UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and the Environment

Transportation Research Group

Developing data collection methods to inform the quantitative design of cycle infrastructure

by

Christopher John Osowski BSc(Hons) MSc

Thesis for the degree of Doctor of Philosophy

June 2017

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT
Transportation Research Group

Doctor of Philosophy

DEVELOPING DATA COLLECTION METHODS TO INFORM THE QUANTITATIVE DESIGN OF CYCLE INFRASTRUCTURE

by Christopher John Osowski BSc(Hons) MSc

Increased share of urban travel by bicycle is widely desired as a cost-effective and environmentally-beneficial means of travel, and one which has the potential to reduce road congestion and improve health outcomes. Recent rapid cycling growth in cities such as London has served to highlight the lack of robust empirically-backed quantitative literature to inform the practitioner, and the consequential barrier to the delivery of enabling infrastructure of the scale required to meet that demand.

Even simple measures vary by orders of magnitude in the literature and some depend on intuitively naïve assumptions, so a simulation (based on the Social Force Model) was defined and implemented to test the key underpinning (non-interaction) assumption of the Highway Capacity Manual's quantitative definition of cycle level of service. The simulations indicate that an assumption of non-interaction between cyclists results in an outcome intrinsically at odds with fundamental traffic flow theory. Both the literature and simulation process serve to highlight the lack of existing appropriate empirical data and behavioural understanding. Furthermore, collecting such data is difficult, expensive and not easily scalable using current methods.

Consequently, a methodology for the collection of key cyclist parameters from generic video data was created, and can be applied to bespoke video surveys and existing CCTV capture, across a variety of modes, and at a fraction of the cost of human operators.

In addition, a bicycle simulator is developed which can test cyclist behaviour in a replicable manner and in a range of circumstances. The design and construction process is detailed, and a proof-of-concept, validated against real data, is presented. Subject to some minor improvements identified, the simulator can now be used more widely for the collection of behavioural data.

These methodologies provide new and practical capabilities for the collection and application of cyclist data, and a greater understanding of cycle behaviour.

Contents

List of Figures

List of Tables

Author's Declaration

I, Christopher John Osowski, declare that the thesis entitled *Developing data collection methods to inform the quantitative design of cycle infrastructure* and the work presented in it are both my own, and have been generated by me as the result of my own original research. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University;
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- Where I have consulted the published work of others, this is always clearly attributed;
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- Parts of this work have been published as: Osowski & Waterson (2015) and Osowski & Waterson (2017).

Signed: .....

Date: .....

Acknowledgements

This work was financially supported by an EPSRC Doctoral Training Centre grant (EP/G03690X/1) through the University of Southampton Institute for Complex Systems Simulation. The funders had no role in study design, data collection and analysis, or preparation of this document.

The author also acknowledges the University of Southampton through: the use of the Iridis High Performance Computing Facility (and its associated support services), upon which many of the presented simulations were run; and Jason Noble and John Preston for their various supervision, critique and direction.

Chapter 5 could not have been completed without the assistance of the National Oceanography Centre Southampton's Estates & Facilities and Security teams; nor the contribution of time and leg-power of the participants. Similarly, the author's gratitude goes to the participants in the bicycle simulator pilot study (Chapter 8); not least for their iron stomachs.

Final acknowledgements go to Jono Gray and Vincent Marmion for maintaining my sanity throughout this endeavour; and to Ben Waterson for his supervision prior to and during this process, and for the substantial amount of time and support given throughout.

Abbreviations

3D	Three Dimensional
AADT	Annual Average Daily Traffic
AASHTO	American Association of State Highway and Transportation Officials
ABS	Acrylonitrile Butadiene Styrene
AI	Artificial Intelligence
API	Application Programming Interface
ASCII	American Standard Code for Information Interchange
ASL	Advanced Stop Line
bpmph	Bicycles per metre width per hour
CA	California
CAD	Computer-Aided Design
CCTV	Closed-circuit television
CO_2	Carbon dioxide
COBA	Cost Benefit Analysis
CSH	Cycle Superhighway
CTC	Cyclists' Touring Club (NB. Cycling UK since April 2016)
DARPA	Defence Advanced Research Projects Agency
DfT	Department for Transport
EPSRC	Environmental and Physical Sciences Research Council
EU	European Union
EWLCSH	East–West London Cycle Superhighway
FH(W)A	Federal Highway Administration
FoV	Field-of-View
FV	Financial Year
GUI	Graphical User Interface
HCM	Highway Capacity Manual
HDTV	High-Definition Television
HEAT	Health Economic Assessment Tool
HFET	Hall Field-Effect Transistor
HMD	Head-Mounted Device
IAN	Interim Advice Note
IR	Infrared
ITU	International Telecommunications Union
LCD	Liquid Crystal Display
LCDS	London Cycling Design Standards
LED	Light-Emitting Diode

LoS	Level(s) of Service
MSQ	Motion Sickness Questionnaire
MUTCD	Manual on Uniform Traffic Control Devices
NOC	National Oceanography Centre
NY	New York
NYC	New York City
PC	Personal Computer
PCU(/PCE)	Passenger Car Unit (/ Passenger Car Equivalent)
PhD	Doctor of Philosophy
PTV	PTV Planung Transport Verkehr AG
px	Pixel
QQ	Quantile-Quantile
QR	Quick-Release
RCT	Randomised Controlled Trial
RGB	Red-Green-Blue
RR67	Research Report 67 (Kimber et al., 1986)
SFM	Social Force Model
SS	Simulator Sickness
SSQ	Simulator Sickness Questionnaire
SUV	Sports Utility Vehicle
TCP	Transmission Control Protocol
TfL	Transport for London
TSRGD	Traffic Signs Regulations and General Directions
TX	Texas
UK	United Kingdom
US(A)	United States (of America)
USB	Universal Serial Bus
VDU	Video Display Unit
VR	Virtual Reality
WA	Washington

Chapter 1

Introduction

1.1 The General Case for Cycling

Increasing the mode share contribution of cyclists is a widely held transport policy aspiration across a range of international jurisdictions (e.g. City of Copenhagen, 2011b; CROW, 2007; Dekoster & Schollaert, 1999; Greater London Authority, 2013; L'Ile-de-France service de presse, 2011; New York City Department of Transport, 2008). An increase in cycling levels is held to reduce motor vehicle congestion by "increasing the peak flow of vehicles" on a given route (British Cycling, 2014), and also to cost-effectively improve public transport capacity (Greater London Authority, 2013). For example, the East-West London Cycle Superhighway on Victoria Embankment in London had a proposed capacity "equivalent to almost four entire trainloads of people (based on seating capacity) on the District and Circle lines beneath the same street[, which] could increase effective capacity on this stretch of the Underground by as much as 10 per cent – and for relatively minimal outlay." (Greater London Authority, 2013).

Environmentally, a transfer to cycling is held to reduce CO_2 pollution (as a 'low-carbon' alternative to private motor vehicles) and "improve air quality" (through emission reduction; Department for Transport, 2008). Economically, cycle scheme implementations have been shown to bring more spending in to local businesses than in other non-cycle friendly areas. For example, a scheme to implement a protected bicycle lane on 9th Avenue in Manhattan, New York City, found a 49% increase in retail sales after implementation compared to only 3% borough-wide (New York City Department of Transportation, 2013).

Individuals also see health benefits (as cycling is an 'active' transport mode; World Health Organization, 2011) and these also have wider societal effects such as improved employee productivity (e.g. "employees who cycle to work are more alert when arriving at work and 15% more productive" Cambridge Cycling Campaign, 2014) and reduced inactivity-related burden on health care services. Jarrett et al. (2012)estimated a cost saving to the UK National Health Service of some 17bn through to 2030 – 2010 prices – if assumed long-term

high-trend future growth in active travel were brought forward.

There is therefore, a strong case for the promotion of cycling.

It is generally agreed that the provision of high-quality cycle infrastructure is both a key driver: in reversing the general decline in cycle use in the face of the positive feedback between motor vehicle use and associated land-use planning (e.g. Department for Transport, 2008); and of building upon existing cycle use corridors (e.g. CROW, 2007). However, whilst a general aspiration of 'high quality' cycle infrastructure is perhaps universally agreed, what exactly constitutes such infrastructure is not; even amongst representative user lobby groups (e.g. Amigo, 2016; Cycling Embassy of Great Britain, 2014).

Indeed, ostensibly 'good' infrastructure – including in those circumstances where cycle use already provides a substantial mode share contribution – can still suffer from similar issues to that of highway traffic; e.g. congestion (City of Copenhagen, 2011a). Often, the issue can be distilled down to the disconnect between the qualitative aspiration and the quantitative spatial, political and economic realities which must be addressed in order to deliver (and maintain) a scheme.

Robust quantitative measures for cycle infrastructure can help to direct delivery to achieve the desired outcomes. Generally cycle schemes and infrastructure are delivered at the municipal or sub-regional level but relatively immeasurable benefits – such as cost savings to national health care services – are essentially externalities and thus are not well-captured in any local attempt to justify a scheme.

Establishing a quantitative measure by which highway (and/or pedestrian realm) schemes can be objectively measured and valued at the local level is dependent on the understanding of the capacity for and behaviour of constituent vehicles (and/or pedestrians). The continued development of modelling tools such as Vissim (PTV, 2013a) and SATURN (Atkins, 2013), their validation through use, and the increases in available computer power, have allowed the use of quantitative tools on all scales of highway scheme design. Similarly, pedestrian modelling tools such as Pedroute (Halcrow, 2010), VisWalk (PTV, 2013c) and Legion (Legion Ltd., 2013) are widely used to simulate and ex ante test the design of pedestrian spaces, and quantitatively assess their operational and/or safety-critical aspects.

Quantitative considerations are key to the UK transport planning process at the local scale (e.g. Transport for London, 2010c) and, on a larger scale, through processes such as COBA (Department for Transport, 2004). For bicycles however, the scale of data required to establish robust measures is simply not available. The lack of existing widespread deployed automated counters for cycle traffic (or the lack of reliably capturing such data from non-cycle specific equipment already deployed), and the difficulty of accounting for homogeneous infrastructure and traffic flow, remains (amongst other reasons) a substantial

barrier to the cost-effective collection of such data.

1.2 Motivation

1.2.1 Issues of Scale

To some extent, quantitative issues are considered secondary to the broader arguments surrounding the implementation of cycle infrastructure issues such as policy, cost, deliverability and politics. However, quantitative aspects underpin many of these issues.

As a case in point, London has a population population of London 10.313 million (in 2015; CIA, 2016b) versus that of Amsterdam (held to be an exemplar of a world-leading cycling city) at 1.091 million (CIA, 2016a). London's target is that cycling comprises 5% of *all* trips by 2025 (Golbuff & Aldred, 2011) but, even as of 2012, despite only a 2% London-wide mode share, 5% of Londoners working in the City of London and 6% of those working in the City of Westminster already commute by bicycle (Transport for London, 2012); these two Boroughs covering the majority of the central business district. Some inner-city areas such as Hackney have even higher shares (7%) and the overall number of cyclists on London roads has risen by 32% since 2012 (Transport for London, 2016). The concept of a 5% cycling mode share in London implies an actual number of cyclists roughly equivalent to those in Amsterdam (albeit they are likely distributed over a wider area) if cycle mode share there is assumed to be 50%, itself a high-end (but not unrealistic) estimate.

With cycle mode share forming the plurality (if not the outright majority) of peak vehicle movements across natural funnel-points already (such as the central London bridges, with Blackfriars Bridge having a 70% cycle vehicle share in 2016; Reid, 2016; Transport for London, 2012), achieving greater mode shares than this 'low' 5% level in London (or anything even approaching the 5% target) may require the consideration of infrastructure on a scale not previously conceived.

For example, if one assumes that AM peak cycle flow is uniformly distributed over the 3 hour 0700–1000 peak (an unlikely but conservative assumption), in 2013 some 880 bicycles per hour were surveyed travelling northbound on Blackfriars Bridge and a similar number over Waterloo Bridge (Transport for London, 2013). As established above, this will likely have increased since, in line with underlying cycle mode growth and that Blackfriars Bridge is also part of the route for the recently delivered (opened 2016) North–South Cycle Superhighway in central London (Transport for London, 2015b) which will also have likely attracted cyclists from other less welcoming parallel routes, including Waterloo Bridge.

This two-way section of cycle track has been delivered at approximately a 4m width. AM-peak southbound flow on Blackfriars Bridge (Transport for London, 2013) is some 30% of northbound flow indicating a total two-way flow of around 1200 bicycles per hour, itself almost certainly now a low estimate for

reasons discussed above. Such a flow is an order of magnitude larger than the 150 cyclists per hour that, elsewhere such as The Netherlands, constitutes a "comfort capacity" for a track of this dimension (CROW, 2007).

Within only a few months of opening, potential issues of capacity were demonstrated. In early October 2016, a cycle crash involving at least 3 cyclists on the East–West Cycle Superhighway (also at Blackfriars, but in the underpass of the bridge discussed above) was captured on video (4ChordsNoNet, 2016) and widely reported (Gillett, 2016; Kelly, 2016). Though other events had almost certainly happened that did not capture media attention, a recurring theme in the surrounding commentary and coverage was that the facility was already "at capacity". What exactly capacity means in this circumstance is perhaps not clear, however what is clear is that user perception is already shifting to an opinion that there is a capacity limitation in the peaks.

London is by no means unique; regularly surveyed cyclists in Copenhagen (considered one of the world-leading cycle cities) report a generally falling satisfaction with "cycle path width" over the last two decades (City of Copenhagen, 2013a), indicative of a growing problem of congestion. Indeed, a recent reported increase in those satisfaction scores are linked to a programme of cycle track expansion in some of the busiest places. 'Rush-hour' bicycle congestion is now a common feature in Amsterdam (Van Mead, 2016).

Given then that current qualitative best practice (e.g. CROW, 2007) is unable to scale to rush-hour demands even in the places for which it was created, it is unlikely to scale to the order of magnitude of cyclists possible should a world-city, such as London or New York, truly adopt cycling as a serious modal share.

1.2.2 Economic Valuation

Approaching the subject from a different angle, there exists no method of economically valuing cycle infrastructure properly. Should a stakeholder query: "What is the value of this cycle lane, crossing, etc.? What is the business case for it?", the practitioner, at best, can only respond with valuations of the potential for increased cycling leading to reduced healthcare costs and/or other intangible social benefits. General tools such as 'HEAT' (Health Economic Assessment Tools) created by the World Health Organization (2011) attempt to quantify these high-level items, but make no attempt to consider the operational level. Such a lack of quantification makes it difficult to justify spend at the local level.

In the UK, where there is a desire to increase cycling but where cycle spending is only £2 per head as opposed to general road spending of £75 per head of population (British Cycling, 2014), there is a standing need for robust economic evaluation tools. Even in those places where spend has locally been substantially higher and schemes have already been (or are being) delivered, there

is a standing need to be able to evaluate those schemes which have been implemented to inform ongoing spend, and spend elsewhere.

Department for Transport (2014) presents the value for money assessment of the cycle schemes delivered (and being delivered) through the Cycle City Ambition Grant and Cycling in National Parks Grant programmes (some £94m total funding, awarded in 2013; 2010 prices); with benefit:cost ratios of some 5:1 and 7:1 identified respectively (combined, 5.5:1). Of these, the overwhelming majority (approximately 80%) are attributable to 'physical fitness' (calculated based on World Health Organization, 2011) and 'journey ambiance'. However actual assessments are quantified based on a simple 'better'/'neutral'/ 'worse' qualitative (and subjective) determination, and whilst these are important benefits to any scheme delivery, neither are operational. This assessment is demonstrative of the lack of the ability to quantify the operational impacts of cycle interventions on the cyclists themselves.

The majority of the remaining values in Department for Transport (2014) are based around the impact on road traffic and commensurate reduction in noise, CO_2 emissions etc. However the inability of such a consideration to capture scheme value is a standing item of friction between local transport authorities and national authorities in the UK: "Bristol [City Council] asserted that the DfT's modelling massively under-estimates the benefits of cycling walking, claiming that meeting the 2025 targets would result in benefits of £61bn, 'yet the DfT envisages £500m'." (Baker, 2016).

Returning to London, the East–West London Cycle Superhighway (EWLCSH) scheme across central London (first announced as part of "The Mayor's Vision for Cycling in London"; Greater London Authority, 2013) was controversial with some motoring groups and groups comprised of a high proportion of road users (e.g. the London Taxi Drivers' Association and the Canary Wharf Estate; Lydall, 2015) mounting a series of strong (including legal) objections at all stages of consultation and delivery. Despite strong public and political backing (some 84% of consultation respondents; Transport for London, 2015a), TfL's motor vehicle traffic modelling (using Vissim) indicated delays to motor traffic, particularly in the east of London, as a result of the reduced road capacity through the centre of the city. Modifications to the scheme were made to reduce these, however modelling for the final scheme (as approved and delivered) still predicted delays which generated net present costs sufficient to create a negative benefit:cost ratio for the scheme (Transport for London, 2015b).

