. Again, once the familiarisation run was completed, the test was repeated three times, resulting in four runs in total.

When all drivers had completed the experiment, the data was analysed and driver behaviours compared.

3.4.4 Drive Cycle Results

Figure 3-21 illustrates the recorded mean speed errors from the drivers. Generally, the drivers maintain a speed quite close to the target speed profile. After high accelerations, however, some overshoot tends to occur. At lower acceleration rates, the drivers were better at anticipating the future speed demand. At speed peaks, the human drivers tended to keep accelerating until the maximum speed was reached, and then coasted down until they reached the target speed again.

Figure 3-22 illustrates the drivers' control of the accelerator pedal. For the vast majority of the time, the drivers are able to accurately follow the target speed, without depressing the pedal beyond 50%. There are definite trends apparent in the pedal position against time plot. Examining the pedal position more closely, the drivers' intermittent control becomes apparent (Johns et al., 2015). This is where drivers update their control actions at discrete points and maintain near constant pedal force between those times.

Comparing the pedal position with pedal forces in Figure 3-23 shows that there are higher frequency elements to the driver pedal force. The friction in the pedal force damps out a lot of these higher frequencies, resulting in the more apparent noise in the pedal position. It was noted post experimentation that there was potentially a small cable-tie intruding over the force transducer, and that the forces measured, may not be accurate over all pedal positions below approximately 15%. For this reason, later analysis will focus on the pedal position, rather than pedal force.

Figure 3-24 demonstrates the costs for the drivers in the drive-cycle task. Surprisingly, there is a near linear trend of increasing RMS speed errors with increasing RMS pedal forces. However, the variation in RMS pedal forces, only has a small impact on the RMS speed errors, so this results may be due to the different drivers being better or worse at locating the optimal weightings of their cost functions.

In this scenario, there is little to distinguish the professional HGV drivers from the non-professionals.

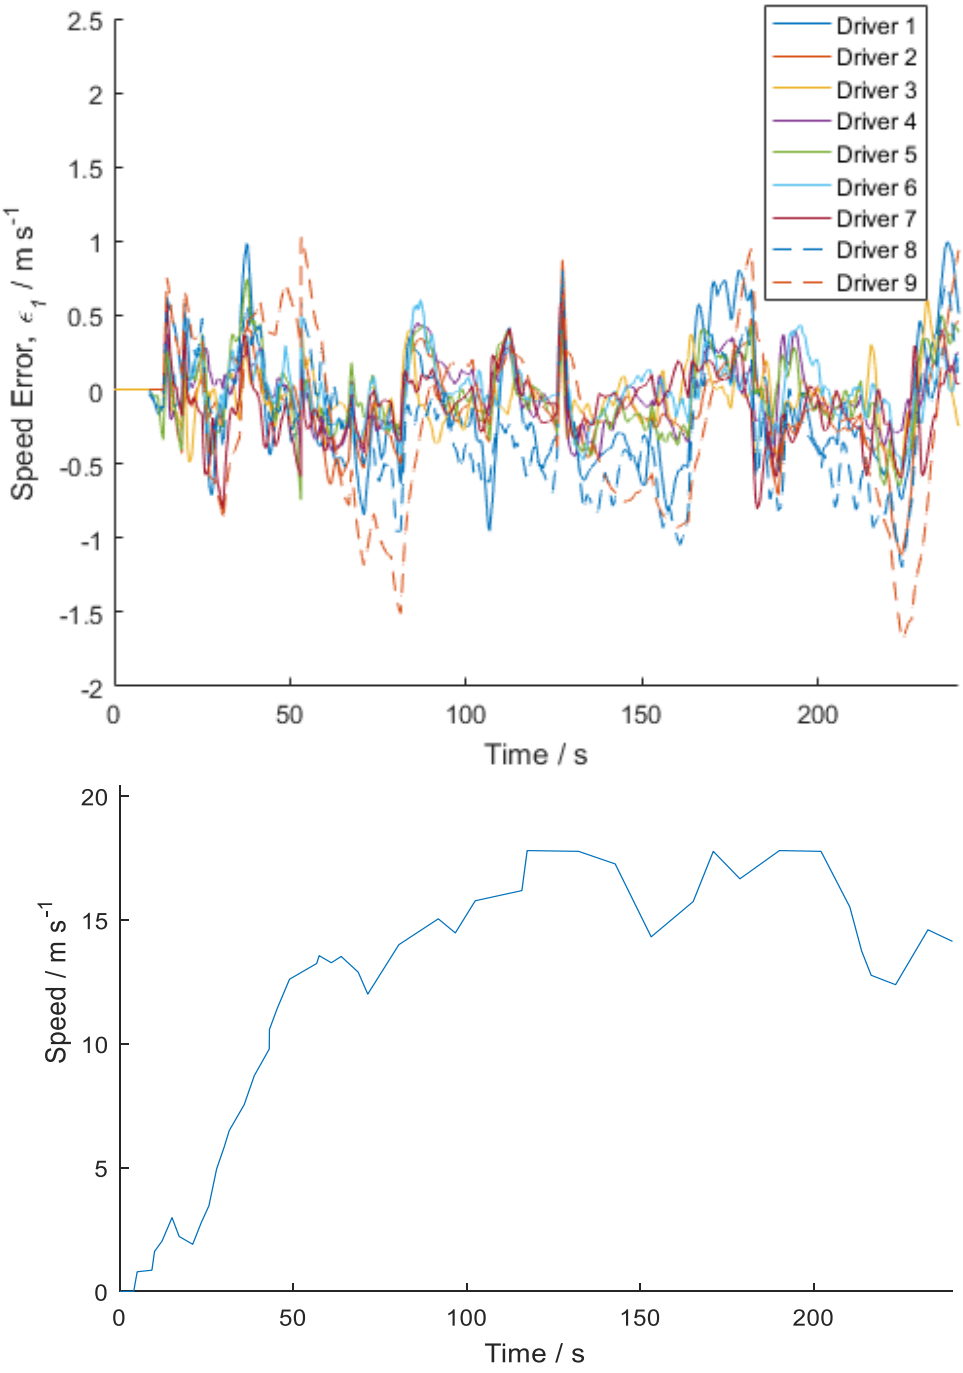

Figure 3-21 – Mean vehicle speed error against time for the drive cycle following task (top) and drive cycle speed profile for reference (bottom)

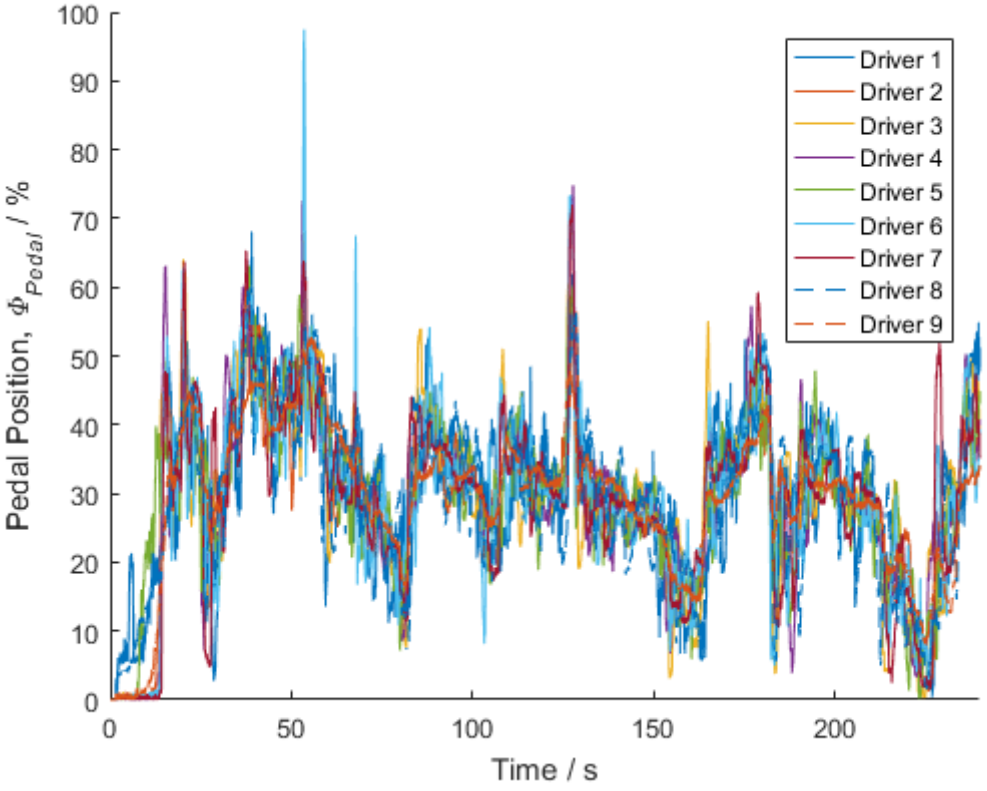

Figure 3-22 – Mean pedal position against time for the drive cycle following task

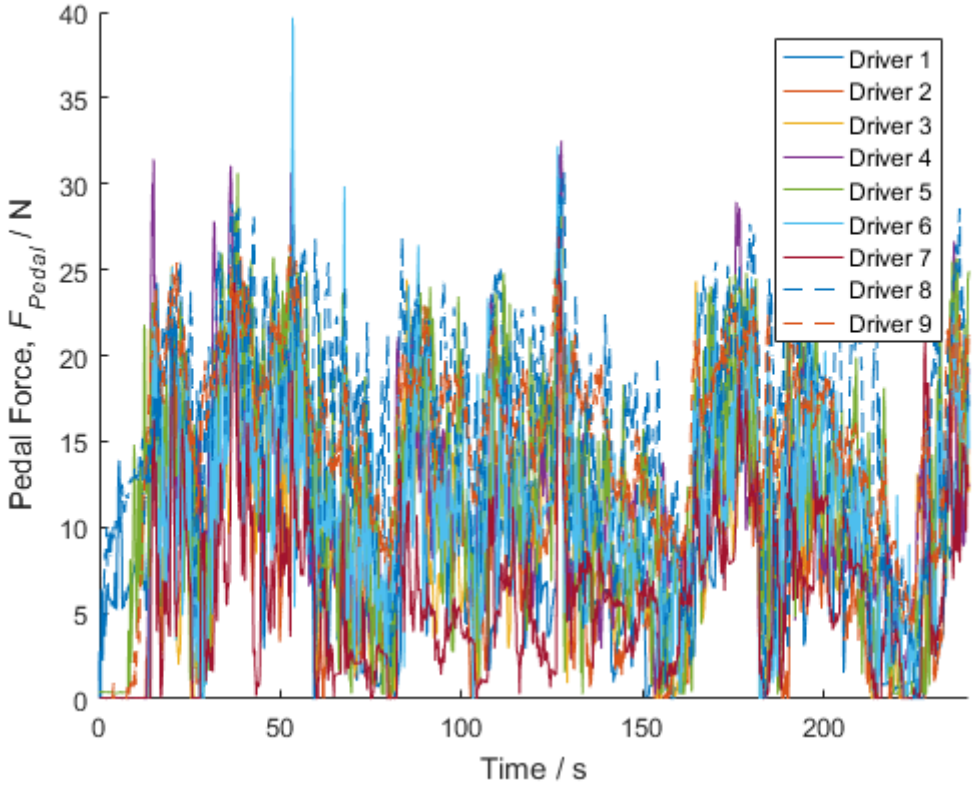

Figure 3-23 - Mean pedal force against time for the drive cycle following task

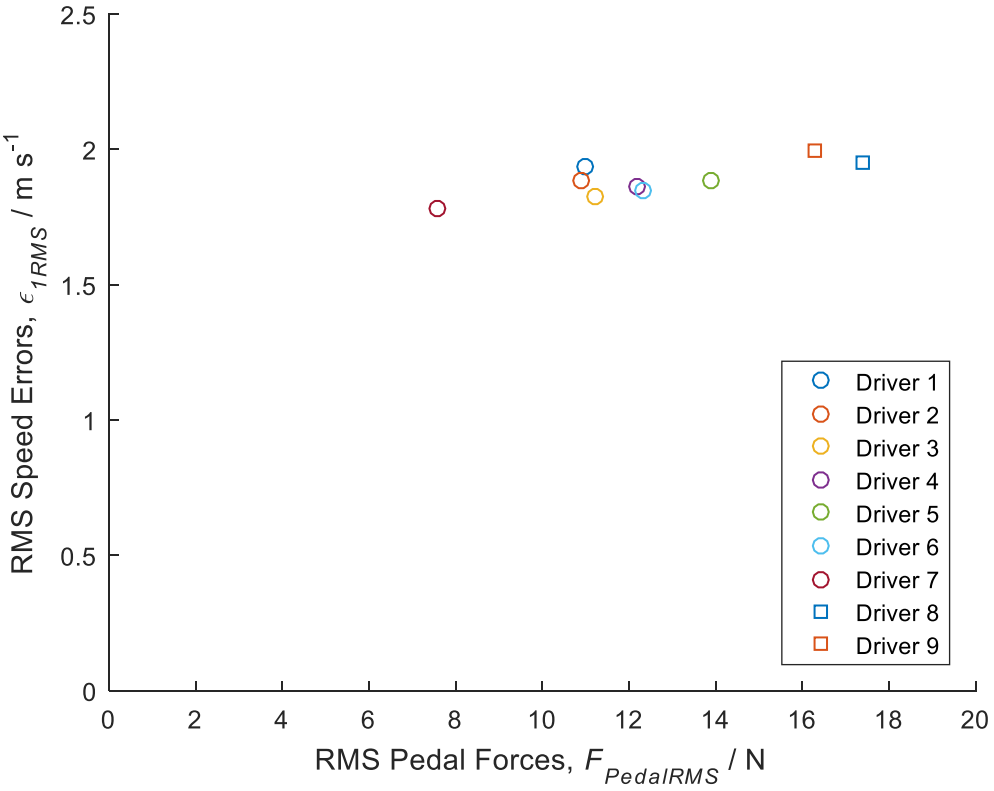

Figure 3-24 – Cost diagram for drive cycle scenario

It is interesting to note that it is the professional drivers that have the highest RMS pedal forces, $F_{PedalRMS}$ of all the drivers. This could reflect a difference in styles observed between truck and car drivers, or may be coincidental in a small sample size such as this one.

As a final consideration, the mean vehicle speed for each driver is plotted with one standard deviation of the individual's runs shaded either side in Figure 3-25. This clearly demonstrates that the standard deviations amongst each driver is similar to the standard deviation across drivers in this scenario, meaning that there is not a statistically significant difference between most of the drivers across most of the test.

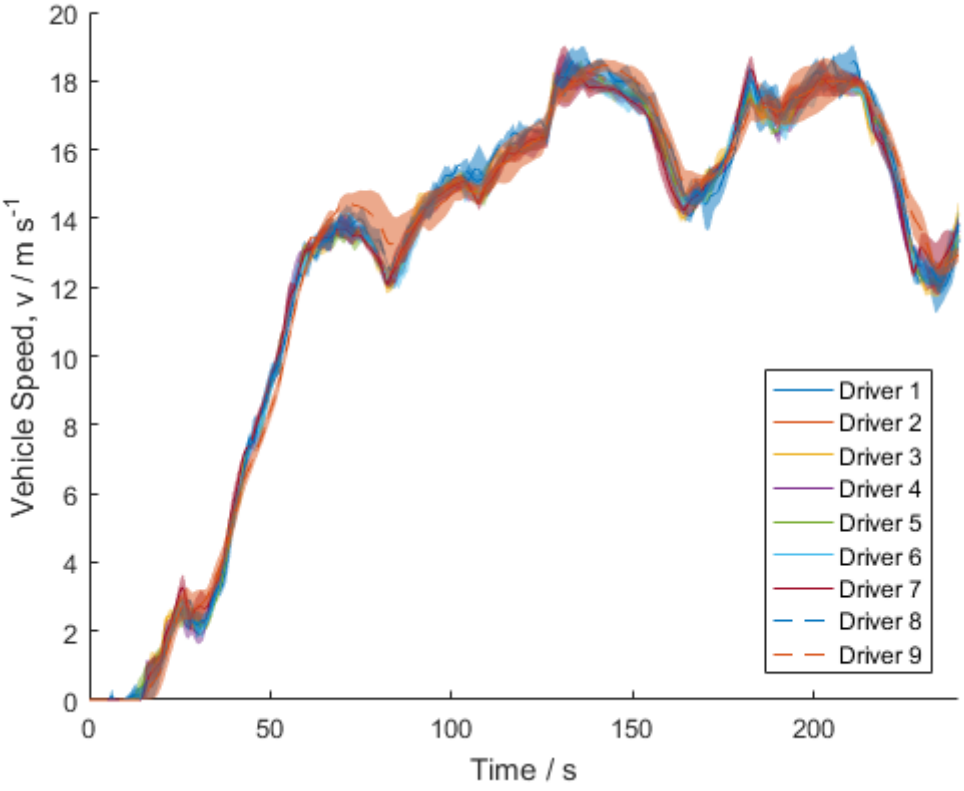

Figure 3-25 – Driver mean vehicle speeds with one driver standard deviation shaded either side

3.4.5 Car-Following Results

Figure 3-26 illustrates the recorded mean (across all three runs) speed profiles from the drivers alongside the speed of the target vehicle. Generally, the drivers maintain a similar speed to the target vehicle. There is more variation in the speed between drivers in this task than there was in the drive-cycle task, due mostly to the fact that the drivers were asked to follow at a safe distance, not match speeds. However, characteristic overshoots in speed are still present, as they were in the drive-cycle task.

Figure 3-27 illustrates the drivers' control of the accelerator pedal. It is apparent that the change of task from drive-cycle to car-following leads drivers to demand more from the vehicle as higher pedal positions are witnessed. The maximum pedal position in this scenario reaches 80% in Driver 6's case. Once again, general trends are still apparent between drivers.

In this scenario, the drivers' intermittent control is even more apparent. With perception thresholds impacting the control actions, and much more visual stimulus to the driver, large step changes in pedal displacement (Figure 3-27) and pedal force (Figure 3-28) are very apparent.

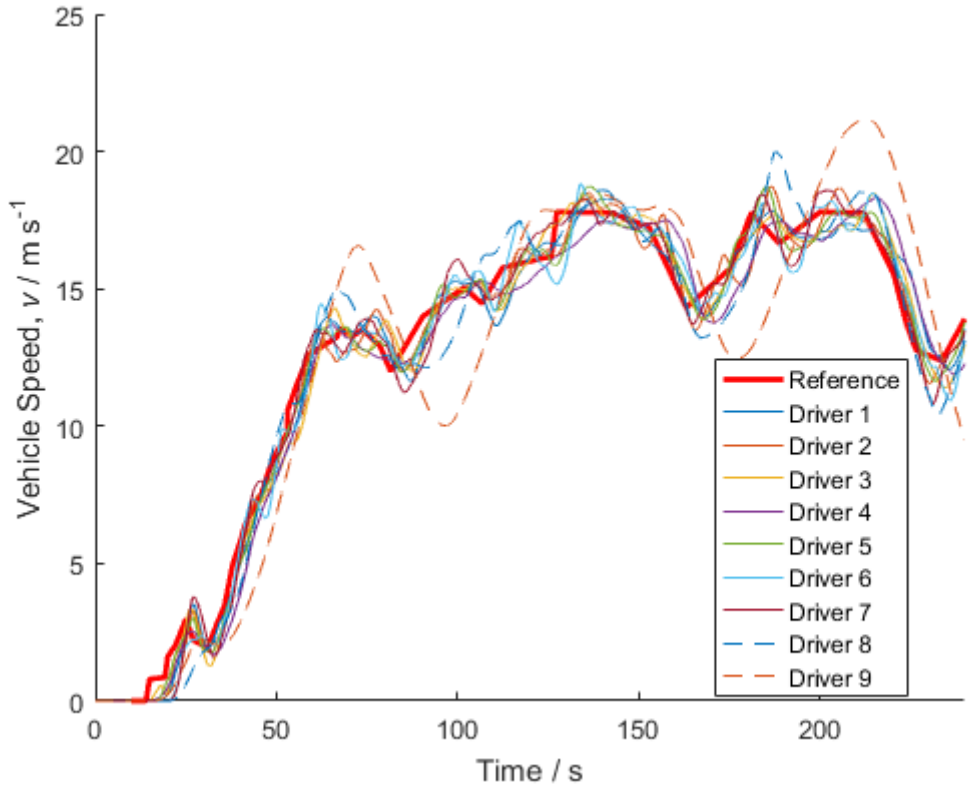

Figure 3-26 - Mean vehicle speed against time for the car-following task

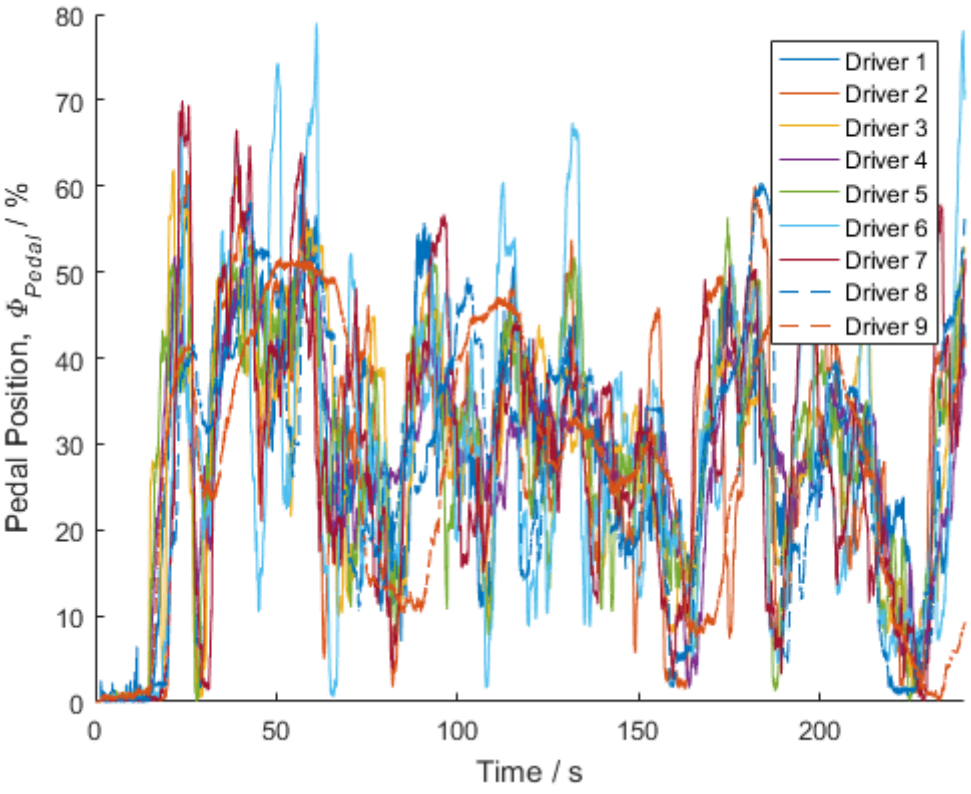

Figure 3-27 - Mean pedal position against time for the car-following task

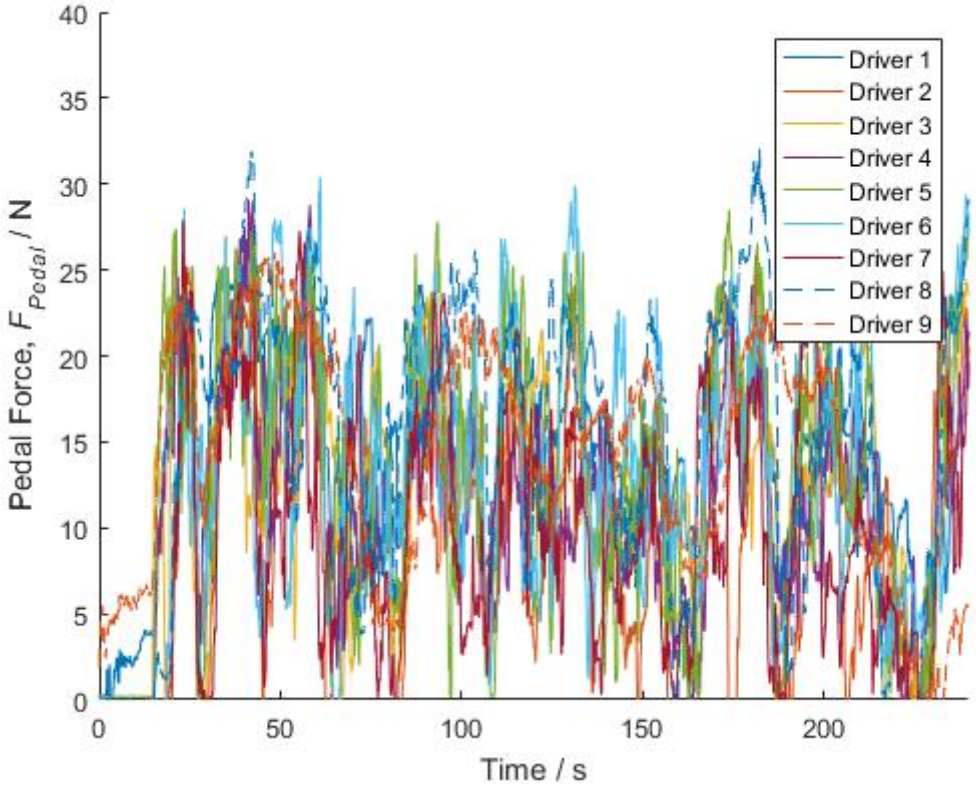

Figure 3-28 - Mean pedal force against time for the car-following task

The additional variable measured in the car-following scenario is distance between the two vehicles. This is simply illustrated in Figure 3-29. Most drivers remain somewhere between 20m and 40m behind the target vehicle throughout most of the experiment. However, drivers 4 and 9 are the exception to this case, falling up to 120m behind the target vehicle. On communicating with the subjects, it became apparent that Driver 4 felt a greater distance was required. As the driver did not have a brake pedal, deceleration rates would be limited, and hence the safe distance was considered much greater. Driver 9, a professional driver, was uncomfortable driving without a brake, and found it very difficult to maintain any sort of constant following distance. This is clearly demonstrated by the significant oscillations in following distance observed.

An alternative way of looking at the following distance is to look at the Time Headway (THW) (3.38). The THW for the seven drivers is in Figure 3-30. All drivers, bar Drivers 4 and 9, maintain near constant THW over the final 180s of the experiment. This corresponds to when the vehicle has accelerated from very low speeds to medium speeds.

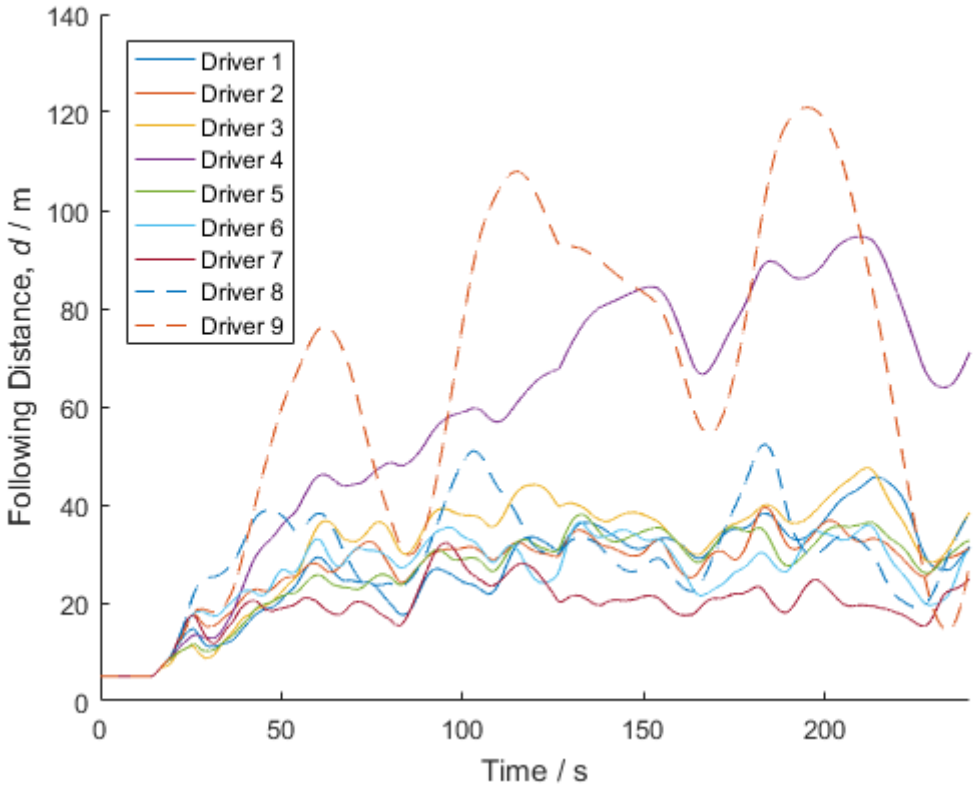

Figure 3-29 - Mean following distance against time for the car-following task

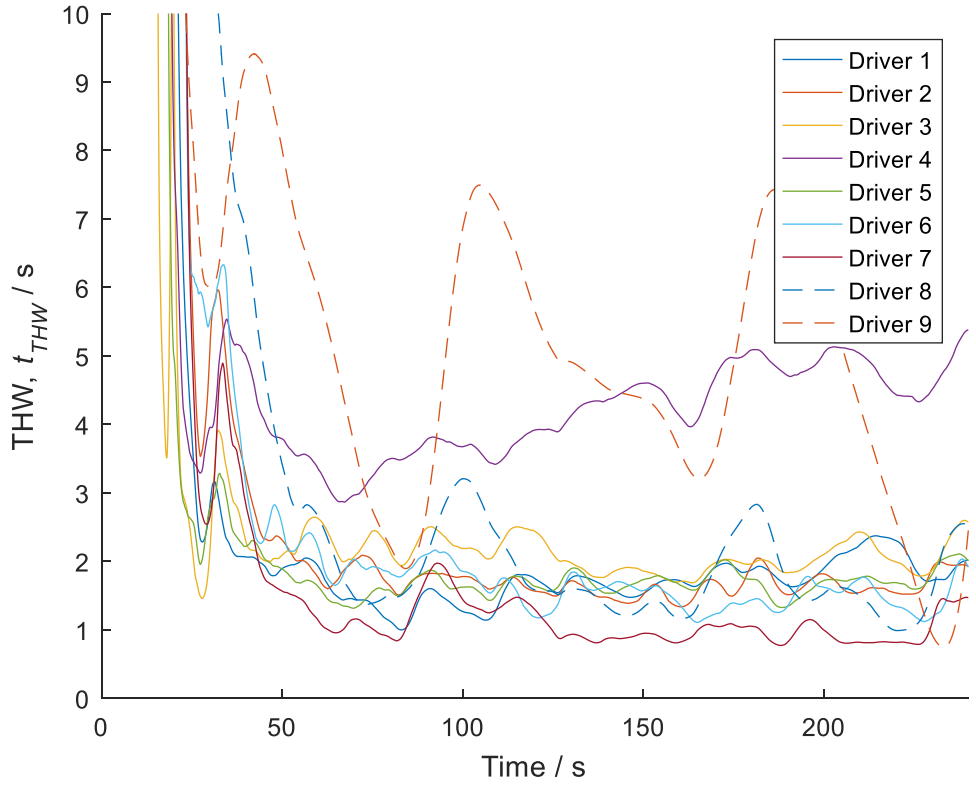

Figure 3-30 - Mean Time Headway (THW) against time for the car-following task

Figure 3-31 demonstrates the costs for the drivers in the car following task. In this comparison, the drivers target following distance is set to the mean THW observed from each driver over the final 180s of the task (3.45). This ensures that all starting effects are not included.

$$\bar{t}_{THW} = \frac{t_s}{240 - 60} \sum_{t=60}^{t=240} t_{THW} \qquad (3.45)$$

As drivers were not set a specific THW target, this approach ensures all drivers are compared fairly about whichever value of THW they followed. Driver 9 stands out once again as due to the variability of their following distances. The remainder of the drivers demonstrate a slight trend for increasing pedal forces with increasing following distances error. This trend suggests that the drivers are of different ability levels.

In this scenario, once again, there is little to distinguish the professional HGV drivers from the non-professionals, but the HGV drivers are at the higher end of the RMS pedal force scale.

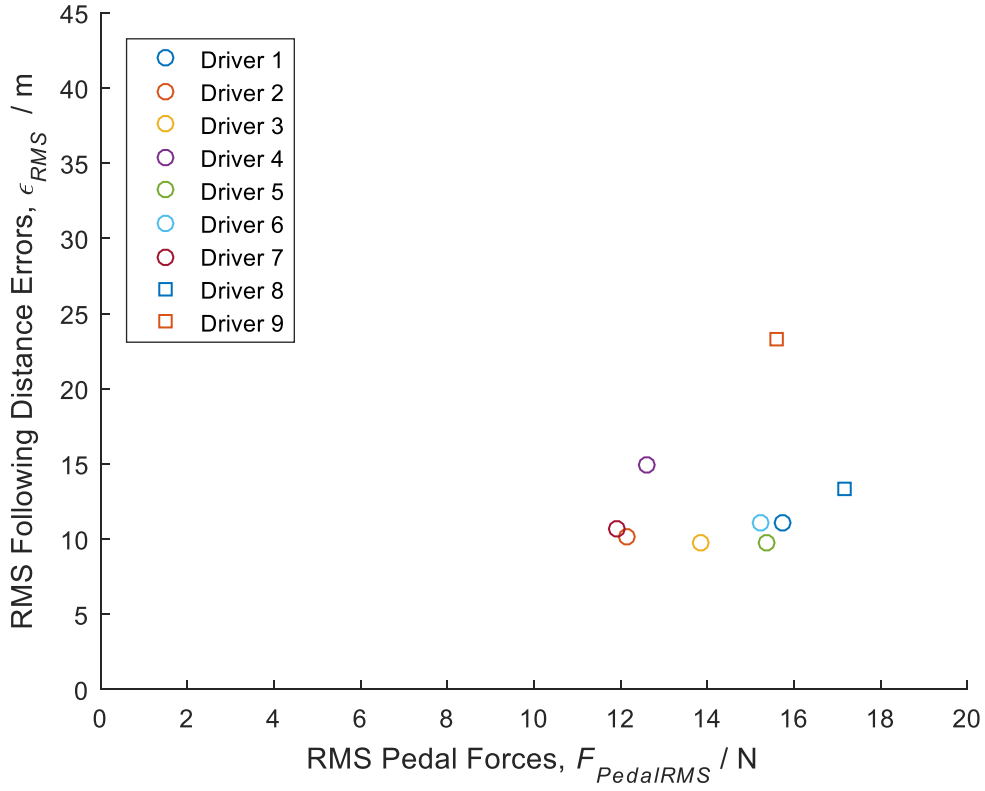

Figure 3-31 – Cost diagram for car following scenario.

Like with the drive cycle case, as a final consideration, the mean following distance for each driver is plotted with one standard deviation of the individual's runs shaded either side in Figure 3-32. It is very noticeable that Driver 9, one of the professional drivers, has a very high

following distance standard deviation, indicated by the large shaded region either side of the mean line. It is of particular interest to note that at 80s and 230s, the standard deviation is so high that the shading crosses the zero axis, implying a collision might be possible. The following distance behaviour of Driver 9 is clearly very inconsistent, possibly due to a lack of familiarity with the driving simulator, so further conclusions should be made cautiously.

Of the other drivers, Driver 4, also has a very high standard deviation, but not as high as Driver 9. The remaining drivers all have significantly lower standard deviations across the cycle, and in most cases, overlap is seen in the following distances adopted by the drivers.

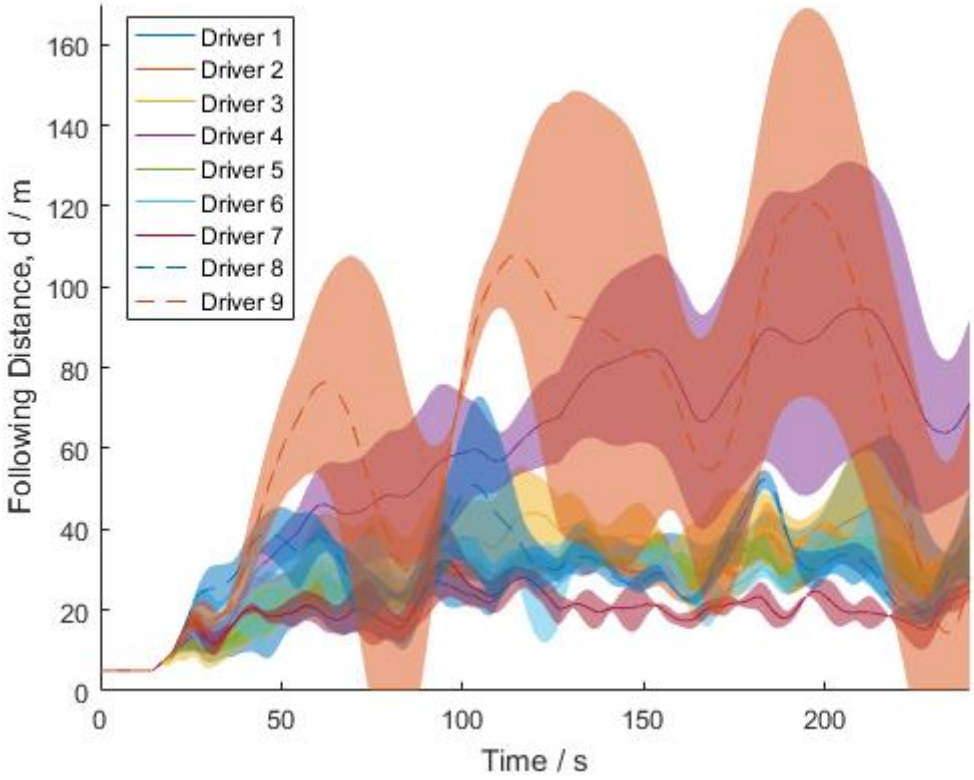

Figure 3-32 - Driver mean following distances with one driver standard deviation shaded either side

3.5 Model Identification

3.5.1 Driver model drive-cycle task

An optimisation algorithm is used to fit the driver model cost function weightings to the recorded data. The Matlab fmincon function is used to find the cost function weightings that minimise the sum of squared speed errors and pedal displacement errors when the driver model is applied to a drive cycle following task. This is called the fitting function:

$$\Gamma_{DC}(q_{1s}, q_{1\phi}, q_{1\dot{\phi}}) = \frac{1}{2\left|\sum_{t=0}^{t=240}(\phi_{PedalMeasured} - \phi_{PedalModelled})^2\right|_0} \sum_{t=0}^{t=240} (\phi_{PedalMeasured} - \phi_{PedalModelled})^2 + \frac{1}{2\left|\sum_{t=0}^{t=240}(v_{measured} - v_{modelled})^2\right|_0} \sum_{t=0}^{t=240} (v_{measured} - v_{modelled})^2 \quad (3.46)$$

where $|...|_0$ denotes a term to be assessed with all q values set to zero (i.e. the only component of the driver cost function is pedal force, so the model will respond by keeping the vehicle stationary throughout each task). The component weightings are such that the value of Γ_{DC} will be 1 for the case of all q values set to zero. It is worth noting that there are two cost functions in this process: the driver cost function, V_1, which determines how the driver model behaves, and the fitting function, Γ_{DC}, which determines how the driver model is fitted to the data.

For the optimisation, each driver cost function weighting is constrained to be positive and the pedal force weighting, r_1, is set to 1. In order to assess the significance of each term in the driver's cost function, different cost function combinations were assessed to see the effect on the fitting function. The results are summarised below in Table 3-4. The numbers included here are the average values over all drivers.

Table 3-4 – Average fitting function values for different driver cost functions. A dash indicates that the corresponding q value was not included in that particular cost function.

q_{1s}	$q_{1\phi}$	$q_{1\dot{\phi}}$	Γ_{DC}
4.3×10^4	2.3×10^{-1}	4.1×10^8	0.0135
3.8×10^2	5.0×10^{-4}	-	0.0134
6.5×10^6	-	4.2×10^8	0.0135
3.8×10^2	-	-	0.0134

It can be seen from Table 3-4 that there is little difference in the fitting function value if neither of the pedal displacement or pedal speed weightings are included. As extra weightings add extra complexity to the driver model, it is proposed to adopt the simplest cost function assessed, the cost function that includes only speed errors, q_{1s}, and pedal force, r_1. As the pedal force weighting, r_1, is set to a value of 1, the value of speed error weighting, q_{1s}, is hence the relative weighting on speed errors compared with pedal force.

A sensitivity study is completed to assess how the fitting function values vary with speed error weighting, q_{1s}, for Driver 2: Figure 3-33. In this figure the components of the fitting function are also included. It can be seen that at low values of speed error weighting, q_{1s}, the Fitting Function, Γ_{DC}, approaches 1. This is where the vehicle remains stationary and the difference between measured and modelled forces are large as the model exerts zero force. There is a critical weighting at which point the modelled vehicle is able to accelerate from stationary, and this is detected from the kink in the speed curve. At high speed error weightings, q_{1s}, the speed errors reach near zero as both modelled and measured drivers are close to the speed target, but there is some steady pedal force error required. As the speed error weighting, q_{1s}, increases, the modelled pedal force increases to achieve smaller speed errors relative to the target, but this results in greater speed and pedal force errors relative to the measured driver.

The resulting curve (Figure 3-33) identifies a clear lower bound to the value of speed error weighting, q_{1s} at approximately $q_{1s} = 5$. A minimum is found at approximately $q_{1s} = 300$. Examining how the driver's cost function summed over the task varies with speed error weighting, q_{1s}, as well (Figure 3-34), reveals that the driver's cost also varies little above the minimum detected from the fitting function technique. This suggests that the fitting function technique is a good method for finding a lower bound to the value of speed error weighting, q_{1s}, and that exceeding the lower bound has little in the way of impact on driver cost function performance. Similar conclusions are drawn from the other drivers, although their sensitivity studies are not included here for conciseness.

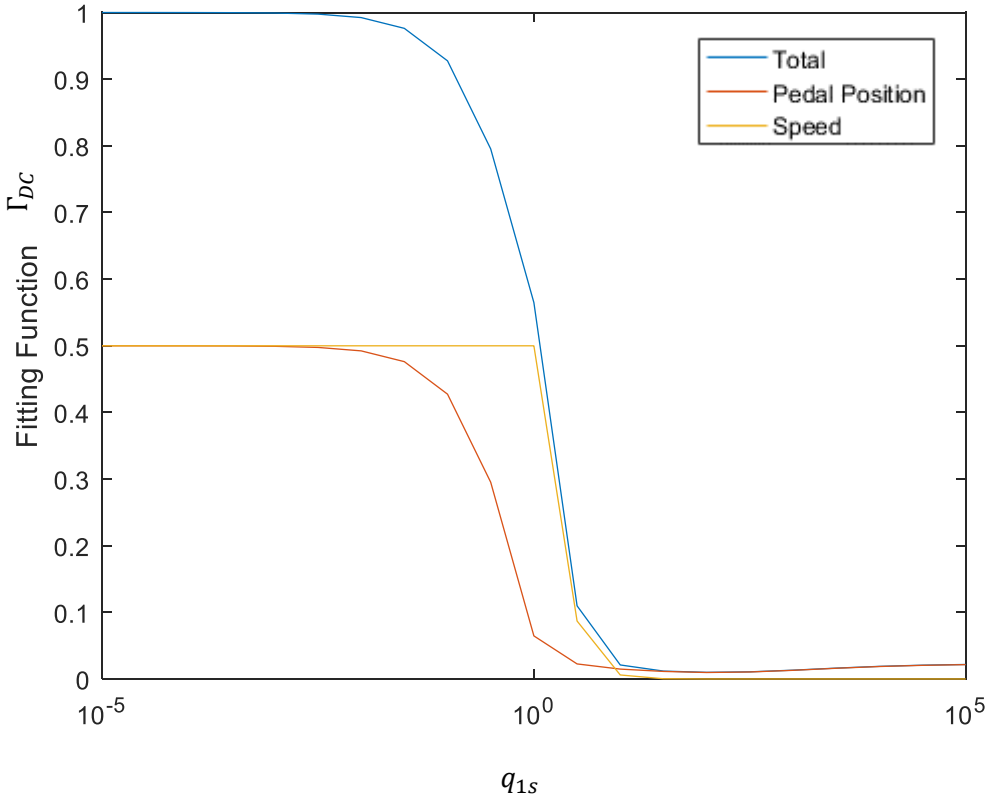

Figure 3-33 – Fitting function sensitivity study for the drive cycle task

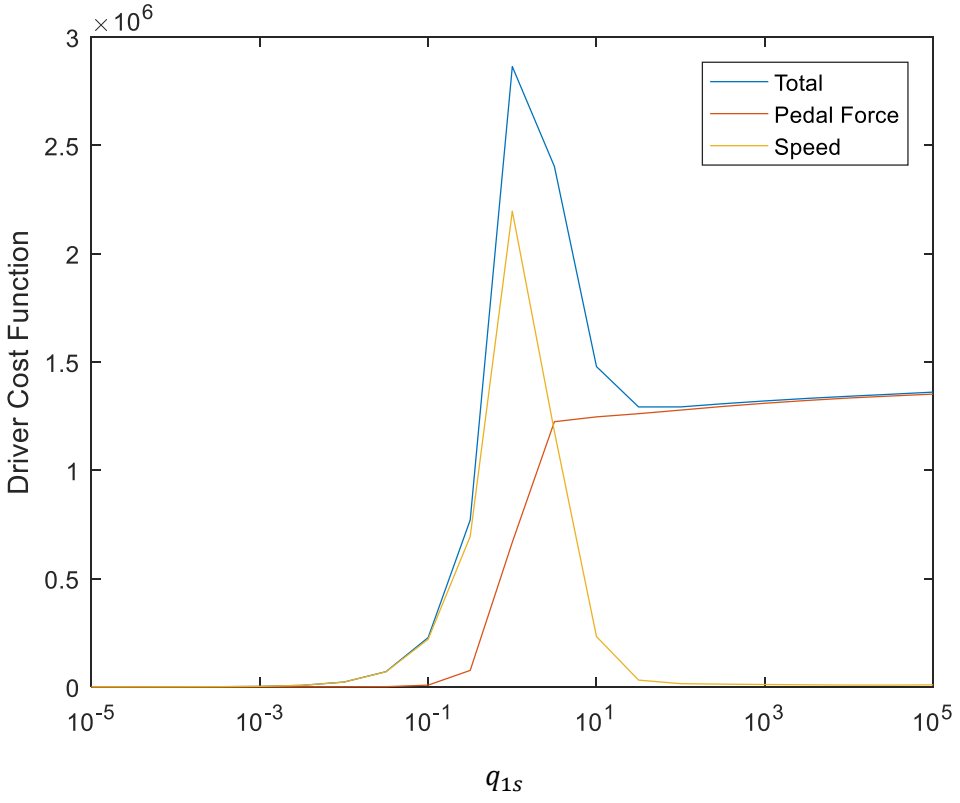

Figure 3-34 –Summed driver cost function sensitivity study for drive cycle task

The resulting cost function weightings for all nine drivers are displayed in Figure 3-35. The fitting function values are included in Figure 3-36 as an indicator of goodness of fit.

It can be seen that the driver model had a best fit with Driver 9, and the worst fit with Driver 7. Driver 7's speed profile is dominated by large overshoots when acceleration demands change – a behaviour that the driver model will not replicate because it demonstrates a better anticipation of future speed demands. The poor fit is also exaggerated because the driver model does not account for the driver's imperfect perception of speed and limited understanding of the vehicle dynamics. Driver 9 on the other hand has much smaller overshoots and achieves speeds closer to the target. The driver model is able to recreate this behaviour much more closely.

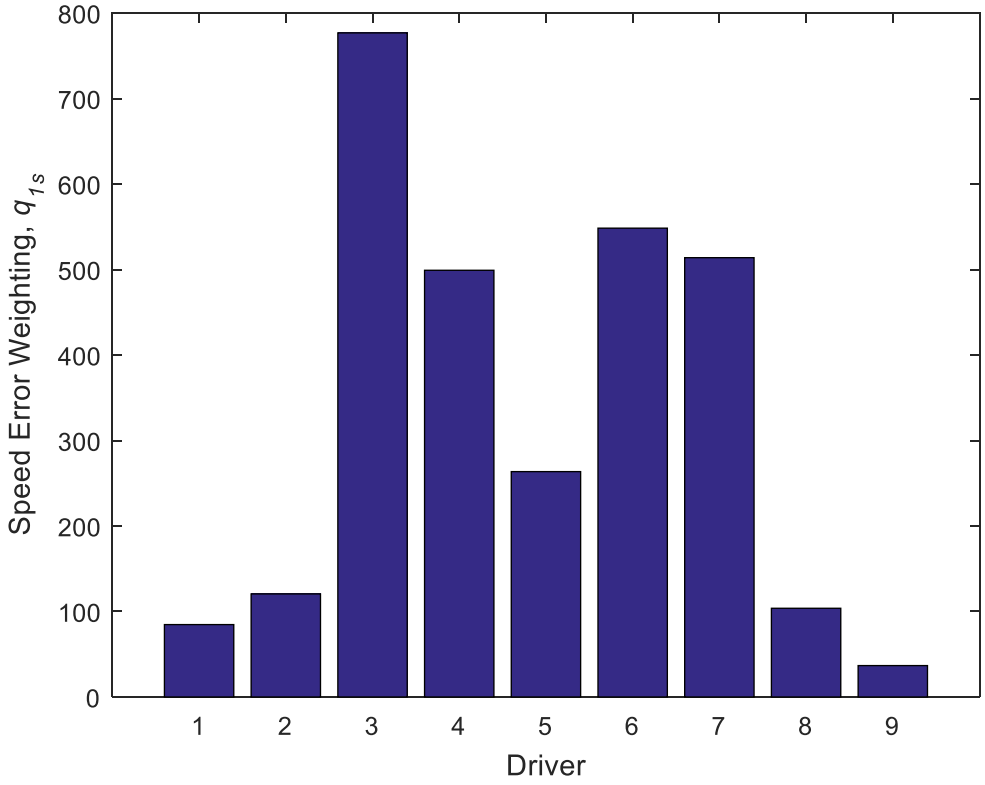

Figure 3-35 – Fitted cost function weightings for the Drive Cycle scenario

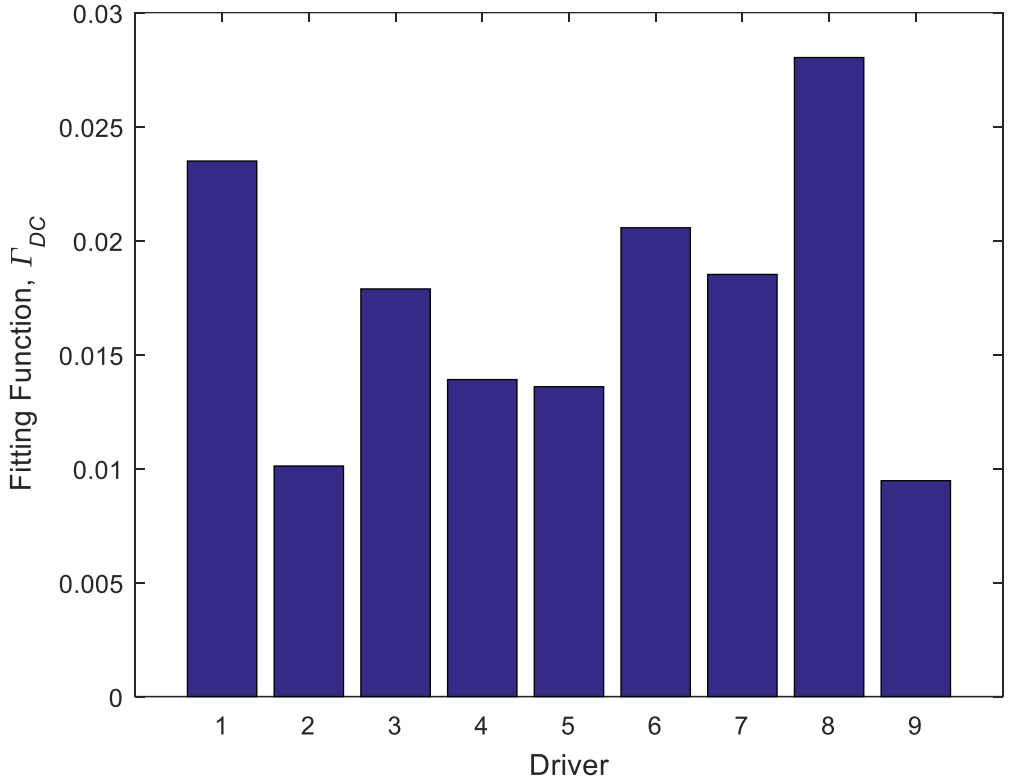

Figure 3-36 – Total driver-model driver squared errors for Drive Cycle scenario

Figure 3-37, Figure 3-38 and Figure 3-39 compare the recorded speeds, pedal positions and pedal forces of Driver 9 going through the drive cycle scenario with the modelled Driver 9. (Driver 9 is selected as the best agreement between human driver and model driver occurs for this driver). The speed profiles of the two compare very well, with the most significant difference being the lack of small overshoots in the model driver's performance. The model driver sticks very close to the target speed profile, except for rounding off speed troughs and peaks where accelerations change sharply.

The modelled pedal position and pedal forces differ slightly more from the recorded values. In the pedal position case, as the human driver is driving at a slightly different speed and acceleration to the model driver, the demanded torque from the engine is slightly different. As the vehicle models are identical in both the driving simulator and the model case, this means that different pedal positions are required. These differences are passed onto the pedal forces as well. However, as discussed in Section 3.1.2, the linear pedal model does not represent the actual pedal dynamics completely accurately. It is also expected that the human driver also introduces significant noise to the system that is not included in the driver model. The result is that the modelled pedal positions and forces are much smoother than the recorded data from the simulator.

The high pedal force and displacement witnessed at the start of the drive cycle are explained by the one non-linearity programmed into the driving simulator – the constant opposition forces

(rolling resistance and gradient component of mass) drop to zero at zero speed. In the driver model, this is not the case as the full model is linear.

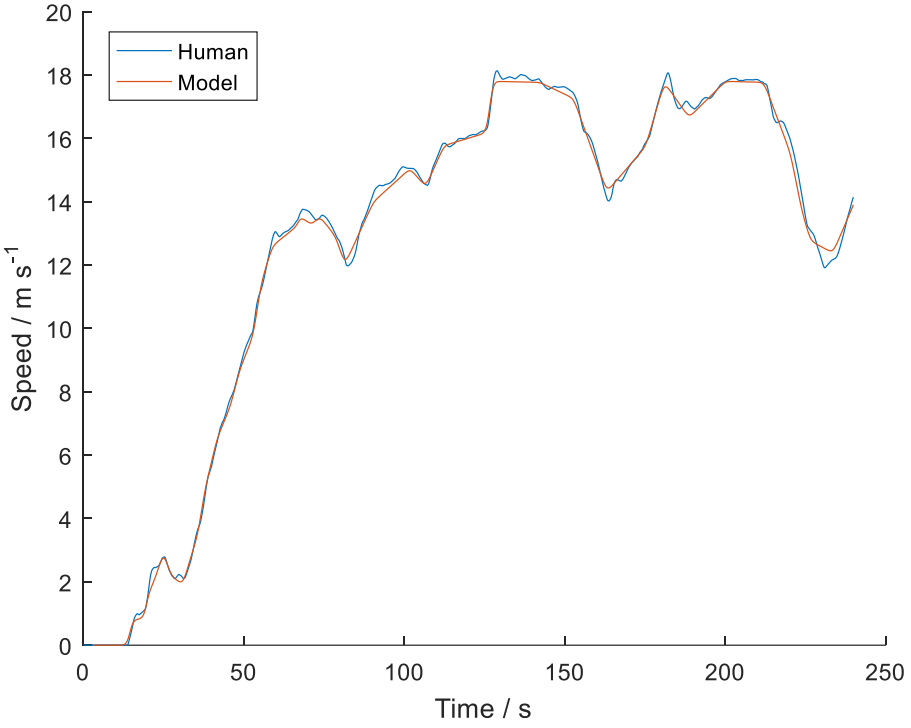

Figure 3-37 – Measured human driver and modelled driver speed profile for Driver 9

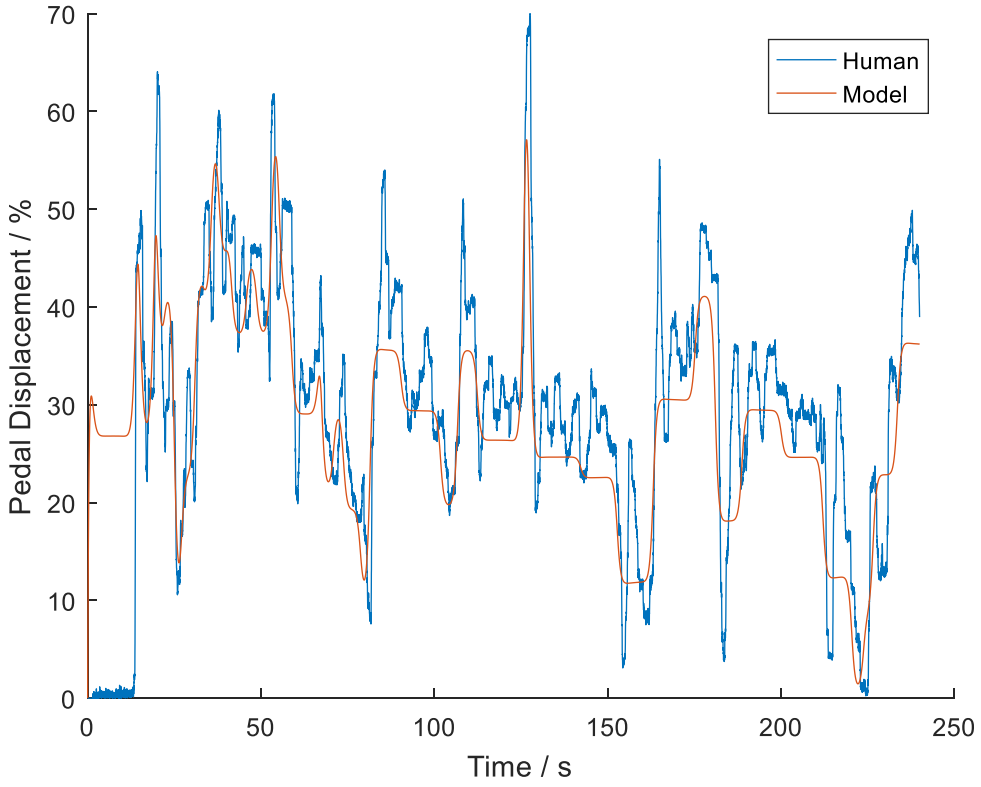

Figure 3-38 - Measured human driver and modelled driver pedal position for Driver 9

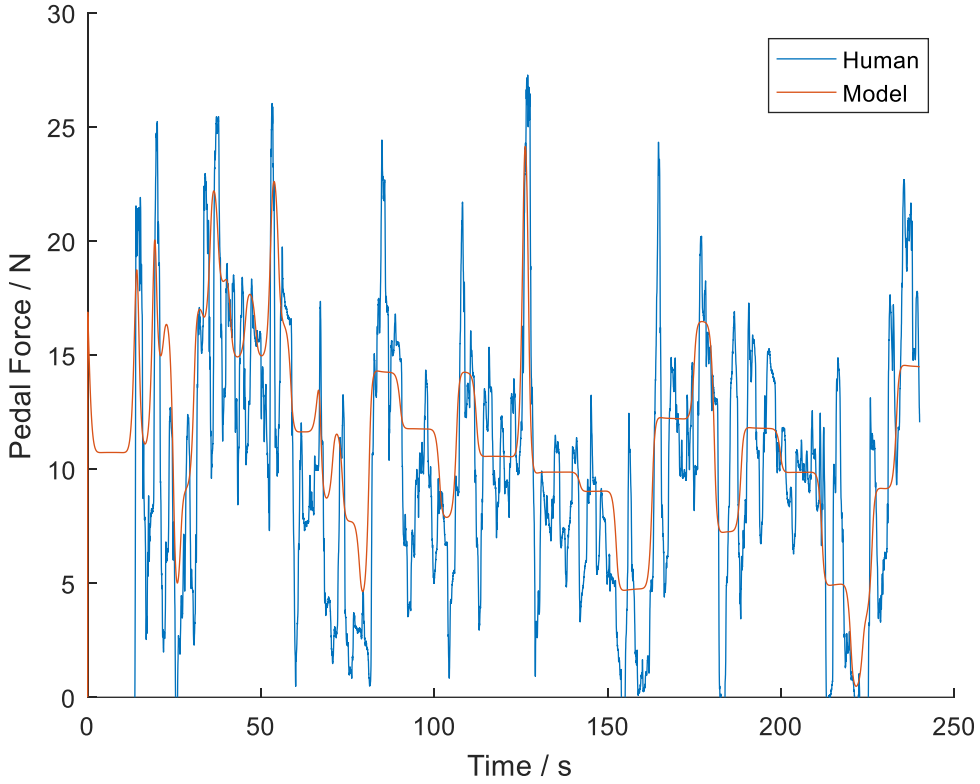

Figure 3-39 - Measured human driver and modelled driver pedal force for Driver 9

Further analysis of the driver model is completed through a parameter study examining how by varying the cost function weightings, differences in driver model behaviour can be seen. Figure 3-40 demonstrates the trade-off between different costs for different driving styles. The driver model is given a drive-cycle following task using the modified Millbrook cycle and the RMS pedal force and RMS speed error are calculated.

Looking at the RMS speed error and RMS pedal force plot (Figure 3-40), several key features of the shape are noted. Firstly, generally, as the weighting on speed error (q_{1s}) increases, the RMS pedal force increases.

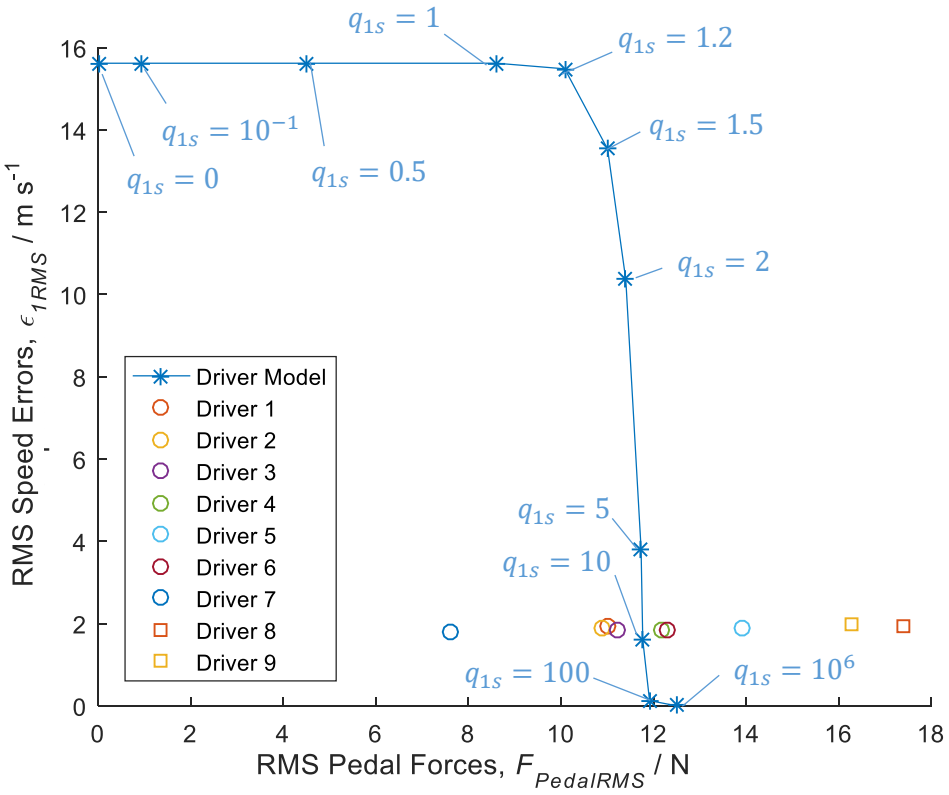

Figure 3-40 – Trade off graph for drive-cycle scenario

The horizontal line at approximately $16\ \mathrm{ms}^{-1}$ in the speed errors plot is a result of the weightings meaning that the pedal is not depressed enough to overcome the drag forces and accelerate the vehicle. As the vehicle cannot accelerate, the RMS speed error is therefore equal to the RMS value of the drive-cycle speeds – a constant value.

At some critical speed error weighting, approximately $q_{1s} = 1.2$, the driver model is able accelerate the vehicle from stationary, and the corresponding RMS speed error decreases. There is a limit to this behaviour, however, as to keep reducing the speed errors further, high accelerations of the vehicle are required. For high accelerations, high pedal displacements are required, and hence high pedal forces. This results in a higher RMS pedal force for the drive cycle.

It is interesting to compare the data points from the individual drivers with the curve for the driver model. Although the human drivers all have a similar RMS speed error to each other, their RMS pedal forces vary more, and it is this variation that draws the individuals away from the driver model plot.

3.5.2 Driver model car-following task

In the car following driver model, the safe displacement of the target vehicle is included as a state. The safe displacement is defined as a safe following distance added to the following vehicle's displacement. Section 3.5.1 has demonstrated that in most cases the THW approach is a good

approximation to the observed driver behaviour. To incorporate this into the driver model, it is therefore required to assess the target THW for each driver. To do this, the average THW is calculated by taking the mean of the THWs in Figure 3-30 over the time 60s to 240s. These average THWs are displayed in Figure 3-41 with error bars indicating the standard deviation of the THW over the same time segment.

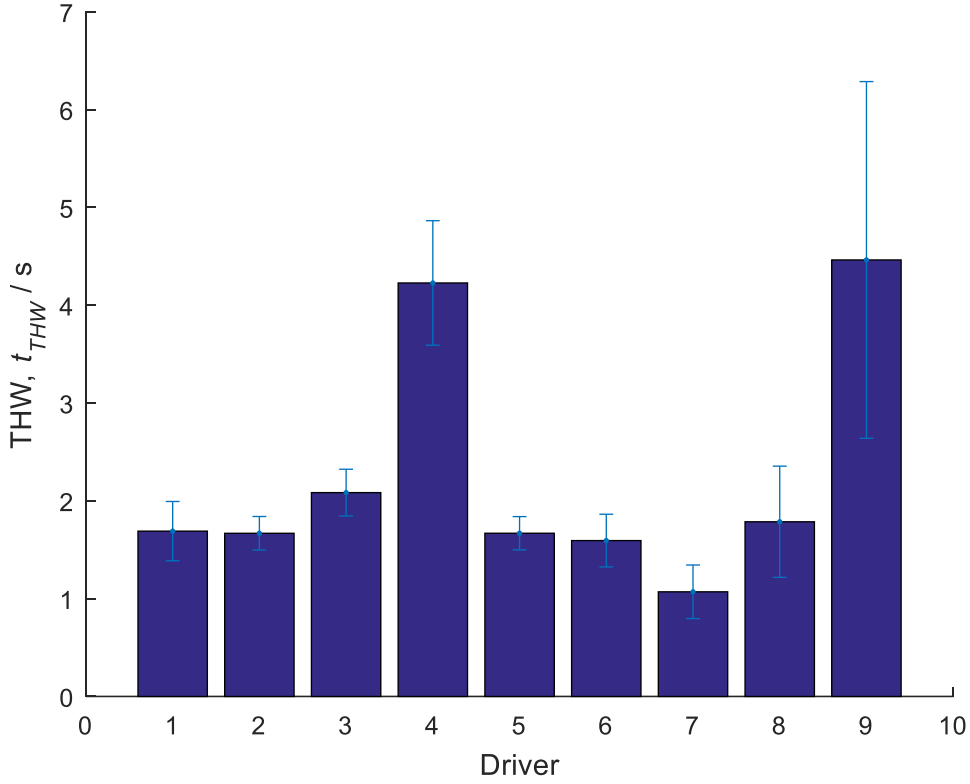

Figure 3-41 – Mean THW by driver including error bars of ± 1 standard deviation

In the car following scenario, the cost functions have an extra variable, following distance, so an extra weighting is included in the fitting optimisation. Using the same approach as before, the fmincon function is used in Matlab to identify the cost function weightings that minimise the fitting function defined in (3.46):

$$\Gamma_{CF}(q_{1s}, q_{1\phi}, q_{1\dot{\phi}}) = \frac{1}{2\left|\sum_{t=0}^{t=240}(\phi_{PedalMeasured} - \phi_{PedalModelled})^2\right|_0} \sum_{t=0}^{t=240} (\phi_{PedalMeasured} - \phi_{PedalModelled})^2 + \frac{1}{2\left|\sum_{t=0}^{t=240}(d_{measured} - d_{modelled})^2\right|_0} \sum_{t=0}^{t=240} (d_{measured} - d_{modelled})^2 \quad (3.47)$$

where $|...|_0$ denotes a term to be assessed with all q values set to zero (i.e. the only component of the driver cost function is pedal force, so the model will respond by keeping the vehicle stationary throughout each task). The weightings are such that the value of the fitting function, Γ_{CF}, will be 1 for the case of all other weighting, q, values set to zero.

For the optimisation, each driver cost function weighting is constrained to be positive and the pedal force weighting is set to 1. In order to assess the significance of each term in the driver's cost function, different weightings were set to zero to see the effect on the fitting function. The results are summarised below in Table 3-5. The numbers included here are the average values over all drivers.

Table 3-5 – Average fitting function values for different driver cost functions. A dash indicates that the parameter is not included in a combination

q_{1s}	q_{1d}	$q_{1\phi}$	$q_{1\dot{\phi}}$	Γ_{CF}
4.3×10^{-4}	7.2×10^{2}	4.0×10^{4}	8.4×10^{-2}	0.0808
6.9×10^{-4}	8.1×10^{2}	4.3×10^{4}	-	0.0808
4.7×10^{-4}	7.3×10^{2}	-	7.2×10^{-2}	0.0806
5.6×10^{-4}	7.4×10^{2}	-	-	0.0808
-	7.5×10^{2}	-	-	0.0805
-	8.1×10^{2}	4.4×10^{4}	-	0.0805
-	7.5×10^{2}	-	3.0×10^{-2}	0.0807

In a similar result to the drive cycle case, it can be seen from Table 3-5 that there is little difference in the fitting function value if none of the vehicle speed, q_{1s}, pedal displacement, $q_{1\phi}$ or pedal speed, $q_{1\dot{\phi}}$ weightings are included. As extra weightings add extra complexity to the driver model, it is proposed to adopt the simplest cost function assessed, the cost function that includes only following distance error, q_{1d} and pedal force r_1. As r_1 is set to a value of 1, the value of q_{1d} is hence the relative weighting on following distance errors compared with pedal force.

A sensitivity study is completed to assess how the fitting function values vary with following distance error weighting. q_{1d} for Driver 3: Figure 3-44. In this figure the components of the fitting function are also included. It can be seen that at low weightings on following distance error, q_{1d}, the Fitting Function, Γ_{CF}, approaches 1. This is where the vehicle remains stationary and the difference between measured and modelled forces are large as the model exerts zero force. There is a critical weighting at which point the modelled vehicle is able to accelerate from stationary, and this is detected from the kink in the following distance curve. At high values of q_{1d}, the following distance errors reach near zero as both modelled and measured drivers are close to

the target following distance, but there is some steady pedal force error. As the weighting increases, the modelled pedal force increases to achieve smaller errors relative to the target, but this results in greater errors relative to the measured driver.

The resulting curve identifies a clear lower bound to the value of following distance weighting, q_{1d}, at approximately $q_{1d} = 10^{-2}$. A minimum is also found at approximately $q_{1d} = 200$.

Examining how the summed driver's cost function varies with following distance error weighting, q_{1d}, as well, reveals that the driver's cost does vary at weights above the minimum detected from the fitting function. This is because the weighting on following distance errors increases faster than the following distance errors decrease, meaning that the summed driver cost function increases rapidly with high values of q_{1d}. Over most of the weighting ranges, the cost function is dominated by the following distance, and pedal force is a minor contributor. Similar conclusions are also drawn from the data for the other drivers.

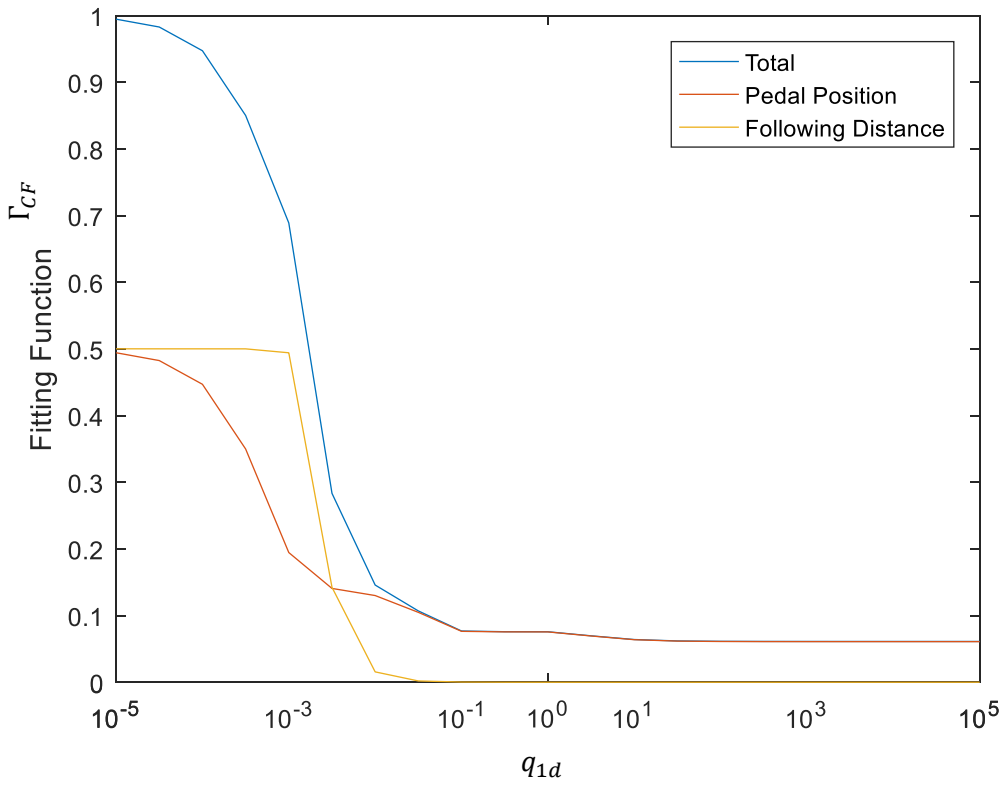

Figure 3-42 – Fitting function sensitivity study for the drive cycle task

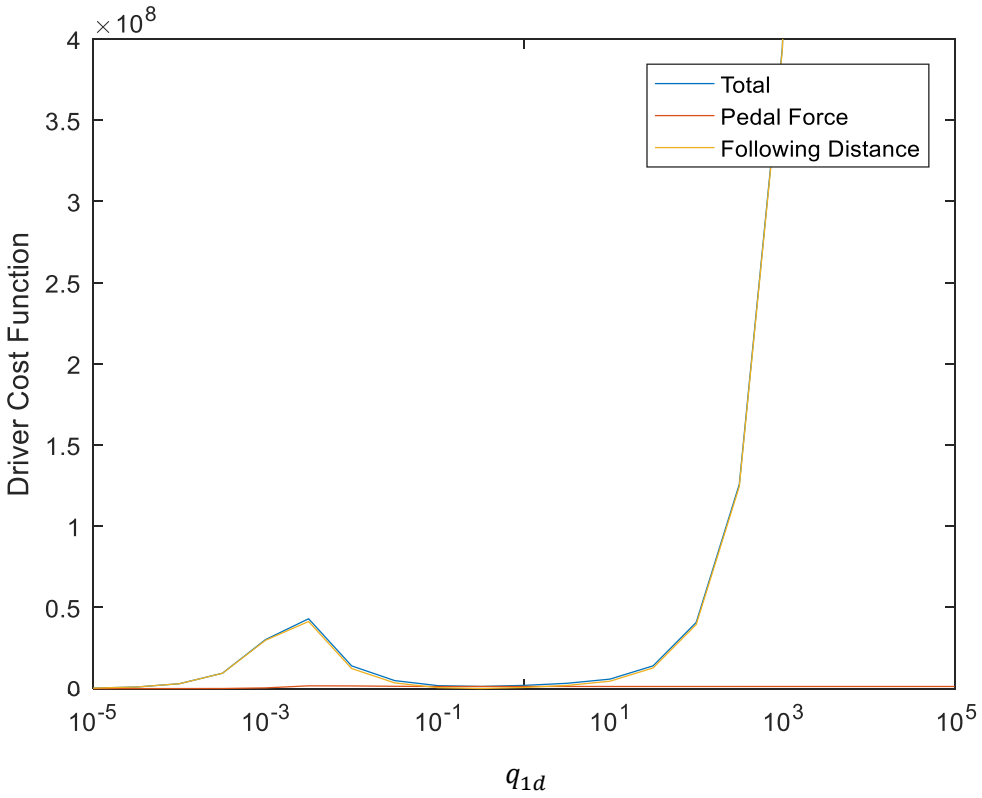

Figure 3-43 – Summed driver cost function sensitivity study for car-following task

Now examining the cost function weighting calculated for the other drivers, it is noted that there are significant differences - Figure 3-44. Although the differences in weightings are significant, all weightings are well above the lower bound observed in Figure 3-42 and the equivalents for other drivers. These weightings do not therefore cause major differences in the behaviour of the model. Figure 3-45 indicates there is good agreement between for modelled and measured drivers for all cases.