Whilst the underlying business case calculations that are presented in summary in Transport for London (2015b) are not available, the Business Case Development Manual (Transport for London, 2010a) provides comprehensive detail for TfL's methods and the likely underpinning of this business case. Inspection of said Manual highlights the limited capture of cycle effects in either being based in derivations from the 'HEAT' cycling tools (World Health Or-

ganization, 2011) or in a series of specific yet uncited (in what is otherwise a heavily cited document) values in pence per journey (or minute of journey; see Table E4.4.5 therein). Those values which in reality are time-based (which includes the categorical infrastructure width items) are fundamentally sensitive to factors that may influence journey time (like congestion), yet insensitive to them in any business case based on this Manual. In a similar vein, the per-journey figures cover a number of items which in reality are likely distance-dependent to the user (such as surface quality, segregation or standing water) and yet have binary numerical application to the TfL Business Case process.

In any case, the eventual delivery decision for the EWLCSH hinged on political and other value which is not captured in the approved business case. Indeed the briefing document that went to the approval board (Transport for London, 2015b) makes a strong case for the value of the scheme, despite what the business case might imply:

"Whilst the benefit to cost ratio is not positive for the proposed East-West Cycle Superhighway based on the monetised benefits that can be captured, this does not mean that the project is poor value for money. This scheme is an essential part of the wider cycling network and re-allocates space from motorised modes to cycling, allowing more people to travel on the system as a whole."

Implicit in the above quote is an uplift in overall road network user capacity that can be expected from the reallocation of motor vehicle space. Furthermore, this has been a theme since the scheme was announced in 2013:

"Our new segregated East-West Superhighway along the Victoria Embankment will have the capacity for about 1,000 cyclists an hour, each way. That is equivalent to almost four entire trainloads of people (based on seating capacity) on the District and Circle lines beneath the same street. We could increase effective capacity on this stretch of the Underground by as much as 10 per cent - and for relatively minimal outlay." (Greater London Authority, 2013)

It is evident that there is a desire to quantify this (and other) cycle schemes and to frame their value in those quantities, as is the case with motor vehicle schemes. Even if quantification is only a component part of the case for (or against) a scheme, or is only to rank schemes relative to one another, and even if in the extreme case one were not to consider quantification to have any intrinsic benefit, there remains an unmet demand for quantification from delivery practitioners and stakeholders.

1.3 Aims and Objectives

There is then, a need to develop the tools necessary to inform the design of cycle infrastructure in a quantitatively robust manner. This document therefore seeks to do the following:

1. Demonstrate that the state-of-the-art of the quantitative cycle infrastructure capacity literature is insufficient to meet practitioner needs;
2. Identify the appropriateness of quantitative modelling techniques as applied to cycle infrastructure;
3. Improve understanding of methodologies for data collection pertaining to cyclists; and
4. Develop and apply those methodologies to determine if there is scope for their further development and/or use. In particular:
   (a) Identify the scope for the use of video observation as a data collection method for bicycle parameters.
   (b) Identify the scope for, design, construct and test an immersive simulator as a data collection method for bicycle parameters.

1.3.1 This Document

In order to address the above aims, this document is structured as shown in Figure 1.1. Chapter 2 starts the main body of the document by drawing together the state of the current literature (both practitioner and the academic literature that can inform it) that is available to inform the quantitative delivery of cycle infrastructure, by the practitioner.

Chapter 3 then moves on to consider the state-of-the-art with regard to quantitative modelling of cyclists before moving on to directly modelling cyclists by taking the current literature and applying it to an implementation of the Social Force Model pertaining to bicycles. A result of this modelling is that the range of parameters necessary to deliver such a model is not robustly backed by the literature. The standing issue is therefore the need to collect data on a scale not currently achieved.

To enable this, Chapter 4 reviews the potential methodologies available to collect cyclist data on a large scale. For the rationale discussed therein, instrumentation of bicycles is declared out of the scope of this project. Chapter 5 then follows with the implementation and testing of a process to collect cyclist data by remote observation on video, and in an automated manner. In addition to reducing the labour-intensity of data collection for cyclists, the process also has the capability of being applied to existing data already held and non-dedicated data sources such as CCTV, whilst also providing a transparency

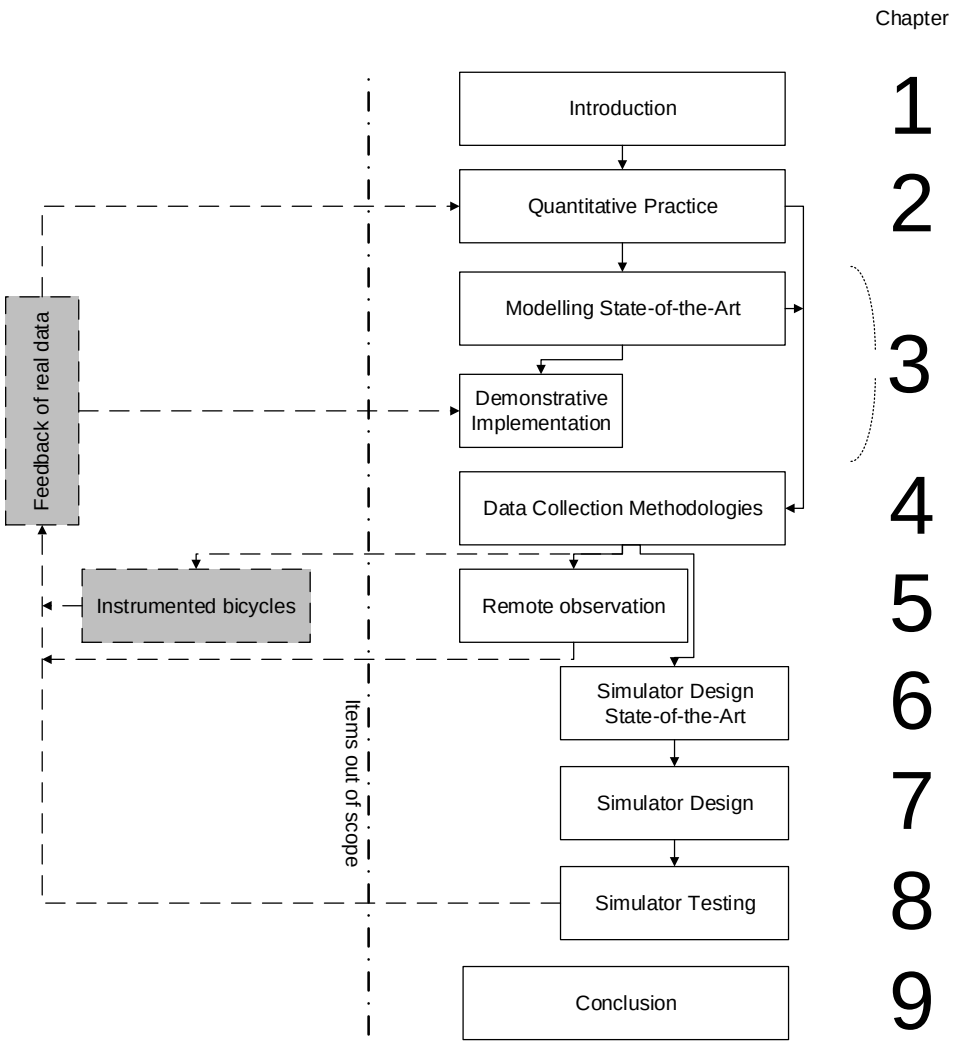

Figure 1.1: Document Structure

of process not available from proprietary products which have emerged in the time since that methodology was developed.

Chapter 6 then considers the literature necessary to deliver a cycle simulator and establishes a specification against which one can be delivered. Chapters 7 and 8 then follow with the design, testing and review of a bicycle simulator that meets that specification. This simulator allows users to cycle in virtual environments and provides the opportunity to collect data from cyclists in experimentally-controlled situations, in repeatable circumstances, and in scenarios which do not need to already physically exist.

Chapter 9 then concludes this document.

Chapter 2

Numerical Measures of Cycle Infrastructure Capacity

2.1 Introduction

The first aim of this document (Section 1.3) is to demonstrate that the state-of-the-art of the quantitative cycle infrastructure capacity literature is insufficient to meet practitioner needs. As was also noted in the previous chapter, the decision to provide (or not provide) a given form of cycle infrastructure is often a political battle more than one for the engineer. However, the engineer requires a solid underpinning of engineering and scientific principles. The following therefore identifies that literature state-of-the-art and highlights the limitations. This will then provide a basis for going forward (in the chapters to follow) to improve the theoretical foundation of cycle infrastructure delivery.

2.2 Summary of Literature

Few attempts have been made in the literature to establish numerical capacity for cycle infrastructure. Therefore, Table 2.1 presents the result of a review by this author of both practitioner and academic literature with particular regard to the explicit specification of the numerical flow capacity of unidirectional cycle infrastructure. The data is presented in chronological order of publication and grouped as either an academic or practitioner source.

Flow rates are often expressed in a range of units depending on the prevailing specification for cycle infrastructure dimensions. As a result to enable intercomparison, here and throughout this document, bicycle capacity measures are converted to a bicycles per metre width per hour measure which normalises these differing sources to a common flow rate baseline unit.

Table 2.1: Literature capacity measures of unidirectional cycle-only infrastructure

Source	Country of Origin	Capacity (bicycles per metre per hour)
Academic Literature		
De Leuw, Cather & Company (1972)	USA	1000
Homburger (1976)	USA	2598
Liu et al. (1993)	China	1836–2088
Navin (1994)	Canada	4000
Botma (1995)	USA	650
Jin et al. (2015)	Various (majority China)	1250–2600
Z. Li et al. (2015)	China	3300–3375
Practitioner Literature		
CROW (2007)	Netherlands	2611–3300
Vejdirektoratet (2012)	Denmark	1000–1500
Transport for London (2014a)	UK	133–363
Department for Transport (2016)	UK	60–187.5

2.3 Academic Literature Review

2.3.1 Cycling in Davis, California

Davis, CA, has long been a leader in the United States for the implementation of cycle infrastructure. Davis is a city of approximately 64000 residents (2005 data) and is home to a University of California campus with an enrol-

ment of approximately 30000, of which approximately a third are city residents and 48% (2002 data) cycle to the university (City of Davis Bicycle Advisory Commission, 2009).

In the face of rapid city growth through the 1960s, in 1972, De Leuw, Cather & Company were commissioned by the City of Davis to produce a "Bicycle Circulation and Safety Study" to inform the ongoing delivery of the city's cycling network. In the absence of US design guidance, the authors drew together a range of design guidance sources from Europe (specifically the Netherlands, West Germany, Switzerland, Sweden and Norway) and interpolated to produce the graph replicated in Figure 2.1.

The manner, weighting and methodology for interpolation is not specified and the authors note that the original sources are in wide disagreement, however, a broadly linear capacity increase with width capacity is indicated (for reasonable ranges) of approximately 1000 one-way bicycles per metre width per hour. No evidence for a linear relationship is indicated and the authors note "relatively limited pathway widths afford capacity in excess of that which might be required even in the most heavily traveled areas. This basic maneuvering space, level of service and other considerations will determine bikeway width rather than capacity." In other words, they assert capacity is unlikely to be a major factor in a determination of bikeway width.

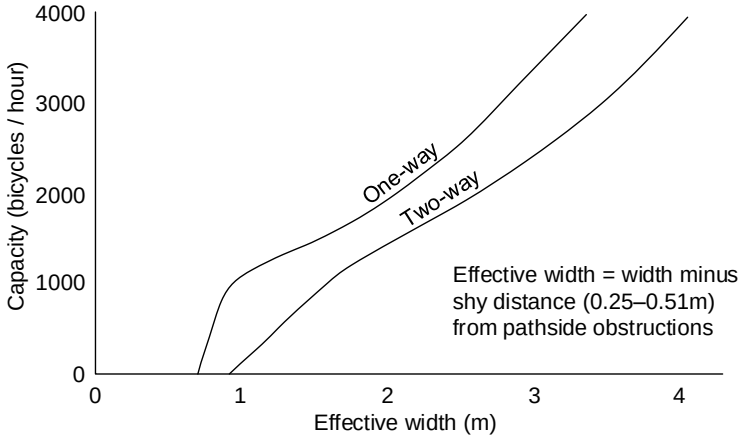

Figure 2.1: Cycle path width-capacity relationship taken from De Leuw, Cather & Company (1972)

A few years later, Homburger (1976) noted that "Bicycle flow patterns are governed by the same basic relationships between flow rate, density and speed which are characteristic of the flow of motor vehicles or pedestrians." but notes that "Capacity and jam density conditions occur seldom, if ever." Despite this, the document also presents an adaptation of (preliminary) empirical data (Miller, 1976) collected in Davis which uses flow density and speed to define levels of service.

'Levels of Service' (LoS; Transportation Research Board, 2010), are qualitative gradings for the operation and user quality of service of a given piece of highway infrastructure. They range from LoS A which represents free-flow highway (with a low traffic density) and is the 'best' grade, to LoS F which represents a standstill traffic jam (and is the worst grade). Peak flow is generally achieved at a LoS of C to D.

The values from Miller (1976) are (part-)replicated in Table 2.2 and, assuming that flow breakdown at the LoS E/F boundary is representative of capacity, indicate capacity of almost 2600bpmph, some 2.6 times that identified by De Leuw, Cather & Company (1972).

Table 2.2: LoS for bicycle one-way flow (adapted from Homburger, 1976)

LoS	Flow (bpmph)	Density (bikes.m^{-2})	Speed (kmh^{-1})
A	864	< 0.050	≥ 17.7
B	1302	< 0.075	< 17.7
C	1968	< 0.125	< 16.9
D	2340	< 0.180	< 15.3
E	2598	< 0.270	< 12.9
F	Variable	> 0.270	< 9.7

2.3.2 An Off-Road Experimental Study

Conducted at University of British Columbia in Vancouver, Canada, Navin (1994) instructed a group of students under instruction to ride around a large radius loop "as fast as possible". Maximum flow was found at 14kmh^{-1} and a flow-density curve was derived and is replicated in Figure 2.2. The graph demonstrates empirical resemblance to theoretical traffic flow theory relationships (e.g. Figure 3.1, and discussed further in Chapter 3). Attention should also be drawn to a lower speed bound to the curve relating to the ability to maintain balance on the bicycle at 3.3kmh^{-1}.

Navin (1994), being empirically derived, makes consideration of the variable Level of Service which exists between the isolated cyclist and a crowded one. However, a maximum capacity is derived based essentially on extrapolated close-contact 'packing' of the cyclists. Whilst this could represent a 'jamming density' in perhaps the absolute sense, real cyclists would be unlikely to arrange themselves in such a dense manner, even in times of complete flow breakdown. Therefore, the capacity derived – 4000bpmph – is perhaps unrealistically high. Also of note is the lower speed bound determined from the study; such a bound is not found elsewhere in the literature.

2.3.3 Chinese Sources

Bicycles are a common transportation mode in China and as such, a number of academic papers provide data on cycleway capacity.

Liu et al. (1993) presented a synthesis of a 1986 "traffic study of 178 intersections/interchanges in Beijing". Capacities are quoted of 1836bpmph for an one-way facility which is unsegregated from traffic, and 2088bpmph for a segregated facility. Of relevance to note, the survey consisted of a large number of partially (or entirely) grade-separated interchanges and, given that even in the intervening sub-period 1995–2005, the number of motor vehicles in Beijing increased by three-times from 0.89–2.58 million vehicles (Yao et al., 2009), a similar survey completed more recently may well yield different results. One should also note the recent increases in electric-assisted bicycles (see also Jin et al., 2015, below) which may also have affected utilisation of the given facilities should a more recent version of this study be completed. Electrically-assisted bicycles are not considered further within this thesis as the purpose is to seek to establish the state-of-the-art with regard to ordinary bicycles, and whilst of perhaps of general interest, equivalence to ordinary bicycles would only stand if there were sufficiently robust data pertaining to them.

Jin et al. (2015) synthesised a number of sources to produce a table of estimated cycleway capacities. Table 2.3 lists a number of (mostly Chinese language) sources presented therein and not otherwise broken out in this document. Those sources range from 1250–2549bpmph. The remainder of that paper focused upon relationship fitting to empirical data to derive a bicycle equivalent unit for electric-assisted bicycles.