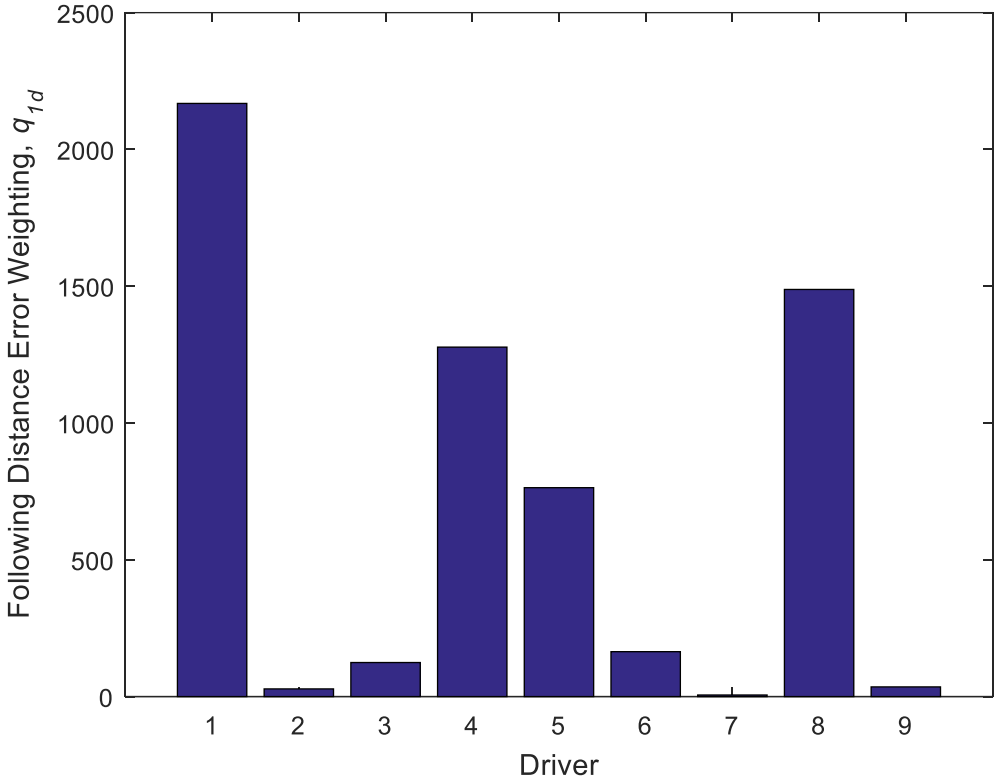

Figure 3-44 - Fitted cost function weightings for the car-following scenario

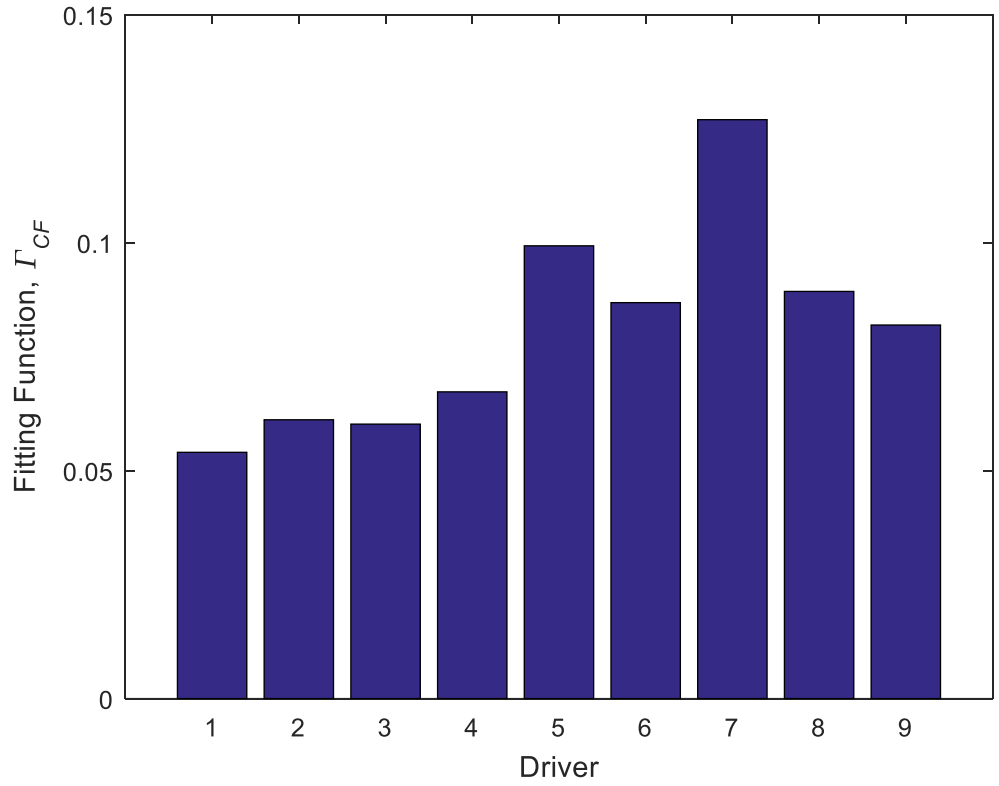

Figure 3-45 – Total driver-model driver squared errors for car-following scenario

As the driver with best agreement between model and driver, Driver 1 is selected to demonstrate the performance of the driver model. Figure 3-46 to Figure 3-49 illustrate the recorded and modelled vehicle speed, following distance, pedal displacement and pedal force. Due to accurate knowledge of the motion of the vehicle in front, the driver model is very quick to respond to changes in speed of the target vehicle. The result of this is that in the first minute of the test, the driver model is considerably closer to the target vehicle.

Throughout the rest of the scenario, the driver model maintains a much smoother speed, with little overshoot from the reference vehicle speeds (Figure 3-46).The driver model is also able to balance the speed and following distance targets to reach a more consistent following distance than the oscillating distances recorded from the driver(Figure 3-47). Again, this is due to the driver model's better knowledge of the vehicle in front as perceptual limitations are not included in the model.

Mostly, the driver model produces pedal forces that agree well with the measured human driver (Figure 3-48). There are several impulses in the driver model pedal forces which occur when the acceleration of the target vehicle changes. As the driver model is unable to anticipate this (it works on a constant acceleration assumption for preview information), when changes in acceleration are detected, it results in sharp changes to its pedal force to meet the new following distance and speed targets. The noise that the driver adds to the system is also not included in the driver model, so the pedal force appears much smoother from the driver model than the human driver (with the exception of the impulses as already discussed).

Finally, the pedal displacement plot shape (Figure 3-49) is well represented by the driver model, but the displacement values frequently vary by $\pm$ 20% when compared to the human driver. This is a direct consequence in the difference in speed profiles between the human driver and the driver model, as the vehicle models are identical.

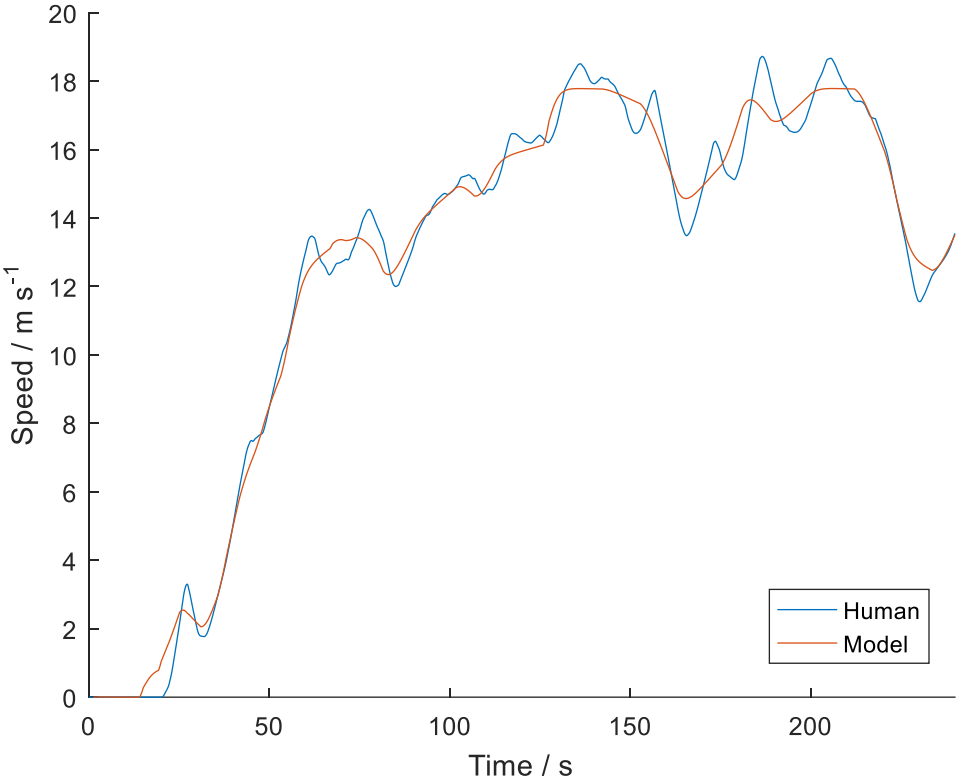

Figure 3-46 – Measured human driver and modelled driver vehicle speed for Driver 1

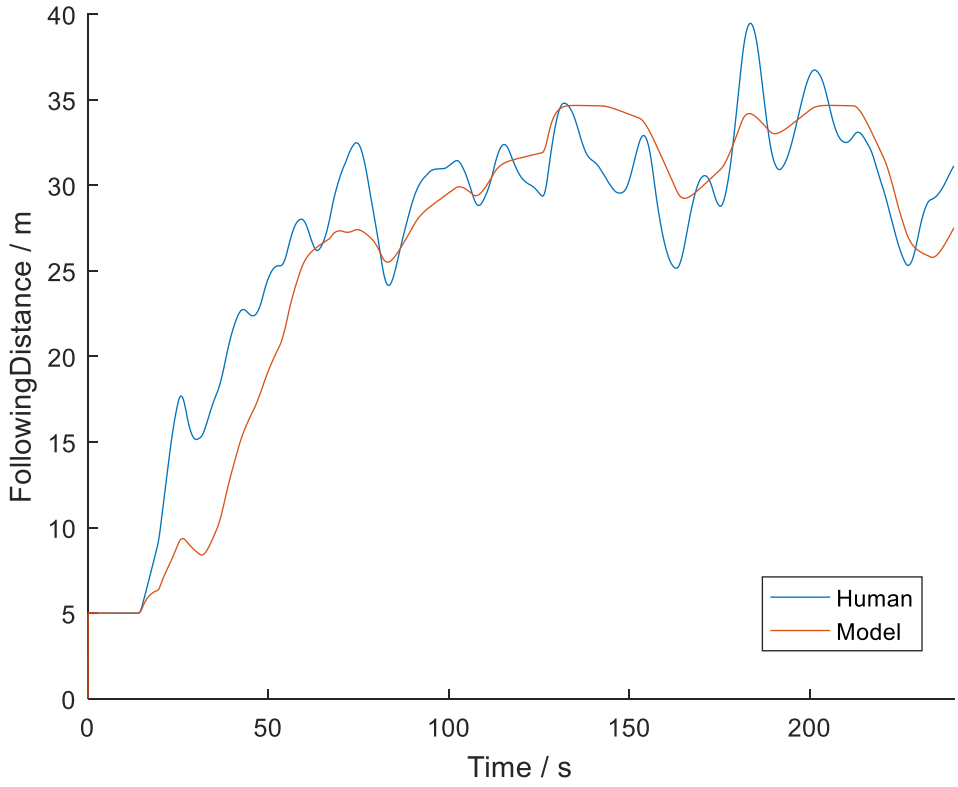

Figure 3-47 - Measured human driver and modelled driver following distance for Driver 1

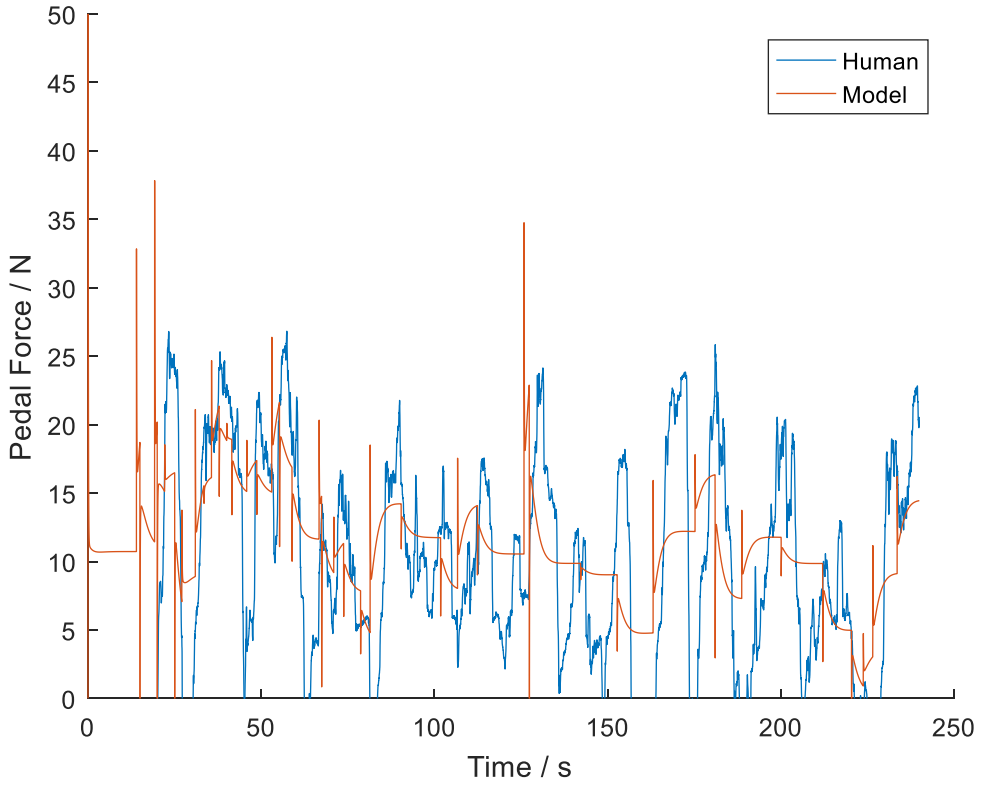

Figure 3-48 – Measured human driver and modelled driver pedal force for Driver 1

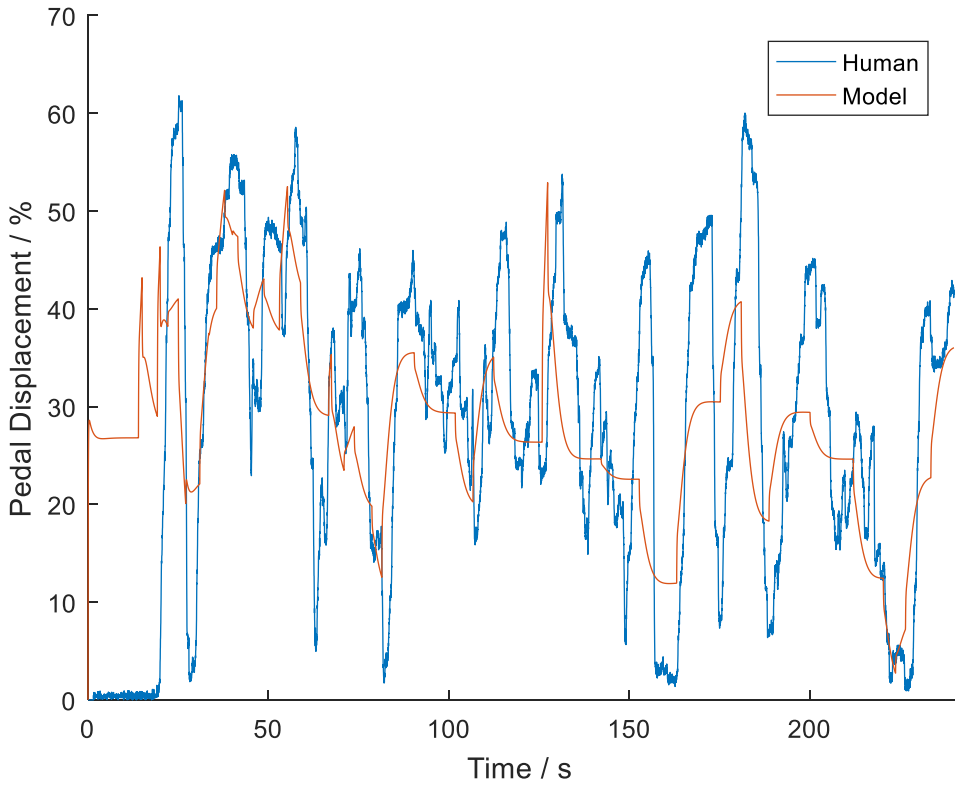

Figure 3-49 – Measured human driver and modelled driver pedal displacement for Driver 1

Finally, a parameter study on the effects of the cost function weighting q_{1d} during a car following task, where the target vehicle follows the modified Millbrook cycle described earlier, reveals the driver model behaviours in Figure 3-50.

Outside the range illustrated in Figure 3-50, at low values of following distance error weighting, q_{1d}, the vehicle is stationary, and will therefore have constant RMS speed and following distance errors. At some critical weighting, the vehicle will accelerate away from rest. As the vehicle is able to accelerate up to speed the speed and following distance errors decrease. As the following distance errors get smaller, large changes in weightings are required to achieve further improvements in following distances.

The limits of Figure 3-50 are set to allow comparisons between the recorded data from the human drivers and the driver model. Over the range of following distance error weightings (q_{1d}) illustrated, it is interesting to note that there is very little change in the RMS pedal force, but the RMS pedal force does tend to increase at both extremes of the figure. This is because the controller undergoes a slight regime change from smooth steady forces to generally lower forces, but with high impulses to achieve changes in acceleration (as illustrated in Figure 3-48). Although the time histories themselves will look very different, the RMS value of pedal force is not significantly affected.

Outside the range of q_{1d} covered in the parameter study covered in Figure 3-50, the RMS pedal force tends to zero as q_{1d} approaches zero. At this point, the RMS following distance error is equal to the RMS distance travelled by the leading vehicle, as the host vehicle will not accelerate beyond rest, if pedal forces are zero, and hence pedal displacements zero.

It is noted that once again there is a much greater variation in the RMS pedal forces from the human drivers than is displayed by the driver model. Driver 7 is of particular note as they appear to lie on the driver model contour. However, as Figure 3-45 illustrates, the fitting function for this driver is not the minimum across all drivers. This is because although the RMS pedal forces are near equal, the pedal force time histories themselves are different.

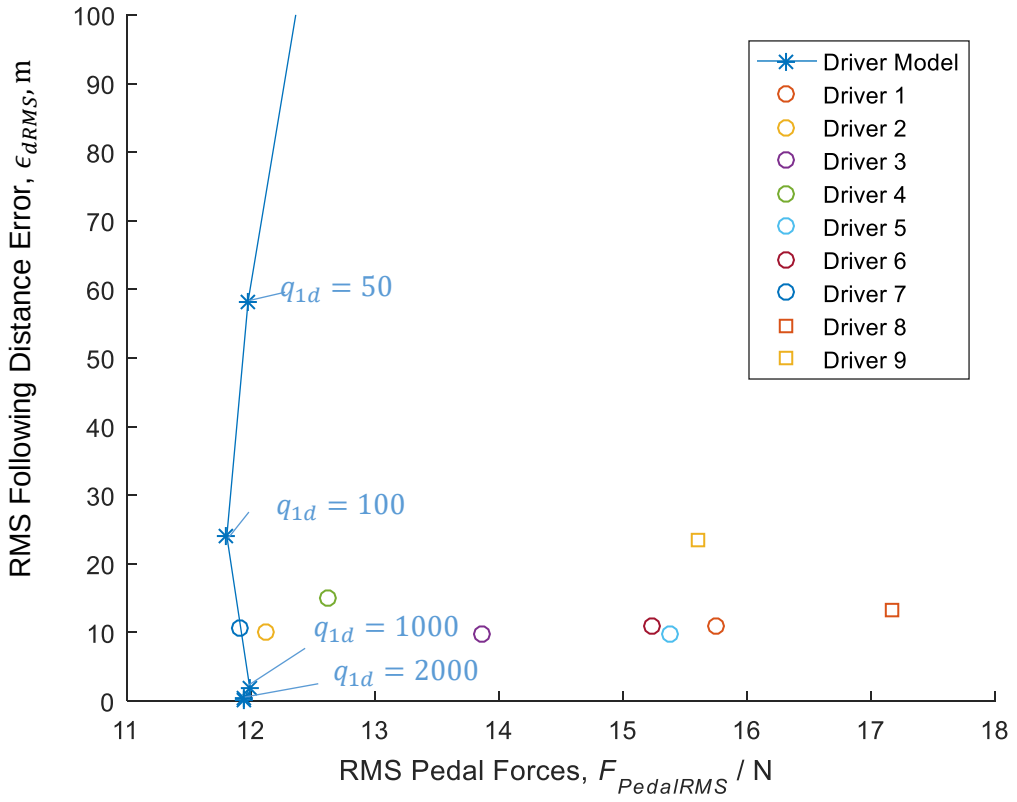

Figure 3-50 – Trade off graph for the car following scenario

3.6 Conclusions

In this chapter, a new mathematical model of the human driver's longitudinal control of a heavy goods vehicle through use of the accelerator pedal is proposed. Model Predictive Control theory, commonly applied to the driver's lateral control of the vehicle, is used as the basis for the model. This model is similar to Sharp's (2007), but extends it to explicitly incorporate the pedal control, and uses MPC, instead of Linear Quadratic Regression (LQR).

A linear vehicle model was first developed using parameters suitable for a heavy goods vehicle and pedal properties are taken from a pedal assembly installed in a driving simulator. The vehicle model is the key central component to the driver model (as the driver's internal vehicle model) and the simulation of the vehicle response to the driver's control actions.

The driver model was validated using a fixed base driving simulator and a range of human drivers. The driving simulator was extensively upgraded and developed to improve the realism of

the driving experience. Each driver completed a drive-cycle following and car-following scenario multiple times.

It was found that the human drivers had a range of driving styles and that this was reflected in the results. Using an optimisation to minimise the squared errors between simulated and measured variables, a suitable driver cost function of pedal force and speed error or following distance error (depending on task) was identified. The sensitivity of the fitting function to changes in driver cost function weightings means that the optimal value calculated can be treated as a lower bound to the weighting value. Weightings above the lower bound had little impact on model performance.

The driver model was tuned to each of the human drivers in turn. Weightings on driver pedal force were set to one and the relative speed error weightings ranged from 100-750 and the following distance error weightings ranged from 25-2200. This means that the human drivers were more consistent in their weighting of speed errors than following distances. The fact that following distance weighting reached approximately triple the maximum speed error weighting signifies that in this scenario, the drivers were placing much more emphasis on performance over effort than in the drive cycle following scenario. Much more variation was witnessed in the car following task in terms of numerical weightings, but the model performed acceptably well, with goodness-of-fit being up to twenty times better than the no-action cases (a fitting function value of 0.05).

It was found that there was acceptable agreement between the human driver and the driver model, with RMS pedal forces within approximately 20% of each other, but that the driver model was unable to represent the noise and disturbances observed in the human drivers. The fully linear model also has limitations at low speeds where the presence of constant drag forces results in pedal depression to maintain a stationary position. A non-linear vehicle model may assist the driver model in better recreating the control forces from the human drivers. A parameter study explored how different parameters in the driver model could be selected to represent the different driving styles of different drivers.

This chapter has applied the work Cole et al. (2006), Qu et al. (2015) and Flad et al. (2013) to the longitudinal control of a vehicle and has extended the work of Sharp (2007), by expanding optimal preview control to the car-following scenario, as well as the drive-cycle scenario. This new application is also applied to the heavy goods industry, which previously had been little touched by similar research. Driving simulator experiments have also validated the modelling approach taken.

Chapter 4 - Pedal Feedback

In the previous chapter, the driver was modelled using Model Predictive Control Theory (MPC). The pedal feedback controller can also be represented in the same manner. When considering the effectiveness and 'likeability' of pedal feedback, it is important to look at how the two controllers (human driver and the separate pedal controller) interact. The literature review identified that game theory can be used to model how two different 'players' interact as they try to achieve their own objectives. Mathematically, there are several subtly different game theoretical frameworks that can be taken. In this chapter, three two-player frameworks are proposed for consideration in turn: decentralised feedback in 4.1, cooperative feedback in 4.2 and one-sided cooperative feedback in 4.3 (Na and Cole, 2013) and (Rawlins and Mayne, 2009).

4.1 Two Player Decentralised Control Theory

In the first case to be considered, the driver and pedal feedback controller are said to have no knowledge or understanding of each other's control. This means that both act completely independently and apply forces on the pedal in parallel.

In Chapter 3, the human driver was modelled as a Model Predictive Controller with additional limb dynamics. In this chapter, a second Model Predictive Controller is added to work in parallel with the driver model. This controller is considered the 'feedback pedal controller' and can apply additional forces on the pedal, to feedback to the driver - Figure 4-1.

The feedback pedal controller is derived in the same manner as the driver model was (equations (3.21) to (3.34)), so derivation is not included here. The matrices relating to the feedback controller are denoted with a subscript '2'. In the driver model, the input dependence matrix was given the symbol $\boldsymbol{\Theta}$. As the driver input may also affect the states of the pedal, a second symbol, $\boldsymbol{\Omega}$, is used to denote the dependence on the feedback pedal controller's input.

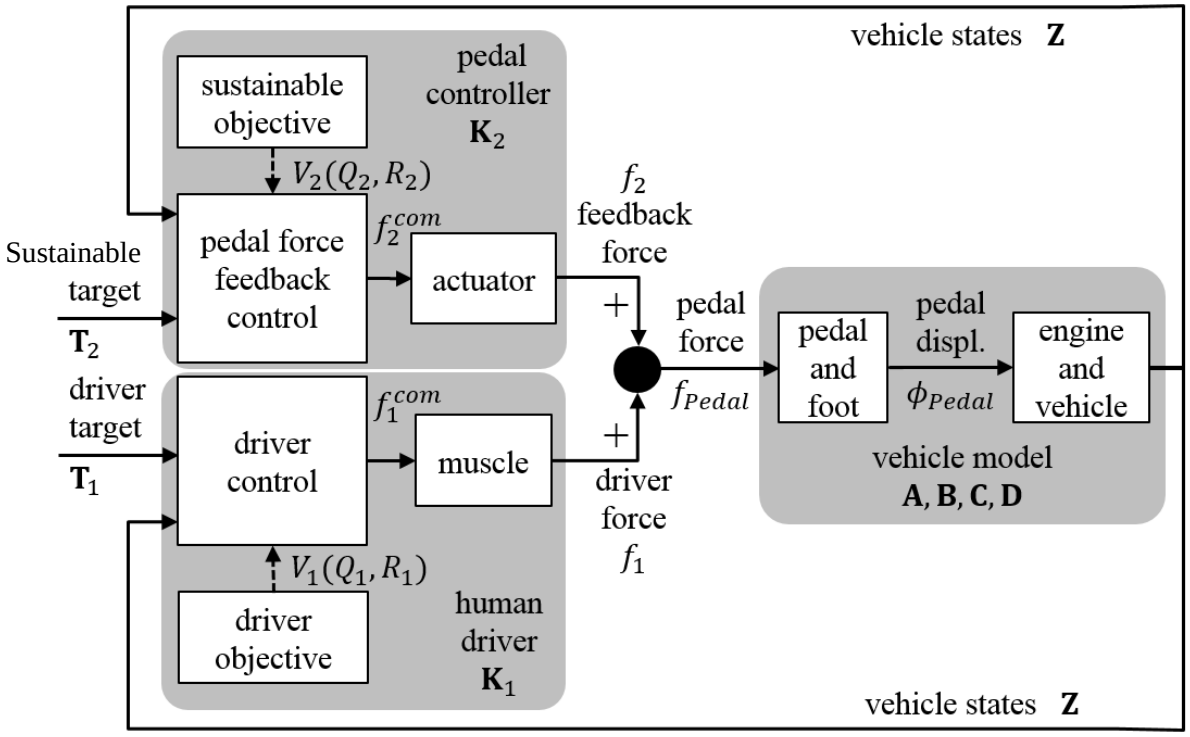

Figure 4-1 – Decentralised Control block diagram

The unassisted driver prediction equation from (3.24) is:

$$\mathbf{Z}_1(\mathrm{k}) = \mathbf{\Psi}_1\mathbf{x}(\mathrm{k}) + \mathbf{\Theta}_1\mathbf{U}_1(\mathrm{k})$$

When modified to include feedback, the prediction equations for both controllers are as follows:

$$\mathbf{Z}_1{}^{\mathrm{dec}}(\mathrm{k}) = \mathbf{\Psi}_1\mathbf{x}(\mathrm{k}) + \mathbf{\Theta}_1\mathbf{U}_1(\mathrm{k}) + \mathbf{\Omega}_1\mathbf{U}_2(\mathrm{k}) \quad (4.1)$$

$$\mathbf{Z}_2{}^{\mathrm{dec}}(\mathrm{k}) = \mathbf{\Psi}_2\mathbf{x}(\mathrm{k}) + \mathbf{\Theta}_2\mathbf{U}_1(\mathrm{k}) + \mathbf{\Omega}_2\mathbf{U}_2(\mathrm{k}) \quad (4.2)$$

Where the superscript ‘dec’ is used to indicate this is the decentralised feedback case. The prediction equations can be combined:

$$\mathbf{Z}^{\mathrm{dec}}(\mathrm{k}) = \mathbf{\Psi}^{\mathrm{dec}}\mathbf{x}(\mathrm{k}) + \mathbf{\Theta}^{\mathrm{dec}}\mathbf{U}_1(\mathrm{k}) + \mathbf{\Omega}^{\mathrm{dec}}\mathbf{U}_2(\mathrm{k}) \quad (4.3)$$

where

$$\boldsymbol{Z}^{\mathrm{dec}}(k) = \begin{bmatrix} \mathbf{Z}_1{}^{\mathrm{dec}}(k) \\ \mathbf{Z}_2{}^{\mathrm{dec}}(k) \end{bmatrix},\ \mathbf{\Psi}^{\mathrm{dec}} = \begin{bmatrix} \mathbf{\Psi}_1 \\ \mathbf{\Psi}_2 \end{bmatrix},\ \mathbf{\Theta}^{\mathrm{dec}} = \begin{bmatrix} \mathbf{\Theta}_1 \\ \mathbf{\Theta}_2 \end{bmatrix} \text{ and } \mathbf{\Omega}^{\mathrm{dec}} = \begin{bmatrix} \mathbf{\Omega}_1 \\ \mathbf{\Omega}_2 \end{bmatrix}$$

Each 'player' has their own cost function to determine performance. The cost functions for the driver and feedback pedal controller, respectively, are below:

$$V_1(k) = \sum_{i=1}^{N_p} \|\boldsymbol{z_1}(k+i) - \boldsymbol{r_1}(k+i)\|^2_{\boldsymbol{Q1}(i)} + \sum_{i=0}^{N_p-1} \|f_1(k+i)\|^2_{\boldsymbol{R1}(i)} \quad (4.4)$$

$$V_2(k) = \sum_{i=1}^{N_p} \|\boldsymbol{z_2}(k+i) - \boldsymbol{r_2}(k+i)\|^2_{\boldsymbol{Q2}(i)} + \sum_{i=0}^{N_p-1} \|f_2(k+i)\|^2_{\boldsymbol{R2}(i)} \quad (4.5)$$

Using the defined cost functions, and the approach outlined in Chapter 3, the forces at any time step are hence given by:

$$\mathrm{f}_1(\mathrm{k}) = \mathbf{K}_{1\mathrm{p}} \begin{Bmatrix} \mathbf{x_1}(\mathrm{k}) \\ \mathbf{T_1}(\mathrm{k}) \end{Bmatrix} \quad (4.6)$$

And

$$\mathrm{f}_2(\mathrm{k}) = \mathbf{K}_{2\mathrm{p}} \begin{Bmatrix} \mathbf{x_2}(\mathrm{k}) \\ \mathbf{T_2}(\mathrm{k}) \end{Bmatrix} \quad (4.7)$$

4.2 Two Player Cooperative Feedback Theory

The second framework to be examined here it is that of cooperative control. In this framework, the two controllers have knowledge of the other's objectives and targets (Figure 4-2 – Cooperative Control with independent reference signals block diagram) and are prepared to work together to reach a common goal. This common goal is a compromise between each individual target's objectives.

The driver's understanding of the vehicle may be very different to the internal model the feedback controller has of the vehicle. For this reason, in the derivation of both controllers, we explicitly consider the case where the internal models are different.

The prediction equations for both controllers are as follows:

$$\mathbf{Z}_1{}^{\mathrm{coop}}(\mathrm{k}) = \boldsymbol{\Psi}_1\mathbf{x}(\mathrm{k}) + \boldsymbol{\Theta}_1\mathbf{U}_1(\mathrm{k}) + \boldsymbol{\Omega}_1\mathbf{U}_2(\mathrm{k}) \quad (4.8)$$

$$\mathbf{Z}_2{}^{\mathrm{coop}}(\mathrm{k}) = \boldsymbol{\Psi}_2\mathbf{x}(\mathrm{k}) + \boldsymbol{\Theta}_2\mathbf{U}_1(\mathrm{k}) + \boldsymbol{\Omega}_2\mathbf{U}_2(\mathrm{k}) \quad (4.9)$$

Where definitions are as in Section 3.2 but subscript numbers denote relevant player: Player 1 is the human driver and player 2, the feedback controller.

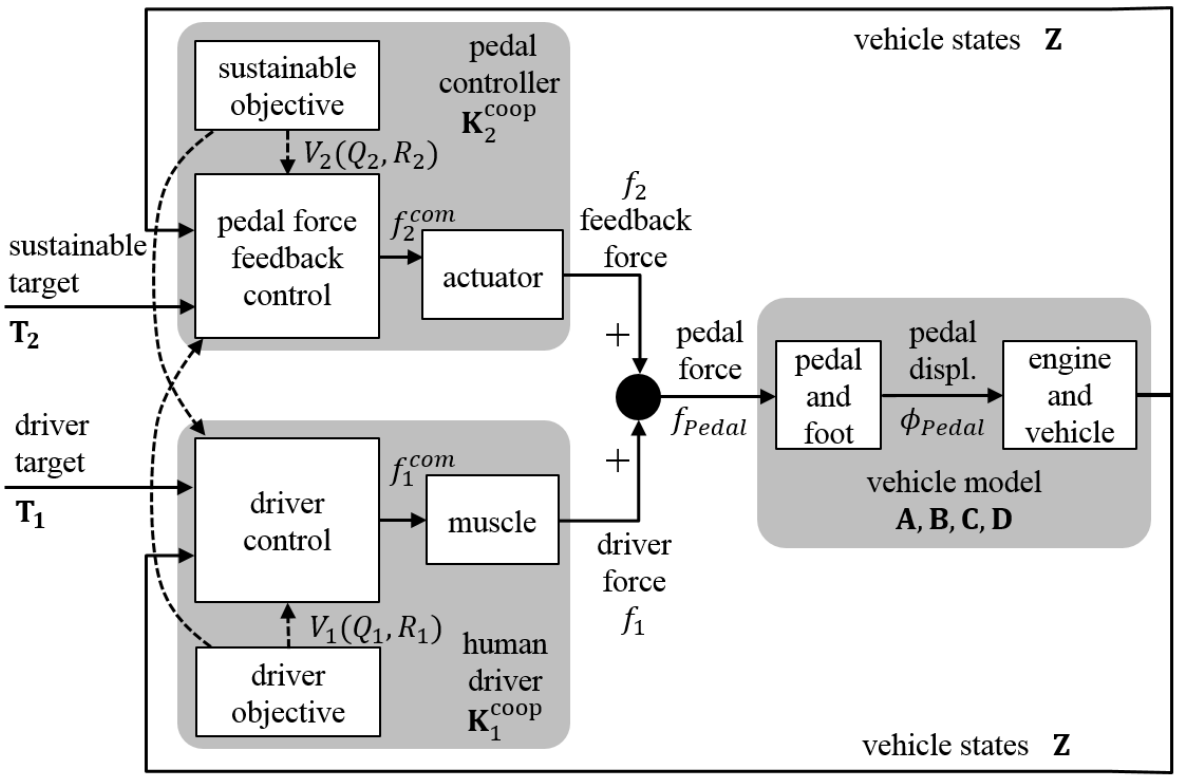

Figure 4-2 – Cooperative Control with independent reference signals block diagram

These two equations are combined as follows:

$$\mathbf{Z}^{\text{coop}}(\text{k}) = \mathbf{\Psi}^{\text{coop}}\mathbf{x}(\text{k}) + \mathbf{\Theta}^{\text{coop}}\mathbf{U}_1(\text{k}) + \mathbf{\Omega}^{\text{coop}}\mathbf{U}_2(\text{k}) \quad (4.10)$$

Where

$$\boldsymbol{Z}^{\text{coop}}(k) = \begin{bmatrix} \mathbf{Z}_1{}^{\text{coop}}(k) \\ \mathbf{Z}_2{}^{\text{coop}}(k) \end{bmatrix},\ \mathbf{\Psi}^{\text{coop}} = \begin{bmatrix} \mathbf{\Psi}_1 \\ \mathbf{\Psi}_2 \end{bmatrix},\ \mathbf{\Theta}^{\text{coop}} = \begin{bmatrix} \mathbf{\Theta}_1 \\ \mathbf{\Theta}_2 \end{bmatrix} \text{ and } \mathbf{\Omega}^{\text{coop}} = \begin{bmatrix} \mathbf{\Omega}_1 \\ \mathbf{\Omega}_2 \end{bmatrix}$$

Following Rawlings and Mayne (2009) analysis, a common objective, or plant objective, is set as a weighted sum of the individual controllers' cost functions. Each controller tries to maximise performance whilst minimising its own control effort by penalising the plant objective cost alongside its control effort:

$$\begin{aligned} V_1^{\text{coop}}(k) = \frac{1}{2}\sum_{j=0}^{N_p} \Big[& \big(\mathbf{z}(k+j) - \mathbf{r}_1(k+j)\big)^{\text{T}} \rho_1 \mathbf{Q}_1 \big(\mathbf{z}(k+j) - \mathbf{r}_1(k+j)\big) \\ & + \big(\mathbf{z}(k+j) - \mathbf{r}_2(k+j)\big)^{\text{T}} \rho_2 \mathbf{Q}_2 \big(\mathbf{z}(k+j) - \mathbf{r}_2(k+j)\big) \Big] \\ & + \frac{1}{2}\rho_1 \sum_{j=0}^{N_u - 1} [f_1(k+j)^{\text{T}} \mathbf{R}_1 f_1(k+j)] \end{aligned} \quad (4.11)$$

$$V_2^{\text{coop}}(k) = \frac{1}{2}\sum_{j=0}^{N_p}\Big[\big(\mathbf{z}(k+j) - \mathbf{r}_1(k+j)\big)^{\mathrm{T}}\rho_1\mathbf{Q}_1\big(\mathbf{z}(k+j) - \mathbf{r}_1(k+j)\big) + \big(\mathbf{z}(k+j) - \mathbf{r}_2(k+j)\big)^{\mathrm{T}}\rho_2\mathbf{Q}_2\big(\mathbf{z}(k+j) - \mathbf{r}_2(k+j)\big)\Big] + \frac{1}{2}\rho_2\sum_{j=0}^{N_u-1}[f_2(k+j)^{\mathrm{T}}\mathbf{R}_2 f_2(k+j)]$$

(4.12)

Where $\mathbf{z}(k), \mathbf{r}_i(k), \mathbf{Q}_i(k)$ and $\mathbf{R}_i$ correspond to definitions given in 6.2 and subscript i indicating the relevant player. ρ_i is a parameter used to specify the relevant weights on the controllers' individual target-following performance in the plant objective.

In matrix form:

$$\mathrm{V}_1^{\text{coop}}(\mathrm{k}) = [\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})]^{\mathrm{T}}\begin{bmatrix}\mathbf{Q}_1^{\text{coop}} & 0\\ 0 & \mathbf{Q}_2^{\text{coop}}\end{bmatrix}[\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})] + \mathbf{U}_1(\mathrm{k})^{\mathrm{T}}\mathbf{R}_1^{\text{coop}}\mathbf{U}_1(\mathrm{k})$$

(4.13)

$$\mathrm{V}_2^{\text{coop}}(\mathrm{k}) = [\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})]^{\mathrm{T}}\begin{bmatrix}\mathbf{Q}_1^{\text{coop}} & 0\\ 0 & \mathbf{Q}_2^{\text{coop}}\end{bmatrix}[\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})] + \mathbf{U}_2(\mathrm{k})^{\mathrm{T}}\mathbf{R}_2^{\text{coop}}\mathbf{U}_2(\mathrm{k})$$

(4.14)

Where

$$\mathbf{T}^{\text{coop}}(k) = \begin{bmatrix}\mathbf{T}_1(k)\\ \mathbf{T}_2(k)\end{bmatrix},\ Q_i^{coop} = \rho_i\begin{bmatrix}\mathbf{Q}_i & 0 & \cdots & 0\\ 0 & \mathbf{Q}_i & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & \mathbf{Q}_i\end{bmatrix} \text{ and } \mathbf{R}_i^{coop} = \rho_i\begin{bmatrix}\mathbf{R}_i & 0 & \cdots & 0\\ 0 & \mathbf{R}_i & \cdots & 0\\ \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & \cdots & \mathbf{R}_i\end{bmatrix}$$

Equivalently:

$$\mathrm{V}_1^{\text{coop}}(\mathrm{k}) = \|\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})\|^2_{\mathbf{Q}^{\text{coop}}} + \|\mathbf{U}_1(\mathrm{k})\|^2_{\mathbf{R}_1^{\text{coop}}} \quad (4.15)$$

$$\mathrm{V}_2^{\text{coop}}(\mathrm{k}) = \|\mathbf{Z}^{\text{coop}}(\mathrm{k}) - \mathbf{T}^{\text{coop}}(\mathrm{k})\|^2_{\mathbf{Q}^{\text{coop}}} + \|\mathbf{U}_2(\mathrm{k})\|^2_{\mathbf{R}_2^{\text{coop}}} \quad (4.16)$$

Where

$$\mathbf{Q}^{\text{coop}} = \begin{bmatrix}\mathbf{Q}_1^{\text{coop}} & 0\\ 0 & \mathbf{Q}_2^{\text{coop}}\end{bmatrix}$$

Define tracking errors as:

$$\boldsymbol{\varepsilon}_1^{\text{coop}}(\text{k}) = \mathbf{T}^{\text{coop}}(\text{k}) - \boldsymbol{\Psi}^{\text{coop}}(\text{k})\mathbf{x}(\text{k}) - \boldsymbol{\Omega}^{\text{coop}}\mathbf{U}_2(\text{k}) \quad (4.17)$$

$$\boldsymbol{\varepsilon}_2^{\text{coop}}(\text{k}) = \mathbf{T}^{\text{coop}}(\text{k}) - \boldsymbol{\Psi}^{\text{coop}}(\text{k})\mathbf{x}(\text{k}) - \boldsymbol{\Theta}^{\text{coop}}\mathbf{U_1}(\text{k}) \quad (4.18)$$

So cost function is

$$\text{V}_1^{\text{coop}}(\text{k}) = \left\|\boldsymbol{\Theta}^{\text{coop}}\mathbf{U}_1(\text{k}) - \boldsymbol{\varepsilon}_1^{\text{coop}}(\text{k})\right\|^2_{\mathbf{Q}^{\text{coop}}} + \|\mathbf{U}_1(\text{k})\|^2_{\mathbf{R}_1^{\text{coop}}} \quad (4.19)$$

$$\text{V}_2^{\text{coop}}(\text{k}) = \left\|\boldsymbol{\Theta}^{\text{coop}}\mathbf{U}_2(\text{k}) - \boldsymbol{\varepsilon}_2^{\text{coop}}(\text{k})\right\|^2_{\mathbf{Q}^{\text{coop}}} + \|\mathbf{U}_2(\text{k})\|^2_{\mathbf{R}_2^{\text{coop}}} \quad (4.20)$$

Alternatively expressed as

$$\text{V}_1^{\text{coop}}(\text{k}) = \left\|\begin{matrix}\mathbf{S}_{\text{Q}^{\text{coop}}}\{\boldsymbol{\Theta}^{\text{coop}}\mathbf{U}_1(\text{k}) - \boldsymbol{\varepsilon}_1^{\text{coop}}(\text{k})\} \\ \mathbf{S}_{\text{R}_1^{\text{coop}}}\mathbf{U}_1(\text{k})\end{matrix}\right\|^2 \quad (4.21)$$

$$\text{V}_2^{\text{coop}}(\text{k}) = \left\|\begin{matrix}\mathbf{S}_{\text{Q}^{\text{coop}}}\{\boldsymbol{\Theta}^{\text{coop}}\mathbf{U}_2(\text{k}) - \boldsymbol{\varepsilon}_2^{\text{coop}}(\text{k})\} \\ \mathbf{S}_{\text{R}_2^{\text{coop}}}\mathbf{U}_2(\text{k})\end{matrix}\right\|^2 \quad (4.22)$$

Where

$$\mathbf{Q}^{\text{coop}} = \mathbf{S}^{\text{T}}_{\mathbf{Q}^{\text{coop}}}\mathbf{S}_{\mathbf{Q}^{\text{coop}}} \text{ and } \mathbf{R}_i^{\text{coop}} = \mathbf{S}^{\text{T}}_{\mathbf{R}_i^{\text{coop}}}\mathbf{S}_{\mathbf{R}_i^{\text{coop}}}$$

To minimise the cost function:

$$\mathbf{U}_1(\text{k}) = \mathbf{K}_{1\text{full}}^{\text{coop}}\boldsymbol{\varepsilon}_1^{\text{coop}}(\text{k}) \quad (4.23)$$

$$\mathbf{U}_2(\text{k}) = \mathbf{K}_{2\text{full}}^{\text{coop}}\boldsymbol{\varepsilon}_2^{\text{coop}}(\text{k}) \quad (4.24)$$

Where

$$\mathbf{K}_{1\text{full}}^{\text{coop}} = \begin{bmatrix}\mathbf{S}_{\text{Q}^{\text{coop}}}\boldsymbol{\Theta}^{\text{coop}} \\ \mathbf{S}_{\text{R}_1^{\text{coop}}}\end{bmatrix} \setminus \begin{bmatrix}\mathbf{S}_{\text{Q}^{\text{coop}}} \\ 0\end{bmatrix}$$

$$\mathbf{K}_{2\text{full}}^{\text{coop}} = \begin{bmatrix}\mathbf{S}_{\text{Q}^{\text{coop}}}\boldsymbol{\Theta}^{\text{coop}} \\ \mathbf{S}_{\text{R}_2^{\text{coop}}}\end{bmatrix} \setminus \begin{bmatrix}\mathbf{S}_{\text{Q}^{\text{coop}}} \\ 0\end{bmatrix}$$

'\' is the Matlab operator for matrix left division by the QR algorithm.

Substitute equations (4.17) and (4.18) into (4.23) and (4.24) respectively:

$$\mathbf{U}_1(\mathrm{k}) = \left[-\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}} \;\; \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\right]\begin{bmatrix}\mathbf{x}(\mathrm{k})\\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{bmatrix} - \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{U}_2(\mathrm{k}) \quad (4.25)$$

$$\mathbf{U}_2(\mathrm{k}) = \left[-\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}} \;\; \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\right]\begin{bmatrix}\mathbf{x}(\mathrm{k})\\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{bmatrix} - \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\mathbf{U}_1(\mathrm{k}) \quad (4.26)$$

Substitute (4.25) and (4.26) into each other:

$$\mathbf{U}_1(\mathrm{k}) = \left[\boldsymbol{\Lambda}_1^{\mathrm{coop}} \; \boldsymbol{\Gamma}_1^{\mathrm{coop}}\right]\begin{Bmatrix}\mathbf{x}(\mathrm{k})\\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{Bmatrix} \quad (4.27)$$

where

$$\boldsymbol{\Lambda}_1^{\mathrm{coop}} = \left[\mathbf{I} - \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\right]^{-1}\cdot\left[-\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}} + \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}}\right]$$

$$\boldsymbol{\Gamma}_1^{\mathrm{coop}} = \left[\mathbf{I} - \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\right]^{-1}\cdot\left[\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}} - \mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\right]$$

and

$$\mathbf{U}_2(\mathrm{k}) = \left[\boldsymbol{\Lambda}_2^{\mathrm{coop}} \; \boldsymbol{\Gamma}_2^{\mathrm{coop}}\right]\begin{Bmatrix}\mathbf{x}(\mathrm{k})\\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{Bmatrix} \quad (4.28)$$

where

$$\boldsymbol{\Lambda}_2^{\mathrm{coop}} = \left[\mathbf{I} - \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\right]^{-1}\cdot\left[-\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}} + \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Psi}^{\mathrm{coop}}\right]$$

$$\boldsymbol{\Gamma}_1^{\mathrm{coop}} = \left[\mathbf{I} - \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Omega}^{\mathrm{coop}}\right]^{-1}\cdot\left[\mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}} - \mathbf{K}_{2\mathrm{full}}^{\mathrm{coop}}\boldsymbol{\Theta}^{\mathrm{coop}}\mathbf{K}_{1\mathrm{full}}^{\mathrm{coop}}\right]$$

Now, by applying the receding horizon approach, the two controllers, $\mathbf{K}_1^{\mathrm{coop}}$ and $\mathbf{K}_2^{\mathrm{coop}}$ are identified:

$$f_{1\mathrm{opt}}(k) = \mathbf{K}_1^{\mathrm{coop}}\begin{Bmatrix}\mathbf{x}(k)\\ \mathbf{T}^{\mathrm{coop}}(k)\end{Bmatrix} \quad (4.29)$$

where

$$\mathbf{K}_1^{\mathrm{coop}} = \left[\boldsymbol{\Lambda}_1^{\mathrm{coop}}(1,:) \;\; \boldsymbol{\Gamma}_1^{\mathrm{coop}}(1,:)\right]$$

and

$$\mathrm{f}_{2\mathrm{opt}}(\mathrm{k}) = \mathbf{K}_2^{\mathrm{coop}}\begin{Bmatrix}\mathbf{x}(\mathrm{k})\\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{Bmatrix} \quad (4.30)$$

where

$$\mathbf{K}_2^{\mathrm{coop}} = \left[\boldsymbol{\Lambda}_2^{\mathrm{coop}}(1,:) \;\; \boldsymbol{\Gamma}_2^{\mathrm{coop}}(1,:)\right]$$

where $\boldsymbol{\Lambda}_1^{\mathrm{coop}}(1,:)$ means the first row of $\boldsymbol{\Lambda}_1^{\mathrm{coop}}$ and so on.

Practically, it would be very difficult to produce a fully cooperative AAP due to the need to detect changes in the driver's target and cost functions. Different drivers have different styles, and the pedal would need to detect which driver was driving at any one time. For this reason, the two player cooperative case is not taken any further forward, but is included as a theoretical case.

4.3 One-sided Cooperative Feedback

The third case is an attempt to acknowledge that it is difficult for a machine to learn and understand human behaviour, but it is plausible that a human driver will develop an understanding of what the pedal feedback controller is aiming for and may well cooperate with it. Figure 4-3 illustrates this framework in block diagram form.

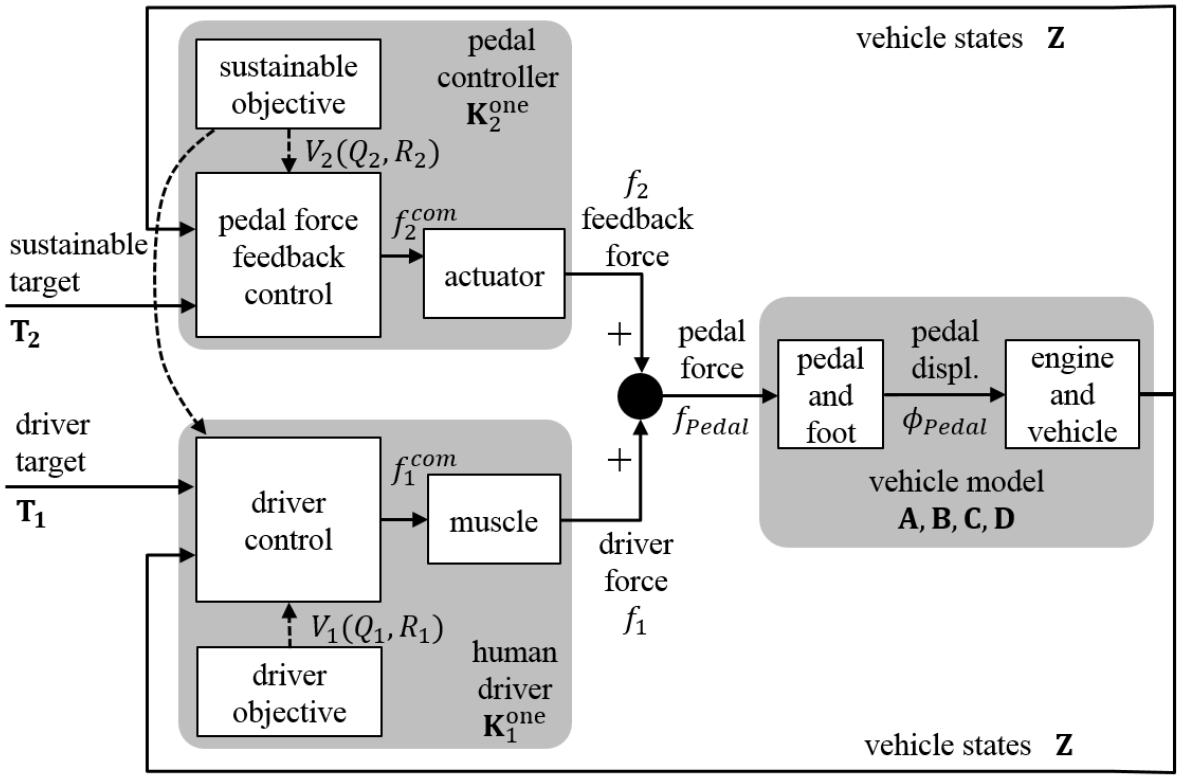

Figure 4-3 – One-sided cooperative control block diagram

Start again with the pair of prediction equations from (4.1) and (4.2):

$$\mathbf{Z}^{\text{one}}(\text{k}) = \mathbf{\Psi}^{\text{coop}}\mathbf{x}(\text{k}) + \mathbf{\Theta}^{\text{one}}\mathbf{U}_1(\text{k}) + \mathbf{\Omega}^{\text{coop}}\mathbf{U}_2(\text{k}) \qquad (4.31)$$

Where

$$\boldsymbol{Z}^{\text{one}}(k) = \begin{bmatrix}\mathbf{Z}_1(k)\\ \mathbf{Z}_2(k)\end{bmatrix} \text{ and } \mathbf{\Theta}^{\text{one}} = \begin{bmatrix}\mathbf{\Theta}_1\\ 0\end{bmatrix}$$

Notice that in place of $\mathbf{\Theta}_2$, we have a zero matrix instead.

The feedback controller is unaware of the driver, so the feedback controller gains can be calculated as the driver model was in 3.2:

$$\mathbf{U}_2(\mathrm{k}) = \mathbf{K}_{2\mathrm{full}}^{\mathrm{one}}\boldsymbol{\varepsilon}_2^{\mathrm{one}}(\mathrm{k}) \quad (4.32)$$

$$\mathbf{K}_{2\mathrm{full}}^{\mathrm{one}} = \begin{bmatrix}\mathbf{S}_{\mathbf{Q}_2}\boldsymbol{\Theta}_2 \\ \mathbf{S}_{\mathbf{R}_2}\end{bmatrix} \backslash \begin{bmatrix}\mathbf{S}_{\mathbf{Q}_2} \\ 0\end{bmatrix} \quad (4.33)$$

By applying the receding horizon approach, the control law becomes:

$$\mathrm{f}_{2\mathrm{opt}}(\mathrm{k}) = \mathbf{K}_2^{\mathrm{one}}\begin{Bmatrix}\mathbf{x}(\mathrm{k}) \\ \mathbf{T}_2(\mathrm{k})\end{Bmatrix} \quad (4.34)$$

where

$$\mathbf{K}_2^{\mathrm{one}} = [-\mathbf{K}_{2\mathrm{full}}^{\mathrm{one}}(1,:)\boldsymbol{\Psi}_2 \;\; \mathbf{K}_{2\mathrm{full}}^{\mathrm{one}}(1,:)]$$

The cost function for the model driver is the same as that used in the cooperative case

$$\mathrm{V}_1^{\mathrm{one}}(\mathrm{k}) = \|\mathbf{Z}^{\mathrm{one}}(\mathrm{k}) - \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\|_{\mathbf{Q}^{\mathrm{coop}}}^2 + \|\mathbf{U}_1(\mathrm{k})\|_{\mathbf{R}_1^{\mathrm{coop}}}^2 \quad (4.35)$$

and hence the controller can be derived in the same way as the cooperative controllers in 4.2 were. The tracking error is defined as

$$\boldsymbol{\varepsilon}_1^{\mathrm{one}}(\mathrm{k}) = \mathbf{T}^{\mathrm{coop}}(\mathrm{k}) - \boldsymbol{\Psi}^{\mathrm{coop}}(\mathrm{k})\mathbf{x}(\mathrm{k}) - \boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{U}_2(\mathrm{k}) \quad (4.36)$$

So the cost function is

$$\mathrm{V}_1^{\mathrm{one}}(\mathrm{k}) = \|\boldsymbol{\Theta}^{\mathrm{one}}\mathbf{U}_1(\mathrm{k}) - \boldsymbol{\varepsilon}_1^{\mathrm{one}}(\mathrm{k})\|_{\mathbf{Q}^{\mathrm{coop}}}^2 + \|\mathbf{U}_1(\mathrm{k})\|_{\mathbf{R}_1^{\mathrm{coop}}}^2 \quad (4.37)$$

Alternatively expressed as

$$\mathrm{V}_1^{\mathrm{one}}(\mathrm{k}) = \left\| \begin{matrix}\mathbf{S}_{\mathbf{Q}^{\mathrm{coop}}}\{\boldsymbol{\Theta}^{\mathrm{one}}\mathbf{U}_1(\mathrm{k}) - \boldsymbol{\varepsilon}_1^{\mathrm{one}}(\mathrm{k})\} \\ \mathbf{S}_{\mathrm{R}_1^{\mathrm{coop}}}\mathbf{U}_1(\mathrm{k})\end{matrix} \right\|^2 \quad (4.38)$$

To minimise the cost function:

$$\mathbf{U}_1(\mathrm{k}) = \mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}\boldsymbol{\varepsilon}_1^{\mathrm{one}}(\mathrm{k}) \quad (4.39)$$

where

$$\mathbf{K}_{1\mathrm{full}}^{\mathrm{one}} = \begin{bmatrix}\mathbf{S}_{\mathrm{Q}^{\mathrm{one}}}\boldsymbol{\Theta}^{\mathrm{one}} \\ \mathbf{S}_{\mathrm{R}_1^{\mathrm{one}}}\end{bmatrix} \backslash \begin{bmatrix}\mathbf{S}_{\mathrm{Q}^{\mathrm{one}}} \\ 0\end{bmatrix}$$

Substituting into (50) into (53):

$$\mathbf{U}_1(\mathrm{k}) = [-\mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}\boldsymbol{\Psi}^{\mathrm{coop}} \;\; \mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}]\begin{bmatrix}\mathbf{x}(\mathrm{k}) \\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k})\end{bmatrix} - \mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{U}_2(\mathrm{k}) \quad (4.40)$$

and then substituting (46) in

$$\mathbf{U}_1(\mathrm{k}) = [-\mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}\boldsymbol{\Psi}^{\mathrm{coop}} \quad \mathbf{K}_{1\mathrm{full}}^{\mathrm{one}} \quad -\mathbf{K}_1^{\mathrm{one}}\boldsymbol{\Omega}^{\mathrm{coop}}\mathbf{K}_{2\mathrm{full}}^{\mathrm{one}}]\begin{Bmatrix}\mathbf{x}(\mathrm{k}) \\ \mathbf{T}^{\mathrm{coop}}(\mathrm{k}) \\ \mathbf{T}_2(\mathrm{k})\end{Bmatrix} \quad (4.41)$$

Following Wang (2015), by dividing $\mathbf{K}_{1\mathrm{full}}^{\mathrm{one}}$ into two components $\mathbf{K}_{1\mathrm{full}1}^{\mathrm{one}}$ and $\mathbf{K}_{1\mathrm{full}2}^{\mathrm{one}}$ that correspond to the two previews $\mathbf{T}_1$ and $\mathbf{T}_2$ respectively, the control law can be written as:

$$f_{1opt}(k) = [\mathbf{K}_{s1}^{one}\ \mathbf{K}_{p1}^{one}\ \mathbf{K}_{p2}^{one}]\begin{Bmatrix}\mathbf{x}(k)\\ \mathbf{T}^{coop}(k)\end{Bmatrix} \quad (4.42)$$

where

$$\mathbf{K}_{s1}^{one} = -\mathbf{K}_{1full}^{one}\boldsymbol{\Psi}^{coop} + \mathbf{K}_{1full}^{one}\boldsymbol{\Omega}^{coop}\mathbf{K}_{2full}^{one}\boldsymbol{\Psi}_2$$

$$\mathbf{K}_{p1}^{one} = \mathbf{K}_{1full1}^{one}$$

$$\mathbf{K}_{s2}^{one} = \mathbf{K}_{1full2}^{one} - \mathbf{K}_{1full}^{one}\boldsymbol{\Omega}^{coop}\mathbf{K}_{2full}^{one}$$

Finally, using the receding horizon method, the control law can be condensed to

$$f_{1opt}(k) = \mathbf{K}_{1}^{one}\begin{Bmatrix}\mathbf{x}(k)\\ \mathbf{T}^{coop}(k)\end{Bmatrix} \quad (4.43)$$

where

$$\mathbf{K}_{1}^{one} = [\mathbf{K}_{s1}^{one}(1,:)\ \mathbf{K}_{p1}^{one}(1,:)\ \mathbf{K}_{p2}^{one}(1,:)]$$

4.4 Driving Experiments and Data Analysis

With two frameworks in place for the modelling of the driver's interaction with pedal feedback, a series of experiments were designed and performed to identify and validate the models.

4.4.1 Experiment Design

The fixed base driving simulator was used once again, and a full description is in Section 3.4. In this experiment, however, the pedal motor was powered, and exerted feedback forces on the pedal.

For the drive-cycle following task, the feedback pedal controller was given speed error weighting, $q_{2s} = 80$ and pedal position weighting, $q_{2\phi} = 50$ to provide a feedback pedal controller driving style more focussed on the reduction in pedal displacements than the human drivers measured in Chapter 3. The car-following feedback pedal controller was given speed error weighting $q_{2s} = 80$, pedal position weighting $q_{2\phi} = 50$ and following distance error weighting $q_{2d} = 40$ (for the same reason as the drive cycle cost function weightings) and a target following distance of 2.5s THW plus a constant 5m offset to avoid the risk of collisions at low speeds. In both the drive cycle and car following scenario, the feedback pedal controllers generate their previews using a constant acceleration assumption – the controller identifies the current target acceleration (either demanded drive cycle acceleration, or acceleration of the vehicle in front) and assumes this remains constant throughout the preview horizon, enabling the calculation of demanded speeds. For the human driver in the drive-cycle scenario, an accurate preview of upcoming speed demands is available. For the car following scenario, the human driver, like the feedback controller, has to make assumptions on how the target vehicle will behave.

4.4.2 Participants

The same nine participants were selected for this experiment as the experiments in 3.4, so all drivers had already completed the unassisted tasks.

4.4.3 Procedure

Once again, the experiment consisted of two scenarios.

In the first scenario, drive cycle following, a four minute segment of the Modified Millbrook Suburban Drive Cycle for HGVs was set as the target speed. The driver was given a practice run through to familiarise themselves with the behaviour of the vehicle, the dynamics of the pedal and the characteristics of the feedback. Drivers were asked to follow the target speed. The test was then completed three times, to improve the reliability of the results.

In the second scenario, car following, the speed of the target vehicle was set as the Modified Millbrook Suburban Drive Cycle for HGVs. Again, the driver was allowed a familiarisation run to get used to the virtual world and the behaviour of the feedback controller. Drivers were asked to follow the vehicle in front at a safe distance, as they normally would on the road. Again, once the familiarisation run was completed, the test was repeated three times.

When all drivers had completed the experiment, the data was analysed and differences in driver behaviours compared.

4.4.4 Drive Cycle Results

Figure 4-4 illustrates the recorded mean speed error profiles from the drivers. Generally, the drivers maintained a speed quite close to the target speed profile, as they were assisted by the pedal feedback. From further analysis, however, it is noted that the higher errors tended to occur after high accelerations. At lower acceleration rates, the drivers were better at anticipating the future speed demand. At speed peaks, the human drivers tended to keep accelerating until the maximum speed was reached, and then coast down until they reached the target speed again.

The human drivers appear to be more consistently closer to the target with the aid of the pedal feedback device. Speed errors in this scenario mostly remain within the range $-1.0\ \mathrm{ms}^{-1}$ to $0.7\ \mathrm{ms}^{-1}$, whereas in the unassisted case of Chapter 3 (Figure 3-21), the speed errors spread within the range $-1.5\ \mathrm{ms}^{-1}$ to $1.0\ \mathrm{ms}^{-1}$. The pedal feedback has contributed to a 40% drop in the range of speed errors.

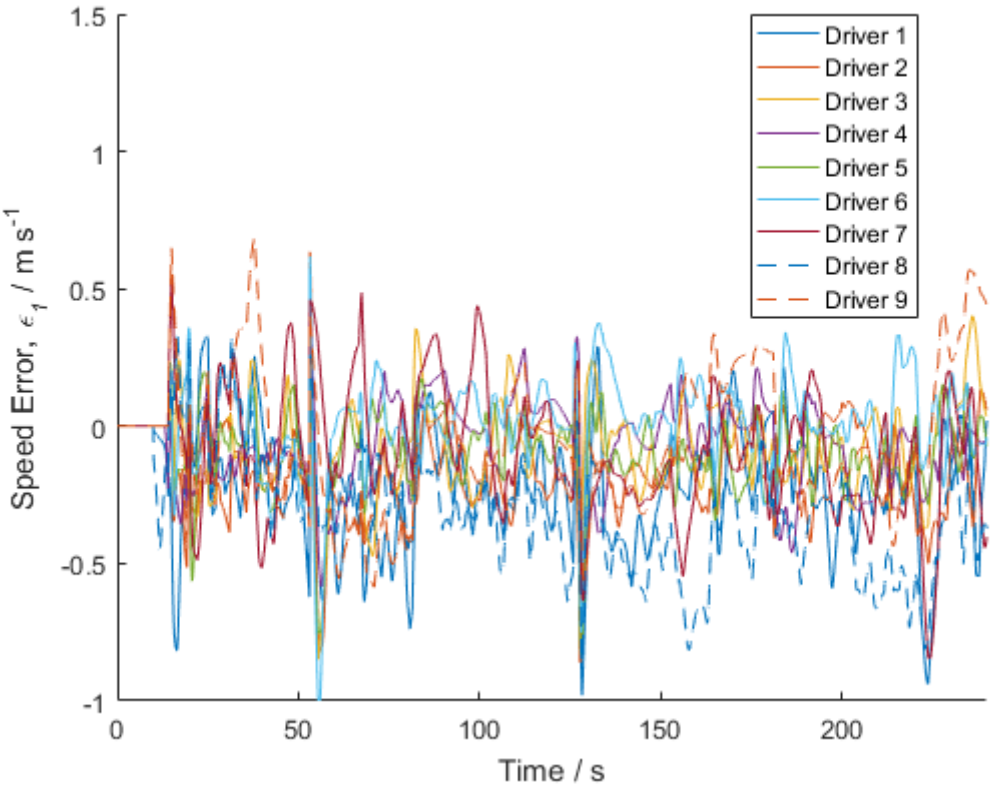

Figure 4-4 – Mean speed error against time for the drive cycle following task

The introduction of pedal feedback has resulted in several high peaks in pedal position over the drive cycle - Figure 4-5. The peaks are more numerous and of greater magnitude than those that were observed in the unassisted case (Figure 3-22). This suggests that the feedback controller is very sensitive to changes in acceleration, and that when it detects one (its previews operate on a constant acceleration assumption, so the controllers are unable to anticipate it) the controller reacts very quickly, resulting in the pedal position peaks. Although not very clear in the figure, the driver's intermittent control is still apparent when they are assisted by the pedal feedback.

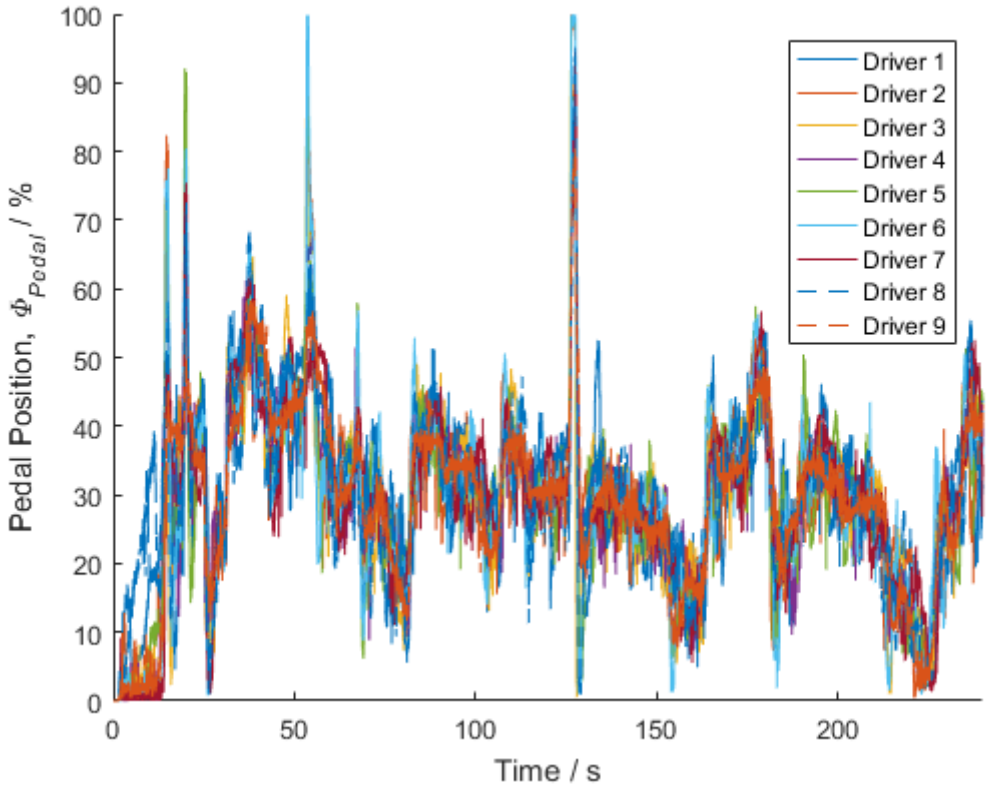

Figure 4-5 – Mean pedal position against time for the drive cycle following task

By comparing the driver pedal forces for the unassisted and assisted cases, (Figure 4-6 with Figures 3-23 show the forces for the full drive cycle), it is noted that the magnitude of the forces from the driver decreases by approximately 50%, suggesting that the pedal feedback has helped decrease speed errors, whilst keeping effort on the driver's part, low. Note that, as mentioned in Chapter 3, there is some concern about the validity of pedal forces at low pedal displacements, caused by a small clip on the pedal. The driver pedal force does not appear in the fitting approaches described in later chapters, as the pedal displacement and feedback force are sufficient.

There is little variation in feedback force across the drivers. A spread of approximately 4N of feedback force covers the range of drivers – Figure 4-7. The feedback force remains near constant during constant accelerations.

It is interesting to note that the addition of pedal feedback has greatly reduced the RMS pedal forces for the drivers, suggesting the workload has decreased, but that the RMS speed errors have not significantly decreased - Figure 4-8 compared with Figure 3-24. This is especially the case for Driver 9, which has switched from one the highest RMS pedal forces in the unassisted case, to the lowest RMS pedal force in this instance. This implies that the pedal feedback has had a big impact on the effort Driver 9 puts into the control of the vehicle.