Table 2.3: Selected literature cycleway capacities (adapted from Jin et al., 2015)

Source	Flow (bpmph)
National Swedish Road Administration (1977)	1250
Ministry of Housing and Urban-Rural Development of China (2012)	1600–1800 (with physical separation) 1400–1600 (without physical separation)
F. Li (1995)	2000
Wei et al. (1997)	2344 (with physical separation)
Wei et al. (1993)	2549 (with physical separation) 2227 (without physical separation)

Z. Li et al. (2015) presented empirical data collected from video surveys of four sites in Nanjing, China, in order to derive fundamental diagrams for cycle flow and establish capacity of separated one-way bicycle facilities. Capacity values of 3960 bicycles per hour for a 1.2m path and 8100 bicycles per hour for a 2.4m path (3300–3375bpmph, respectively) were concluded. Critical density

was found to be 102 bicycles per km per metre of width, or 0.102 bicycles.m^{-2}. However, contrast this with Table 2.2 where such a density corresponds to only a LoS C rating. In any case, these values represent the top-end of the authors' observations and were measured at a bottleneck location and thus they may not be a sustainable capacity derivation for a simple uninterrupted flow facility.

In summary of these Chinese sources, capacities are found to range between 1250bpmph and 3375bpmph, which whilst a wide margin, are within the extremities of values presented above from US and Canadian sources (1000–4000bpmph).

2.4 Practitioner Literature Review

2.4.1 United States

Botma (1995) considers (from first principles) the issue of bicycles interacting by considering the distribution of speeds and mathematically deriving the probabilities of passing another bicycle over a given distance. This work is derived from some of the historic Dutch standards (CROW, 1993) where a passing rate of $\leq 10\%$ is considered acceptable. Botma considers this intuitively low and derives Levels of Service (LoS) based on various increasing rates of "hindrance" up to a maximum of 100% at the LoS E/F threshold; i.e. the threshold between low speed congested movement and complete standstill. Botma also makes the assumption that bicycles do not impede one another, even at the highest rates of passing. The results of this computation are presented in Table 2.4.

Table 2.4: Analytically-derived LoS for a one-way path of 2.0m width (from Botma, 1995)

LoS	% hindrance	Service volume	Passing frequency (event.s^{-1})
A	0-10	130	$< 1/150$
B	10-20	260	$< 1/75$
C	20-40	520	$< 1/35$
D	40-70	910	$< 1/20$
E	70-100	1300	$< 1/15$
F	100	-	$> 1/15$

Based on these calculations, the threshold from LoS E to F represents a frequency of only one passing in 15 seconds (on average). With an average passing found (in earlier work by the same author) to take 10 seconds (Botma & Papendrecht, 1991), this is representative of being in a practically continual state of passing. However, on a 'two-lane' track, such a situation could represent a substantial impediment to passing flow as no more than one lateral passing can occur at one point at any given time. Both this, and the corollary assumption that the bicycles do not impede one another, is noted as a potential limitation of the technique by Botma, and that it should be tested.

Despite a lack of such testing, the LoS measure in the Highway Capacity Manual (HCM; Transportation Research Board, 2010) is informed by Barker et al. (2008) which aggregates various US-centric research programmes – including Botma, 1995 – into individual measures for bicycle facilities, pedestrian facilities and mixed-mode circumstances. Generally these consist of a weighted combination of street-segment and intersection experiences. Such definitions are calibrated in mixed-traffic circumstances and mainly relate to cycle/motor interactions, explicitly individual factors (such as travel time) or general user perception (see Landis et al., 1997). In none of these cases (save for Botma, 1995 itself), is the interaction of multiple bicycles with each other considered.

In an effort to test the impact of these assumptions, Z. Li et al. (2012) constructed a cellular automata implementation (N.B. Botma, 1995 was analytical) from the same first principles to provide estimations of the number and types of passings between cyclists and how these change with increased flow volumes. This was then calibrated with real data from 9 sites in Nanjing, China. Generally the modelling shows close agreement with reality for unimpeded and simple passing manoeuvres (i.e. no speed change was necessary). However, because it was difficult to observe in the video – both because of the manual process used for data collection from the video, and a lack of obvious bicycle congestion at the sites chosen – the empirical work excluded those passing events where the speed of the passing cyclist needed to change (so called, 'delayed passing'). This is a key component of the behavioural profile of a cyclist and therefore a not inconsequential determinant of capacity and practical level of service.

The validity of the non-interaction assumption is also raised in Hummer et al. (2006) which uses the same principles, combined with field surveys, to develop a practitioner tool for establishing LoS on shared-use facilities. The resulting quantitative model is converted to a LoS measure by use of focus group surveys. The result is an effective uprating of the (pseudo-quantitative) 'capacities' in the 2000 version of the HCM to those subsequently found in the 2010 version of the HCM.

However, Hummer et al. experienced a similar limitation to Z. Li et al. (2012) in that quantifying 'delayed passing' from field data was difficult. In addition, a series of assumptions were incorporated into the analytical model relating to:

1. An assumed respect for opposing traffic such that an overtaking cyclist would not ever use the outermost 'virtual lane' (e.g. to 'double overtake');

2. The majority of the surveys being performed over the winter months, and on a set of paths/trails which were simply 'known to the team' as opposed to being systematically sampled; and

3. For the model calibrations, the highest use paths did not fit their linear model.

Given these limitations, the Levels of Service derived are only likely to be suitable for low use scenarios. Indeed: "most applications of this model would be outside of such an urbanized environment", where "such" refers to the Chicago Lakefront Trail which is repeatedly noted as a high quality example that was extremely busy, and yet was originally excluded from the variable selection (but not the calibration) for the model that resulted.

In other US federal design guidance, the Federal Highway Administration Separated Bike Lane guidance (Federal Highway Administration, 2015) specifies width with a dependence on capacity: for example, "Width considerations include expected bicycle volumes," and "Wider separated bike lanes [...] should be considered where a high volume of bicyclists is expected." Additionally, American Association of State Highway and Transportation Officials (1999) recommends a two-way cycle path width of 3.0m, allowing reduction where (amongst other scenarios) "bicycle traffic is expected to be low". Alternatively, wider tracks are suggested "due to substantial use by bicycles" (amongst other possible reasons). Neither source contains quantification with regard to definitions of "high volume" or equivalents.

These design guidance documents are often codified into design standards (i.e. requirements on the practitioner, rather than recommendations) at the state and local level. For example:

- Seattle Department of Transportation (2007) on-road standards codify state-level standards, which themselves codify American Association of State Highway and Transportation Officials (1999). For example, "[Shared-use paths] ... should be a minimum of 10-feet [3.0m] wide. Minimum width may be reduced to eight feet [2.4m] where physical or right-of-way constraints are severe. Trail widths of 12, 14, and even 16 feet [3.7m, 4.3m, 4.9m, respectively] are appropriate in high-use urban situations."

- City of Los Angeles Department of City Planning (2011) is synthesised "from the Caltrans Highway Design Manual, the California MUTCD [(Manual for Uniform Traffic Control Devices; a document that specifies signage, road markings, signals and other traffic control equipment)], and existing City of Los Angeles design practice." and states that "Wide bicycle lanes are also appropriate in areas with high bicycle use. A bicycle lane width of 6 to 7 feet [1.8–2.1m] makes it possible for bicyclists to pass each other without leaving the lane, increasing the capacity of the bicycle lane."

- City of Davis Bicycle Advisory Commission (2009) states that "The city's guideline width for on-street bike lanes is 8 feet [2.4m] ". Furthermore, this goes so far as to note that "There is a consensus among bicycle planning and safety experts that bike lanes constructed to the Davis guidelines are appropriate."

- Klotz Associates (2007) states that "Depending on the usage and whether the path is one or two-directional, the recommended width varies." and notes that "AASHTO [(American Association of State Highway and Transportation Officials, 1999)] also recommends that wider shoulder widths are used when high bicycle traffic is expected."

- Oregon Department of Transportation (2012) states that "The standard width for bike lanes is 6 feet [1.8m]". It is also stated that "Bike lanes may also be wider than the standard 6 feet [1.8m] in areas of very high use".

- New York State Department of Transportation (2015) codifies AASHTO guidance discussed above and notes that "incidence of [shared-use] conflicts will increase with reduced width"

Despite generally noting that increased use necessitates increased widths, none of these cases contain a quantification as to low/high flows that would inform such provisions. Furthermore, given that a number of the literature sources presented are themselves the result of international review, it is almost certain that if any reasonable quantitative sources did exist then these may have been captured.

2.4.2 Northern European 'Best Practice'

The Netherlands and Denmark have a high rate of cycling. For example, over a third of journeys to work or education are completed by bicycle in Copenhagen (36%; City of Copenhagen, 2013a) and over two thirds of journeys in Amsterdam city centre (68%; Van Mead, 2016). By contrast, only 2% of journeys in London are cycled (Transport for London, 2012). Consequently, these jurisdictions are often considered to be world-leaders.

'Record 25: Design Manual for Bicycle Traffic' (CROW, 2007) – and previously its earlier version 'Sign Up for the Bike' (CROW, 1993) – is the prevailing design standard for the provision of bicycle infrastructure in The Netherlands (with CROW (1998) stipulating some engineering specifics.

In terms of quantifying simple-unidirectional cycle infrastructure, CROW (2007) provides a few capacity indications. Firstly, a 'comfort capacity' is specified of 75–187.5 bicycles per metre per hour (width depending; p.173 therein) for a one-way cycle track on or adjacent to the highway. This range is derived from Table 2.5 which is reproduced from CROW (2007). Elsewhere in that document, an absolute capacity of 2611.1–3300 bicycles per metre per hour is quoted CROW (2007, p. 204); specifically, 3300 bicycles per hour for a path of 1.0m, and 4700 bicycles per hour for a path of 1.8m.

CROW (2007) is widely considered as the 'gold-standard' for design guidance for cycling infrastructure (e.g. Parkin & Rotheram, 2010). Certainly, the document is cohesive, comprehensive and backed with decades of practice from

the Netherlands. However, some aspects of the document are less robust; infrastructure capacity is an example of this.

Flows are stipulated for rush hour peaks, reflecting both the widespread use of cycles as a commuting mode and the need to consider the peak capacity needs of the route. Additionally, different widths are stipulated based on flows, indicating a specific consideration of the capacity need of the route; however, the flow rates beyond which capacity is considered are low.

Two-way flows are specified as low as 20 bicycles per metre per hour for a 2.5m track. Isolated cycle track is explicitly specified elsewhere in CROW (2007, p. 144) with similar capacity constraints to those for a two-way track in Table 2.5 but with widths 0.5m narrower and a requirement for crossable adjacent verge to compensate for the loss in width. No guidance is provided for situations where capacity requirements demand a width of more than 4.0m.

Note too that the values drawn out in the preceding are given as ranges because the capacities of given sections of infrastructure (at least as presented) are implicitly non-linear with increasing width. However this is not quantified or explored in any of the literature.

Juxtaposing the two ranges of values relating to 'comfort' versus 'absolute' capacity, the implication is that the range between the two sets of values is 'uncomfortable' but not necessarily over capacity. If the premise that cycle traffic is likely to exhibit non-linear behaviour (as motor and pedestrian traffic does) is accepted, then this transition zone – which accommodates a wide range of infrastructure – demands consideration as the user experience of infrastructure carrying, say, 2500 bicycles per metre per hour, is likely to be so uncomfortable as to be potentially unusable. This also assumes that such high capacities are themselves realistic, which later in this document, is demonstrated to not be the case (Chapter 3).

As a final point with regard to this source, a one-way peak flow of >750 bicycles per hour is specified for a path of 4.00m width, with flows under 750 bicycles per hour covered by the next lowest width. No guidance is available as to the appropriate width provision for flow rates substantially in excess of that rate.

In Denmark, Vejdirektoratet (2012) provides the design standards in use

Table 2.5: Cycle track adjacent to highway design dimensions (from CROW, 2007, p. 173)

One-way Track		Two-way Track	
Rush hour intensity (one direction; bikes/hr)	Width (m)	Rush hour intensity (two direction; bikes/hr)	Width (m)
0–150	2.00	0–50	2.50
150–750	3.00	50–150	2.50–3.00
>750	4.00	>150	3.50–4.00

and dictate a cycle track should be 2.0m wide at minimum with a 2.2m recommendation providing for 2000 cyclists per hour in an implied 2no 'virtual' lanes. These recommendations date from 1943 road standards and have been retained in broadly that form in subsequent revisions and are also referred to in guidance documents such as Cycling Embassy of Denmark (2012). It is stated that each additional metre of width can provide for an additional 1500 cyclists per hour.

No information is available as to how user comfort or service quality degrades in the approach towards this capacity, nor how sensitive it is to non-homogeneity in the flow or directionality of cycles; although, the standard suggests that capacity is unlikely to be the driving design factor. A minimum of 2.5m for a two-way cycle track and guideline 2.2m for a one-way track (with an absolute minimum of 1.7m) are stipulated from a design perspective but, for a one-way facility, a link to capacity is not drawn. The base recommendation implies 1000 cyclists per hour per metre and a non-linear proportional *increase* in flow with width, as opposed to the reverse in CROW (2007) where increased width provides proportionately diminishing capacity benefits.

In summary then with regard to accepted 'best practice', there is no common agreement as to the capacity of even the simplest bicycle infrastructure (figures range from 1000 to 3300 bicycles per metre width per hour); no agreement as to whether flow increases proportionately with increasing width, provides diminishing capacity increases or provides non-linear increases with width; and no agreement as to minimum widths (either practically or in an absolute sense). Capacity indications, if one assumes them valid, do not extend to the scale of use that is likely to be seen in a very large urban area and little if any indication is given as to how the highway design practitioner could realistically handle cyclists in the order of many thousands per hour on a given route. As discussed in Section 1.2.1, flows of this scale are potentially already seen in cities such as London despite what is usually considered to be a low cycling contribution and immature cycling infrastructure.

2.4.3 United Kingdom

Cycle infrastructure capacity issues are generally absent from UK design standards with national design standards in the 'Design Manual for Roads and Bridges', such as 'TA 90/05: The Geometric Design of Pedestrian, Cycle and Equestrian Routes' (Department for Transport, 2005), and 'TA 91/05: Provision for Non-motorised Users' (The Highways Agency, 2005), containing no consideration of capacity. 'Local Transport Note 2/08: Cycle Infrastructure Design', the prevailing design guideline document, also contains no consideration of capacity; nor do other general guidance documents such as the 'Manual for Streets' (Department for Transport, 2007), or the 'Sustrans Design Manual: Handbook for Cycle-Friendly Design' (Sustrans, 2014). 'Local Transport

Note 1/12: Shared Use Routes for Pedestrians and Cyclists' (Department for Transport, 2012) does note a limited literature with regard to shared-use capacity (ranging 25–180 users per hour per metre width) but noting there is "little consistency", declines to make a recommendation.

In October 2016, an Interim Advice Note (a document short of a full standard such as TA90/05 above, but still carrying the force of such) was published which superseded a number of parts of the national standards (Department for Transport, 2016). Of note is a consideration of capacity which states that "Infrastructure shall provide sufficient capacity to accommodate growth in volumes of cycle traffic"; and elsewhere related to surfacing, "Cycle lanes with widths of more than 2.0m, [are] used for high demand flows". A table replicates the content of Table 2.5 above (and is implicitly sourced from CROW, 2007). Further evidence for the choice of these values is not provided, nor is further qualification of "sufficient capacity" or "high demand".

It should also be noted that whilst it is a requirement on the practitioner to abide by the content of the Interim Advice Note (IAN) for design works on trunk road schemes (i.e. schemes on roads of national importance), the IAN carries only the weight of recommendation for schemes on other local highways, upon which the overwhelming majority of cycle schemes will be delivered.

In London, the current iteration of the London Cycle Design Standards (LCDS) states: "Lanes/tracks should be designed with adaptability and growth in cycling numbers in mind. [...] Indicatively, high cycle flows – over 800 cycles per hour at peak one-way, or 1,000 two-way – will require widths of 2.5 metres one-way or 4.0 metres two-way," (Transport for London, 2014a). Recommended minimum widths in the LCDS are reproduced in Table 2.6 and the associated peak hour flow categories are also reproduced in Table 2.7. It is not stated in the LCDS where these figures were sourced from, but it should be noted that whilst these are not explicitly capacity measures, "failure to meet recommended minima represents a low level of service". Despite shared terminology, the TfL cycle levels of service are not related to the LoS in the HCM and are qualitative in nature.