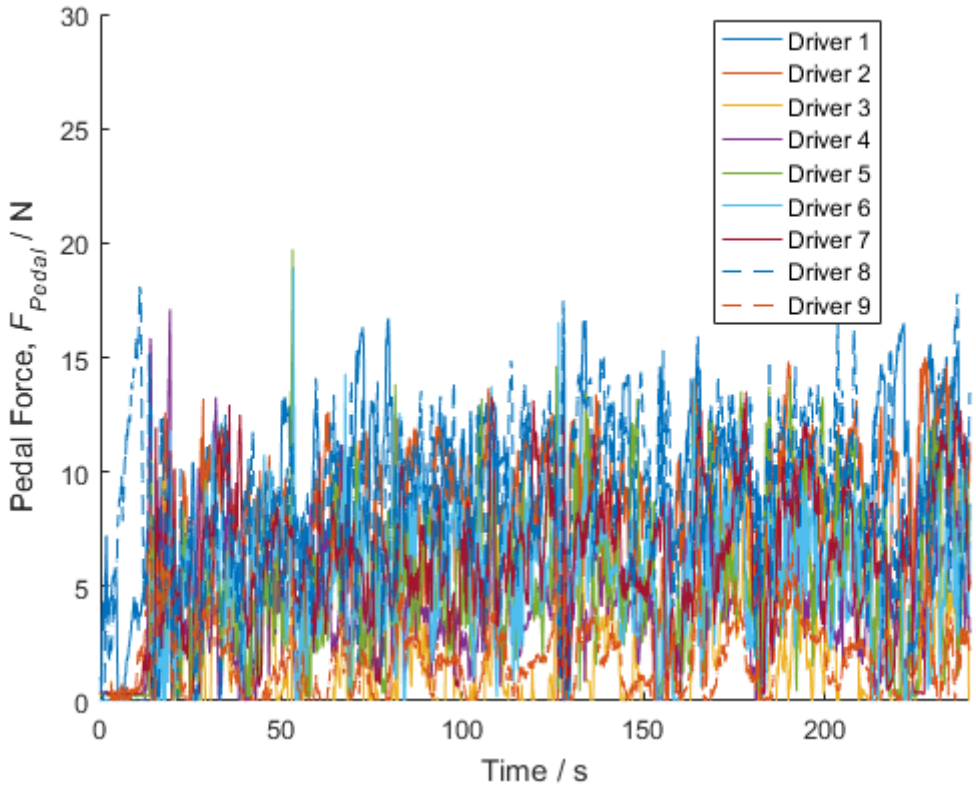

Figure 4-6 – Mean driver pedal force against time for the drive cycle following task

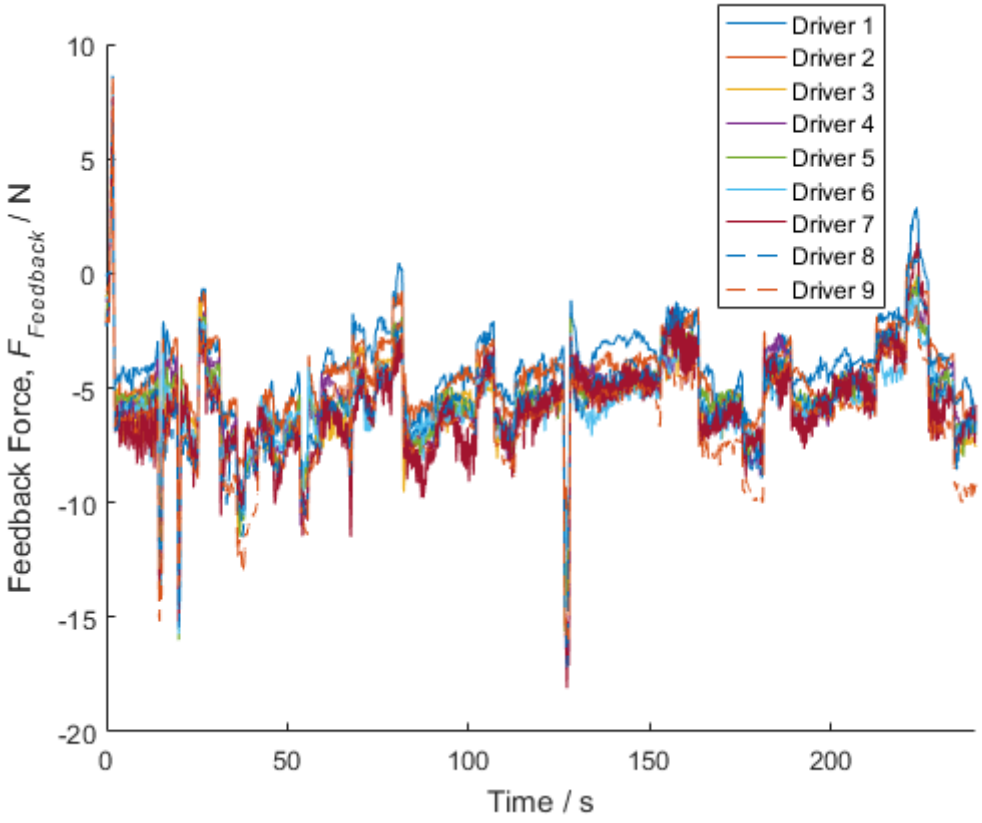

Figure 4-7 – Mean feedback force against time for the drive cycle following task

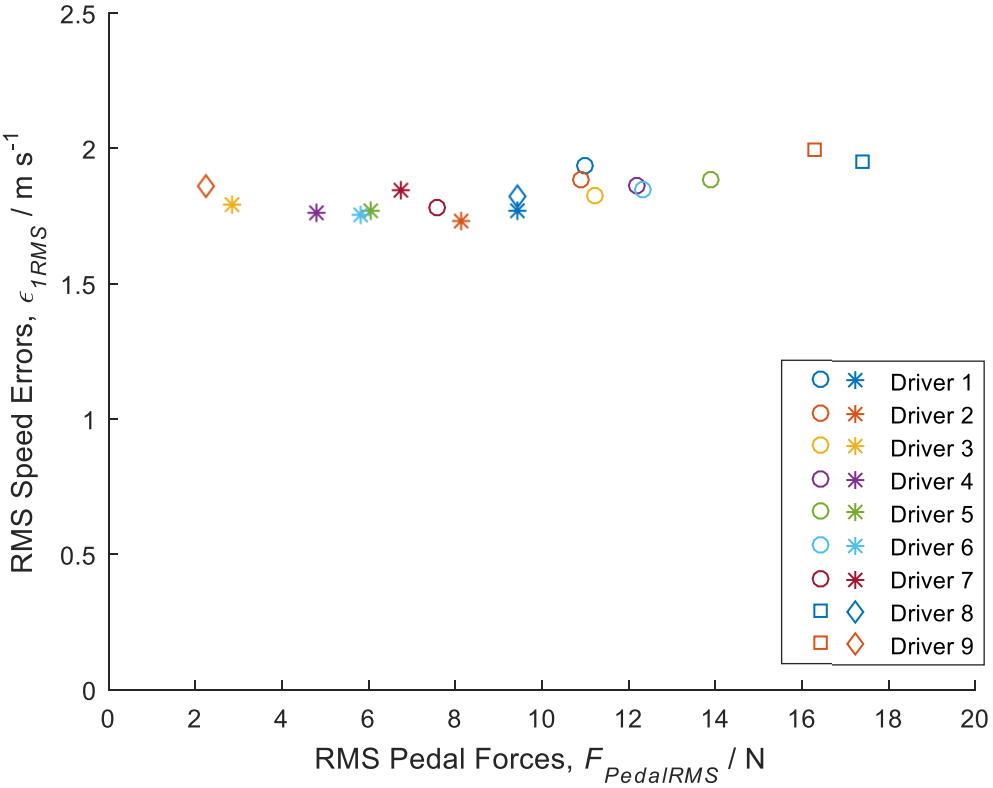

Figure 4-8 – Cost diagram for the drive cycle scenario with feedback – stars (non-professionals) and diamonds (professionals) illustrate data from the assisted drivers. Their original data points from the unassisted case are illustrated with circles (non-professionals) and squares (professionals). Colours are unchanged.

As a final consideration, the mean vehicle speed for each driver is plotted with one standard deviation of the individual's three runs shaded either side in Figure 4-9. Once again, this clearly demonstrates that the standard deviations within each driver is similar to the standard deviation across drivers in this scenario, meaning that there is not a statistically significant difference between most of the drivers across most of the test. It is interesting to note that the addition of pedal feedback has greatly reduced the standard deviations in vehicle speed compared with the unassisted case in Figure 3-25.

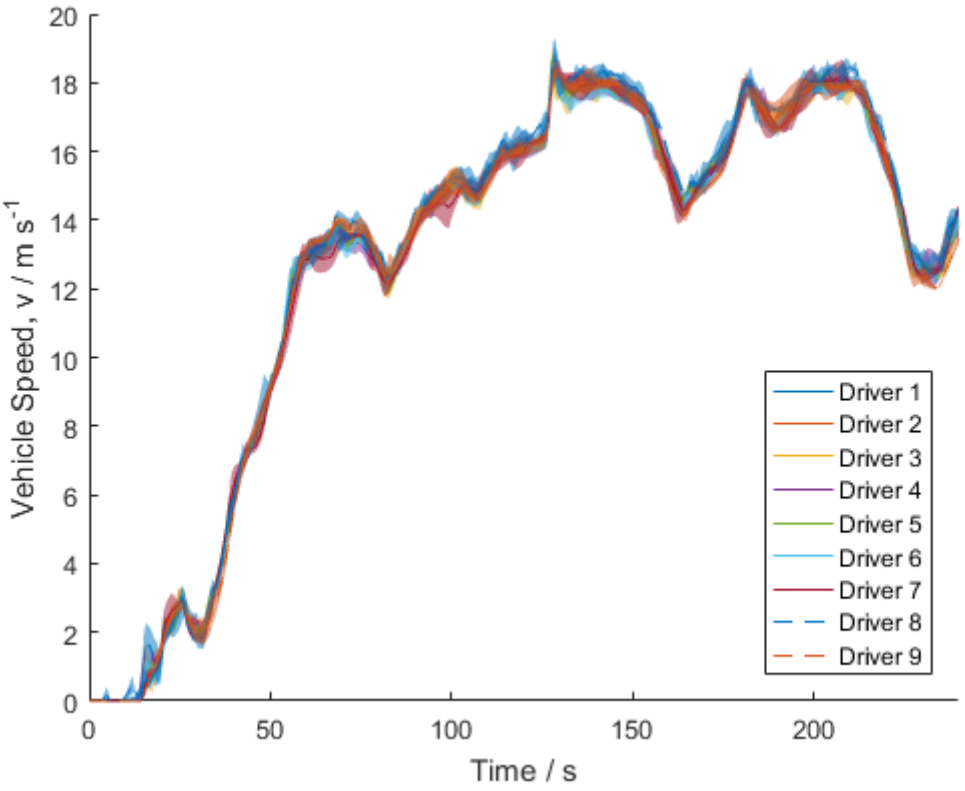

Figure 4-9 - Driver mean vehicle speeds with one driver standard deviation shaded either side

4.4.5 Car-Following Results

The mean speed profiles for the assisted drivers in the car-following scenario are displayed in Figure 4-10. All drivers remain quite close to the speed of the target vehicle, but it is noted that Drivers 1 and 9 drift further from the target vehicle speed than other drivers.

Now observing the pedal position in Figure 4-11, it can be seen that the addition of pedal feedback has smoothed the pedal position curve when compared to Figures 3-26. This is due to the fact that the pedal feedback assesses and acts at every time step of 0.02s, whereas the driver, when implementing intermittent control, can go for much longer (up to a second) without updating. The addition of pedal feedback has also reduced the peaks in pedal position. The unassisted driver would frequently depress the pedal to 60% and occasionally up to 80%. The assisted driver, on the other hand only occasionally depresses above 60%. The reduced pedal positions are an indicator of improved fuel usage.

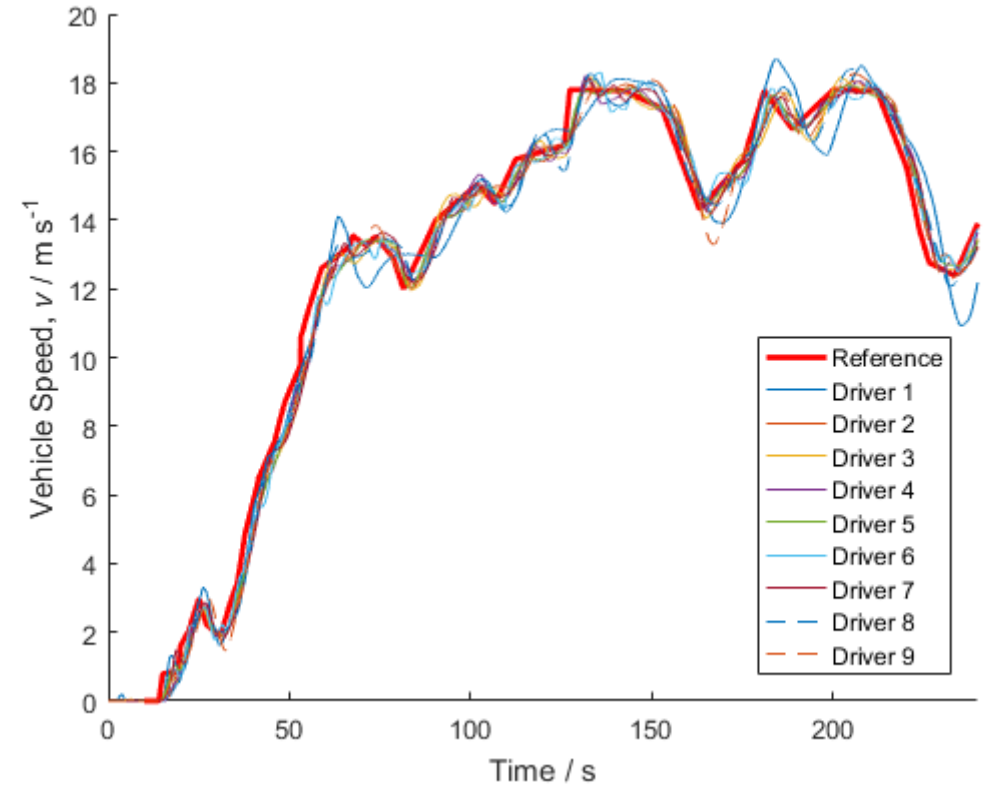

Figure 4-10 – Mean vehicle speed against time for the car-following task

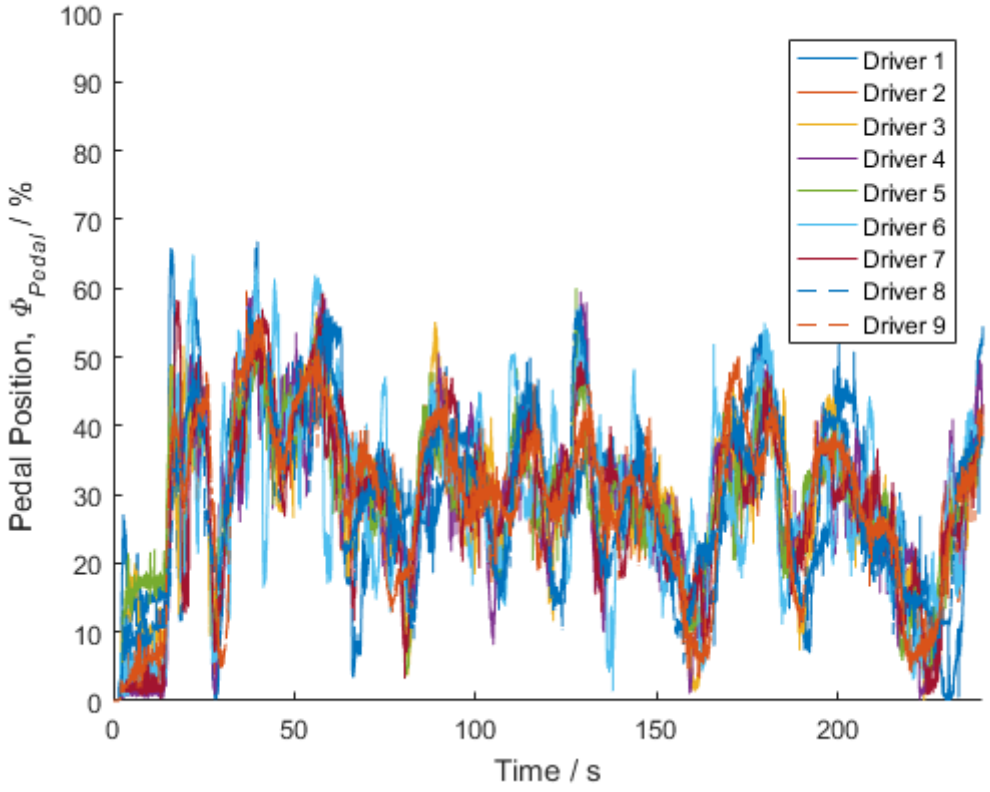

Figure 4-11 – Mean pedal position against time for the car-following task

As already mentioned, the speed profile for Driver 1 looks different to the other drivers'. This is matched by the driver's pedal force (Figure 4-12). The other drivers maintain forces in the 0-35N bracket, whereas Driver 1 is exerting forces up to 45N. This implies that the driver is not prepared to accept the guidance of the feedback controller and therefore tries to oppose it. The pedal feedback force (Figure 4-13) for Driver 1 is also significantly higher than the other drivers.

The result of this conflict between Driver 1 and the feedback controller is a vehicle that has poor vehicle following performance (Figure 4-14 and Figure 4-15) whilst having a highly fluctuating pedal position and vehicle speed, which are likely to have negative effects on vehicle fuel consumption. The other drivers on the other hand have utilised the guidance from the feedback controller to adopt a more consistent vehicle following strategy, and now accurately follow the target vehicle at a 2.5s THW gap.

Unlike the drive cycle scenario, the addition of pedal feedback here has increased both the drivers' RMS pedal forces and RMS following distance errors in some cases - Figure 4-16. This is because the RMS following distance error is based on the target estimated from the driver in the unassisted case, as explained in 3.4.5. The pedal feedback guides the driver towards a different THW target, resulting in greater errors relative to the driver's original target. This measure is used, however, as a direct comparison with the unassisted case in Chapter 3. The RMS pedal forces also dramatically increases for Driver 1 as the driver tries and oppose the pedal feedback to remain close to their original target.

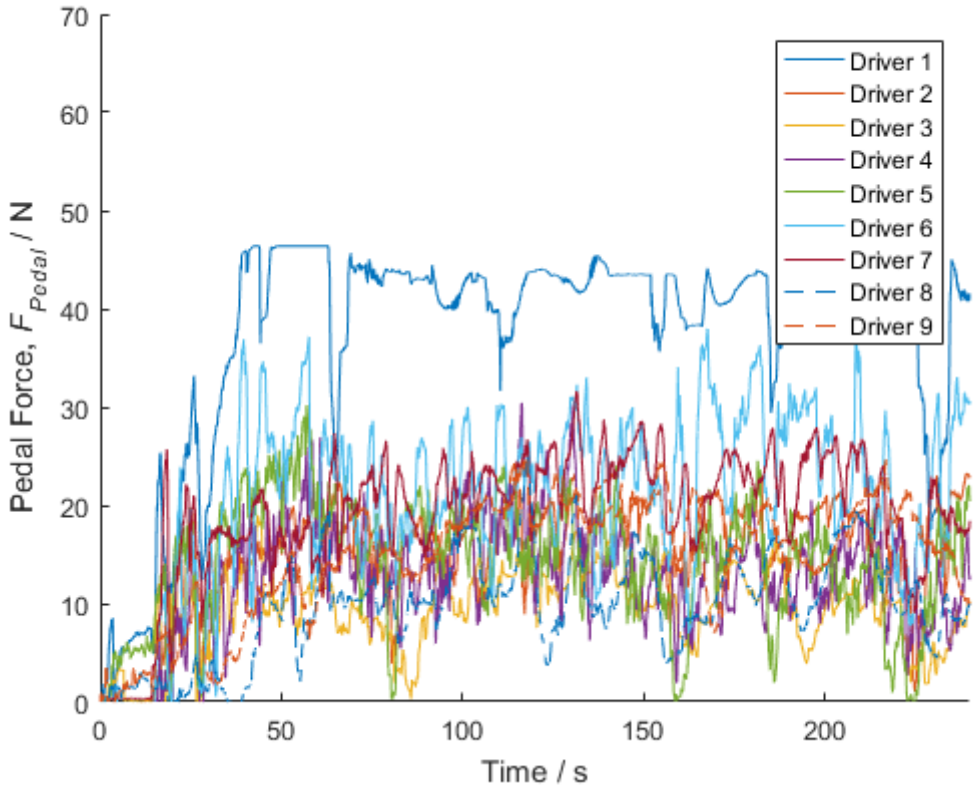

Figure 4-12 – Mean driver pedal force against time for the car-following task

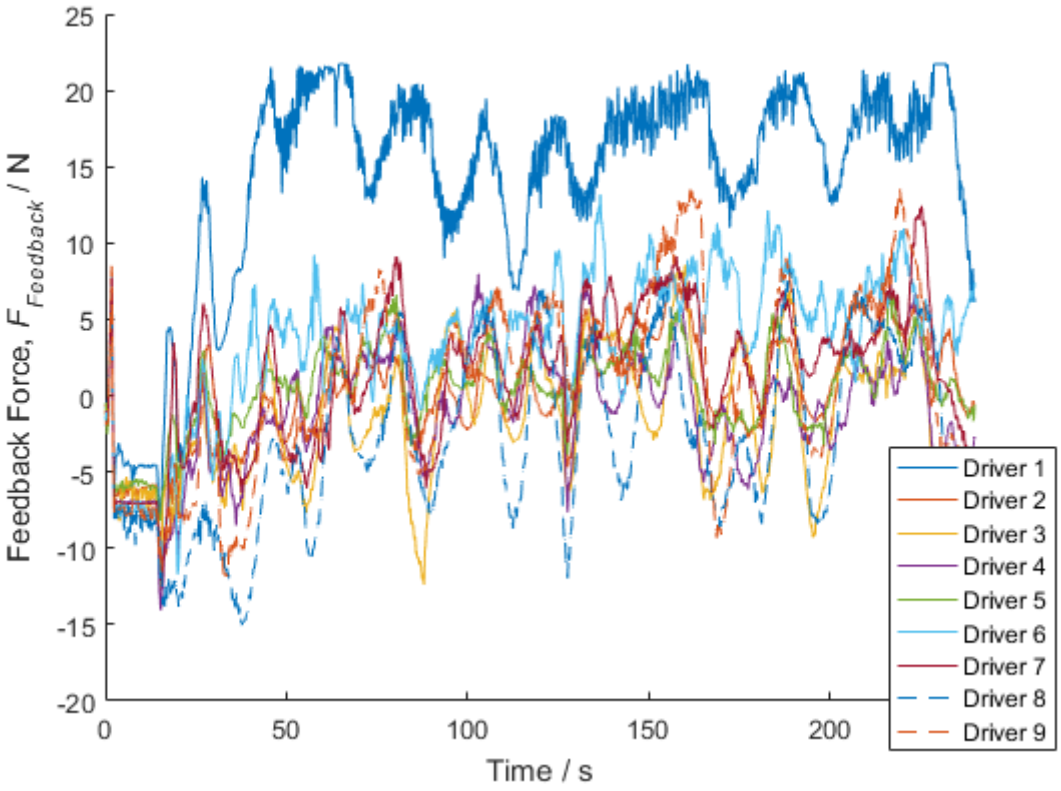

Figure 4-13 – Mean feedback force against time for the car-following task

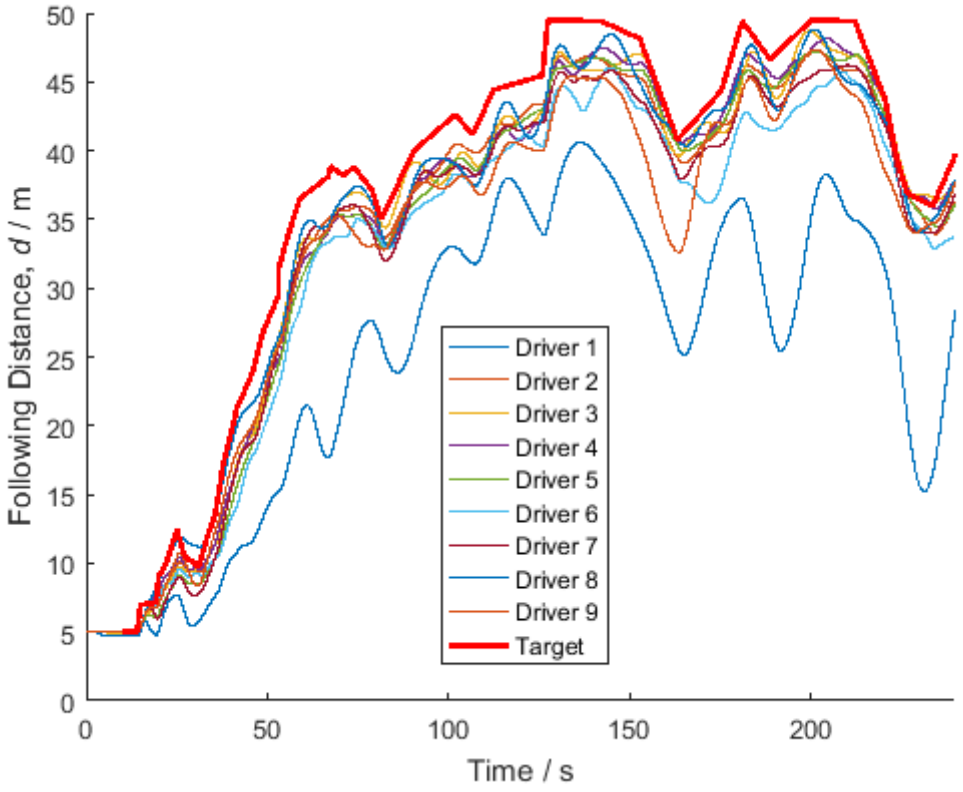

Figure 4-14 – Mean following distance against time for the car-following task

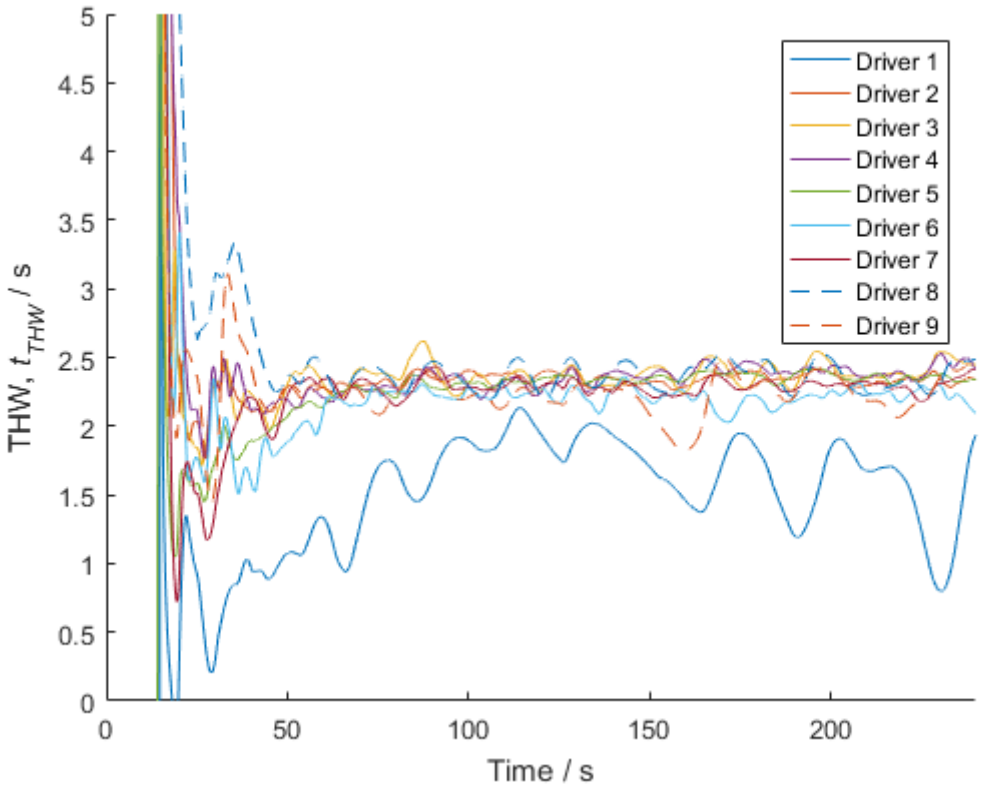

Figure 4-15 – Mean THW against time for the car-following task

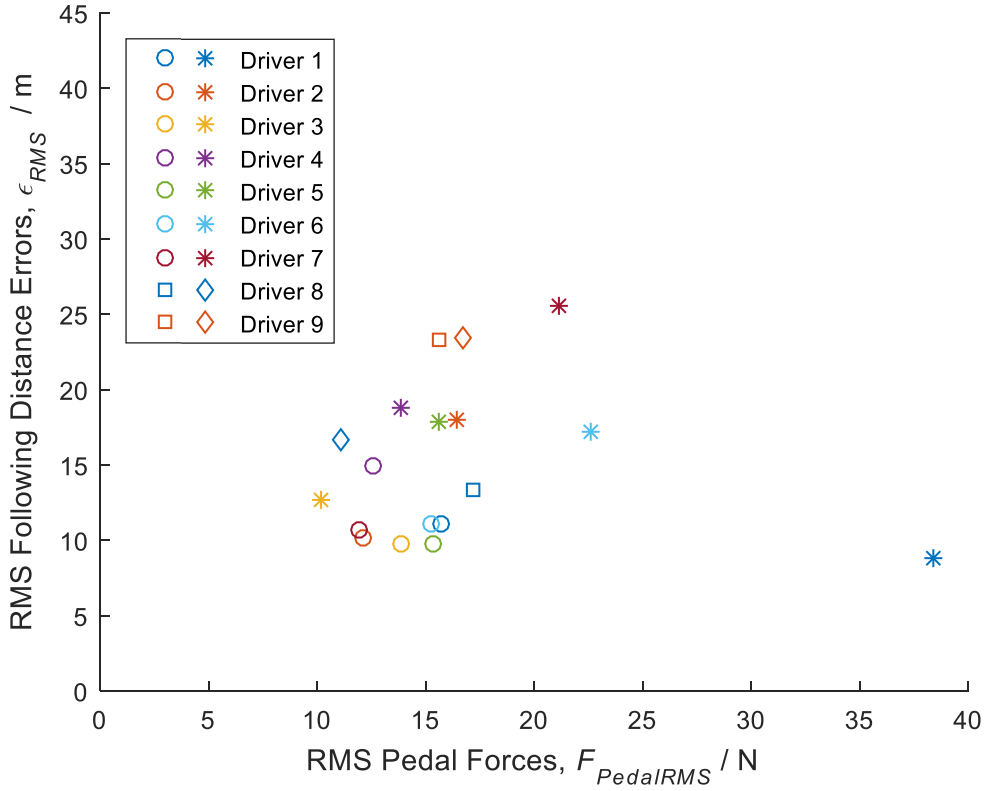

Figure 4-16 – Cost diagram for the car-following scenario with feedback – stars (non-professionals) and diamonds (professionals) illustrate data from the assisted drivers. Their original data points from the unassisted case are illustrated with circles (non-professionals) and squares (professionals). Colours are unchanged.

Finally, the following distance standard deviations are shaded alongside the means for each driver in Figure 4-17. It is very noticeable that the addition of pedal feedback has not only greatly reduced the following distances for a few drivers (mainly Drivers 4 and 9), but has also greatly reduced the standard deviation in following distance (compare with Figure 3-32). The feedback has enabled the drivers to be much more consistent in their following distance behaviour.

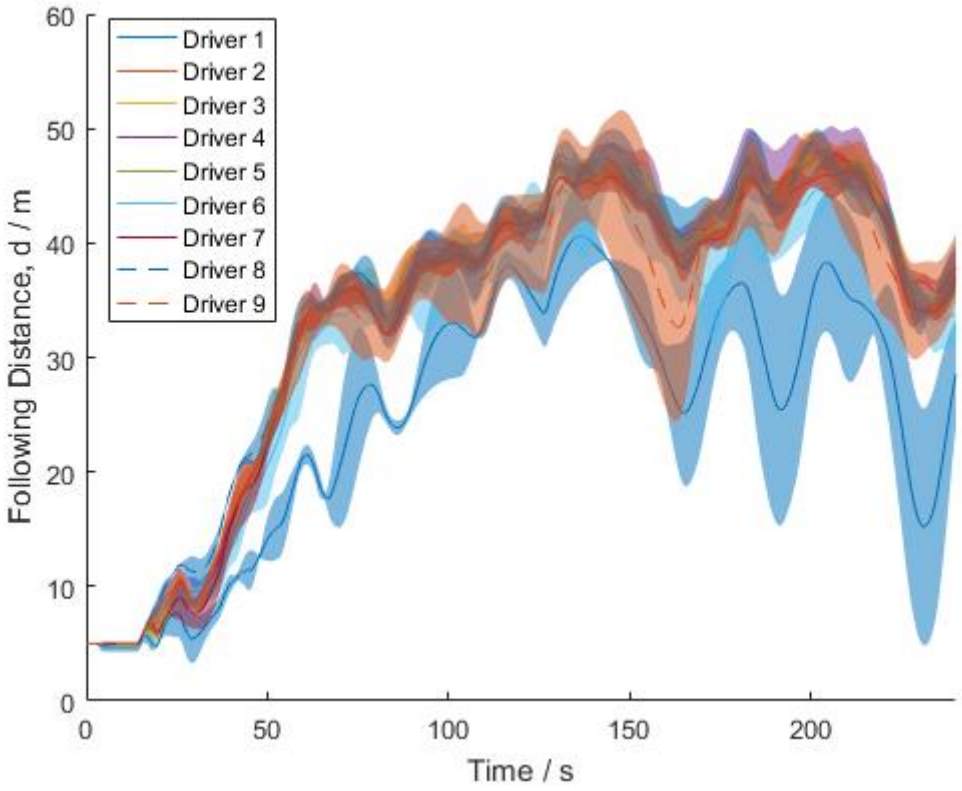

Figure 4-17 - Driver mean following distances with one driver standard deviation shaded either side

4.5 Driver model identification

The aim of this section is to determine which of the two frameworks, decentralised or cooperative, best represents the measured human behaviour.

4.5.1 Driver model drive-cycle task with feedback

The Matlab fmincon function is used in an optimisation algorithm to fit the driver model frameworks to the recorded data from the simulator. The driver cost function weightings on speed error are varied to identify the combination resulting in the minimum fitting cost function value. The fitting function is as defined in (3.46). Constraints are initiated to ensure that the weightings are positive. In the one-sided cooperative case, the driver dominance, ρ_1 is also included in the optimisation.

Figure 4-18 demonstrates the effect of varying the driver dominance, ρ_1, on the fitting function for Driver 6 over a range of values of q_{1s}. It is noted that the global minimum is located on the contour of $\rho_1 = 1$, meaning the optimisation will match the decentralised model of the driver.

This is also the case for the other eight drivers, although the minima are located at slightly different values of q_{1s}. Hence, for this task, one-sided cooperative control is no longer considered. It is also noting that the contour for $\rho_1 = 0$ is a horizontal line. This is because, for $\rho_1 = 0$, the driver's own target is no longer taken into account, and there is no cost associated with the driver's pedal force. The result is that the driver exerts the forces necessary to reach the sustainable target.

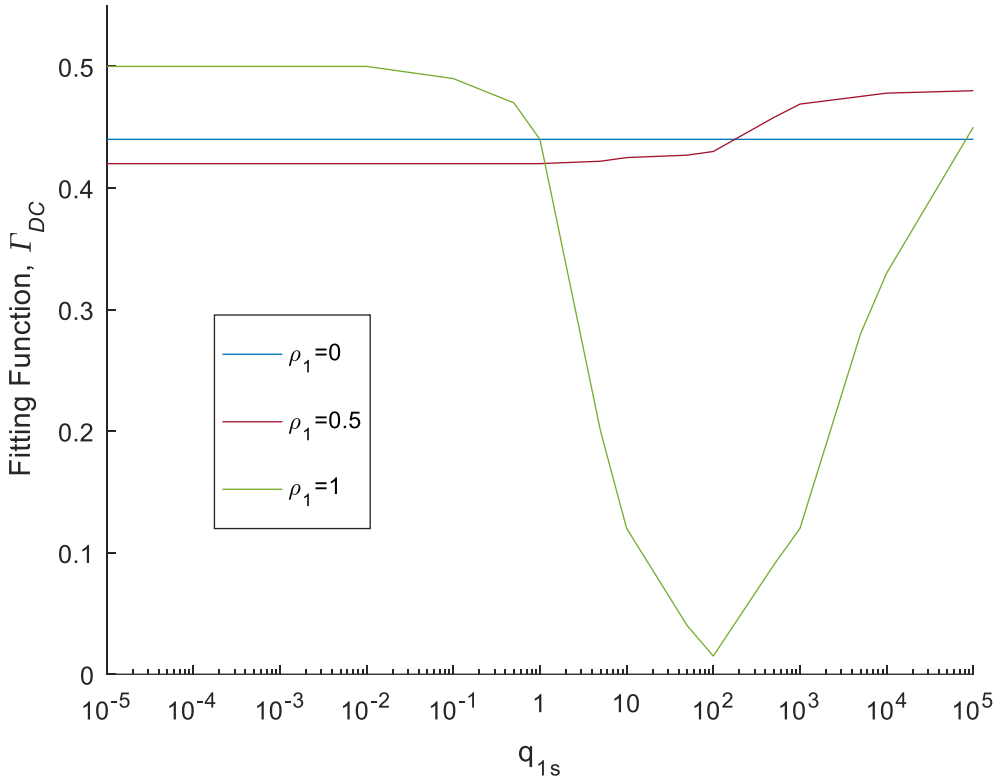

Figure 4-18 – Fitting function sensitivity for the one sided cooperative for a variety of values of driver dominance, ρ_1, and decentralised ($\rho_1 = 1$) cases

Figure 4-19 illustrates the optimal cost function weightings for each of the nine drivers. From an initial comparison, it is noted that the weightings on speed error between the unassisted driver case (from 3.2) and the decentralised controller model are of similar magnitude, supporting the idea that drivers might interact with the feedback controller in a decentralised manner.

In terms of 'goodness of fit', Figure 4-20, shows that the fitting function for the decentralised case is less than that for the unassisted case for six of the nine drivers. This means that, in most cases the driver model is better able to predict the behaviour of a decentralised driver with feedback, than the driver unassisted. The framework is well suited to modelling the human driver-active pedal interaction.

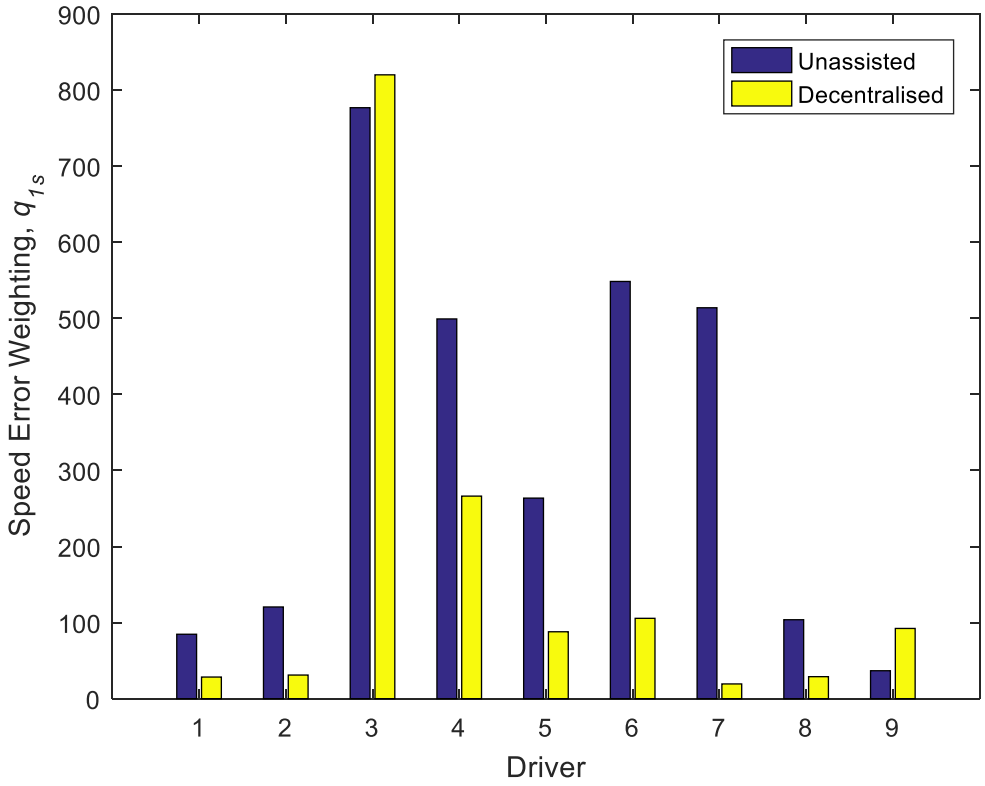

Figure 4-19 – Fitted cost function weightings for the Drive Cycle task

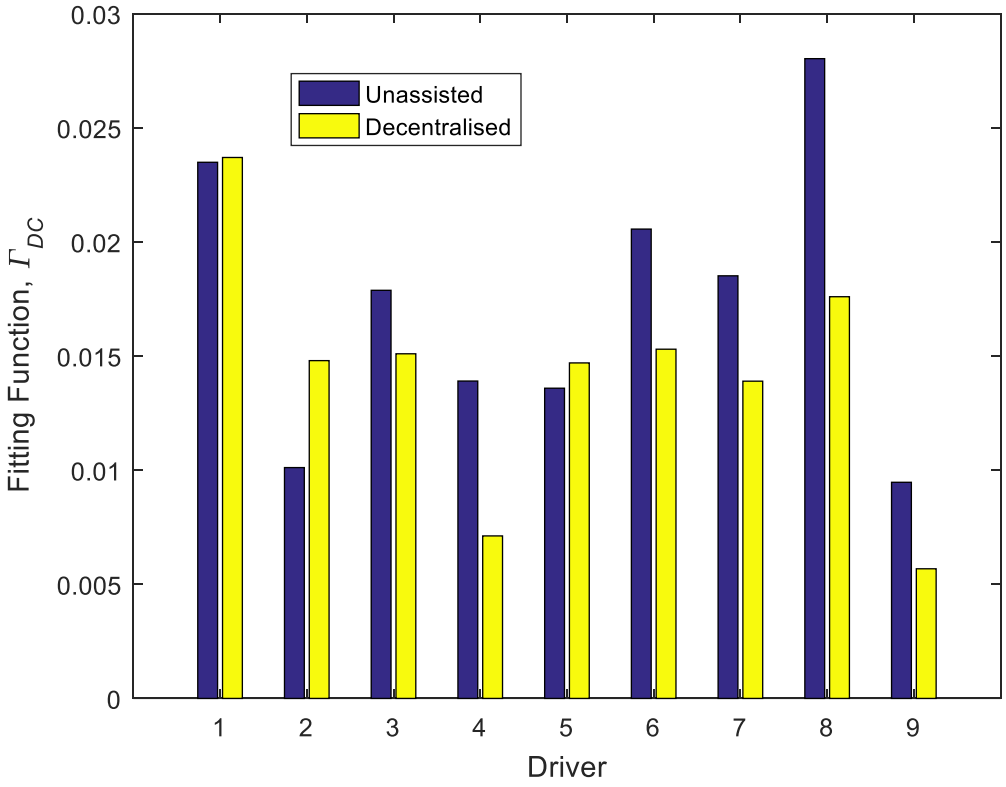

Figure 4-20 – Total driver-model driver squared errors for the Drive Cycle task

The behaviour of the assisted driver model is now demonstrated over the next set of figures. Driver 9 is selected for demonstration as the driver models best fitted with this driver over both cases of unassisted and assisted. Figure 4-21 illustrates the recorded and modelled vehicle speed for the drive cycle following scenario. The real driver, does not have perfect knowledge of the behaviours of the vehicle and therefore is unable to control the vehicle to the same high standard as the very capable driver model. Another possible explanation for the discrepancy is that there is noise in the human sensorimotor system and an inaccurate simulation model (as intermittency is not included for example). With this is mind, the driver model does quite closely resemble the shape of the driver speed profile.

Significant differences between the human driver and the driver model occur in the driver pedal forces -Figure 4-22. The decentralised driver model underestimates the forces from the driver.

As the driver pedal forces are inaccurate, the pedal feedback forces will also be inaccurate -Figure 4-23. In the model situation, the pedal feedback is having to exert much higher forces to reach its objective, as the driver model is exerting lower forces than the human driver did. The resultant of these uneven pedal forces between the driver and feedback controller even out to produce a pedal position that closely resembles the measured pedal position - Figure 4-24.

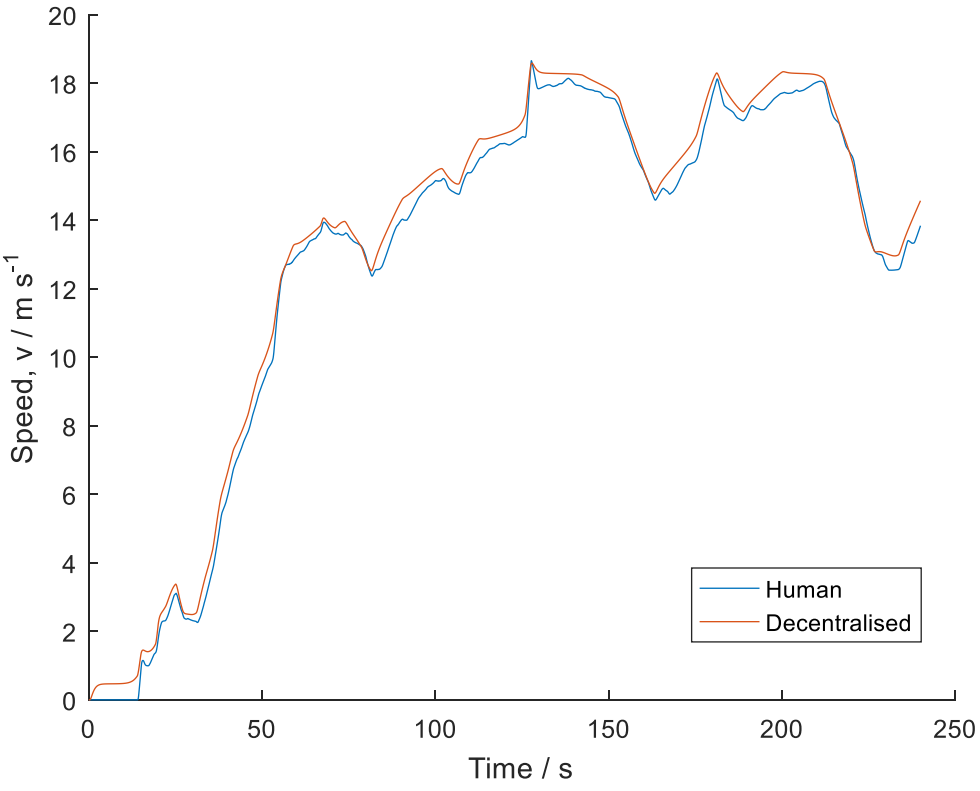

Figure 4-21 - Measured human driver and modelled vehicle speed for Driver 9

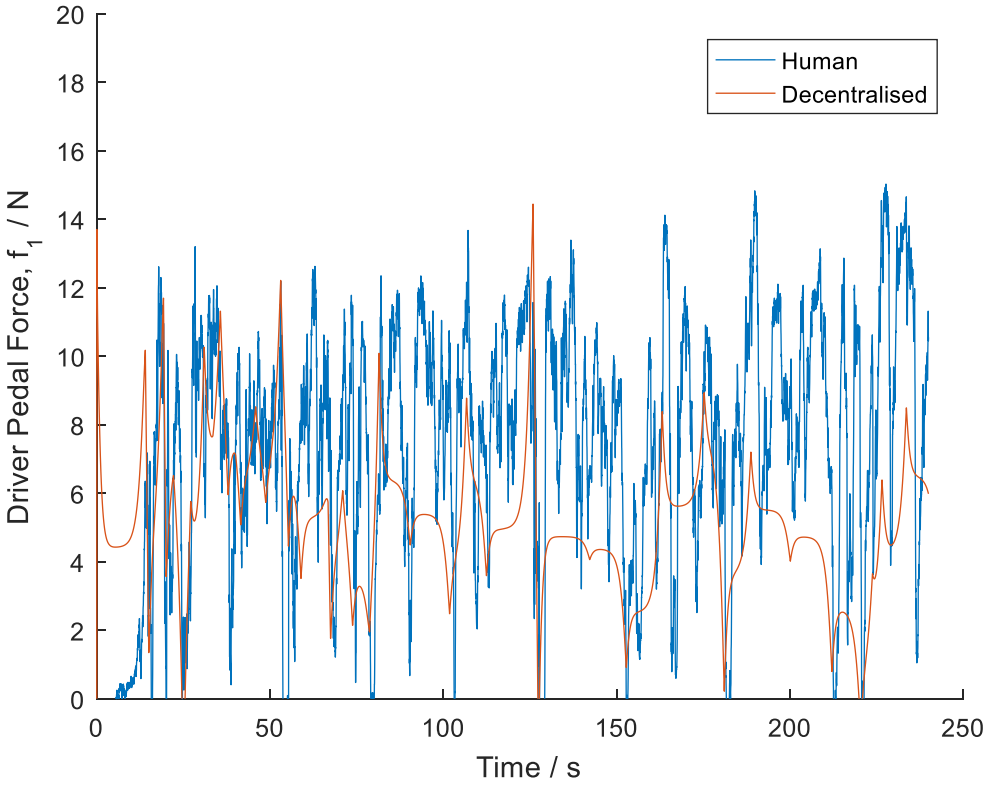

Figure 4-22 – Measured human driver and modelled pedal force for Driver 9

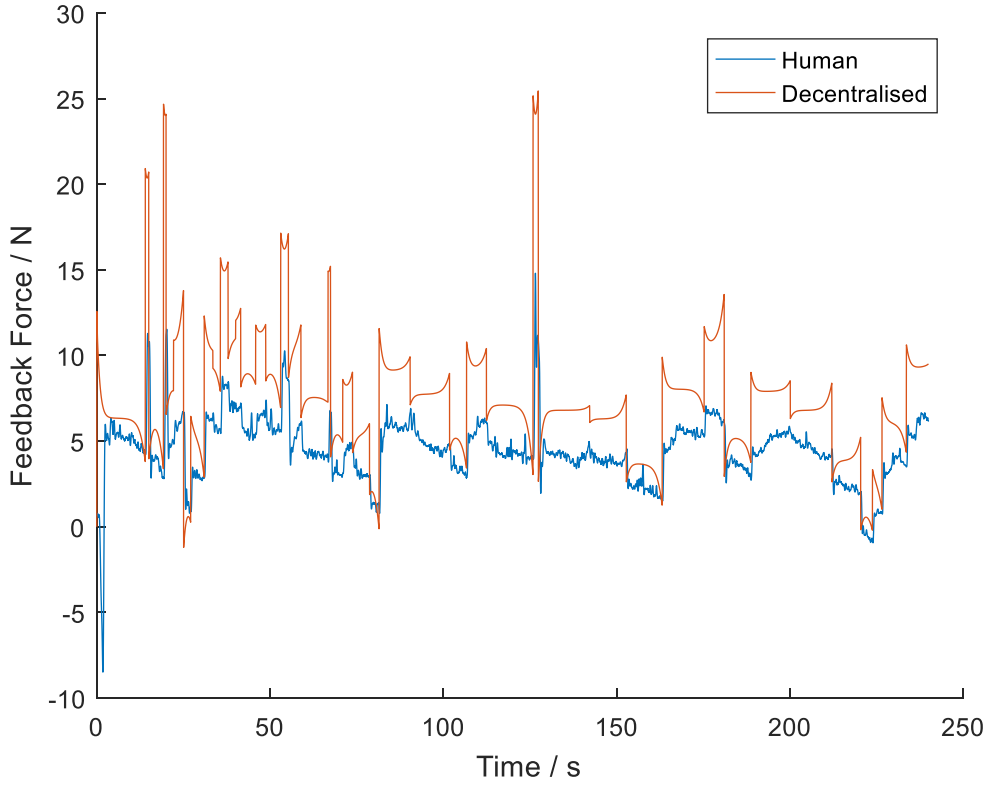

Figure 4-23 – Measured human driver and modelled feedback force for Driver 9

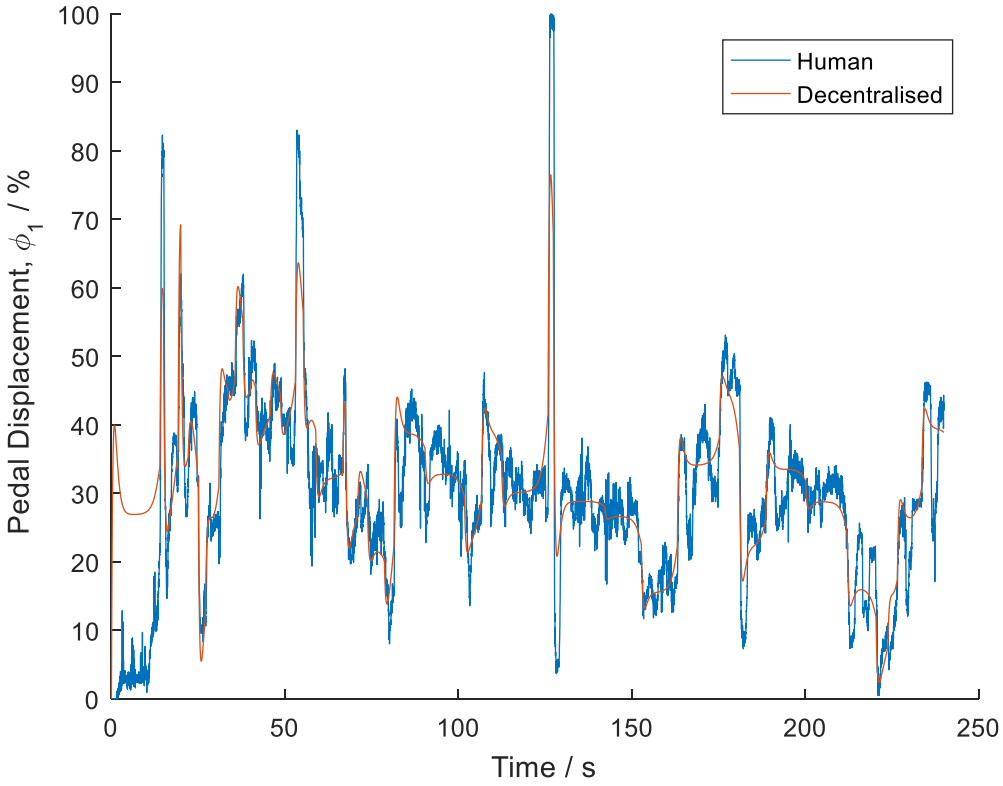

Figure 4-24 – Measured human driver and modelled pedal displacement for Driver 9

Further analysis of the driver model with feedback is completed through a parameter study examining how, by varying the cost function weightings, differences in driver model behaviour can be seen. Figure 4-25 demonstrates the trade-off between different costs for different driving styles (determined by different driver cost function weightings on speed error). The driver model and feedback controller are given a drive-cycle following task using the modified Millbrook cycle in both the decentralised and one sided cooperative framework (note that the one-sided cooperative case is introduced for comparison again, and the value of ρ_1 is fixed to 0.75.

Looking at the RMS speed error and RMS pedal force plots (Figure 4-25), several key features of the shapes are noted. Firstly, it is noted that the one-sided cooperative control framework produces lower (or near equal) speed errors over all driver cost function weightings than the decentralised controller. This is because the driver is modelled to have understanding of the feedback controller's objective, and will hence make some accommodation to this. It is also noted that in the decentralised case, that by increasing the driver's weighting on speed errors, q_{1s}, it is possible to increase the resultant speed errors, before a continued increase in speed error weighting results in reduced speed errors. This is due to the transition between the driver being overpowered by the feedback controller's strategy at low weighting of speed error and the driver being very dominant over the feedback controller at high weightings on speed error. In the region in between the two, the controllers behave similarly to each other, which results in non-optimal performance.

It is also noted that the experimental data points for each human driver are significantly higher than the modelled values, suggesting the human drivers behave in a non-optimal way with the pedal feedback – this could be an effect of the presences of sensorimotor noise not included in the simulation.

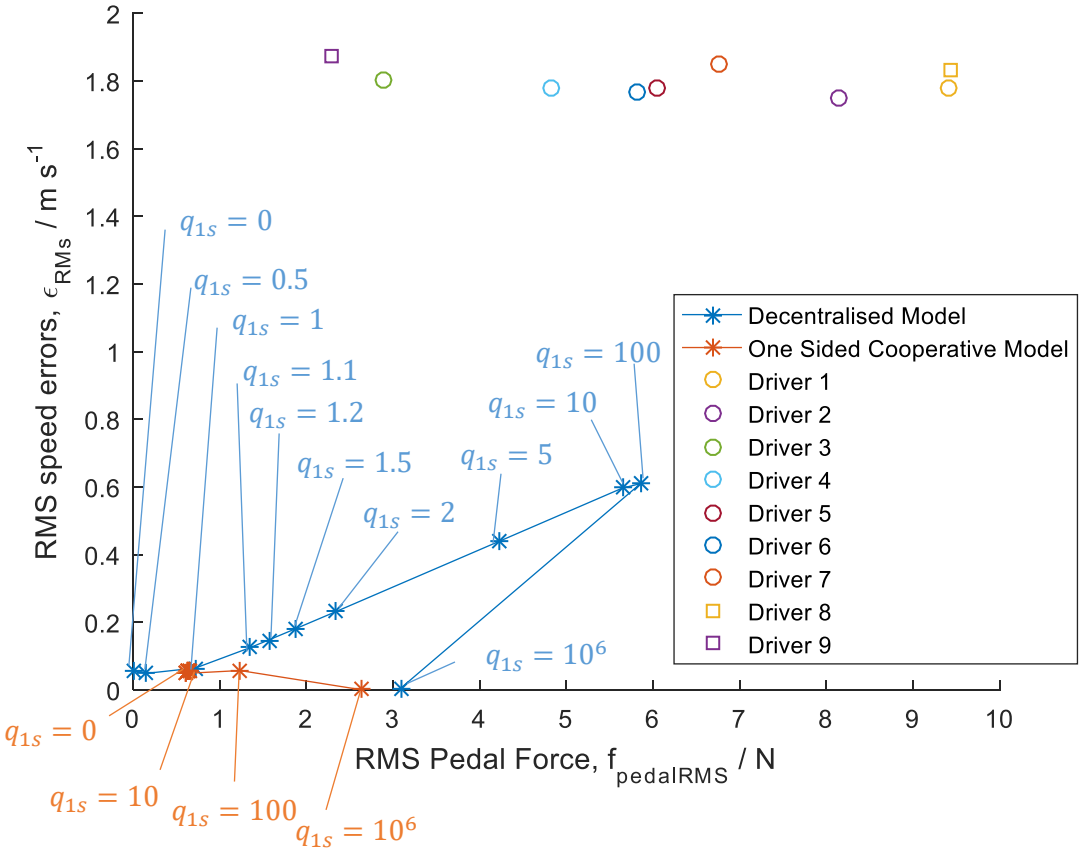

Figure 4-25 –Trade off graph for the drive cycle with feedback scenario

4.5.2 Driver model car-following task with feedback

The Matlab function fmincon is once again used to find the optimal set of cost function weightings to fit the driver model to the recorded simulator data. Each driver is given a target THW equal to the observed value from the unassisted case in Chapter 3. The optimisation is then run to find the cost function weightings for each driver in the decentralised and one-sided cooperative frameworks that minimise the fitting function from (3.45).

The calculated cost function weightings are displayed in Figure 4-26 alongside the cost function weightings from the unassisted case in Chapter 3. For the one-sided cooperative case, the additional parameter values, ρ_1 is displayed in Figure 4-27.

There are several significant differences in the cost function weighting between the three frameworks – unassisted, decentralised and one-sided cooperative. In most cases, the decentralised

cost function weighting is lower than the unassisted case, and the one-sided cooperative case frequently produced weightings in between that of the unassisted and decentralised cases.

The one-sided cooperative model has identified two groups of drivers – those who are prepared to cooperate with the feedback, and those who aren't. In terms of parameter values, the split is between those with high values of driver dominance, ρ_1 (for those who don't cooperate) and those with low driver dominance, with ρ_1 values in the region of 0.5 for those who do cooperate. Only Driver 6 falls in the low ρ_1 group, classifying them as cooperative. Drivers 2 and 7 are labelled as non-cooperative, but demonstrates a significant change in car-following performance between decentralised and one sided cooperative. The change in performance is caused by a change in driver cost function weightings, rather than by the introduction of a cooperative strategy.

The fitting function values for each driver are displayed in Figure 4-28 as an indicator of goodness of fit. In all cases, bar Driver 1, the assisted driver achieves a better fit with the model than the unassisted model achieved. This is because the feedback has aided the driver in becoming more predictable. For drivers 1 and 7, very high values of ρ_1 mean that the decentralised and one-sided cooperative case are near identical. For the others, the one-sided cooperative case provides a better fit to the driver data.

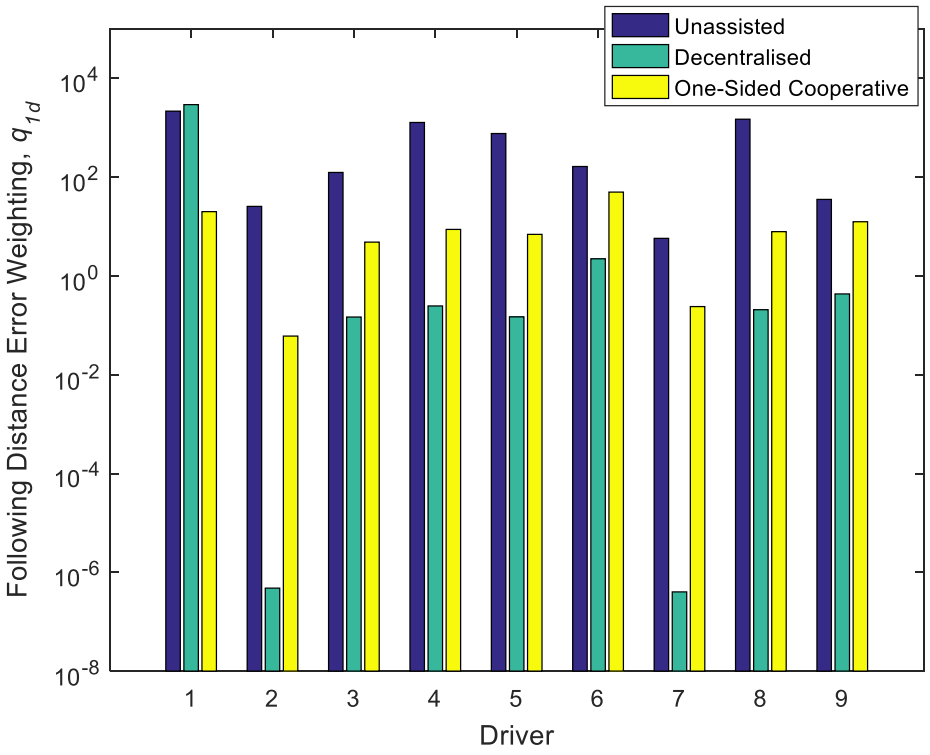

Figure 4-26 – Fitted cost function weightings for the car-following task

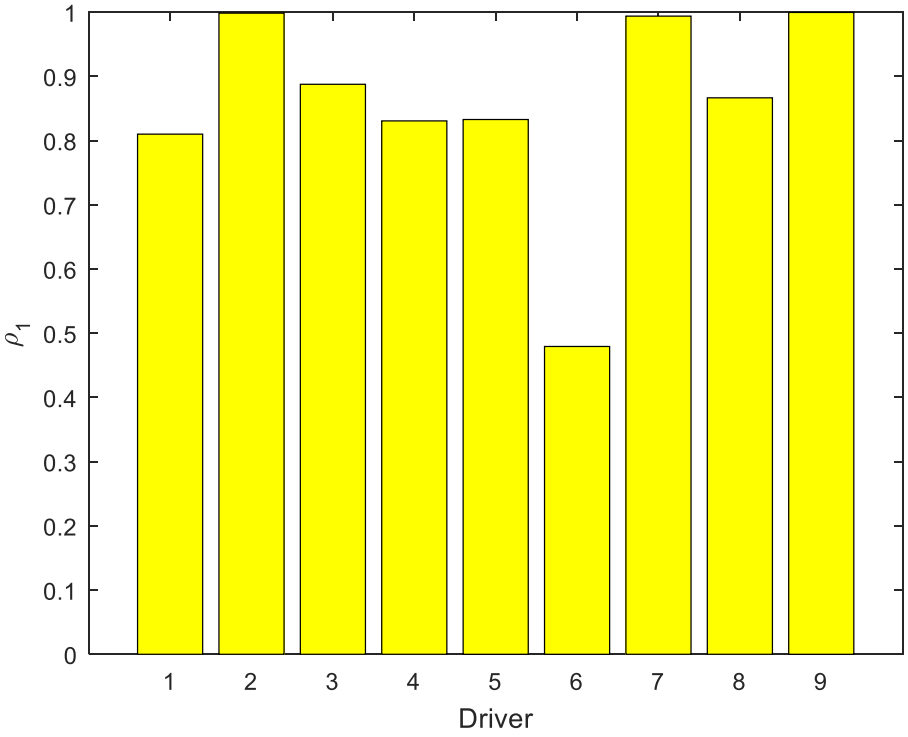

Figure 4-27 – Fitted ρ_1 values for the one sided cooperative car-following task

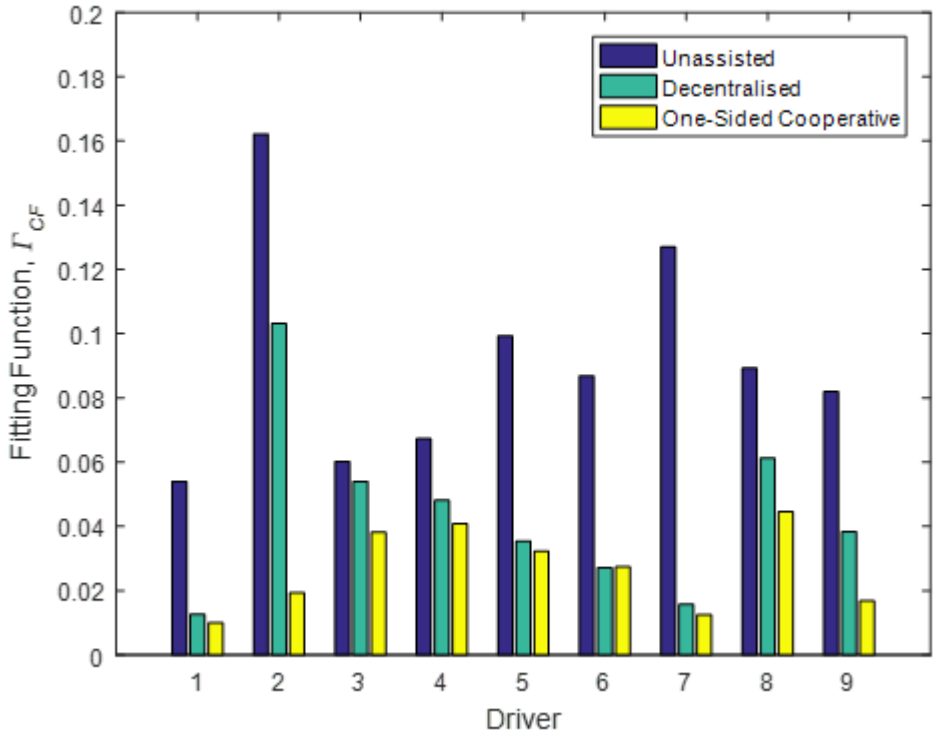

Figure 4-28 – Total driver-model driver squared errors for car-following scenario

To demonstrate the performance of the driver model in both decentralised and one-sided cooperative frameworks, the modelled car following scenario is displayed over the next set of

figures for Driver 1 – the driver with the best fit to the model. In Figure 4-29, the vehicle speeds are compared. Both modelled scenarios have very similar speed profiles, but the human driver has higher accelerations. Similarly, with following distance, the driver models produce a smoother output than the human driver. This is because the driver models do not completely accurately reproduce the pedal force from the driver (Figure 4-31).

It is in the pedal force (Figure 4-31) that differences are seen between the two frameworks. Unlike the drive cycle following task, in this scenario the driver and the pedal feedback controller have sufficiently different objectives that the driver model is forced to work and apply forces to the pedal. The one-sided cooperative driver model produces a more constant pedal force which aligns with the human driver allowing the pedal feedback to do a lot of the work. The decentralised driver model on the other hand fluctuates more as the model is unaware of what the pedal feedback will achieve. Due to the discrepancies between the measured and modelled pedal forces, the modelled feedback forces are also different to the measured forces, even though the controllers are identical. Generally, the model exerts lower forces on the pedal, as the driver model underestimates the driver's pedal forces.

With the pedal forces summed, the pedal displacement is determined, and broadly, the measured pedal position and modelled pedal positions are similar. The one-sided cooperative driver model achieves marginally better results as it more closely resembles the human driver.

Overall, the driver model represents the human driver reasonably well. Although the model makes no attempts to model the noise from the human driver, agreement between the driver model and human driver is good. In all cases, the one-sided cooperative controller is able to produce a better fit to the real human driver than the decentralised controller, suggesting that the human driver is able to learn and adapt their control strategy to reflect the feedback given.