Table 2.6: Recommended width minima for cycle tracks or segregated lanes (adapted from Transport for London, 2014a)

Flow	One-way width (m)	Two-way width (m)
Very low or Low	1.5	2.0
Medium	2.2	3.0
High or Very high	≥ 2.5	≥ 4.0

Table 2.7: Peak hour flow category definitions (adapted from Transport for London, 2014a)

Flow	One-way flow (bicycles)	Two-way flow (bicycle)
Very low	< 100	< 100
Low	100–200	100–300
Medium	200–800	300–1000
High	800–1200	1000–1500
Very high	≥ 1200	≥ 1500

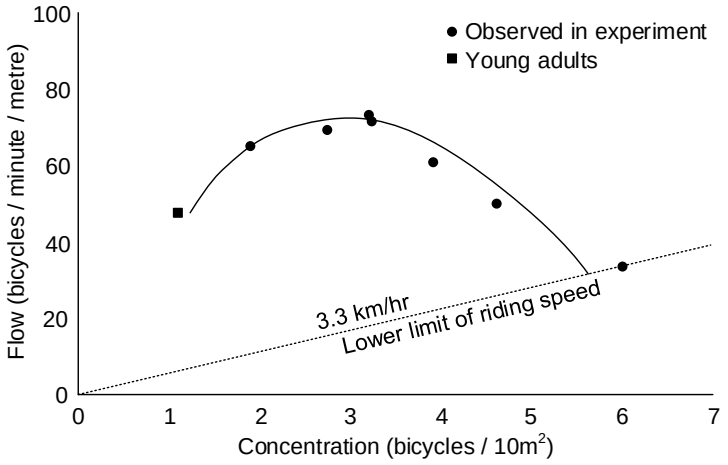

Figure 2.2: Empirically-extrapolated flow/density curve for cycle flow (Navin, 1994)

2.5 Conclusions

The foregoing has explored the state of the literature from both a practitioner and academic perspective, per Aim 1 in Section 1.3. Despite a longstanding desire to identify capacity constraints of cycle infrastructure, the literature remains sparse and in wide disagreement. Basic capacity is defined widely across two orders of magnitude for even the simplest circumstances, such as uni-directional isolated cycle-only links. Given this, attempting to establish capacity constraints of more complex infrastructure arrangements is not possible.

Furthermore, where some literature does identify quantitative aspects of infrastructure capacity, the underlying assumptions do not generally stand up to any realistic scrutiny such as being based on unrealistic expectations of the ability of cyclists to pack close to (or pass) one another without interacting. In the best case, this could lead to over-engineered infrastructure which is simply poor value for money, but more seriously, could reduce the number and scale of schemes by unnecessarily diverting resources to a more limited number of projects, undermining the economic valuations of those facilities; or alternatively resulting in facilities that are potentially inadequate for their demanded use. The latter case also creating safety implications, not simply economic ones.

It is therefore reasonable (and necessary) to return to first-principles traffic theory – which is well-established – to consider if it is possible to establish a more robust estimate for the capacity of cycle infrastructure though the use of traffic modelling tools developed from well-understood first principles.

Chapter 3

Bicycle Microsimulation Modelling

The previous chapter demonstrated the general lack of literature with regard to quantitative design of cycle infrastructure and the point was also made that the quantitative literature that does exist, does not stand up to quantitative interrogation. Consequently, it is prudent to return to the established fundamentals of traffic flow theory, establish the current state-of-the-art of cycle modelling it informs, and thereby examine if the state-of-the-art has simply not filtered through to practice, or indeed is also not fit-for-purpose.

The majority of this chapter is published in the form of a journal paper in the International Journal of Sustainable Transportation (Osowski & Waterson, 2017).

3.1 Traffic Flow Theory

The determination of highway capacity is not simply a matter of establishing the physical spatial capacity of a highway to accommodate vehicles. The interaction of vehicles with one another means that as the density of traffic increases, speeds fall as drivers are less able to choose their speed, change lane or overtake. At the lower extreme, traffic flow is zero as there are no vehicles; at the upper extreme, the 'jamming density' of the highway is reached (i.e. stationary traffic). In both cases, the flow of traffic is nil. As traffic flow increases then, flow rates increase until an intermediate point that beyond that density, flow is sensitive to breakdown (a 'jam') due to stochastic interactions where, even though the resulting traffic density might then be high, the flow rate is now not. Capacity then is usually considered to be the maximum rate of flow of vehicles for a given section of highway, and is generally termed 'saturation flow'.

This complex relationship between density, speed and flow is usually idealised as shown in Figure 3.1, and is commonly referred to as 'the fundamental diagram(s)' of traffic flow theory. In the idealised case, as traffic density

increases, speed decreases linearly as shown in the top left part of Figure 3.1 (N.B. in reality, this is not usually a linear relationship). Given this, and the boundary cases defined above – where flow is zero in an absence of traffic or a complete standstill (the jamming density; D_j) – flow therefore peaks at an intermediate value. This peak flow is termed the saturation flow and is indicated by F_s on Figure 3.1 (top right and bottom). The speed at which the saturation flow is found is therefore termed the 'optimal speed' (S_O) and the density at which that flow occurs, the 'optimal density' (D_O).

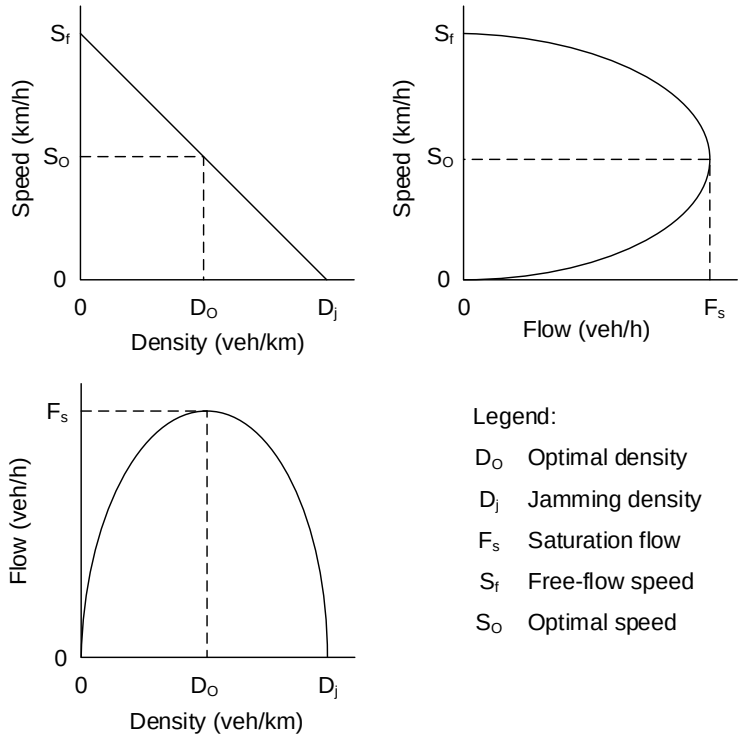

Figure 3.1: Fundamental traffic diagram for an idealised highway (from Transportation Research Board, 2010)

This non-linear behaviour of traffic is an emergent result of the aggregation of the interactions of the many vehicles in the traffic stream. Flow breakdown also displays hysteresis behaviour in that a small change in one direction of a variable causes a state change but to reverse that state change a larger reversal in that variable is required. In the case of traffic, once flow has increased such that breakdown (a jam) occurs, the inflow to that area must reduce substantially before the jam frees up again. Moreover, the stochastic interactions of the vehicles (specifically the lag of the drivers in responding to the situation around them) is enough in of itself to cause the observed hysteresis behaviour (Resnick, 1994). Empirical behaviour widely bears these principles out (e.g. Arasan & Koshy, 2005; Brackstone et al., 2009; Salter & Hounsell, 1996). Indeed road traffic flow is a canonical example of a Complex System.

3.1.1 Pedestrian and Cycle Flow

The principles of traffic flow theory are defined and well understood for motor vehicles, per the above. They are also more widely applicable.

In the late-1960s, the design of pedestrian space was largely informed by arbitrary capacity measures based upon the physical dimensions of the pedestrian. Fruin (1971) was one of the first pieces of work to identify that planning in such way was not an effective manner of producing comfortable, inviting or safe infrastructure:

"[...] many authorities are using maximum capacity ratings for dimensioning pedestrian space. No evaluation or consideration of human convenience has been made in developing these design standards. The flow curves [...demonstrate] that the maximum capacity of a pedestrian traffic scheme is attained only when there is a dense crowding of pedestrians."

Through a programme of empirical surveys, a series of Level of Service (LoS) measures were developed (paralleling those for highway traffic) and results yielding similar fundamental diagrams to highway traffic were observed. Despite the development of potential alternate measures (e.g. Duives et al., 2015), these measures remain in common use today (e.g. Transport for London, 2010c).

Despite not being constrained to lanes or following behaviour, the quantitative analysis of pedestrian spaces was established independently along the same principles as highway traffic was some decades earlier. Mode-specific emergent behaviour – such as as dynamic (or 'virtual') lane formation – is accommodated coherently within these measures (e.g. Fruin, 1971; Helbing & Johansson, 2013a; Hoogendoorn & Daamen, 2005; Still, 2000). Indeed, the operation of pedestrians in continuous space and at substantially lower speeds does not invalidate those interaction principles of vehicular traffic flow; if anything, it reinforces their wider applicability.

However, given the lack of defined lanes and 'rules-of-the-road' in the majority of pedestrian situations, a number of collective behaviours develop which are readily observed in pedestrian situations. The concept of 'virtual lanes' already introduced, extends to more general intersection behaviour – i.e. between two different streams of pedestrians which are not directly opposed – where the phenomenon of 'striping' occurs where the 'stripes' move dynamically through the intersecting flow and "extend sidewards into the direction that is perpendicular to their direction of motion" (Helbing et al., 2005).

Additionally, at bottlenecks such as doorways, flow oscillates between the two directions of flow and 'clogs' – i.e. the outflow is less than the theoretical capacity of the doorway as congestion occurs around and before pedestrians have entered the bottleneck – even in the absence of an opposing flow (Hoogendoorn & Daamen, 2005).

Certain aspects of these pedestrian behaviours are applicable to cyclists. At

the most fundamental level, Zhang et al. (2013) demonstrated that when appropriately scaled, pedestrian, cyclist and motor vehicle fundamental diagrams can be unified over the range of data collected in their study. Given that cyclists display some pedestrian behaviours, and some vehicular behaviours, this is perhaps to be expected.

More dynamic behaviours such as virtual lane formation are also readily observable in cyclists. It is plausible that the more complex behaviour such as striping and bottleneck-oscillation might apply to cyclists, but infrastructure is generally constructed with prioritisation such that occurrences of those behaviours are avoided. Shared-spaces with high cycle flows might plausibly provide such a demonstration, but observations of these will be limited by the rarity of such circumstances and densities. In any case, it is resoanble to assume that there is no inherent reason that these emergent behaviours would not also apply to cyclists.

3.1.2 Passenger Car Units for Cyclists

When dealing with the practical implications of capacity in the delivery of infrastructure, as a practitioner, one would normally be inclined to argue for a conversion of cycle flow to passenger car units (PCU).

Sometimes referred to as 'passenger car equivalent' (PCE), passenger car units provide a mechanism for converting a traffic flow comprised of mixed vehicle types (e.g. where there are heavy vehicles present) to a common baseline measure. As implied by the term, a standard passenger car is defined a 1.0 with other vehicles being assigned a factor relative to that. The exact PCU factors used on a given project are a matter for the practitioner but factors in the order of 2.0 for a bus and 0.2 for a pedal cycle are in common use (Transport for London, 2010b). A larger PCU value indicates a vehicle with an impact on the traffic flow equivalent to that multiple of standard passenger cars.

A pedal cycle is usually assumed (in the UK) as 0.2pcu (Department for Transport, 2006) which has remained unchanged as a figure since its inclusion in Research Report 67, commonly referred to as RR67 (Kimber et al., 1986). However, tracking the value back to RR67 shows that for bicycles, that value was taken from Kimber et al. (1982) and consulting that document shows that the value derived from their study was in fact 0.25. The focus of RR67 was intended upon motor traffic therefore it seems possible that the value was truncated in the process of inclusion as values in RR67 are to 1 decimal place, but those in Kimber et al. (1982) are presented to 2 decimal places. Furthermore, the other value from Kimber et al. (1982) included in RR67 is a PCU value of 0.4 for motorcycles; however, the original value is 0.42. Unfortunately, this minor truncation for the bicycle value has a potential impact of some 25% on the PCU value. In addition, the original value in Kimber et al. (1982) resulted from a study of actual UK roads, which almost certainly means that

bicycles were a marginal component of the flow volume at most if not all the sites (though specific sites are not noted therein). RR67 is widely considered the definitive text for saturation flow values (at least in the UK) and it is possible therefore that the foundation of cycle PCU values in UK practice is based (at least in part) on a typographical truncation.

Further fuel for this view is found in the other comparison made in RR67 to the value of 0.2 found in Webster & Cobbe (1966); also considered a definitive document. That document contains a value of $\frac{1}{5}$ but also notes that the value was not found by that work of the authors but was itself sourced from elsewhere; specifically Holroyd (1963). Consulting that source, it can be seen that it refers (non-specifically) to past work by that organisation indicating PCU values for pedal cycles of "approximately $\frac{1}{3}$", and the results of a study conducted across 9 sites that found PCU values ranging between -0.03 and $+0.25$ for the first 6 seconds of saturation flow (non-weighted average to 0.10), and values between 0.08–0.37 for the remainder of the saturation period thereafter (non-weighted average to 0.21). Aside from the wide range of values including some negative (indicating that the presence of cyclists actually *increases* the capacity of the signals, likely due to filtering in the red period), the first 6 seconds is then disregarded by the authors and the remaining value explicitly rounded down to $\frac{1}{5}$. The potential inclusion of the complexity of the capacity effect of cyclists did not propagate into the literature that followed.

In current practice, Transport for London recommend PCU values equal (though not cited as such) to those in RR67 (Transport for London, 2010b), including the value of 0.2 for pedal cycles. However, the document notes that:

"Where cyclists are present, their volume can have an impact on the calibration and validation of traffic models. As the volume of cyclists change, their impact on traffic behaviour varies in a non-linear manner. It is not appropriate to assign a common PCU value to cyclists where a significant proportion of cyclists and powered two-wheelers are present, as where their volume exceeds approximately 20% of the total volume on any one approach this may have a disproportional effect on modelling results."

No suggestion is made as to what would instead be appropriate, and indeed as is clear from evidence such as Holroyd (1963), potentially a substantive understatement. Homburger (1976) also contains a similar warning as to the need for larger PCU values than the "British method" value of 0.2 "where turning or crossing conflicts are involved". Meanwhile, Jin et al. (2015) summarises a number of south-east Asian sources which provide PCU values in the range 0.20–0.41.

CROW (1998) specifies a PCU value of 0.3 for a cyclist, ranging between 0.4 on a main road and 0.2 on a side road where the flows are "distributed disproportionately". However, no rationale or source for these numbers is presented and it should be noted that a PCU value of 0.4 is *double* that in RR67.

Elsewhere, research relating to heterogeneous traffic flow in India found empirical pedal cycle PCU values of between 0.3 and 2.8, depending on the link type (Tiwari et al., 2000); and back in the UK, Carrignon (2009) used Vissim to develop PCU values between 0.19 and 0.30 depending on (mixed traffic) lane width.

This wide range of bicycle PCU values, and lack of agreement, highlights the lack of widespread quantitative research backing for assigning PCU values to bicycles. Certainly, the literature ranges by an order of magnitude and even around the values clustered in the 0.2–0.3 bracket, the practical significance is a difference of some 50% in capacity. From a practitioner perspective, and almost by definition, PCU values are used where the primary concern of the practitioner relates to motor vehicles. In such cases, and where cycle flows are marginal to begin with, differing outcomes relating to the choice of PCU are likely within the bounds of variability of the data and methods used. However, if cycle flows are a significant component (as they are likely to further become), then this will no longer be something which can be ignored (a point noted by Transport for London, 2010b). At the least, it illustrates that using PCU values for evaluating the quantitative aspects of cycle-only infrastructure is potentially not acceptable without substantial further research on the topic.

3.2 Simulation Modelling

Transport, as a discipline, was a relatively early adopter of modelling and simulation as a means of informing design. As far back as the early 1950s, comprehensive (and numerical) models were being applied in the US (McNally, 2007). Simulation modelling tools have become widely adopted by practitioners and the wide acceptance in the industry of most mainstream transport models (when properly calibrated), means that for the majority of transport projects, a level of modelling is not only desirable, but a necessity to achieving the relevant permissions and funding necessary to deliver a scheme (e.g. Transport for London, 2010c).

A range of highway traffic models exist and are (generally) structurally derived from the basic Four Stage Model (Figure 3.2). The stages are generally iterative with the result ideally achieving convergence, or a stochastic output which can be better understood through Monte Carlo processes. Often, stages are collapsed as required: e.g. a single junction model might condense trip generation and distribution to simple factoring to future flow forecasts from existing traffic flow counts; mode choice would not feature and assignment would simply be existing measured turning proportions. Condensing stages in this manner allows the application of analytical solutions such as multiple linear regression (as used in the roundabout modelling package ARCADY; TRL Software, 2012), or more complex analytical models (such as those used to compute the operation of signals in LinSig; JCT Consultancy, 2011).

In any case, simulation models are usually considered within two broad categories:

Macroscopic models These are models that consider the movement of traffic on a whole-of-flow basis. Parallels can be drawn to fluid dynamics and mathematical network modelling. Commonly used highway traffic software modelling packages that operate on this scale are SATURN (Atkins, 2013) and Visum (PTV, 2013b). Link operations are (non-exhaustively) determined based on input parameters including speed-flow curves (i.e. applying the principles of traffic flow fundamentals, see Section 3.1), numbers of lanes and link lengths. The models then typically iterate through an assignment algorithm until the traffic is distributed over the modelled network in equilibrium. Various other packages can be said to operate at the microscopic level, given they operate on the junction scale yet are essentially an applied form of macroscopic model. Examples include LinSig (for signals; JCT Consultancy, 2011) and ARCADY (for roundabouts TRL Software, 2012), as these are based on time-averaged flow capacities of link components of the junction, rather than the specifics of vehicle-on-vehicle interactions.

Microscopic models In contrast to macroscopic models, microscopic models operate based on the interactions of the individual vehicles with one another and therefore the aggregate traffic flow fundamental behaviours developing from this. Vissim (PTV, 2013a) and Paramics (SIAS, 2000) are commonly used packages in industry. With appropriate calibration the results are similar to macroscopic models, although the computation time required to achieve these results can be substantial. However, the disbenfits of runtime and the scale of input data required is counterbalanced by the ability to interrogate the operation of the model at a high level of detail allowing the assessment of vehicle-vehicle and vehicle-infrastructure interactions, and to present results in a detailed, visual and intuitive way to stakeholders.

There are also a developing category of models referred to as 'mesoscopic' which combine parts of macro and microscopic models such as using microsimulation junction models within a network where traffic on links is treated at the macroscopic level, or treating intra-urban links microscopically and inter-urban links macroscopically. The disadvantage of the loss of some details of the model operation is balanced with the advantages in terms of run-time reduction, parameter complexity and consequently, cost to run the model. Examples of this type of model include the combined use of Vissim and Visum, or SATURN and DRACULA (Institute for Transport Studies, 2007).

In terms of applicability to bicycles, the state of highway modelling highlights the barriers to the direct application of the principles to bicycles. In or-

der to properly run and calibrate a highway model, the model must be calibrated with suitable data and underpinned by a knowledge of the general properties of traffic over a range of parameters; for example, to define the speed-flow curves (see Section 3.1) necessary to establish the macroscopic properties of link flows. Alternatively, the individual interactions and their impact on flow and capacity must be understood to inform a microscopic model. A mesoscopic model, demands both. As established in Chapter 2, such data is generally unavailable with regard to bicycles.

In all of the highway modelling cases, the existence of lanes and defined rules of the road in reality, provide useful constraints on the model which lend some simplification to their design. This is not necessarily the case with pedestrian modelling although the underlying principles are similar. Consequently, pedestrian modelling suffers from some complications when compared to highway traffic modelling:

- Unlike vehicles which exhibit lane-keeping behaviour, pedestrians operate in continuous space where, even if explicit demarcated lanes exist (e.g. a segregated shared-use pedestrian/cycle route), the behaviour of pedestrians is such that compliance is not high enough to be uniformly assumed. Yet, crowds of pedestrians tend to self-organise into dynamic 'virtual lanes' so following/overtaking behaviours may not simply be able to be ignored.

- Pedestrian behaviour is more erratic than vehicles. For example, pedestrians slow or stop to make decisions (even at places other than decision points), travel in groups or ignore one-way systems. These inherent complexities are generally not applicable in highway modelling and are certainly not captured by current models.

- Pedestrian behaviour changes depending on the circumstance. Pedestrians move with different properties in the case of a fire evacuation in comparison with an average weekday in a rail station.

At the macroscopic level, these lower level issues tend to be 'smoothed out' in space and time. For example, PEDROUTE (Halcrow, 2010) effectively combines macroscopic generalised link cost principles with a cellular automaton network to simulate constrained pedestrian environments. The behaviours of the individual pedestrians are not considered.

By contrast, microscopic pedestrian models consider the complexities of the behaviour of individual pedestrians. A range of commercial packages are available but perhaps the widest used is Legion (Legion Ltd., 2013). Legion considers individual pedestrian behaviour, interactions and wayfinding in a three-dimensional environment. Industrially, Legion (and more generally pedestrian microsimulation) is widely considered to be the state-of-the-art.

To return to cycling, there is a need for a proper quantitative understanding of cycle flow. Cyclists (generally) are bound by highway 'rules-of-the-road', yet within those constraints, operate in continuous two-dimensional space with dynamics that emerge from the interactions of the cyclists with one another. It is reasonable then that – with appropriate modifications – some existing tools pertaining to highway traffic, and particularly pedestrian flows, may have transferability to cycles.

3.2.1 Cycle Modelling State-of-the-Art

Cycle modelling at the individual interaction level is required to consider the aspects of cycle behaviour that differ substantively from motor vehicle and pedestrian modelling whilst simultaneously adopting wider best practices for modelling which have developed over many decades. For example, cyclists display different kinematics to both motor vehicles and pedestrians, and cyclists tend to avoid speed changes in preference for trajectory changes so as to minimise their exertion (Twaddle et al., 2014).

3.2.1.1 Existing Modelling Packages

Some limited attempts have been made to accommodate cyclists within existing highway models. Carrignon (2009) used Vissim to establish PCU values for cyclists (Section 3.1.2) based on a new vehicle type being accommodated in the model. Whilst the model modifications accommodated some in-lane filtering, the outputs were limited by some core limitations of the model (such as an inability to incorporate red-light jumping) and by a general validity for only low proportions of cycles in the vehicle flow.

City of Copenhagen (2013b) details a more recent attempt to comprehensively include cyclists in a Vissim model. A set of five different bicycle types were developed and parametrised across eight different situational speed distributions. Also included is more complex behaviour such as red light violation and the basis for the various parameters are generally indicated to be video data measurements and experience. However, the interaction of cyclists with one another remains, at its core, a basic following/overtaking vehicle model, albeit that it can occur within lanes. Explicit results of the implementation are not provided, so it is unclear as to whether the proposals were generally successful or not; although promotional material indicates it was (PTV, n.d.). Further, whilst Vissim provides the capability to incorporate cycles in the most recent version of the Vissim software, calibration and validation are still very much the responsibility of the practitioner and depend on modified motor vehicle models, rather than one developed from first-principles to address the unique properties of cyclists.

3.2.1.2 Cellular Automata

The aforementioned Vissim modifications operate in continuous space within the model of individual lanes. More generally, a limited number of cycle models have been developed in the academic literature that utilise various amounts of spatial discretisation to simplify the arrangements. Broadly, all such models are classified as 'cellular automata' where a bicycle agent occupies a cell (or multiple cells) and navigates the spatial lattice per a defined rule set.

A 'pure' rule-based cellular automata follows global rules applied to all cells. An example would be simple 'if the cell ahead is empty, move forward one cell' rule. Such rules are generally limited by the nature of the cellular lattice (e.g. Cartesian, rectilinear and hexgrid) and are sensitive to its dimensions. This clearly of limited practicality when considering the behaviour of traffic agents and so more complex rules that depend on the situation of the agent and more complex lattices can be applied to achieve more specific ends.

Yao et al. (2009) assumed a Cartesian grid with cell dimensions of 1m×1m. A cycle lane here was assumed to have a width of 3 cells (i.e. 3m) and cyclists were given dimensions of 1×3 cells (and vehicles in the adjacent lane, 3×5 cells). A five step process is iterated to update each bicycle:

1. "Determine the desired velocity"
2. "Judge whether lane changing is needed and whether adjacent lanes could be changed to"
3. "Determine bicycle's velocity of y-coordinate"
4. "Determine bicycle's velocity of x-coordinate"
5. "Update bicycle's position"

The result is outputs that broadly conform to established flow theory with decreased speed against increasing flow and an 'inverted-U' shaped density-speed curve, similar to Figure 3.1. However a lateral arrangement such as this formalises 'virtual lane' behaviour by nature of the grid, rather than it developing in an emergent manner from the rules of the system. Additionally, the rule set essentially formalises a following/overtaking behaviour – which as a cyclist behaviour, is not well-understood, if indeed it is the case at all – and establishes an order of operation of speed change then direction (lane) change, which does not correspond to the observation that cyclists prefer to change direction before slowing down.

Gould & Karner (2010) presented a multiple-lane cellular automata model with cells of 2.1m length and cyclists occupying one cell. Again, a multiple layer speed consideration then lane consideration rule set is applied. The output of the model corresponds well to empirical data in that paper which includes a range of low density cycle flows from Davies (CA, USA) but the empirical data does not extend to the higher densities of the model, which beyond

a peak, indicate a fall in flow with increasing density (i.e. as expected from traffic flow theory).

Vasic & Ruskin (2012) use two parallel one-dimensional 'tracks' of cells for each of motor vehicles and cyclists on a parallel cycle lane. Cells for the motor vehicle lane are defined as twice the length of those in the parallel cycle lane, thus capturing the size difference between the two vehicle types. Each vehicle occupies a single cell and multiple rule sets are also defined to allow junction interactions (i.e. two 'tracks' that share a common cell) and lateral interaction between the motor vehicles and cyclists (i.e. the rules also capture interactions in the laterally-adjacent cell). Whilst the model indicates promising results for the circumstance modelled, it is incapable of capturing the true complexity of cycle behaviour. Cyclists cannot overtake one another, nor can they mix directly with the motor traffic, as might be expected if their contributory proportion of flow were high.

3.2.1.3 Continuous Space Models

As a wider alternative to discretisation, continuous space models allow movement of non-arbitrary length across the modelled space. Generally, such models are agent-based in that processing such a model involves iterating through all the agents in the model space, and updating each based on the agent's rules which may be (but need not be) common to many agents. Such models are obviously closer to how real cyclists behave (if properly defined), but are more computationally complex and may require more parametrisation to operate.

Venkatesan et al. (2008) developed a continuous space model where heterogeneous (motor vehicle and cyclist) agents move forward along a 'mid-block' (i.e. isolated) link by standard kinematic equations of motion unless they approach the rear of another agent, where they then pass on the side of the forward vehicle to which they are nearer, if there is sufficient space between that vehicle and the boundary. If there is insufficient space on either side, then explicit following behaviour is implemented.

The result of the model implementation yields a speed-flow curve generally in keeping with traffic flow theory (Figure 3.1), but is not explored in detail further in the paper as the focus of the paper is the implementation architecture. Gowri et al. (2009) expands this model further to include signal junction operation and more diversity in vehicle types; although the effect on the validity of the model from these expansions is unclear. What is clear is that even a simple movement/passing model operating in continuous space is capable of outputs in keeping with traffic flow theory.

Outside that mentioned, the other form of continuous space model that can be found in the literature with regard to cycling is the Social Force Model (SFM). Originally constructed by Helbing & Molnár (1995), the Social Force Model (SFM) has become an established tool for the modelling of pedestrian

movements. The two major commercial pedestrian microsimulation packages in wide use, PTV Viswalk (PTV, 2013c, part of the Vissim package) and Legion (Legion Ltd., 2013) both depend on the principles established by Helbing & Molnár (Viswalk directly, and Legion by way of distillation in Still, 2000). Their wide use in industry has established the validity of the core models when, as with any model, properly calibrated. Given the demonstrated value of the SFM, it has seen some application to bicycles in the literature as, amongst other benefits, there exists a potential to unify models across modes if the SFM can be used as model common to multiple modes (e.g. Anvari et al., 2015; Pascucci et al., 2015).

The SFM (as defined in Helbing & Molnár, 1995) is based upon a vectoral 'force-like' quantity named the social force. As opposed a physical force, the social force is a metaphorical force which exists to quantify the "motivation to act". The result, as is observable, is that pedestrians behave as if they were subject to an invisible force; namely, the social force. Social force is 'generated' by the agents in the model, and by boundaries. The result (generally) is that agents seek to avoid boundaries and one another, whilst maximising their speed towards their destination. Other aspects such as the ability of a force to act as a 'push' to the agents, allows the modelling of crowd surges and attractive forces allow the modelling of agent grouping.

M. Li et al. (2011) used an implementation of the SFM for cars and cyclists in an adjacent cycle lane. Line demarcation is treated as boundaries in the model so, for example, cars stay in lane as a function of the repulsive force generated by the edges of the lane. The model allows cyclists to move out of the cycle lane (i.e. the boundary road marking is less repulsive than the kerbside) and to mix with the motor traffic should the density of bicycles in the lane be sufficiently high. However the model is not calibrated with empirical data and the focus of the models run is on the resulting behaviour of the motor vehicles (whether or not they are forced to suddenly brake), rather than that of the bicycles. In particular, the interaction of bicycles with one another is only of limited concern in the model and aspects such as the likely propensity of real cyclists to move out of the cycle lane and mix with motor traffic (rather than slow down for example), are not explored.

Liang et al. (2012) created a bicycle-only 'psychological-physical force model' which was based upon the SFM, and then augmented with: an additional "physical force model" to simulate contact and sliding friction between cyclists; and a "trajectories choice model" which chooses a steering direction to either side so as to minimise the cyclist's exposure to force density. The simulation data is shown as agreeing well with experimental data, although the source of the experimental data is not defined. Furthermore, the physical force model is included to simulate contact, but bicycles in contact are perhaps not realistic in that they would essentially be crashing.

Finally, Schönauer et al. (2012) used the SFM to simulate a 'shared-space' where cars, bicycles and pedestrians mix. Bicycles were treated the same way as cars (albeit with different parameters) in their model in that they navigate a 'guiding field' overlay which augments the social field that emerges from the agents. Further to this, conflict resolution (between all modes) is captured at a tactical level, not through the SFM (which it is suggested only fits well for slow movement, such as that of the pedestrians), but through a non-cooperative game theoretic interaction which considers conflicting agent strategies of zero change, accelerating, decelerating or 'dodging' either left or right. The focus of the model upon this game theoretic interaction and aside from showing a good agreement between chosen paths in the simulation and in reality (as might be expected given the 'guiding field' exists to ensure so), also appears to capture bicycles only incidentally.

3.2.1.4 Summary

A common thread across the state-of-the-art presented above is the preliminary nature of the work. Cycle modelling is in its earliest stages academically and has not filtered into general practitioner activities where wide and open use allows the refinement and validation of the tools and principles necessary.

Attempts at incorporating bicycles in existing motor vehicle modelling packages suffer from the limitations inherent in those models, such as lane-based movement and existing vehicle interaction models, and as such do not capture the nuance of how bicycles behave where that behaviour is distinct from motor vehicles. Modifications to incorporate such differences may demand modifications to the simulations at a low (and proprietary) level, and thus may remain a matter for those developers for the foreseeable future.

The foregoing also detailed a number of cellular automata-based models from the literature. These cover scales from a bicycle being defined equal to the size of a cell, to those where a bicycle occupies a number of cells and thus the discretisation of space is less granular. In a number of cases, there is clearly some merit, given they appear to replicate known fundamental traffic behaviour. However in all cases, those models are at the earliest stages of development, are backed with limited (if any) empirical data, and in some cases make assumptions which do not conform to known bicycle behaviour (e.g. changing speed before changing direction). Furthermore, it is not clear whether behaviours such as virtual lane formation are incorporated as a function of the 'gridded' nature of these models, rather than emerging from the simple interaction of the agents, as is the case in reality.

Continuous-space models, generally represented in the literature by variations of the SFM, also display potential merits. They can replicate aspects such as virtual lane formation in an emergent manner and as a category are already widely validated and used for pedestrian models, and in the early stages of ex-

ploration relating to motor vehicle and the sharing of space. In particular, the shared-space aspects are of great interest as they provide the opportunity to develop models that unify all modes in a common model, as opposed to a consolidation of separate but connected ones. Unfortunately, as is the case with all the literature highlighted, the potential for these types of model are not yet well explored.

As evidenced by the literature noted, there is potential merit in many of these methods. However, going forward, focus in this document will be on a microscopic continuous-space agent-based model. These models have already displayed the early stages of validity pertaining to cyclists and of the basis for the future incorporation of other modes – something which is arguably of greater importance for cyclists, than for the majority of circumstances for other modes such as pedestrians and motor vehicles. Furthermore, they capture emergent behaviour such as virtual lane formation as a function purely of the interaction of the agents, and not one of the hard-coded environment they are navigating.

The SFM, upon which the majority of the continuous-space literature above is based, provides a useful and established basis for the formalisation of a basic continuous-space model – in that case as pertaining to the movement of pedestrians. It captures the basic behaviours of pedestrians such as virtual lane formation and striping as functions of the simple interactions of the agents in the simulation (e.g. Helbing et al., 2005; Helbing & Molnár, 1995; Helbing et al., 2001). Importantly, the SFM is simply a mathematical formalisation of interactions in continuous-space that are framed narratively in the concept of a 'social force'. Alternative formalisations of the core concepts, such as charged-particle movement in a field or multiple-body gravity interactions, are mathematically similar, if not narratively so.

Given: that the basic SFM is well-understood and defined; that the use of the core SFM as a starting point for a continuous-space model does not demand an acceptance of the concept of social force (although the metaphor is likely to be applicable to cyclists); and that the concept of the SFM is merely a formalisation of general continuous space interactions; the following will use the SFM as a starting basis for the development of a bicycle model.