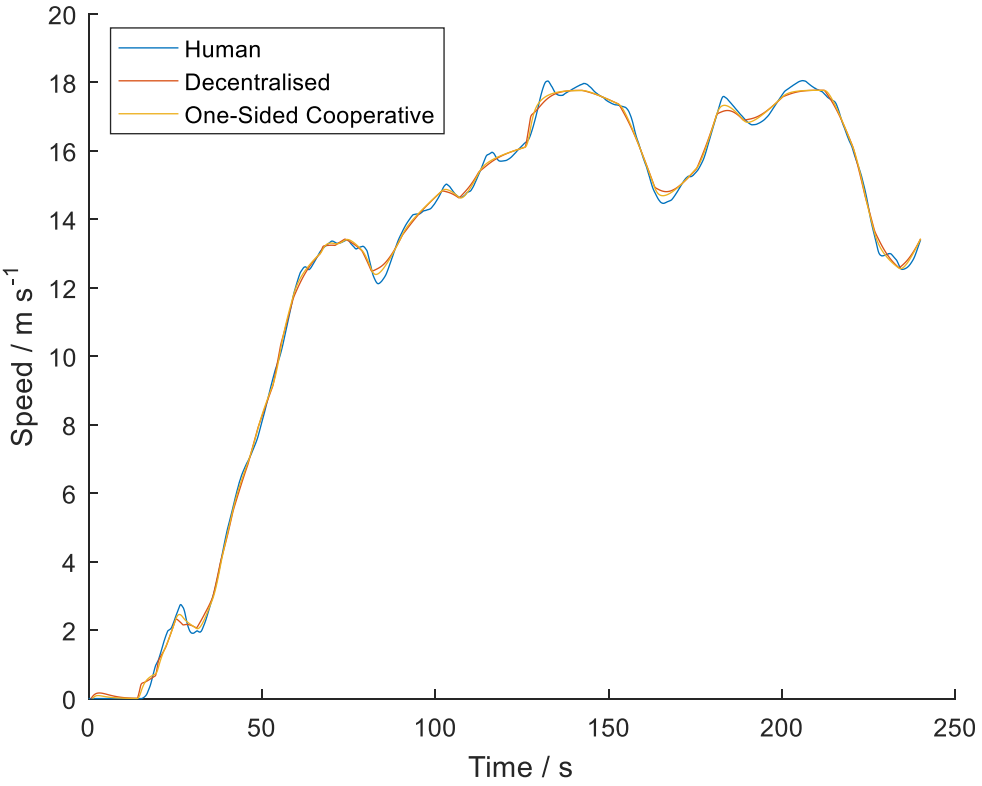

Figure 4-29 – Measure human and modelled vehicle speed for Driver 1

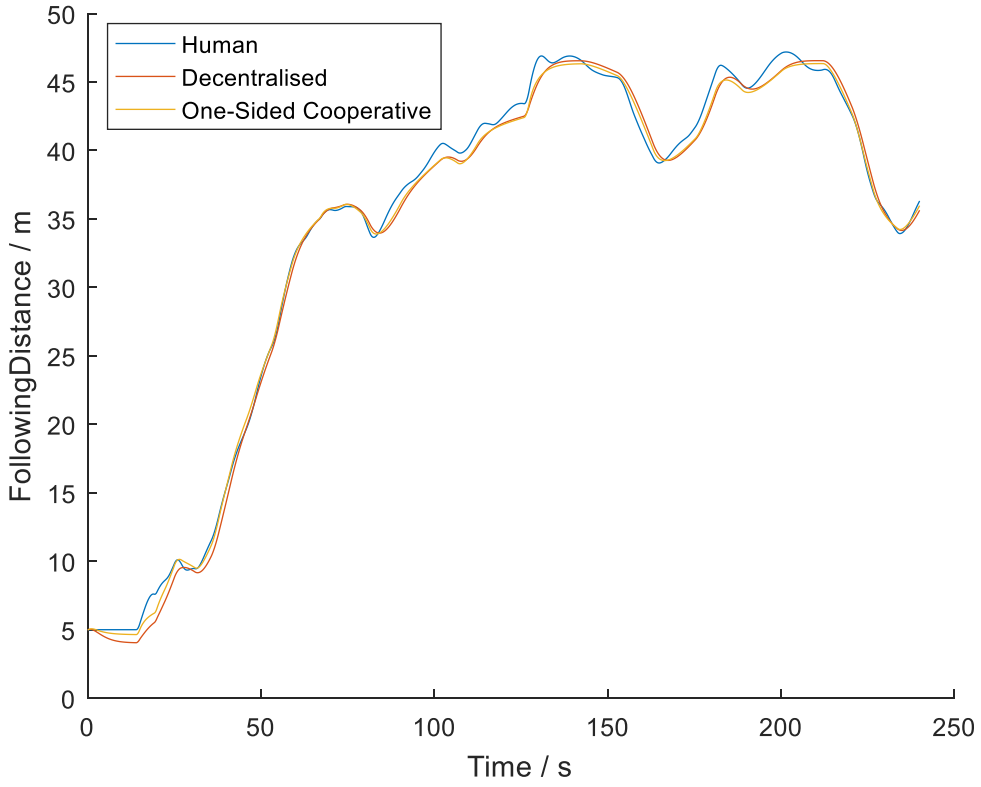

Figure 4-30 – Measured human and modelled following distance for Driver 1

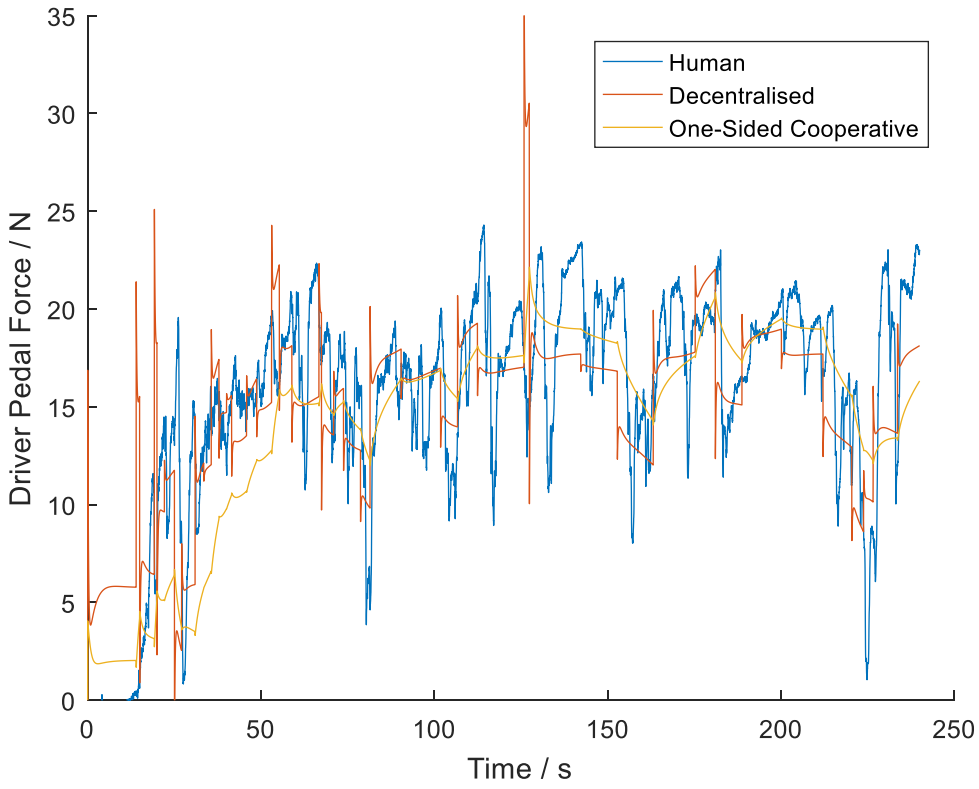

Figure 4-31 – Measured human and modelled driver pedal force for Driver 1

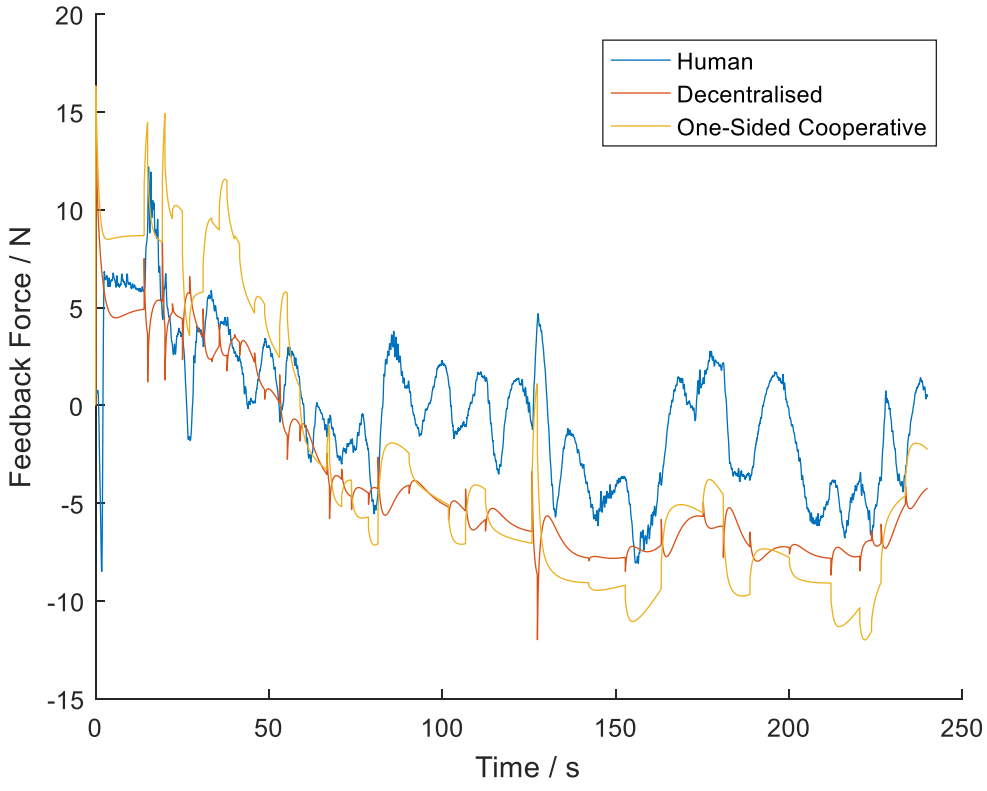

Figure 4-32 – Measured human and modelled feedback force for Driver 1

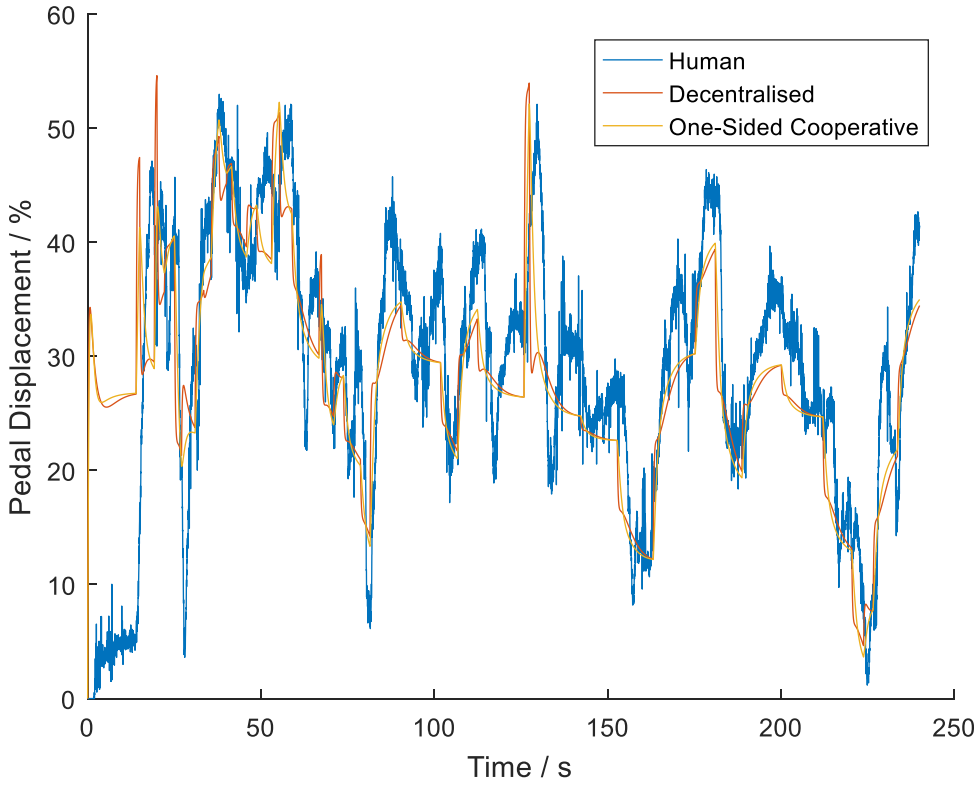

Figure 4-33 – Measured human and modelled pedal displacement for Driver 1

Further analysis of the driver model with feedback is completed through a parameter study examining how, by varying the cost function weightings, differences in driver model behaviour can be seen. Figure 4-34 demonstrates the trade-off between different costs for different driving styles (determined by different driver cost function weightings on following distance error) for both the decentralised and one sided cooperative models. The driver model and feedback controller are given a car-following task using the modified Millbrook cycle in both the decentralised and one sided cooperative framework (note that the one-sided cooperative case is introduced for comparison again, and the value of ρ_1 is fixed to 0.75).

Looking at the RMS following distance error and RMS pedal force plots (Figure 4-34), several key features of the shapes are noted. Firstly, the decentralised controller points appear to lie very close to a straight line. On closer inspection, it is also noted that this line is in fact a very sharp peak. When the cost function weighting on following distance errors, q_{1d}=0, the driver model takes no action, meaning a zero RMS driver pedal force, and leaves all the work to be performed by the pedal feedback controller. As the value of q_{1d} is increased, the driver model starts to contribute, increasing the RMS pedal force. The net effect of this contribution is in fact to worsen the resulting car following performance – the RMS following distance error increases as well. In this region, the contribution from the driver can be seen as an interference to the pedal feedback controller.

However, at a critical point, located with an approximate value on speed error weighting of $q_{1d} = 10$, this relationship reverses, and the driver model starts to decrease the force contributions in order to reduce the acceleration of the vehicle. This behaviour comes from the difference in cost functions between the driver model and the pedal feedback controller. The pedal feedback controller has cost function components including following distance, speed errors and pedal position and force, meaning that the controller is trying to match speeds as well as achieve safe distances. The driver model, on the other hand, is only trying to balance a cost function of following distance errors and pedal forces. Due to the relationship between vehicle speed, and safe following distances (3.38), this produces a slight conflict as vehicle speeds change – the following vehicle cannot perfectly match vehicle speed and achieve a changing following distance. In this scenario, the driver model would prefer slower accelerations, and hence applies lower pedal forces. The result is decreasing RMS following distance errors, and decreasing RMS pedal forces from the driver. This creates the sharp peak.

The one-sided cooperative case is also plotted in Figure 4-34. However, the change in driver cost function weightings has a very small impact on RMS pedal force and RMS following distance errors when compared with the span of errors observed from the human drivers in the simulator experiments, and the decentralised controller case. This is because by incorporating a level of understanding between the driver model and the pedal feedback controller (reflected in the value of $\rho_1 = 0.75$) the driver model accommodates the demands from the feedback controller as well.

Finally, it is worth noting again that the span of RMS errors from the human drivers is significantly larger than the span of RMS errors modelled in the two frameworks. This may be due to the apparent random component of human driver control caused by the noise not incorporated into the driver model.

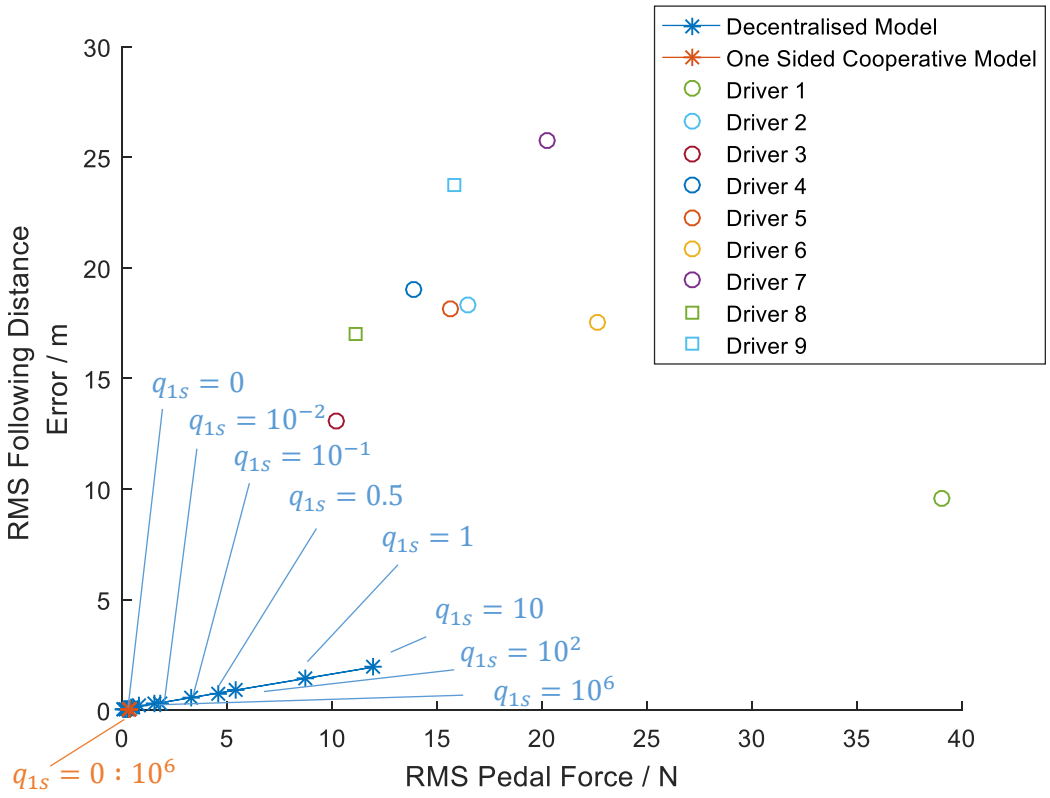

Figure 4-34 –Trade off graph for the drive cycle with feedback scenario

4.6 Conclusions

In this chapter, three mathematical game theoretical frameworks (Rawlings and Mayne, 2009) have been proposed for modelling the interaction between a human driver and pedal force feedback controller: decentralised, cooperative and one-sided cooperative.

After reflecting on the practical implementation of such a system, it was deemed that a fully cooperative pedal feedback device is not feasible with current technology due to the difficulties in continually assessing the drivers target and cost function weightings. The other two frameworks were carried forward and a series of driving simulator experiments were designed and carried out in a fixed base driving simulator with a range of drivers.

Two driving scenarios were examined: drive cycle following and car-following. Each driver completed multiple runs of the scenario to average out any random behaviour. By optimising the variables explored in the parameter study to minimise the errors between modelled and measured data, the driver models were fitted to the observed human drivers. This optimisation has quantified the difference in driving styles between drivers, but has also demonstrated that the driver models are able to represent mean human driver behaviour to a reasonable level.

In the first scenario to be examined, the drive cycle following case, the decentralised control framework best represented the human driver, as the one-sided cooperative human driver model exploited the behaviour of the pedal feedback device to minimise its own workload, to the point of exerting near zero forces on the pedal. This is a behaviour that was not observed in the human driver.

In the second scenario, the car-following scenario, the one-sided cooperative control framework produced the best fit to the driver data as it was found that the human drivers were adapting their driving styles to reflect the guidance given and adopt the time headway target set by the feedback controller.

In both cases, a decrease in peak (by approximately 20%) and average pedal position was detected between the unassisted and the assisted drivers, suggesting (but not proving) that driver pedal force feedback has potential for reducing the energy consumption of HGV's. Drivers also demonstrated higher performance and maintained more consistent following distances in the car following scenario.

Using a fully linear vehicle model, as developed in the previous chapter, a parameter study also explored how, by varying different parameters, different driving styles could be represented by the model.

Looking forward, this chapter has demonstrated that pedal force feedback to the driver does have potential to modify the driving styles of human drivers, and it is expected that these behaviour changes may lead on to fuel reductions as well.

Chapter 5 – Conclusions and Future Work

5.1 Conclusions

5.1.1 Chapter 2

In the literature review, it was identified that there is little data available on the behaviour of Heavy Goods Vehicle (HGV) drivers in real-world environments. To overcome this, two identical heavy goods vehicles were fitted with data logging equipment and monitored for a period of four months. The loggers recorded signals from the vehicle CAN bus through the FMS port and supplemented this information with data from built in sensors in an Android smart phone. The bank of data covers a wide variety of routes from urban and city roads in central London, to long distance motorway driving. Each vehicle is driven by a sole driver, so a good picture of each driver's driving styles is developed.

Firstly, though, another research objective was to create an accurate mathematical model relating the driver's longitudinal control input (accelerator pedal position) to the vehicle states in order to develop a greater understanding of the behaviour of human drivers. The bank of data was used to develop a mathematical model of the vehicles fitted with the data logger. The gross vehicle mass was estimated by iterating the reported vehicle mass from the FMS port. It was found, however, that this mass estimate was poor (errors up to 25%), but that it could detect a change in mass during a journey (due to a payload drop). Weigh bridge data was made available to measure the mass of the vehicle. Drag characteristics of the vehicles were estimated using the SRF Mapper tool. Traction force and fuel rate maps were also generated from the recorded data and performed well. A large degree of uncertainty made the maps unreliable when combined, and it was deemed that suitable conclusions about changes in driving style could not be drawn from the combined map. The model was able to estimate cumulative fuel use to within 20% when cruise control was not engaged.

Secondly, the driver behaviour was analysed. To do this, two independent groups of driving tasks were defined to separate the speed states of the vehicle from the driver control states. This analysis showed that a large proportion of the distance covered by these HGVs was done so under

Cruise Control and that a predictable driving manoeuvre frequently utilising the accelerator pedal was the navigation of roundabouts and junctions.

Analysis over several roundabouts revealed differences in the driving styles of the two drivers and the potential for one of the drivers to achieve 7% fuel savings by using slower speeds and accelerations, and by engaging cruise control through the roundabout. Statistical analysis over the bank of data as a whole identified that one driver made much greater use of cruise control than the other driver and also noted that this driver was less likely to use very high pedal depressions.

In Chapter 2, several of the key research objectives of this project were achieved. Measurements of vehicle and driver states were obtained from real world driving environments and a mathematical model of the vehicle was developed. However, in validating the mathematical model, results were not as promising as hoped for, so further work is suggested in 5.2 to improve the vehicle model. The chapter was also able to explain how different methods had been used to successfully investigate the differences in driver behaviour on the road.

5.1.2 Chapter 3

In Chapter 3, a new mathematical model of the human driver's longitudinal control of a HGV was proposed. Model Predictive Control Theory, commonly used to model the driver's lateral control of the vehicle, was used as the basis for the model.

The base of this model was the vehicle model developed in Chapter 2, along with linearised pedal dynamics, derived experimentally from a Volvo truck accelerator pedal. The pedal was then incorporated into a fixed base driving simulator.

The driver model was validated in a series of experiments in the driving simulator. Nine human drivers (two of which were professional HGV drivers) were given two driving tasks: drive cycle following and car following. Each task was repeated multiple times and an average taken.

Differences in human behaviour were observed, particularly in the case of choice of safe following distances in the car following scenario. Observed mean time headways ranged from 1s to 6s. The driver model was fitted to each human driver in turn by tuning cost function weightings in the model. For the drive cycle task, the cost function components were speed error and pedal force, whilst for the car following scenario, the components were following distance error and pedal force.

After fitting, it was found that there was good agreement between the driver model and the human drivers in both scenarios. The simulated RMS pedal forces lay within 20% of the measured simulator values. The driver model was, however, unable to represent the noise and disturbances observed in the human driver behaviours.

Chapter 3 was successfully able to accomplish one of the research objectives from 1.6. A new driver model, incorporating the driver’s preview of upcoming demands, was developed. This

model was then validated in a series of driving simulator experiments, specifically tailored to the heavy goods vehicle application.

5.1.3 Chapter 4

The driver model developed in the previous chapter was expanded to incorporate interactions with an Active Accelerator Pedal (AAP). Three Mathematical Game Theoretical frameworks were considered: decentralised, cooperative and one-sided cooperative. Practically, a cooperative system was deemed unachievable, so was not examined beyond a preliminary derivation. The difficulty in a cooperative framework is producing a pedal feedback device that would always be capable of understanding the driver's targets and style.

A pedal feedback controller was proposed for both the drive cycle and car following scenario and the same nine drivers from Chapter 3 were given both tasks whilst being assisted by the pedal feedback. As with the unassisted case, a range of driving styles was observed, and the driver model was fitted to each driver for each task. The cost function components were the same as the unassisted case but the weighting values were identified specifically for the assisted tasks.

It was found that for the drive cycle task, the decentralised framework best represented the driver behaviour, whilst in the car following scenario, the one-sided cooperative framework was best suited. It was also noted that the addition of pedal feedback reduced the resultant RMS pedal position, suggesting that fuel savings may also be observed.

Chapter 4 was able to address the final research objective of this thesis. The driver model was successfully extended to incorporate interactions with an AAP, and driving simulator experiments were able to support the models. However, the work carried out in this chapter only considered the interaction of human drivers with one fixed feedback controller. Further work is suggested in 5.2 to examine the interactions with a variety of controllers.

Overall, it was concluded that pedal feedback to the driver has the potential for dramatically improving driver behaviour and speed control, and it is proposed to be worth continual investigation.

5.2 Future Work

5.2.1 Mass Estimation

In Chapter 2 it was noted that accurate mass information allows better estimations of the traction forces, and hence better prediction of the vehicle response to driver control actions. Although mass is reported through the FMS port, this information has been shown to be inaccurate compared with weighbridge data. A more accurate mass estimation would reduce the uncertainty in the traction force maps.

5.2.2 Road Gradient

Another important aspect of the estimation of traction forces is the road gradient. Although two sources of road elevation are considered in this thesis, both have flaws. The GPS receiver in the phone did not accurately record the road elevation when compared with mapped data. The Google Maps data did better represent the prevailing land height but could not account for cuttings and embankments meaning that the road height was not similar to the local land height. A suggested alternative approach could involve barometric pressure measurements.

5.2.3 Non-linear vehicle and pedal dynamics

The introduction of traction force and fuel rate maps was a first step towards a non-linear vehicle model. However, as they were subject to significant uncertainty, further work is required to improve the accuracy of the simulations performed. Improved road gradient and mass data (as discussed in 5.2.1 and 5.2.2) are likely reduce the uncertainty in the maps.

If the non-linear vehicle maps are improved, then it is also possible to incorporate other non-linear properties into the vehicle model to improve the rest of the model as well. The pedal dynamics were linearised for the inclusion in the Model Predictive Controllers in this thesis, but the pedal itself has clearly non-linear behaviour. This could be better reflected in a fully non-linear model.

A non-linear vehicle model will result in non-linear control. Initial considerations for a non-linear model predictive controller are considered in Appendix 2.

5.2.4 Limb dynamics

A further improvement to the driver model could be achieved by incorporating the driver's limb dynamics. The current model does not reflect this, as it focusses solely on the driver's cognitive control, and not any reflex control actions.

5.2.5 Pedal Feedback Strategy Optimisation

In this thesis, the interaction between the driver and one feedback pedal controller has been examined. It is now possible to examine the interaction between the human drivers and other pedal controllers in order to identify an optimal pedal feedback controller for the purposes of reducing fuel consumption in heavy goods vehicles.

References

Abbink, D., 2006. Neuromuscular Analysis of Haptic Gas Pedal Feedback During Car Following (PhD Thesis). TU Delft.

Abbink, D.A., Boer, E.R., Mulder, M., 2008. Motivation for Continuous Haptic Gas Pedal Feedback to Support Car Following, in: 2008 IEEE Intelligent Vehicles Symposium, Vols 1-3. Presented at the IEEE Intelligent Vehicles Symposium, IEEE, Eindhoven, pp. 1201–1208.

Abbink, D.A., Mulder, Mark, van der Helm, F.C.T., Mulder, Max, Boer, E.R., 2011. Measuring Neuromuscular Control Dynamics During Car Following With Continuous Haptic Feedback. IEEE Trans. Syst. Man Cybern. Part B-Cybern. 41, 1239–1249.

Adell, E., 2010. Acceptance of Driver Support Systems, in: Drivers' Needs and Acceptance of Assistance Functions. Presented at the European Conference on Human Centred Design for Intelligent Transport Systems, Berlin, pp. 475–486.

Adell, E., Varhelyi, A., 2008. Driver comprehension and acceptance of the active accelerator pedal after long-term use. Transp. Res. Pt. F-Traffic Psychol. Behav. 11, 37–51.

Adell, E., Varhelyi, A., Fontana, M. dalla, 2011. The effects of a driver assistance system for safe speed and safe distance – A real-life field study. Transportation Research Part C: Emerging Technologies 19, 145–155.

Adell, E., Várhelyi, A., Hjälmdahl, M., 2008. Auditory and haptic systems for in-car speed management – A comparative real life study. Transportation Research Part F: Traffic Psychology and Behaviour 11, 445–458.

Allen, R.W., Chrstos, J.P., Aponso, B.L., Lee, D., 2002. Driver/Vehicle Modeling and Simulation.

Angkititrakul, P., Ryuta, T., Wakita, T., Takeda, K., Miyajima, C., Suzuki, T., 2009. Evaluation of driver-behavior models in real-world car-following task, in: 2009 IEEE International Conference on Vehicular Electronics and Safety (ICVES). Presented at the 2009 IEEE International Conference on Vehicular Electronics and Safety (ICVES), pp. 113–118.

Azzi, S., Reymond, G., Mérienne, F., Kemeny, A., 2011. Eco-driving performance assessment with in-car visual and haptic feedback assistance. Journal of Computing and Information Science in Engineering 11, 041005.

Bageshwar, V.L., Garrard, W.L., Rajamani, R., 2004. Model predictive control of transitional maneuvers for adaptive cruise control vehicles. IEEE Transactions on Vehicular Technology 53, 1573–1585.

Bao, S., LeBlanc, D.J., Sayer, J.R., Flannagan, C., 2012. Heavy-truck drivers' following behavior with intervention of an integrated, in-vehicle crash warning system: a field evaluation. Hum Factors 54, 687–697.

Barlow, T., Latham, S., McCrae, I., Boulter, P., 2009. A reference book of driving cycles for in the measurement of road vehicle emissions (No. PPR354). TRL Ltd.

Bertolazzi, E., Biral, F., Da Lio, M., Saroldi, A., Tango, F., 2010. Supporting Drivers in Keeping Safe Speed and Safe Distance: The SASPENCE Subproject Within the European Framework Programme 6 Integrating Project PReVENT. IEEE Trans. Intell. Transp. Syst. 11, 525–538.

Birrell, S.A., Fowkes, M., Jennings, P.A., 2014. Effect of Using an In-Vehicle Smart Driving Aid on Real-World Driver Performance. IEEE Transactions on Intelligent Transportation Systems 15, 1801–1810.

Birrell, S.A., Young, M.S., Weldon, A.M., 2013. Vibrotactile pedals: provision of haptic feedback to support economical driving. Ergonomics 56, 282–292.

Bishop, J.D.K., Stettler, M.E.J., Molden, N., Boies, A.M., 2016. Engine maps of fuel use and emissions from transient driving cycles. Applied Energy 183, 202–217.

Boham, S., 2003. Brake Pedal Feedback (Masters Report). University of Cambridge.

Bouchner, P., Novotny, S., 2011. Development of Advanced Driving Simulator: Steering Wheel and Brake Pedal Feedback, in: Proceedings of the 2Nd International Conference on Circuits, Systems, Control, Signals, CSCS '11. World Scientific and Engineering Academy and Society (WSEAS), Stevens Point, Wisconsin, USA, pp. 170–174.

Burnham, G.O., Seo, J., Bekey, G.A., 1974. Identification of human driver models in car following. IEEE Transactions on Automatic Control 19, 911–915.

Chandler, R.E., Herman, R., Montroll, E.W., 1958. Traffic Dynamics: Studies in Car Following. Operations Research 6, 165–184.

Chien, C.C., Ioannou, P., 1992. Automatic Vehicle-Following, in: American Control Conference, 1992. Presented at the American Control Conference, 1992, pp. 1748–1752.

Cole, D.J., 2012. A path-following driver–vehicle model with neuromuscular dynamics, including measured and simulated responses to a step in steering angle overlay. Vehicle System Dynamics 50, 573–596.

Cole, D.J., Pick, A.J., Odhams, A.M.C., 2006. Predictive and linear quadratic methods for potential application to modelling driver steering control. Vehicle System Dynamics 44, 259–284.

Corona, D., Lazar, M., De Schutter, B., Heemels, M., 2006. A hybrid MPC approach to the design of a Smart adaptive cruise controller, in: 2006 IEEE Computer Aided Control System Design, Presented at the 2006 IEEE Computer Aided Control System Design, pp. 231–236.

de Rosario, H., Louredo, M., Díaz, I., Soler, A., Gil, J.J., Solaz, J.S., Jornet, J., 2010. Efficacy and feeling of a vibrotactile Frontal Collision Warning implemented in a haptic pedal. Transportation Research Part F: Traffic Psychology and Behaviour 13, 80–91.

Douglas, D.H., Peucker, T.K., 2006. Algorithms For The Reduction Of The Number Of Points Required To Represent A Digitized Line Or Its Caricature. Cartographica: The International Journal for Geographic Information and Geovisualization.

Fitch, G.M., Hankey, J.M., Kleiner, B.M., Dingus, T.A., 2011. Driver comprehension of multiple haptic seat alerts intended for use in an integrated collision avoidance system. Transportation Research Part F: Traffic Psychology and Behaviour 14, 278–290.

Flad, M., Trautmann, C., Diehm, G., Hohmann, S., 2013. Experimental Validation of a Driver Steering Model based on Switching of Driver Specific Primitives. 2013 Ieee International Conference on Systems, Man, and Cybernetics (smc 2013) 214–220.

Geitner, C., Birrell, S., Skrypchuk, L., Krehl, C., Mouzakitis, A., Jennings, P., 2015. Good Vibrations: Driving with a Haptic Pedal, in: Adjunct Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '15. ACM, New York, NY, USA, pp. 100–105.

Gipps, P.G., 1981. A behavioural car-following model for computer simulation. Transportation Research Part B: Methodological 15, 105–111.

Gonder, J., 2011. Fuel saving effectiveness of various driver feedback approaches. National Renewable Energy Laboratory, Golden, CO :

H Cathcart, D. J. Cole, H Stoffels, K Glover, 2013. A mathematical occupant-vehicle model for investigating discomfort during transient longitudinal acceleration. Presented at the 23rd International Symposium on Dynamics of Vehicles on Roads and Tracks, Qingdao, China.

Hajek, H., Popiv, D., Just, M., Bengler, K., 2011. Influence of a Multimodal Assistance Supporting Anticipatory Driving on the Driving Behavior and Driver's Acceptance, in: Kurosu, M. (Ed.), Human Centered Design, Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 217–226.

Hajnal, A., Fonseca, S., Kinsella-Shaw, J.M., Silva, P., Carello, C., Turvey, M.T., 2007. Haptic selective attention by foot and by hand. Neurosci. Lett. 419, 5–9.

Hartley, P., 2012. Brake Pedal Feel (Masters Report). University of Cambridge.

Hjalmdahl, M., Varhelyi, A., 2004. Speed regulation by in-car active accelerator pedal - Effects on driver behaviour. Transp. Res. Pt. F-Traffic Psychol. Behav. 7, 77–94.

Hof, T., Conde, L., Garcia, E., Iviglia, A., Jamson, S., Jopson, A., Lai, F., Merat, N., Nyberg, J., Rios, S., David Sanchez, Steffan Schneider, Philipp Seewald, Caroline van der Weerdt, Remco Wijn, Adrian Zlocki, 2013. ecoDriver: A state of the art review and user's expectations (No. D11.1. Retrieved from www.ecodriver-project.eu).

Huang, J.L., Ford, J.K., 2012. Driving locus of control and driving behaviors: Inducing change through driver training. Transp. Res. Pt. F-Traffic Psychol. Behav. 15, 358–368.

Hunt, S.W., Odhams, A.M.C., Roebuck, R.L., Cebon, D., 2011. Parameter measurement for heavy-vehicle fuel consumption modelling. Proc. Inst. Mech. Eng. Part D-J. Automob. Eng. 225, 567–589.

Jamson, H., Hibberd, D.L., Merat, N., 2013. The Design of Haptic Gas Pedal Feedback to Support Eco-Driving. Presented at the 7th International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design, New York.

Johns, T., n.d. A Model of Driver Compensatory Steering Control Incorporating Cognitive Limitations. Presented at the AVEC'12.

Johns, T.A., Cole, D.J., 2015. Measurement and mathematical model of a driver's intermittent compensatory steering control. Vehicle System Dynamics 53, 1811–1829.

Larsson, H., Ericsson, E., 2009. The effects of an acceleration advisory tool in vehicles for reduced fuel consumption and emissions, in: Transportation Research Part D-Transport and Environment. pp. 141–146.

Lefevre, S., Sun, C., Bajcsy, R., Laugier, C., 2014. Comparison of parametric and non-parametric approaches for vehicle speed prediction, in: American Control Conference (ACC), 2014. Presented at the American Control Conference (ACC), 2014, pp. 3494–3499.

Li, S., Li, K., Rajamani, R., Wang, J., 2011. Model Predictive Multi-Objective Vehicular Adaptive Cruise Control. IEEE Transactions on Control Systems Technology 19, 556–566.

Maciejowski, J., 2002. Predictive Control: With Constraints. Prentice-Hall: London.

Mastinu, G., Ploechl, M. (Eds.), 2014. Road and Off-Road Vehicle System Dynamics Handbook. CRC Press.

McGordon, A., Poxon, J.E.W., Cheng, C., Jones, R.P., Jennings, P.A., 2011. Development of a driver model to study the effects of real-world driver behaviour on the fuel consumption. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering 225, 1518–1530.

McIlroy, R.C., 2017. Good vibrations: Using a haptic accelerator pedal to encourage eco-driving

McIlroy, R.C., Stanton, N.A., Harvey, C., 2014. Getting drivers to do the right thing: a review of the potential for safely reducing energy consumption through design. IET Intelligent Transport Systems 8, 388–397.

Micah Steele, R Brent Gillespie, 2001. Shared Control Between Human and Machine - Using A Haptic Steering Wheel to Aid in Land Vehicle Guidance.

Mortimer, R., 1974. Foot Brake Pedal Force Capability of Drivers. Ergonomics 17, 509–513.

Mosbach, S., Flad, M., Hohmann, S., 2017. Cooperative longitudinal driver assistance system based on shared control, in: 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC). Presented at the 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 1776–1781.

Mulder, M., 2007. Haptic gas pedal feedback for active car-following support (PhD Thesis). TU Delft.

Mulder, M., Abbink, D.A., Boer, E.R., 2012. Sharing control with haptics: seamless driver support from manual to automatic control. Hum Factors 54, 786–798.

Mulder, M., Kitazaki, S., Hijikata, S., Mulder, Max., van Paassen, M.M., Boer, E.R., 2004. Reaction-time task during car-following with an active gas pedal. Presented at the 2004 IEEE International Conference on Systems, Man and Cybernetics, The Hague, pp. 2465–2470 vol.3.

Mulder, M., Mulder, M., van Paassen, M.M., Abbink, D.A., 2008. Haptic gas pedal feedback. Ergonomics 51, 1710–1720.

Mulder, M., Mulder, M., van Paassen, M.M., Abbink, D.A., 2005a. Effects of lead vehicle speed and separation distance on driver car-following behavior, in: 2005 IEEE International Conference on Systems, Man and Cybernetics. Presented at the 2005 IEEE International Conference on Systems, Man and Cybernetics, pp. 399-404 Vol. 1.

Mulder, M., Mulder, M., van Paassen, M.M., Abbink, D.A., 2005b. Identification of driver car-following behaviour, in: 2005 IEEE International Conference on Systems, Man and Cybernetics. Presented at the 2005 IEEE International Conference on Systems, Man and Cybernetics, pp. 2905-2910 Vol. 3.

Mulder, Mark, Abbink, D.A., van Paassen, M.M., Mulder, Max, 2011. Design of a Haptic Gas Pedal for Active Car-Following Support. IEEE Trans. Intell. Transp. Syst. 12, 268–279.

Mulder, Mark, Pauwelussen, J.J.A., van Paassen, M.M., Mulder, Max, Abbink, D.A., 2010. Active Deceleration Support in Car Following. IEEE Trans. Syst. Man Cybern. Paart A-Syst. Hum. 40, 1271–1284.

Na, X., Cole, D., n.d. Game Theoretic Modelling of a Human Driver's Steering Interaction with Vehicle Active Steering Collision Avoidance System.

Na, X., Cole, D.J., 2015. Game-Theoretic Modeling of the Steering Interaction Between a Human Driver and a Vehicle Collision Avoidance Controller. IEEE Transactions on Human-Machine Systems 45, 25–38.

Na, X., Cole, D.J., 2013. Linear quadratic game and non-cooperative predictive methods for potential application to modelling driver–AFS interactive steering control. Vehicle System Dynamics 51, 165–198.

Newell, G.F., 2002. A simplified car-following theory: a lower order model. Transportation Research Part B: Methodological 36, 195–205.

Nishiwaki, Y., Miyajima, C., Kitaoka, N., Itou, K., Takeda, K., 2007. Generation of Pedal Operation Patterns of Individual Drivers in Car-Following for Personalized Cruise Control, in: 2007 IEEE Intelligent Vehicles Symposium. Presented at the 2007 IEEE Intelligent Vehicles Symposium, pp. 823–827.

Odhams, A.M.C., Roebuck, R.L., Lee, Y.J., Cebon, D., 2008. Energy Use in Heavy Vehicles (No. CUED/C-Mech/TR.94). University of Cambridge.

Odhams, A.M.C., Roebuck, R.L., Lee, Y.J., Hunt, S.W., Cebon, D., 2010. Factors influencing the energy consumption of road freight transport. Proc. Inst. Mech. Eng. Part C-J. Eng. Mech. Eng. Sci. 224, 1995–2010.

Okuyama, Y., Murakami, T., 2012. Model predictive control-based drive assist control in Electric vehicle - An application to inter distance control considering human model, in: 2012 13th Int'l Workshop on Mechatronics (MECATRONICS) , 2012 9th France-Japan 7th Europe-Asia Congress on and Research and Education in Mechatronics (REM). Presented at the 2012 13th Int'l Workshop on Mechatronics (MECATRONICS) , 2012 9th France-Japan 7th Europe-Asia Congress on and Research and Education in Mechatronics (REM), pp. 153–160.