3.2.2 The Social Force Model

The following describes the formulation of the SFM per Helbing & Molnár (1995).

3.2.2.1 Social Force

The primary component of the motion of a pedestrian is their desired velocity. As a vector quantity, velocity requires both a magnitude (speed) and direction.

Direction is defined through the pedestrian seeking the shortest route to their destination (or next waypoint). Mathematically:

$$\vec{e}_\alpha(t) := \frac{\vec{r}_\alpha^{dest} - \vec{r}_\alpha(t)}{\|\vec{r}_\alpha^{dest} - \vec{r}_\alpha(t)\|} \tag{3.1}$$

where $\vec{r}_\alpha(t)$ is the position of pedestrian α at time t (N.B. All positions are defined as vectors in Euclidean space), $\vec{r}_\alpha^{dest}$ is the position of the destination (or waypoint if part of a larger route) of pedestrian α and therefore $\vec{e}_\alpha(t)$ is the desired direction of pedestrian α to its desired destination (note that this is a unit vector).

In the absence of any social force, the pedestrian travels in this direction with speed v_α^0. The desired velocity at time t is therefore defined as:

$$\vec{v}_\alpha^0(t) := v_\alpha^0 \vec{e}_\alpha(t) \tag{3.2}$$

If the pedestrian has deviated from their desired velocity then they seek to return to it subject to a relaxation time denoted by τ_α. The pedestrian's current velocity at any given time is therefore subject to a desire force ($\vec{F}_\alpha^{dest}(\vec{v}_\alpha, v_\alpha^0 \vec{e}_\alpha)$) defined as:

$$\vec{F}_\alpha^{dest}(\vec{v}_\alpha, v_\alpha^0 \vec{e}_\alpha) := \frac{1}{\tau_\alpha}(v_\alpha^0 \vec{e}_\alpha - \vec{v}_\alpha) \tag{3.3}$$

where $\vec{v}_\alpha$ is the actual velocity.

In addition to the desire force, the pedestrian is subject to repulsive forces from other pedestrians. The repulsive force is proposed as a monotonic decreasing function which applies in a direction perpendicular to the isopotentials. The isopotentials are proposed to have the form of an ellipse which is directed in the direction of motion of the other pedestrian β with the foci being pedestrian β's current position and future position after a time period Δt based on current speed v_β and direction $\vec{e}_\beta$. The semi-minor axis (b) of that ellipse is therefore defined as follows:

$$b := \frac{\sqrt{(\|\vec{r}_{\alpha\beta}\| + \|\vec{r}_{\alpha\beta} + v_\beta \Delta t \vec{e}_\beta\|)^2 - (v_\beta \Delta t)^2}}{2} \tag{3.4}$$

where $\vec{r}_{\alpha\beta} := \vec{r}_\alpha - \vec{r}_\beta$.

The spatial dimension of the elliptical isopotential is related to a numerical quantity by way of an exponentially repulsive potential defined:

$$V_{\alpha\beta}(b) = V_{\alpha\beta}^0 \exp^{-b/\sigma} \tag{3.5}$$

where $V_{\alpha\beta}^0$ is a measure of intensity and σ a measure of spread. The quantity is then inverted for repulsion and multiplied by the local maximum gradient so that the direction of application applies perpendicular to the isopotential. These operations are summarised mathematically as:

$$\vec{f}_{\alpha\beta}(\vec{r}_{\alpha\beta}) := -\nabla_{\vec{r}_{\alpha\beta}} V_{\alpha\beta}\big(b(\vec{r}_{\alpha\beta})\big) \tag{3.6}$$

Helbing & Molnár (1995) include the possibility of similarly-defined attractive pedestrian forces (e.g. due to groups) but, for brevity, these are not reproduced here.

The final consideration for the repulsive pedestrian force is the ability of pedestrian α to perceive pedestrian β. Perception is considered to be constrained to the field of view ($\pm\varphi$) in the direction of desired motion (i.e. the pedestrian always faces its destination). If outside this field of view, then the force is simply multiplied by a factor c (where $0 < c < 1$) yielding a force:

$$\begin{gathered} \vec{F}_{\alpha\beta}(\vec{e}_\alpha, \vec{r}_\alpha - \vec{r}_\beta) := \vec{f}_{\alpha\beta}(\vec{r}_{\alpha\beta}) \\ \text{or} \\ \vec{F}_{\alpha\beta}(\vec{e}_\alpha, \vec{r}_\alpha - \vec{r}_\beta) := c\vec{f}_{\alpha\beta}(\vec{r}_{\alpha\beta}) \end{gathered} \tag{3.7}$$

as appropriate, and iteratively, for each other pedestrian β.

Finally, consideration must be given to boundary effects. In simple cases, these are imparted by walls but are translatable to fixed obstacles, path edges or other obstructions. The principles of their force generation (in Helbing & Molnár, 1995) are similar to pedestrians (i.e. monotonically decreasing with distance), save that they apply from a point on the border nearest to the pedestrian α. For a given boundary B, the repulsive force is quantified as:

$$\vec{F}_{\alpha B}(\vec{r}_{\alpha B}) := -\nabla_{\vec{r}_{\alpha B}} U_{\alpha B}(\|\vec{r}_{\alpha B}\|) \tag{3.8}$$

with the function $U_{\alpha B}(\|\vec{r}_{\alpha B}\|)$ defined (by Helbing & Molnár, 1995) as:

$$U_{\alpha B}(\|\vec{r}_{\alpha B}\|) = U^0_{\alpha B} \exp^{-\|\vec{r}_{\alpha B}\|/R} \tag{3.9}$$

where again $U^0_{\alpha B}$ is a measure of intensity and R a measure of spread.

The total social force ($\vec{F}_\alpha(t)$) experienced by the pedestrian α is therefore the sum of the attraction/desire to move to the destination/waypoint, pedestrian forces, and boundary forces, or, mathematically:

$$\vec{F}_\alpha(t) := \vec{F}^{dest}_\alpha(\vec{v}_\alpha, v^{dest}_\alpha \vec{e}_\alpha) + \sum_\beta \vec{F}_{\alpha\beta}(\vec{e}_\alpha, \vec{r}_\alpha - \vec{r}_\beta) + \sum_B \vec{F}_{\alpha B}(\vec{r}_{\alpha B}) \tag{3.10}$$

For unit mass, this force equates to the acceleration experienced by the pedestrian α (i.e. Newton's Second Law) with the possibility for the addition of random fluctuations to avoid deadlock situations or to approximate errors made by pedestrians in reality:

$$\frac{d\vec{w}_\alpha}{dt} := \frac{\vec{F}_\alpha(t)}{m} + \text{fluctuations} \tag{3.11}$$

where m is the pedestrian's mass (and which can be excluded if all pedestrians share an assumed equivalent unit mass and as indeed the formula is defined in Helbing & Molnár, 1995) and $\vec{w}_\alpha$ is a preferred velocity which is then finally adjusted by the constraint of the maximum speed possible for the pedestrian ($\vec{v}_\alpha^{max}$).

This acceleration (Equation 3.11) is applied to the pedestrian's current velocity with a final step to constrain the new velocity such that if it exceeds the maximum speed the pedestrian is capable of, then it is reduced to that speed

$$\vec{v}_\alpha(t) := \vec{w}_\alpha \times g(\frac{\vec{v}_\alpha^{max}}{|\vec{w}_\alpha|}) \tag{3.12}$$

and where:

$$g(\frac{\vec{v}_\alpha^{max}}{|\vec{w}_\alpha|}) := 1 \text{ if } |\vec{w}_\alpha| \leq \vec{v}_\alpha^{max} \tag{3.13}$$

or:

$$g(\frac{\vec{v}_\alpha^{max}}{|\vec{w}_\alpha|}) := \frac{\vec{v}_\alpha^{max}}{|\vec{w}_\alpha|} \text{ otherwise.} \tag{3.14}$$

This process is repeated for every pedestrian and repeated at every time step. The result is, as discussed, a simulation of pedestrians moving in multi-dimensional continuous-space which replicates the emergent properties of large numbers of pedestrians such as virtual lane formation, bottleneck-oscillations and striping. More generally the model formalises the concept of point objects in motion through Euclidean space interacting with one another, and the constraints of that space.

3.2.2.2 Refinement and Applicability to Cyclists

The above details the original implementation of the Social Force Model as defined in Helbing & Molnár (1995). Various refinements to the core SFM have been made in subsequent years to address certain issues with the model as it pertains to pedestrians:

- Expansion of the pedestrian from a 'point object' to include a radius, 'compressibility' and friction forces: Augmentation of the core formulae for force generation allows crowding and crush behaviours to be more closely replicated (Helbing et al., 2000);
- Addition of a lambda parameter and replacement of the zonal 'view cone' with respect to force effects: Later versions of the model use a function which is continuous with angle and provides a closer approximation to real pedestrian motion (Helbing & Johansson, 2013a).

The version of the SFM presented above is one which does not include these later refinements as they are specifically targeted at improving the applicability of the SFM to pedestrian circumstances. When considering cyclists, the same refinements may not be logical. For example, cyclist contact, friction and compressibility is (by definition) a crash, and would not be a reasonable replication to seek to make in a general-use cycle model.

The following therefore takes the original version of the SFM per the above and Helbing & Molnár (1995) as a basis, and makes necessary adjustments to ensure transferability to cyclists. The core principles of the model (e.g. continuous two-dimensional individuals moving according to Newtonian physics) are leveraged and augmented where necessary to ensure applicability. The principles, their acceptance or rejection, and what changes are (or are not) made, are detailed in Table 3.1.

Table 3.1: Comparison, acceptance and rejection of underlying properties of the pedestrian SFM

Principle	Acceptance	Rationale
Agents move through, and due to, a socially generated force field	Accept	Principle is transferable to cyclists if additional kinematic constraints are addressed.
Force field is vectoral	Reject	Vector forces result in potential 'pushing' force. Unlikely to be experienced by cyclists in usual operation. See Section 3.3.1.
Agents are attracted to a waypoint/destination	Partial accept	Agents in the model will continue in their direction of travel until perturbed otherwise. This does not necessitate a destination, though is not incompatible with one.
Agent speeds 'relax' to their desired velocity.	Reject	Relaxation implies unbounded acceleration/deceleration. The time frame for speed changes is non-trivial for cyclists. Acceleration is applied based on standard kinematics. See Section 3.3.4.
Social force spatial profiles are 'near-future' distributions of possible location.	Reject	Cyclist speeds and thus time considerations are longer. If accepted then spatial force distributions would be too diffuse to be useful. Future locations are therefore computed by the agents. See Sections 3.3.1 and 3.3.4.
Agents are compressible	Reject	Cyclists in contact would likely have crashed. Proximity of cyclists therefore exhibits a step-change. See Section 3.3.1.
Agents can be attracted to one another (e.g. family/friend groups)	Accept	Equally valid for cyclists as for pedestrians.
Agent perception is variable with angle of view	Partial accept	Concept in original SFM is variable with angle but always non-zero. Rearward visibility for cyclists is much more difficult. See Section 3.3.3.
Agent speeds vary continuously down to zero	Reject	Below a certain speed, maintaining balance on a bicycle is more difficult, imposing a de facto lower cap on speed with a step down to zero.
Agent speeds are bound at a maximum	Accept	Given the lack of pushing forces noted above, this is accepted and is equal to the cyclist's desired speed. See Section 3.3.4.

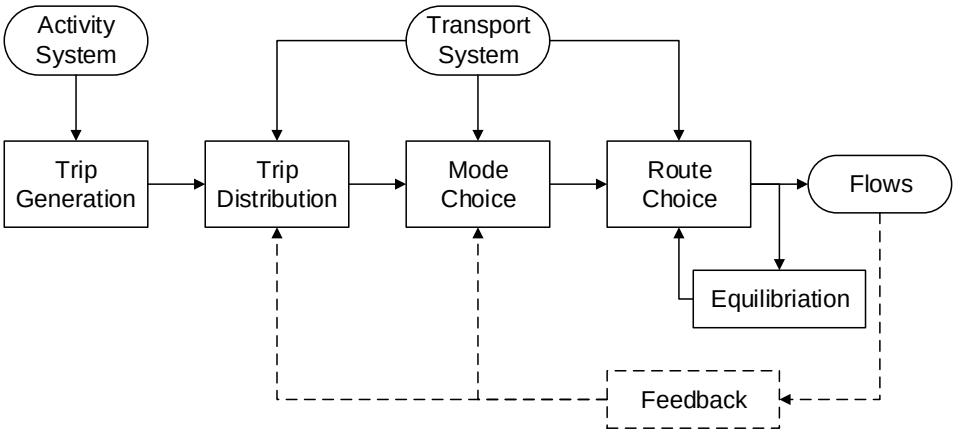

Figure 3.2: General structure for a Four Stage Transport Model (from McNally, 2007)

3.3 Model Development

The model presented here adopts the core principles of the SFM (as proposed originally by Helbing & Molnár, 1995) whilst simultaneously making the necessary modifications to ensure applicability to bicycles. Given there are modifications being applied, a number of different models could have been chosen as a base. The SFM was chosen in the absence of any current accepted model for cycle behaviour and because it is a good candidate for replicating the behaviours observable in real cycle flows. That said, the contribution of this chapter is the principle of the systematic modelling of cycle flows at all, not simply the methodology.

The SFM operates based on agents navigating a force-field like quantity generated by the presence of the agents (the 'social field', Lewin, 1951; and earlier Gibson & Crooks, 1938). The force represents a metaphorical "motivation to act"; for example, to move toward your destination, to maintain your optimum velocity, to avoid other agents, to stay with your group, or to avoid fixed objects in the environment.

As noted in Chapter 2, a literature backing for the fundamentals of cyclist interactions simply does not exist. The following section therefore discusses and presents the theoretical basis upon which modifications to the core SFM have been made (also see Table 3.1).

3.3.1 Bicycle Force Generation

The spatial arrangement of the force generated by pedestrians is assumed in the SFM to take the form of an ellipse projected in the direction of travel with foci separation equivalent to one walking step length. Pedestrians themselves are considered point objects and make decisions based on an instantaneous perception of the current Social Force field which notionally reflects the probable future arrangement in the spatial distribution. However, changing speed or direction is essentially trivial for the pedestrians; i.e. acceleration is unconstrained. For sufficiently long time steps (in the order of a second or more) this may be essentially true, but if time steps are small (especially if substantially <1s) then this starts to become less realistic.

For bicycles, the comparable definition of 'sufficiently long' is likely to be longer than that for pedestrians (given higher speeds) and thus longer than is reasonable for the purposes of modelling the interactions. This time frame is limited (in the absolute) by the ability of the rider to deliver power but, given this time length is proportionally long compared to the time frame of interactions, the constraint of the rate of speed change (i.e. acceleration) must be considered for bicycles.

Furthermore, given the higher speeds of cyclists (relative to pedestrians), it is necessary to separate the future situation from instantaneous perception as

cyclists need to realistically consider a greater time into the future than is the case for pedestrians. Consequently in this model, bicycles do not consider their exposure to forces at the point which they currently occupy, but consider the forces at a range of projected potential future locations based on extrapolation of where other bicycles will be at the relevant time (given an assumed continuation of other bicycles' current behaviour until that time in the future), and the kinematic constraints of the considering agent's current speed and possible range of turn prior to that future point. Future force returns at those potential locations are discounted with increasing time to account for the uncertainty in future decisions of the other bicycles; essentially producing a net present value for the force on each considered path choice (in essence, a 'net present force'; see Section 3.3.4). Force returns include whichever repulsive and attractive forces are present in the given simulation, although for simplicity all attractive forces are set equal to zero (i.e. there is no grouping of cyclists or waypointing) in the demonstration that follows in Section 3.5 onwards. Cyclists navigate the force field, seeking vectors that minimise their overall exposure to repulsive force as calculated in each time step.

Simply modifying the SFM's spatial force projections, which embody the uncertainty about the next few steps of the pedestrian, to fit the above parameters (i.e. projecting through an extended time step), would not produce acceptable outcomes owing to the extensive and diffuse spatial distribution of the projected force that would result given the speed, extended time consideration and potential range of motion of the cyclist.

Bicycle force generation is proposed to be generated with an exponential distribution akin to the SFM but with an additional overlay reflecting the cyclist envelope. Pedestrians in the SFM are treated as point objects (Helbing & Molnár, 1995); by contrast, bicycles cannot compress one another (Helbing et al., 2005) nor rub against one another (Helbing et al., 2000) and simultaneously maintain the ability to cycle (compared to pedestrians which to some extent, are 'compressible'). Physical contact between cycles is almost certainly, by definition, a crash and is not recoverable in trivial time. The equivalent in a pedestrian model would be pedestrians physically falling over one another.

Clearly, such a mode of operation is not appropriate for 'usual' operation and thus is not proposed to be transferred to the bicycle model here. As an alternative, it is proposed to overlay a force equal to the bicycle maximum; i.e. that which would be calculated at a 0.0m from the location of the bicycle) over the spatial area of the bicycle (i.e. the rectangular area in Figure 3.3). For simplicity, this is approximated to a rectangle with dimensions equating to a bicycle. Such a dimension incorporates the rider, clearance for steering and potential panniers etc. For the remaining bicycle force, elliptical spatial force profiles are proposed centred on the cyclist (and orientated longitudinally; Figure 3.3b) reflecting the inappropriateness of a circular approximation for a bicycle/cyclist

combination (see Section 3.4.2).

3.3.2 Boundary Force Generation

Boundary exponential force distributions (as per the original SFM) result in a model in which the ultimate calibrated balance of forces is sensitive to path width as an exponential distribution has infinite extent. To avoid this, and given that cyclist boundary effects are currently poorly understood, a simple linear function for force is proposed (Section 3.4.4) which decreases with distance from the boundary and reaches zero at a given distance (i.e. approximating edge-shyness; Department for Transport, 2008). Thus, the balance of the interactions with other bicycles and with the boundaries is broadly decoupled from the path width for the purpose of this model, with the exception of those in close proximity to the path edge where avoidance behaviour is as expected.

(a) Section through bicycle force profile

(b) Plan view of 'isopotentials'

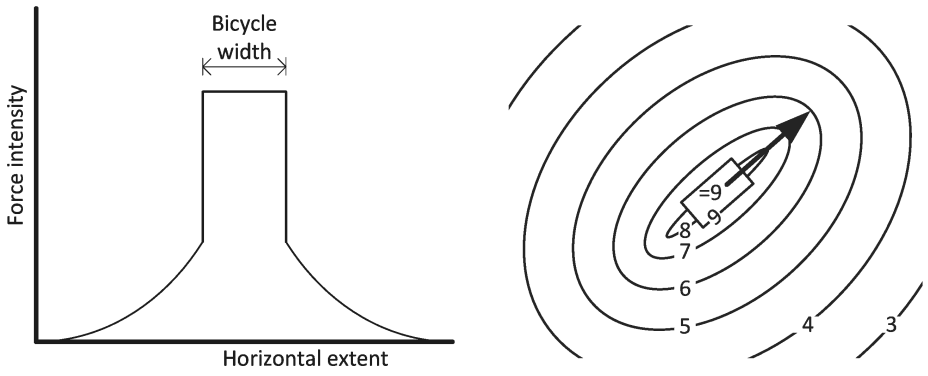

Figure 3.3: Illustrative force profiles (arbitrary scales/values)

3.3.3 Force Perception

The SFM considers the potential of a view-cone and Helbing & Molnár (1995) proposed a reduction factor of 0.5 for pedestrians outside of it (i.e. to the rear; see Figure 3.4). However, combined with the directionality of the force, the result is a 'pushing' force from a pedestrian approaching from behind. Helbing & Molnár were primarily concerned with crowd crushes wherein such pushing forces are explicitly non-trivial. Future development to the SFM (e.g. Helbing et al., 2005) revised this to a continuous anisotropic λ parameter, however, for the purposes used here, a simple augmented zonal view-cone is sufficient when considering the issue from first-principles.

For bicycles in 'usual operation', a simple reduction by half of the perceived effect of other cyclists to the rear is not realistic for two reasons. Firstly, a reduction factor in the order of a half for anywhere outside the view-cone is suggested for pedestrians, but it requires substantially more effort to look directly

behind safely whilst cycling than is the case for walking, thus reducing the ability to perceive such a force. Secondly, it is debatable whether a cyclist would feel any non-trivial requirement to cycle substantially faster due to a perceived social pressure of cyclist(s) behind them. The second issue is alleviated through the use of a scalar force field which, given the proposed directional choice algorithm, does not fundamentally change the operation of the model. With regard to the reduction factor, three different perception profiles are proposed (compare Figures 3.4 and 3.5):

- If the perceived bicycle is within the forward view-cone, then the full force effect of that bicycle is perceived. This area requires no effort to view on behalf of the perceiving cyclist.
- If the perceived bicycle is within one of the side view-cones then notionally, this would require the cyclist to turn their head to observe them. This requires effort on behalf of the cyclist and may result in them not being observed if the cyclist is otherwise engaged. This equates with the side blind-spots for a motor vehicle driver. As a result, the effect of bicycles present within these view-sectors is reduced.
- If the perceived bicycle is not within the side or front view-cones, then the cyclist would have to turn their head and torso simultaneously to look behind; a difficult task on a bicycle whilst maintaining full control. Consequently, the effect of these bicycles are perceived less than the side view-cones (if at all).

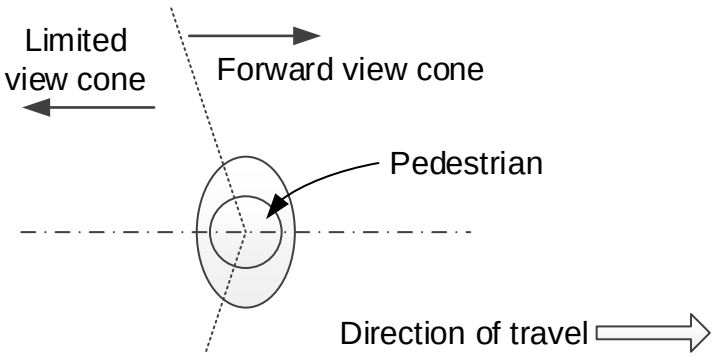

Figure 3.4: Illustrative view cone in original Social Force Model

Perception of the force field ahead of the cyclist is condensed to a 'net present force' on each potential vector choice. The vector with the least repulsive (or most attractive) net present force is selected. This ensures that cyclists react most to obstacles close to them rather than those further away and that the greater certainty of the imminent situation is reflected in the force perception.

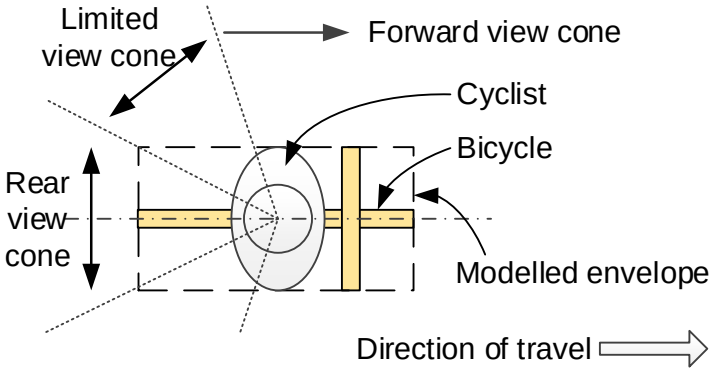

Figure 3.5: Illustrative view cone proposed in this model

As each cyclist makes this consideration once per time step, the cyclist can react appropriately to the changing circumstances around them, as they unfold.

3.3.4 Bicycle Speed Selection

Pedestrian speed in the SFM is only bounded at the maximum; i.e. pedestrians can move at very low speeds. Whilst this may be realistic for pedestrians who can essentially reduce step length and walking cadence continuously to zero, cyclists are limited by a minimum speed below which they cannot maintain balance on the bicycle and (usually) react by lowering one or both feet to the ground.

Acceleration is proposed to be applied to the bicycle using the appropriate kinematic equations of motion (assuming uniform linear acceleration during that time step) with resulting speeds capped at the desired maximum and by the minimum sustainable (which if the speed would fall below, then a 'foot down' stopping state is activated; see Section 3.4.7). For this model, acceleration is simply assumed to be invariant with speed as there is a lack of any empirically-backed literature to the contrary.

Speed selection therefore takes place realistically rather than being abstracted to a "relaxation parameter" as in the original SFM. Acceleration in the model is modified by the cyclists' perception of the force field ahead of them on their chosen path. This too replicates real behaviour where speed selection is achieved through the application of braking or by acceleration.

3.4 Model Implementation

3.4.1 Model Structure

The model for each cyclist agent follows the structure detailed in Figure 3.6 which parallels the SFM as formulated in Helbing & Molnár (1995). However in structural terms, it should be noted that the complexities of cycle behaviour over those of pedestrians mean that the proposed model is not entirely built upon mathematical abstraction (as per the SFM). In common with many models, some aspects are algorithmic and/or iterative in nature; for example, the rule-based lane selection algorithm in PTV (2009), detailed in Barcelo (2010).

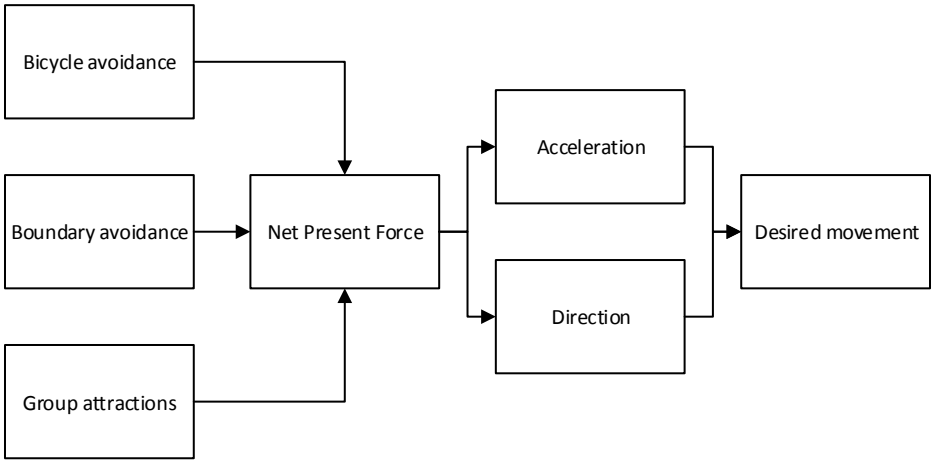

Figure 3.6: Structure of the proposed SFM

At the highest level, the model is designed to be run iteratively, where time is stepped by a fixed amount (t_{step}; seconds) between iterations, every cyclist is iterated through once at each time step (according to the process defined in Figure 3.6), and the location/speed of each cyclist is subsequently updated once per time step (once all other calculations are complete).

3.4.2 Bicycle Repulsive Force Generation

The repulsive force of each bicycle ($force_{bicycle}(pt)$; in kgms^{-2}) is generated according to Equation 3.15:

$$force_{bicycle}(pt) := \begin{cases} forceScale_{bicycle} \times \exp\left(\frac{bikeWidth-b}{forceSpread_{bicycles}}\right) \\ \qquad \text{if } pt \text{ is not within the bicycle spatial envelope; or} \\ forceScale_{bicycle} \times \exp\left(\frac{bikeWidth}{forceSpread_{bicycles}}\right) \\ \qquad \text{if } pt \text{ is within the bicycle spatial envelope.} \end{cases} \tag{3.15}$$

where pt is the point for which force is to be calculated, $bikeWidth$ is the width of the bicycle spatial envelope (in metres), and b is the semi-minor axis length (in metres) of the elliptical 'isopotential' coincident with that point (Figure 3.3b). $forceScale_{bicycle}$ and $forceSpread_{bicycles}$ are scaling factors in kgms^{-2} and metres respectively. b is calculated as per Equation 3.16:

$$b := \frac{\sqrt{(df_1 + df_2)^2 - (df_{inter})^2}}{2} \tag{3.16}$$

where df_1 is the distance (in metres) from the point of interest to the rear focus, df_2 is the distance (in metres) from the same point to the front focus and df_{inter} is the inter-foci distance (also in metres).

3.4.3 Bicycle Attractive Force Generation

The attractive force of each bicycle ($forceAttraction_{bicycle}$; in kgms^{-2}) at a given point (pt) is generated according to Equation 3.17:

$$forceAttraction_{bicycle}(pt) := \begin{cases} related_i \times forceAttractionScale_{bicycle} \\ \times \exp\left(\frac{bikeWidth - b}{forceAttractionSpread_{bicycles}}\right) \\ \quad \text{if } pt \text{ not within bicycle spatial envelope; or} \\ 0 \\ \quad \text{if } pt \text{ is within the bicycle spatial envelope.} \end{cases} \tag{3.17}$$

Attractive forces are defined in a similar manner to the repulsive forces discussed above, however no attractive force is returned for the area within the spatial envelope of the cyclist. As noted above, this is not a desirable area for another cyclist to seek as it would result in a collision, thus no attractive force should be experienced at that point.

Parameters for the spatial distribution of the attractive force ($forceAttractionScale_{bicycle}$ and $forceAttractionSpread_{bicycle}$) are defined similarly to repulsive force, save that if repulsive force is defined positive, then attractive force must be negative, or vice versa.

Appropriate calibration ensures groups of 'related' cyclists remain apart yet still attempt travel in a proximate group. The existence of an attractive force between given bicycles, depends on them being defined in the same group or not; bicycles only experiencing an attractive force from other bicycles in the same group. This behaviour is captured by $related_i$ which is defined in the range 0.0–1.0 inclusive (no units) and depending on the level of relationship that exists between the bicycle concerned and the other bicycle being considered (i.e. no force is returned if the two cyclists do not have a grouping relationship) and allowing for a tunable strength to that relationship. A $related_i = 0.0$ denotes bicycles with no grouping relationship; a $related_i = 1.0$

denotes the strongest attractive relationship. 'Tune-ability' allows for more complex relationships to be quantified: a child–parent desire to stay near one another might be stronger than that between work colleagues for example; the former therefore having a higher value of $related_i$ than the latter.

3.4.4 Boundary Force Generation

Boundary forces ($force(pt)$; in kgms^{-2}) are generated according to a simple linear function (Equation 3.18):

$$force(pt) := \max\Big(forceScale_{boundary} - (forceSpread_{boundary} \times d)\,,\, 0\Big) \quad (3.18)$$

where d is the perpendicular distance (in metres) from the point of interest to the boundary concerned. If repulsive forces are defined to increase in the positive (i.e. 'more' force is repulsive), then the maximum value is chosen as shown here. If repulsive forces are defined to increase negative (i.e. 'more' force is attractive), the reverse is true. Care must be taken to ensure consistency of signing (i.e. boundaries should always be repulsive), however the choice of signing for either repulsive or attractive force as positive, is the choice of the modeller. $forceScale_{boundary}$ and $forceSpread_{boundary}$ are scaling factors (in kgms^{-2} and kgs^{-2}, respectively).

3.4.5 Net Present Social Force

The force returned at a given point and time ($force(pt)$; in kgms^{-2}) is the sum of the repulsive force of all other bicycles ($force_{bicycle}$; in kgms^{-2}), boundaries ($force_{boundary}$; in kgms^{-2}) and any attractive forces from other bicycles ($forceAttraction_{bicycle}$; Equation 3.19; in kgms^{-2}):

$$force(pt) := \sum^{i} force_{bicycle_i} + \sum^{j} force_{boundary_j} + \sum^{i} forceAttraction_{bicycle_i} \quad (3.19)$$

The concept of a net present force is key to this implementation of the SFM. A 'net present' force allows the forces returned at a number of spatial points and for various times in the future to be combined with due regard for the fact that points further in the future have a lower certainty attached to them and should therefore be given less regard.

The discount to a net present value ($netPresentForce$; in kgms^{-2}) is defined by Equation 3.20:

$$\begin{aligned} netPresentForce := \sum_{u=plan_{step}}^{u=plan_{max}} \Bigg(\exp\left(-\lambda \times \left(\frac{u}{plan_{step}} - 1\right)\right) \\ \times \Big(force(pt_{left}) + force(pt_{right})\Big)\Bigg) \end{aligned} \quad (3.20)$$

where u (in seconds) is the current planning step ahead being considered,

$plan_{step}$ is the minimum planning step duration considered and $plan_{max}$ is the maximum time ahead being considered (both in seconds). λ (no units) is the decay constant (not to be confused with the use of λ in later versions of the SFM Helbing & Johansson, 2013b where it relates to the tunable strength of force applied from behind the pedestrian) which is tunable based on the intensity that future force returns are perceived – higher values of λ placing less weight on force perceptions further in the future (i.e. a more rapid decay of force returned with increasing time) and the value chosen being a matter for calibration.

Force at the current position is not considered. $force(pt_{left})$ and $force(pt_{right})$ (both in kgms^{-2}) are the returned forces at that planning time step for a point on the left and right extents of the bicycle envelope on the given vector (see Figure 3.7). This is included so the bicycles do not attempt to force their way through gaps that are too small to pass through, should a single point on that vector happen to fall between two close objects.