Przybyla, J., Taylor, J., Jupe, J., Zhou, X., 2012. Simplified, data-driven, errorable car-following model to predict the safety effects of distracted driving, in: 2012 15th International IEEE Conference on Intelligent Transportation Systems (ITSC). Presented at the 2012 15th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 1149–1154.

Qu, T., Chen, H., Cao, D., Guo, H., Gao, B., 2015. Switching-Based Stochastic Model Predictive Control Approach for Modeling Driver Steering Skill. IEEE Trans. Intell. Transp. Syst. 16, 365–375.

Rajamani, R., 2011. Adaptive Cruise Control, in: Vehicle Dynamics and Control. Springer Science & Business Media, pp. 141–170.

Rawlings, J.B., Mayne, D.Q., 2009. Model predictive control: theory and design. Nob Hill Pub, Madison, Wis.

Samuel, S., Austin, L., Morrey, D., 2002. Automotive test drive cycles for emission measurement and real-world emission levels-a review. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering 216, 555–564.

Schumann, J., Godthelp, J., Hoekstra, W., 1992. An Exploratory Simulator Study On the Use of Active Control Devices In Car Driving. Institute for Perception (Netherlands).

Sentoff, K.M., Aultman-Hall, L., Holmén, B.A., 2015. Implications of driving style and road grade for accurate vehicle activity data and emissions estimates. Transportation Research Part D: Transport and Environment 35, 175–188.

Shakouri, P., Ordys, A., Laila, S., Askari, M., 2011. Adaptive Cruise Control System: Comparing Gain-Scheduling PI and LQ Controllers. Presented at the 18th IFAC World Congress, Milano, pp. 12964–12969.

Sharp, R.S., 2007. Application of optimal preview control to speed-tracking of road vehicles. Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science 221, 1571–1578.

Sivaji, V., Sailaja, M., 2013. Adaptive Cruise Control Systems for Vehicle Modeling Using Stop and Go Manoeuvres. International Journal of Engineering Research and Applications 3, 2453–2456.

Summala, H., Nieminen, T., Punto, M., 1996. Maintaining Lane Position with Peripheral Vision during In-Vehicle Tasks. Human Factors: The Journal of the Human Factors and Ergonomics Society 38, 442–451.

Swaroop, D. v a h g, 1997. String Stability Of Interconnected Systems: An Application To Platooning In Automated Highway Systems. California Partners for Advanced Transit and Highways (PATH).

Tang, T.-Q., He, J., Yang, S.-C., Shang, H.-Y., 2014. A car-following model accounting for the driver’s attribution. Physica A: Statistical Mechanics and its Applications 413, 583–591.

Toffetti, A., Iviglia, A., Arduino, C., Soldati, M., 2014. “ecoDriver” HMI feedback solutions, in: Proceedings of the Human Factors and Ergonomics Society Europe.

van Driel, C., van Arem, B., 2007. Traffic Flow Impacts of a Congestion Assistant, in: Proceedings of TRB 87th Annual Meeting. Washington DC.

Varhelyi, A., Hjalmdahl, M., Hyden, C., Draskoczy, M., 2004. Effects of an active accelerator pedal on driver behaviour and traffic safety after long-term use in urban areas. Accid. Anal. Prev. 36, 729–737.

Varhelyi, A., Persson, A., Kaufmann, C., 2014. In-field evaluation of the effects of Continuous Driver Support on driver behaviour. pp. 1–15.

Vlassenroot, S., Broekx, S., Mol, J.D., Panis, L.I., Brijs, T., Wets, G., 2007. Driving with intelligent speed adaptation: Final results of the Belgian ISA-trial. Transportation Research Part A: Policy and Practice 41, 267–279.

Wahlberg, A., 2002. Fuel efficient driving training - state of the art and quantification of effects.

Yanakiev, D., Kanellakopoulos, I., 1996. Speed tracking and vehicle follower control design for heavy-duty vehicles. Veh. Syst. Dyn. 25, 251–276.

Yang, H.-H., Peng, H., 2010. Development of an errorable car-following driver model. Vehicle System Dynamics 48, 751–773.

Yang, H.-H., Peng, H., Gordon, T.J., Leblanc, D., 2008. Development and validation of an errorable car-following driver model, in: American Control Conference, 2008. Presented at the American Control Conference, 2008, pp. 3927–3932.

Yang, X., Wang, X., Fu, Q., Cheng, Y., Ran, B., 2014. Modeling longitudinal driving behaviors at defective sites on urban expressways. Physica A: Statistical Mechanics and its Applications 410, 359–369.

Zhao, D., Hu, Z., 2011. Supervised adaptive dynamic programming based adaptive cruise control, in: 2011 IEEE Symposium on Adaptive Dynamic Programming And Reinforcement Learning (ADPRL). Presented at the 2011 IEEE Symposium on Adaptive Dynamic Programming And Reinforcement Learning (ADPRL), pp. 318–323.

Appendix 1 – Active Accelerator Pedal Mechanism Analysis

The Active Accelerator Pedal (AAP) analysis included in the main text of Chapter 3 is abbreviated for conciseness. Here, a more detailed explanation, and reasoning behind the approximations made, is included. Note different notation is used here for the purpose of simplifying the derivations. Firstly, a diagrammatic representation of the AAP mechanism with linkage *BC* able to rotate is included in Figure A1-1:

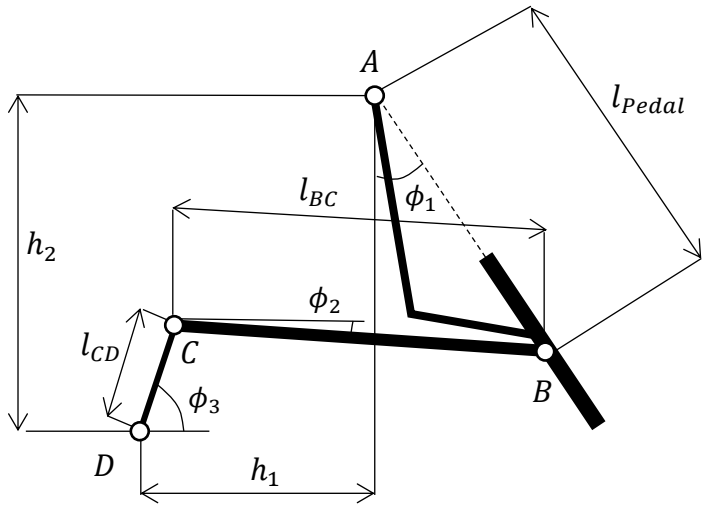

Figure A1-1 – Diagram representation of the AAP mechanism

Firstly, it is worth defining the new reference angles in terms of the ones used in Chapter 3:

$$\phi_1 = \frac{\pi}{2} - \Phi_1 - \phi_{Pedal} \tag{A1.1}$$

$$\phi_3 = \Phi_3 + \phi_m \tag{A1.2}$$

From the diagram, it can be seen that

$$l_{CD}^2 = \left|\overrightarrow{\boldsymbol{AC}} - \overrightarrow{\boldsymbol{AD}}\right|$$

(A1.3)

Now,

$$\overrightarrow{\boldsymbol{AC}} = \begin{Bmatrix} l_{Pedal}\sin\phi_1 - l_{BC}\sin\phi_2 \\ -l_{Pedal}\cos\phi_1 + l_{BC}\cos\phi_2 \end{Bmatrix}$$

(A1.4)

and

$$\overrightarrow{\boldsymbol{AD}} = \begin{Bmatrix} -h_1 \\ -h_2 \end{Bmatrix}$$

(A1.5)

Subtracting (A1.5) from (A1.4) gives

$$\overrightarrow{\boldsymbol{AC}} - \overrightarrow{\boldsymbol{AD}} = \begin{Bmatrix} l_{Pedal}\sin\phi_1 - l_{BC}\sin\phi_2 + h_1 \\ -l_{Pedal}\cos\phi_1 + l_{BC}\cos\phi_2 + h_2 \end{Bmatrix}$$

(A1.6)

The modulus of the resulting vector from (A1.6) is therefore

$$\left|\overrightarrow{\boldsymbol{AC}} - \overrightarrow{\boldsymbol{AD}}\right| = \begin{vmatrix} l_{Pedal}\sin\phi_1 - l_{BC}\sin\phi_2 + h_1 \\ -l_{Pedal}\cos\phi_1 + l_{BC}\cos\phi_2 + h_2 \end{vmatrix}$$

(A1.7)

Combining (A1.3) and (A1.7):

$$l_{CD}^2 = (l_{Pedal}\sin\phi_1 - l_{BC}\sin\phi_2 + h_1)^2 + (-l_{Pedal}\cos\phi_1 + l_{BC}\cos\phi_2 + h_2)^2$$

(A1.8)

Which rearranges to

$$f_2(\phi_1)\cos\phi_2 + g_2(\phi_1)\sin\phi_2 = m_2(\phi_2)$$

(A1.9)

where

$$f_2(\phi_1) = -2l_{Pedal}l_{BC}\cos\phi_1 + 2h_2l_{BC}$$
$$g_2(\phi_1) = -2l_{Pedal}l_{BC}\sin\phi_1 - 2h_1l_{BC}$$
$$m_2(\phi_1) = l_{CD}^2 - l_{Pedal}^2 - l_{BC}^2 - h_1^2 - h_2^2 - 2l_{Pedal}(h_1\sin\phi_1 - h_2\cos\phi_1)$$

Rearranging (A1.9), squaring both sides, and then using trigonometric identities results in:

$$f_2^2(\phi_1)\,(1-\sin^2\phi_2) = m_2^2(\phi_2) - 2m_2(\phi_1)g_2(\phi_1)\sin\phi_2 + g_2^2(\phi_1)\ \sin^2\phi_2$$

(A1.10)

Using the quadratic formula to solve for $\sin\phi_3$:

$$\sin\phi_2 = \frac{2m_2(\phi_1)g_2(\phi_1) \pm \sqrt{4m_2^2(\phi_1)g_2^2(\phi_1) - 4\big(g_2^2(\phi_1)+f_2^2(\phi_1)\big)\big(m_2^2(\phi_1)-f_2^2(\phi_1)\big)}}{2\big(g_2^2(\phi_1)+f_2^2(\phi_1)\big)}$$

(A1.11)

Now ϕ_2 can be calculated by utilisation of the inverse sine function in the appropriate range.

Similarly, from the Figure A1-1, it can be seen that

$$l_{BC}^2 = \left|\overrightarrow{\boldsymbol{AB}} - \overrightarrow{\boldsymbol{AC}}\right|$$

(A1.12)

Now,

$$\overrightarrow{\boldsymbol{AB}} = \begin{Bmatrix} l_{Pedal}\sin\phi_1 \\ -l_{Pedal}\cos\phi_1 \end{Bmatrix}$$

(A1.13)

$$\overrightarrow{\boldsymbol{AC}} = \begin{Bmatrix} -h_1 + l_{CD}\cos\phi_3 \\ -h_2 + l_{CD}\sin\phi_3 \end{Bmatrix}$$

(A1.14)

Subtracting (A1.14) from (A1.13) gives

$$\overrightarrow{\boldsymbol{AB}} - \overrightarrow{\boldsymbol{AC}} = \begin{Bmatrix} l_{Pedal}\sin\phi_1 - h_1 + l_{CD}\cos\phi_3 \\ -l_{Pedal}\cos\phi_1 - h_2 + l_{CD}\sin\phi_3 \end{Bmatrix}$$

(A1.15)

The modulus of the resulting vector (A1.15) is therefore

$$\left|\overrightarrow{\boldsymbol{AB}} - \overrightarrow{\boldsymbol{AC}}\right| = \begin{vmatrix} l_{Pedal}\sin\phi_1 - h_1 + l_{CD}\cos\phi_3 \\ -l_{Pedal}\cos\phi_1 - h_2 + l_{CD}\sin\phi_3 \end{vmatrix}$$

(A1.16)

Combining (A1.12) and (A1.16):

$$l_{BC}^2 = (l_{Pedal}\sin\phi_1 - h_1 + l_{CD}\cos\phi_3)^2 + (-l_{Pedal}\cos\phi_1 - h_2 + l_{CD}\sin\phi_3)^2$$

(A1.17)

Which rearranges to

$$f_3(\phi_1)\cos\phi_3 + g_3(\phi_1)\sin\phi_3 = m_3(\phi_3)$$

(A1.18)

where

$$f_3(\phi_1) = -2l_{Pedal}l_{CD}\sin\phi_1 - 2h_1l_{CD}$$
$$g_3(\phi_1) = 2l_{Pedal}l_{CD}\cos\phi_1 - 2h_2l_{CD}$$
$$m_3(\phi_1) = l_{BC}^2 - l_{Pedal}^2 - l_{CD}^2 - h_1^2 - h_2^2 - 2l_{Pedal}(h_1\sin\phi_1 - h_2\cos\phi_1)$$

Rearranging (A1 .18), squaring both sides, and then using trigonometric identities results in:

$$f_3^2(\phi_1)\,(1 - \sin^2\phi_3) = m_3^2(\phi_3) - 2m_3(\phi_1)g_3(\phi_1)\sin\phi_3 + g_3^2(\phi_1)\ \sin^2\phi_3$$

(A1.19)

Using the quadratic formula to solve for $\sin\phi_3$:

$$\sin\phi_3 = \frac{2m_3(\phi_1)g_3(\phi_1) \pm \sqrt{4m_3^2(\phi_1)g_3^2(\phi_1) - 4\big(g_3^2(\phi_1) + f_3^2(\phi_1)\big)\big(m_3^2(\phi_1) - f_3^2(\phi_1)\big)}}{2\big(g_3^2(\phi_1) + f_3^2(\phi_1)\big)}$$

(A1.20)

Now ϕ_3 can be calculated by utilisation of the inverse sine function in the appropriate range.

With expressions now available for ϕ_2 and ϕ_3 in terms of ϕ_1, a drive cycle following driving simulator experiment (explained fully in Chapter 3) is completed to examine the expected range for the first and second derivatives of ϕ_1, ϕ_2 and ϕ_3. Figures A1-2, A1-3 and A1-4 illustrate the angle, angular speed, and angular acceleration of the three linkages. It can be seen that the motor arm, denoted by the subscript 3, experiences the highest angular displacements, speeds and accelerations by a significant margin. It is also noted that the linkage *BC* remains very close to the horizontal at all times, and has significantly lower angular speeds and accelerations than the other two bars. For these reasons, it is proposed that the effect of the moment of inertia of the linkage *BC* is negligible and that the bar remains horizontal at all times.

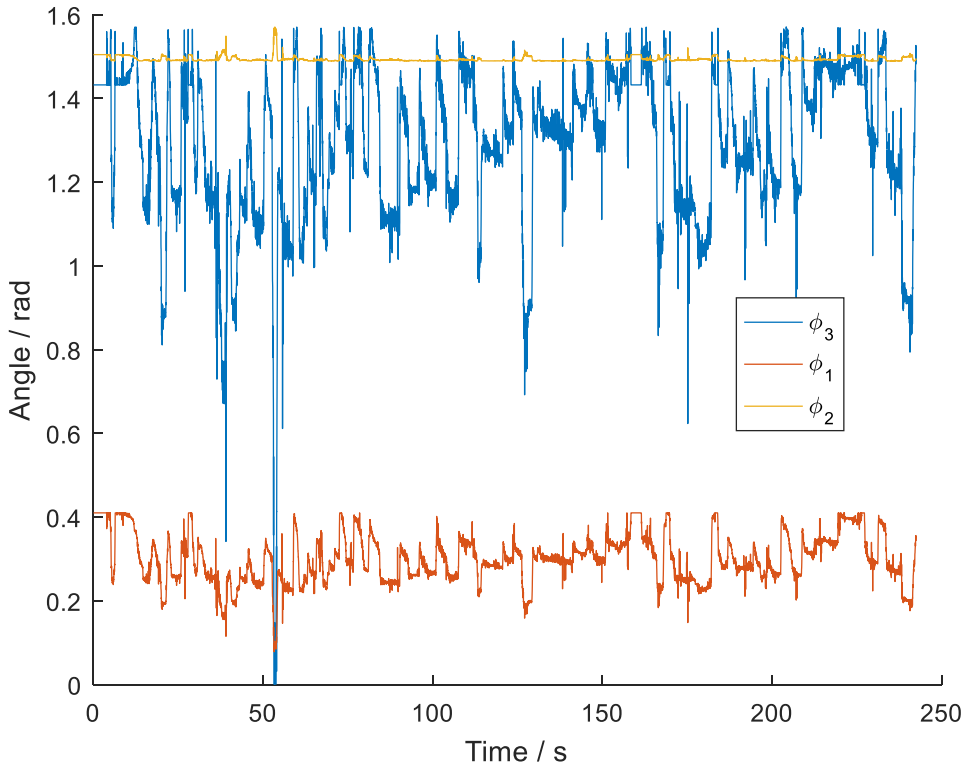

Figure A1-2 – Measured mechanism angles from a drive-cycle driving simulator experiment

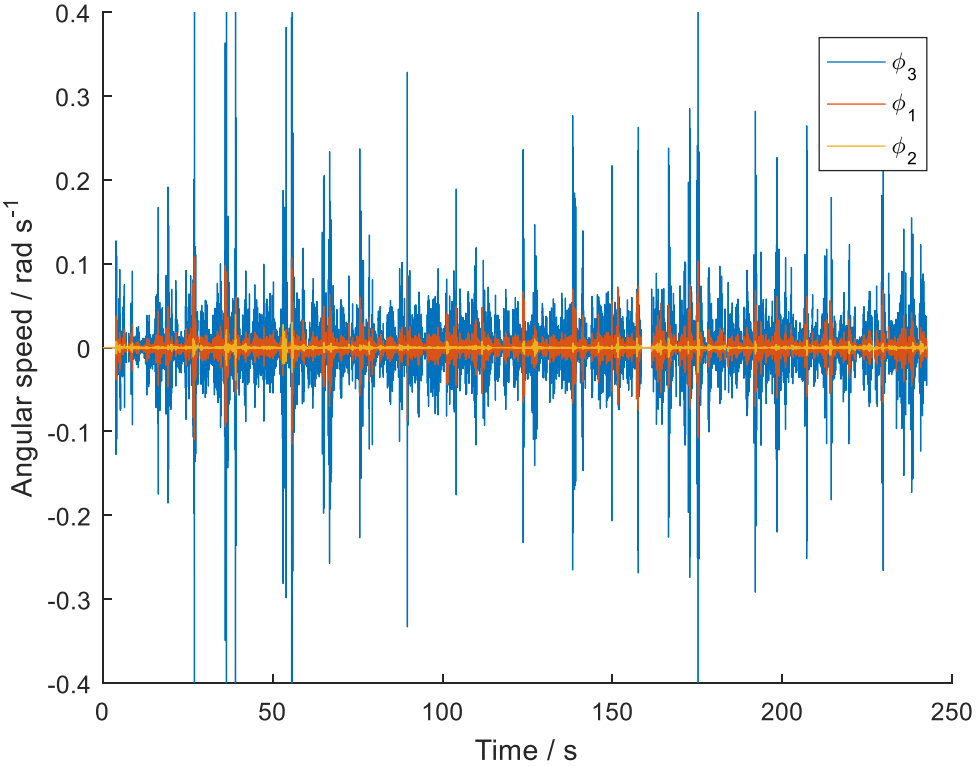

Figure A1-3 – Measured mechanism angular speeds from a drive-cycle driving simulator experiment

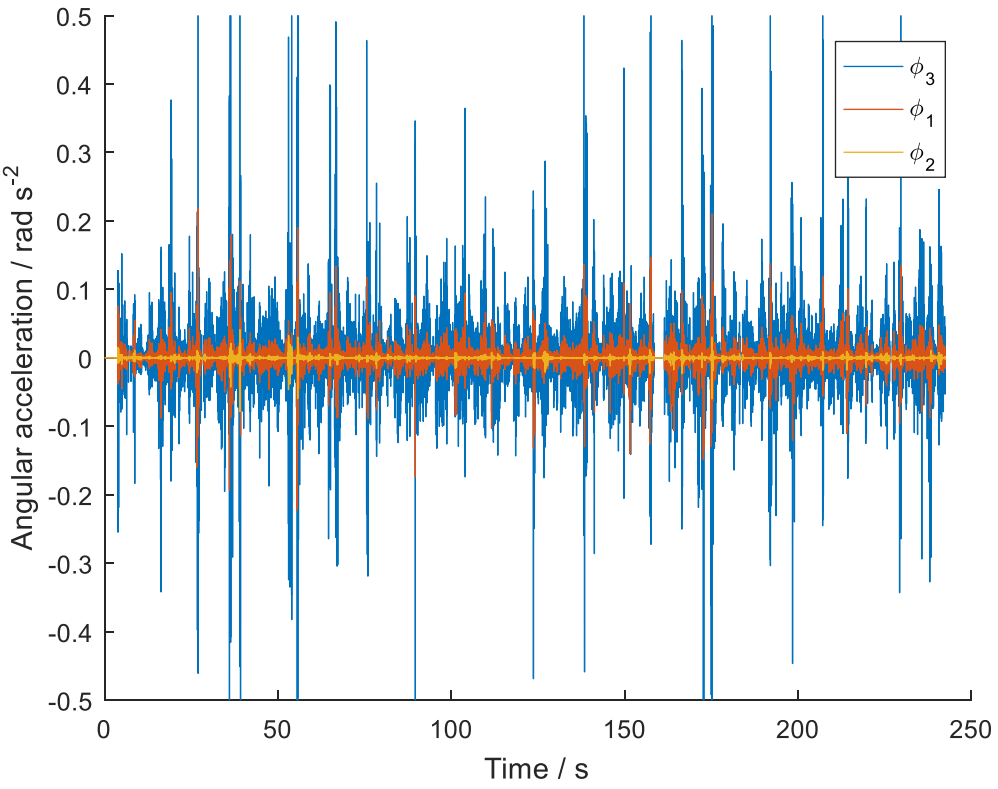

Figure A1-4 – Measured mechanism angular accelerations from a drive-cycle driving simulator experiment

As a further justification to this, the effect of the moment of inertia of *BC* is estimated. From Figure A1-4, at worst case, the maximum angular acceleration recorded was 0.2 $\mathrm{rad\,s^{-2}}$. Modelling the linkage as a rod, of length l_{BC} and mass m_{BC}, the moment of inertia about its centre of mass becomes:

$$I_{BC} = \frac{1}{12} m_{BC} l_{BC}^2$$

(A1.21)

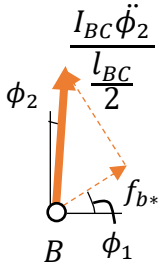

Figure A1-5 – Force components due to the moment of inertia of the linkage *BC*

Now, looking at Figure A1-5, the normal force on the pedal, due to the moment of inertia of BC is given by:

$$f_{B*} = \frac{I_{BC}\ddot{\phi}_2}{\frac{l_{BC}}{2}} \sin\left(\frac{\pi}{2} - \phi_1 - \phi_2\right) \qquad (A1.22)$$

Using the properties from the pedal mechanism in Chapter 3 and assuming maximum values of $\ddot{\phi}_2$ and $\sin\left(\frac{\pi}{2} - \phi_1 - \phi_2\right) = 1$, the effect of the moment of inertia of the connecting rod on the pedal arm itself is expected to be less than

$$f_{B*} = 3.6 \times 10^{-3}\text{N}$$

This is much smaller than the measured forces between the pad of the pedal and the foot (up to approximately 20N), so it is therefore reasonable to assume the effects of the moment of inertia of the linkage *BC* are negligible.

Appendix 2 – Non-Linear Model Predictive Control

In the main text, the objective was always to produce a linear model for operation in the Model Predictive Control Theory (MPC). Linear MPC produces a computationally inexpensive controller as the gains used are time-invariant. However, the restrictions of linear MPC limit the accuracy of the vehicle model used. In order to get around this, and to add a fuel consumption component in the cost function, a non-linear solver can be implemented. This appendix documents some considerations for dealing with a non-linear model.

Two limiting factors observed in the linear approach are: the simplicity of the engine torque model, and the absence of a fuel consumption model. The torque output from an engine is highly dependent on the speed the engine is operating at. It is restricted by the maximum performance of the engine, typically, either a maximum torque rating, or maximum power rating, depending on engine speed. The matching of the engine torque to load is completed by an automatic gearbox controller in many trucks. In order to fully reflect this, a non-linear torque map is required. Secondly, for full analysis of the effectiveness of any pedal force feedback strategies used, a detailed and accurate fuel consumption model is needed to predict fuel consumption rates at a time step by time step basis.

The control problem remains the same – an optimal set of control actions is sought that minimises the cost function. In the linear method, QR decomposition was used to find the optimal set of gains that, when applied to the preview, produced the desired output. In the proposed non-linear controller, the vehicle model is no longer time invariant as there are non-linear dependencies on speed present.

In the non-linear approach taken here, the vehicle model is linearized at an initial start point about its current states and every half second afterwards.

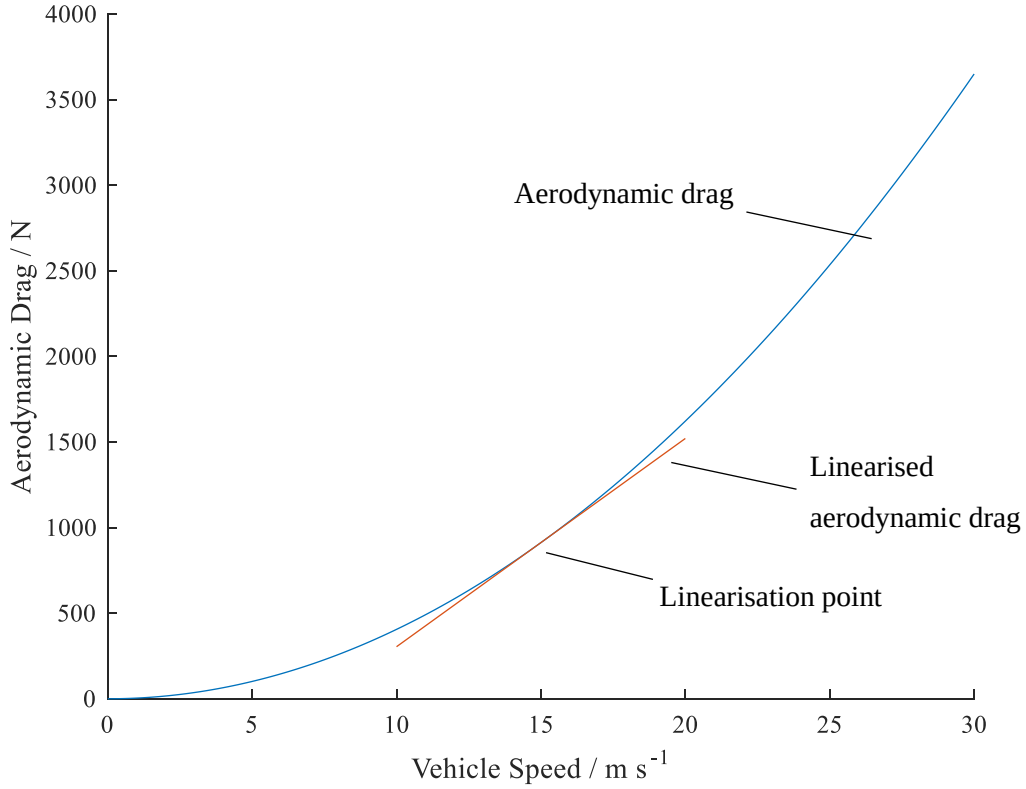

Figure A2-1 – Linearised aerodynamic drag

Conventional aerodynamic drag equations take a quadratic form (A2.1). Here, the aerodynamic drag is linearised about the vehicle speed every half second, by calculating the gradient of the aerodynamic drag against vehicle speed curve: (A2.2) and

Figure A2-1.

$$F_{Aero} = \frac{1}{2}\rho_{air} C_d A v^2$$

(A2.1)

$$F_{AeroL} = \frac{1}{2}\rho_{air} C_d A (2v_0 \times v - v_o^2)$$

(A2.2)

The Matlab function fminsearch was called to find an optimal set of pedal forces that minimised the cost function over the preview. Notice that now the pedal forces are the direct output of the optimisation, rather than the gains applied to the preview information.

As the computational cost of such a controller is so much higher than a linear one, a useful characteristic of human driver control was exploited to reduce the cost: intermittent control. In the linear approach, the preview gains are applied at each time step and only the first control action is

implemented in a receding horizon approach. However, real human drivers to not optimise their control at such high frequencies as there are other tasking occupying the cognitive control mechanism as well in what is known as the central bottleneck (Johns and Cole, 2015).

A driver assesses the current states of the vehicle at a given point and then calculates a set of control actions. This processing takes some time. Once the calculations are complete, the cognitive processing then switches to secondary tasks, if there are any, as the control signal is sent through the limbs and is implemented. When the secondary task is processed, the driving task is then reassessed; the current states are assessed and new actions calculated. This means that there is a finite time between control updates, called the control update time. The control update time is well-studied in the steering control area, and typical update times are approximately 0.5s.

The same control update time is used in the model here, where 0.5s of control are taken and implemented before the next optimisation occurs. This can result in some large discontinuities in the demanded pedal force, and these are filtered out by the limb dynamics in the human driver.

A second step was also taken to reduce the computation cost of the simulation. Instead of optimising the pedal force at every time step over the preview, several points are selected to be optimised and intermediate points are calculated by linearly interpolating between defining points. For example, in a ten second preview, the first second may have defining points at 0.1, 0.2, 0.3 0.4. 0.5 1.0, 2.0, 4.0, 6.0, 8.0 and 10.0 seconds. The spacing increases at the distant end of the preview, as the information is less critical to the control action. The computational resources are focused at the key time frame of the first second.

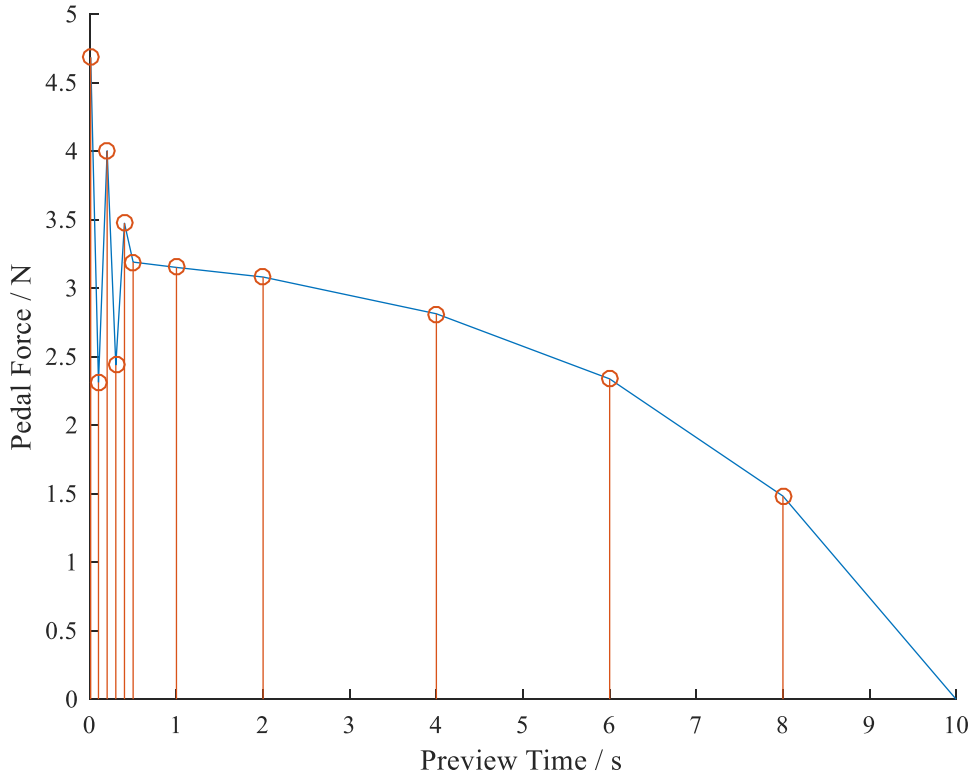

Figure A2-2 – Force optimisation points with a sample set of forces

Figure A2-2 illustrates how the set of defining forces produce a corresponding set of control forces for the timings given in the example. The first half second, sampled at the highest rate is implemented, and the second half second is used as the input to the next optimisation, when the vehicle model has been linearised about its new states.

At this point, the controller is given a simple ramp in speed demand, as illustrated by the red line in Figure A2-3. With a 20 second drive cycle, the controller will produce 10.5 seconds of control actions (the remaining 9.5 seconds are part of the preview information made available to the controller. It can be seen that the initial response oscillates a couple of times before settling in to accurately follow the target.

Figure A2-4 shows the pedal forces corresponding with the ramp speed target. There is a definite trend in the force to increase with time, as when the vehicle accelerates (even with constant acceleration), the drag forces increase meaning that a higher load is being put on the engine.

This Appendix has briefly considered aspects of non-linear control as the basis for further work on the subject.

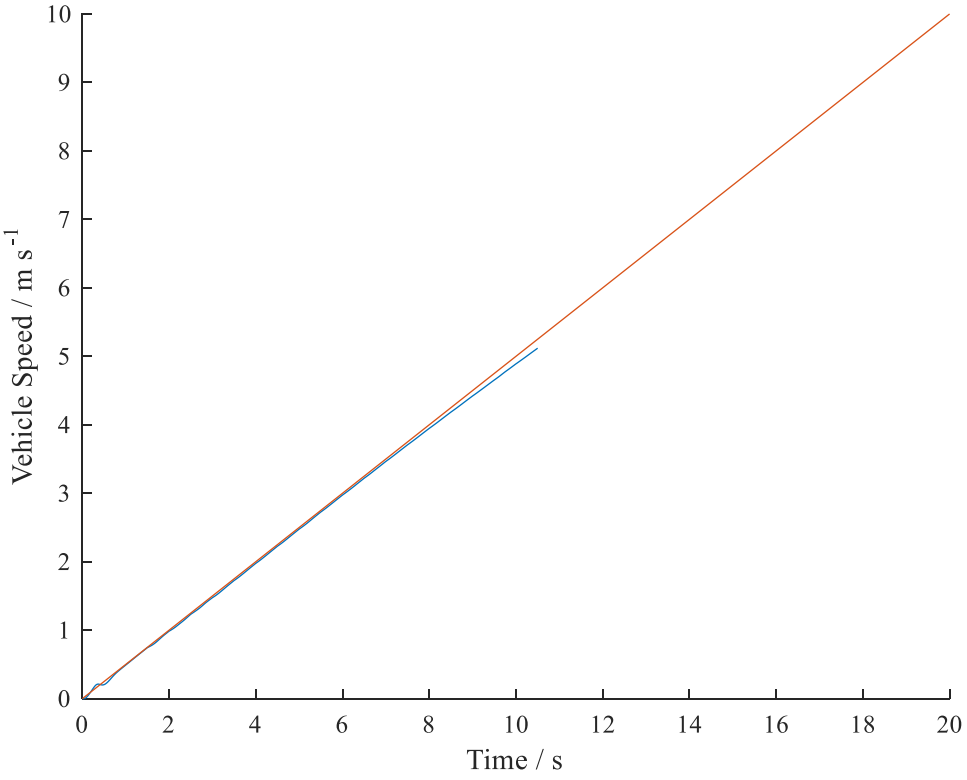

Figure A2-3 – Vehicle speed when set a 0 to $10 ms^{-1}$ ramp target over 20s.

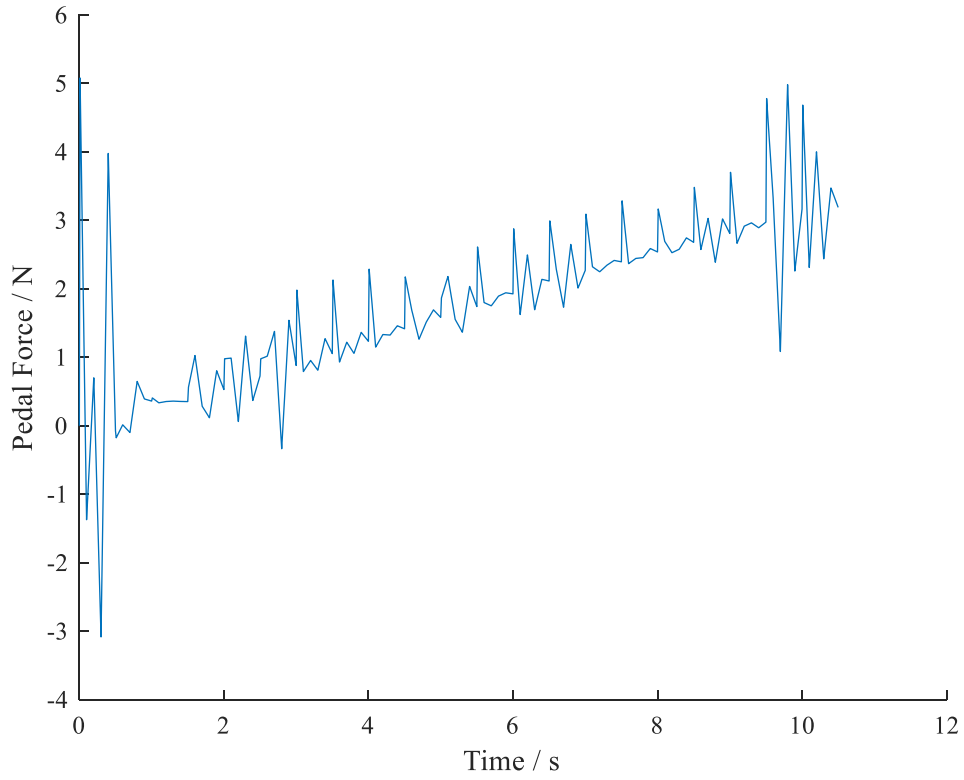

Figure A2-4 – Driver pedal forces for a ramp speed demand