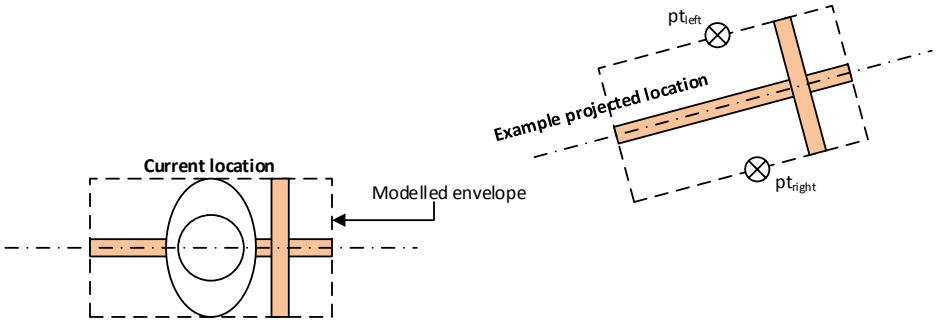

Figure 3.7: Indicative example of arrangement of pt_{left} and pt_{right} used in force computation

3.4.6 Desired Movement

In the SFM, desired movement is defined numerically as a combined function of the desired speed and direction of movement, and/or a relaxation to this. The resultant behaviour of the pedestrian then being affected by the social force. In this model, the concept of desired movement is emergent from the tendency, in the absence of any social force, for bicycles to travel in the desired direction and at the desired speed. Following perturbation, travel direction returns to the desired and speed returns to desired at a rate bounded by maximum acceleration. Speed and direction of travel are independently changed as follows in Sections 3.4.6.1 and 3.4.6.2.

3.4.6.1 Speed

As with the pedestrian SFM, this cyclist SFM gives the cyclist a desired speed. For the purpose of modelling and in the absence of experimental data, a normally-distributed speed distribution is assumed.

Acceleration is equal to the desired acceleration (a_{max}; in ms^{-2}) which is reduced based on the overall net present repulsive force perceived on the chosen direction of travel ($netPresentForce$; in $kgms^{-2}$) divided by the mass (m; in kg) of the cyclist/rider combination (and can be ignored if all are assumed unit mass), and if the remaining acceleration is sufficient, results in deceleration and thus slowing (Equation 3.21). Such a deceleration is bounded by the maximum deceleration parameter (a_{min}; in ms^{-2}). In the absence of any literature to the contrary, acceleration is considered independently of the speed, though it is accepted that this may not be the case in reality.

$$a(t) = \begin{cases} a_{max} - \left(\frac{netPresentForce}{m}\right) \\ \qquad \text{if } > a_{min} \\ a_{min} \\ \qquad \text{if } \leq a_{min} \end{cases} \tag{3.21}$$

Resultant speed is defined as a function of acceleration and speed in the previous iteration time step; itself calculated from the perception of social force on the chosen route. At each time step, acceleration is applied as constant linear acceleration for the duration of the time step, and using the relevant kinematic equation of linear motion (Equation 3.22). Consequently, the need for a critical, yet practically abstract, 'relaxation' parameter is avoided.

$$v(t) = \begin{cases} v(t - t_{step}) + (a(t) \times t_{step}) \\ \qquad \text{if } \leq v_{desired}\text{; or} \\ v_{desired} \\ \qquad \text{if } > v_{desired}. \end{cases} \tag{3.22}$$

where $v(t)$ is the new speed in the current time step (in ms^{-1}), $v(t - t_{step})$ the speed in the previous time step (in ms^{-1}), $a(t)$ (in ms^{-2}) the current acceleration and t_{step} (in seconds) the length of the simulation step.

3.4.6.2 Direction

Desired direction of movement is defined in the SFM as the direction from the current position to the next waypoint/destination. This model proposes similar in the absence of any surrounding social force.

However, the pedestrian SFM modifies the direction of travel vector (i.e. the

velocity) instantaneously according to the ambient social force. Here instead, speed and direction are decoupled and the social force is aggregated into a net present force, which is perceived on each available direction of travel. The direction of travel chosen at each time step is simply the available direction of travel with the most attractive (or least unattractive) social force exposure. In that sense, motion is an agent 'choice' based on planning, rather than an immediate reaction to the surroundings.

3.4.7 Stopping

The model presented thus far produces outputs (when appropriately calibrated) that appear valid at low densities of flow. However, there is an already identified issue (Section 3.3.4) with bicycles which is not yet included: namely, that there is a minimum speed that cyclists can maintain before they need to put one or both feet down, and stop.

A final consideration must therefore be overlaid upon the resulting new speed calculated in each time step. In this regard, the cyclist can be considered as a 'state machine' illustrated in Figure 3.8.

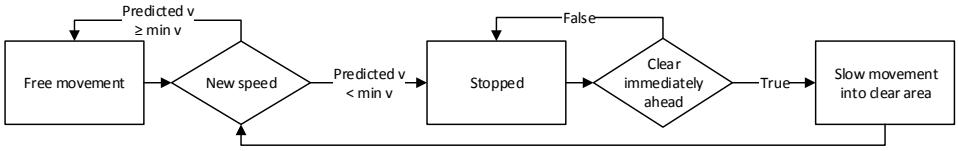

Figure 3.8: Cycle state machine arrangement

A cyclist exhibits movement until such time as the speed of the cyclist falls below the minimum speed threshold. At this time, the cyclist stops (by lowering one or both feet), which given the low speed, can be assumed instantaneous (i.e. stopping occurs within a single time step). Thus far (and in reality), the distance considered ahead is (indirectly) a function of the cyclist's speed — the cyclist considers a given time ahead, a faster cyclist therefore considers a greater distance. When the cyclist stops, such a consideration is no longer appropriate. Instead, the cyclist considers, "Is the area immediately ahead clear?" If this is the case then the cyclist moves off at a minimum speed for as long as it takes to traverse that distance already considered and known to be clear (in this time, the cyclist is regaining their balance). Once that time has passed, consideration is as the standard behaviour where, if the space ahead is clear, the cyclist will accelerate away; if not, then the cyclist may again slow below the minimum threshold and stop again.

Algorithmically, this represents appropriate cycle behaviour relating to moving off or stop-go queuing and is included here in such a manner as this is an implementation aspect of the model, rather than mathematical. The model

retains substantial validity without such a consideration, however as discussed (Section 3.3.4), the ability of a cyclist to travel at potentially infinitesimal speeds is unrealistic.

3.5 Simulation Parameters and Implementation

3.5.1 Parameters

Table 3.2 draws together the parameters required to operate the model and provides suggestions for values based either on literature (as shown), or visual calibration of the implementation of the simulation that follows in the remainder of this chapter. As stated, it should be noted that a number of the parameters require calibration, however this does not impact the underlying validity of the model, as will be demonstrated.

Where available, the literature parameters used in this implementation are taken from CROW (2007) which ensures that they have some element of empirical backing. Though perhaps intuitively low (e.g. a mean of 4.02ms^{-1}), this source also provide values for standard deviation and acceleration/deceleration, so they can be considered to be at least internally-consistent within the given (in this case Dutch) population.

Minimum speed is taken from Navin (1994) as no such value is included in CROW (2007) and no other literature exists. Model parameters for the intensity of the social forces were chosen using visual observation of the simulation running, using informed approximations (such as the angles of sight utilised) in order to produce outputs commensurate with intuitive and empirical knowledge; e.g. forces were calibrated to allow virtual lanes of approximately 1.0m width to form, in keeping with the dynamic envelope of the cyclist specified in Department for Transport (2008) and elsewhere (though it should be noted the empirical basis for such a value is, in itself, unproven). All parameters used in the preparation of the data presented in this chapter are detailed in Table 3.2.

3.5.2 Worked Example

In order to further clarify the basic operation of the model, the following details a worked example for two bicycles. Clearly the value in such a model is in the ability to scale to large numbers of agents and run for large numbers of time steps, if implemented as a computer simulation. However, a worked example is presented here to demonstrate that such a simulation is simply a result of a large number of these calculations.

3.5.2.1 Circumstance

The example worked here concerns two bicycles arranged per Figure 3.9. The workings illustrate the computation for Bicycle 1 and use the same parameters per the implementation in Section 3.5 and Table 3.2.

Table 3.2: Model Parameters

Parameter	Source	Symbol	Value	Requires calibration
Simulation time step		t_{step}	0.1s	No
Simulation length			300.0s	No
Bicycle traversal length			75.0m	No
Boundary force spread	Department for Transport (2008)	$forceSpread_{boundary}$	200kgs^{-2}	Yes
Boundary force scaling		$forceScale_{boundary}$	4000kgms^{-2}	Yes
Bicycle force spread		$forceSpread_{bicycles}$	75m	Yes
Bicycle force scaling		$forceScale_{bicycles}$	150kgms^{-2}	Yes
Bicycle attractive force spread		$forceAttractiveSpread_{bicycles}$	75m	Yes
Bicycle attractive force scaling		$forceAttractiveScale_{bicycles}$	0kgms^{-2}	Yes
Bicycle force ellipse foci separation		df_{inter}	5.0m	Yes
Bicycle width	Transport for London (2014b)		0.75m	Yes
Bicycle length	Department for Transport (2008)		1.8m	Yes
Bicycle fleet speed mean	CROW (2007)		4.02ms^{-1}	Yes
Bicycle fleet speed standard deviation	CROW (2007)		0.21ms^{-1}	Yes
Bicycle fleet maximum acceleration	CROW (2007)	a_{max}	+1.0ms^{-2}	Yes
Bicycle fleet maximum deceleration	CROW (2007)	a_{min}	−1.5ms^{-2}	Yes
Bicycle minimum speed	Navin (1994)		0.92ms^{-1}	Yes
Bicycle maximum steering angle			40.0°	Yes
Bicycle steering step			4.0°	Yes
Bicycle forward planning step		$plan_{step}$	0.25s	Yes
Bicycle forward max planning time		$plan_{max}$	5.0s	Yes
Bicycle forward planning decay constant		λ	1.0	Yes
Bicycle angle of sight			±100.0°	Yes
Bicycle angle of reduced sight			±160.0°	Yes
Side reaction factor			0.1	Yes
Rear reaction factor			0.0	Yes
Bicycle mass			1.0kg	Yes

3.5.2.2 Bicycle Forces

For each of the possible vectors the bicycle could select (any one of the steering steps within the range of the maximum steering angle; illustrated non-exhaustively in Figure 3.10) a 'net present' total force on that vector is computed. This total force is calculated at the location at which the bicycle would be, should it make that steering choice for each time step ($plan_{step} = 0.25s$) in the future, cumulatively to the maximum of 5.0s ($plan_{max}$).

For the purpose of clarity, the example worked here will be for the 0° steering choice. At the first planning step ($u = 0.25s$), the projected location of the bicycle is computed to find pt_{left} and pt_{right}. The projected locations of all the other bicycles (in this case a single bicycle) are also computed based on the current direction and speed of those bicycles; something which a cyclist could reasonably observe. This projection is illustrated in Figure 3.11.

To first compute $force_{bicycle}(pt_{left})$, we require Equation 3.15, which to complete, we require b which is the result of Equation 3.16. Recall b is the length of the semi-minor axis of the elliptical force distribution about the other bicycle. b depends on the distances to the foci about the other bicycle. For illustration, see Figure 3.12.

For this circumstance (Figure 3.12), the distance to the rear focus (df_1) from pt_{left} is 2.335m and the distance to the forward focus is 7.276m. The interfoci distance (df_{inter}) is 5.0m per Table 3.2. b is therefore 4.105m.

By now substituting this into Equation 3.15 with $forceSpread_{bicycles} = 0.075m$ and $forceScale_{bicycle} = 150.0kgms^{-2}$, a force of $5.633 \times 10^{-18}kgms^{-2}$ is calculated for this particular projected location (pt_{left}) at $u = 0.25s$ for a chosen steering angle of 0°. As the other bicycle is ahead of this bicycle, no factoring is required for the harder to view viewcones.

3.5.2.3 Boundary Forces

For the purpose of this example, attractive forces will be ignored (i.e. $related_i = 0.0$ in Equation 3.17). Therefore, the remaining force to be calculated for this circumstance, is that from the boundaries.

By substituting the perpendicular distance d to the 'northern' boundary from pt_{left} at $u = 0.25s$ and a chosen steering angle of 0°, which here has a value of 1.625m, into Equation 3.18, a force of $0kgms^{-2}$ is returned (specifically, the maximum of 0 and -321kgms^{-2}).

The force from the 'southern' boundary is calculated similarly and also results in a force of $0kgms^{-2}$ for this location of pt_{left}.

3.5.2.4 Total Force

Next, these forces are summed to provide the total force at pt_{left}, which is equal to $5.633 \times 10^{-18}kgms^{-2}$. However, forces calculated further in the future

have a greater uncertainty. For example, the other bicycle may have changed direction so will then not be in the position calculated by this bicycle had it continued on its current trajectory. To account for this, future forces are discounted per Equation 3.20.

In this case, $u = 0.25s$, $plan_{step} = 0.25s$, $\lambda = 1.0$, $force(pt_{left})$ was calculated to be $5.633 \times 10^{-18} kgms^{-2}$, and if the calculation is repeated for $force(pt_{right})$, that value is found to be $0.246 \times 10^{-18} kgms^{-2}$. These figures are then substituted into Equation 3.20 to provide a value of $5.879 \times 10^{-18} kgms^{-2}$ for this value of u.

The above is then repeated for each value of u up to (and including) the value of $plan_{max}$ and these values are all summed to yield the total 'net present' force of $3.181 \times 10^{-3} kgms^{-2}$ for a steering choice of 0°.

The implementation of this model then repeats this process for all other possible steering choice angles to result in a list of steering choice angles and total 'net present' force for each of those choices. The least repulsive (or most attractive) of those is then chosen as the next $time_{step}$'s trajectory. In this case, the chosen trajectory is one with a steering angle of +4° (and total net present force of $5.009 \times 10^{-12} kgms^{-2}$), which in this coordinate system, is slightly downwards, away from Bicycle 2. Figure 3.13 illustrates the 'net present' forces on each potential trajectory and illustrates the general range of steering to be constrained to within the pathway. Figure 3.14 then shows the same data plotted on a log scale which highlights the desirable route as one slightly to the right of ahead for Bicycle 1. Note specifically the small peak in the force return to the left (around -4°) which corresponds to the presence of Bicycle 2.

3.5.2.5 Speed Selection

For simplicity, the mass of the bicycle is assumed unitary (i.e. 1.0kg) and thus, per Equation 3.21, the total force on the chosen trajectory (which in this case is very small) is deducted from a_{max} resulting in a positive acceleration of $+1.0 \text{ms}^{-2}$. Note that as the direction was chosen before the adjustment to the speed, the bicycle has a tendency to choose a less obstructed path (if available) rather than to slow down, as is the case in reality.

The current speed of the bicycle is 4.0ms^{-1} (which is less than $v_{desired}$ which is equal to 4.02ms^{-1}) therefore the acceleration is applied as constant linear acceleration for this simulation time step per Equation 3.22 and as an acceleration at this rate would exceed $v_{desired}$ (it would result in a speed of 4.1ms^{-1} over $time_{step} = 0.1s$), the resulting new speed is therefore set equal to $v_{desired}$; i.e. 4.02ms^{-1}.

3.5.2.6 Simulation Update

The new direction of the bicycle is that of the chosen trajectory's angle, here $+4°$ and the new speed has now been calculated (4.02ms^{-1}). This process is computed for each bicycle agent in the simulation in the same way and, once all calculations have been completed, all bicycles are moved based on the new speed and direction by the appropriate distance and angle given the length of the simulation $time_{step}$.

For example, given $time_{step} = 0.1s$, the bicycle calculated here moves at the angle $+4°$ at a speed of 4.02ms^{-1}, which results in a forward motion for this bicycle of 0.401m and a lateral motion ('southwards') of 0.028m.

This procedure is then repeated until the desired length of simulation has been completed. To demonstrate the operation of this example, a simulation was set up with these parameters and the bicycles' locations tracked. The resulting trajectories are shown in Figure 3.15 which illustrates the movement of Bicycle 1 to right of Bicycle 2 where – as it is travelling faster – it is free to overtake. Bicycle 2 continues on its path unimpeded as Bicycle 1 moves ahead and away, and is therefore of minimal influence to Bicycle 2, as would be expected.

3.5.3 Application of the Model

The model above allows the testing of the basic assumption underpinning Botma (1995) which is that passing/overtaking cyclists do not impede/interact with one another, and that any such event is able to be ignored for all but the most extreme rates of flow (see discussion in Section 2.4.1). Furthermore, despite that paper presenting its capacity measures as a "proposal" with "preliminary character", those measures have seen their adoption into the core literature used (internationally) by the practitioner – albeit with some caveats Rouphail et al. (1998) note that the numbers considered are not absolute capacity measures, but their interpretation is the equivalent of considering a highway maximum operational flow to not be the measure of moving vehicle capacity. Indeed, Botma makes the point that such testing is required, however no such consideration has occurred in the literature to date. The remainder of this chapter takes the simulation developed above, and applies it to that standing issue.

3.5.4 Implementation

An agent-based computer simulation was built which implemented the model described. For simplicity, no grouping of cyclists was assumed (i.e. only repulsive forces were present) and unidirectional flow in two-dimensional space was considered on paths of fixed widths (1.0m–10.0m; in steps 0.1m) and parallel sides and a length of 60m. Testing of the model indicated that beyond the

length used in this chapter for this particular scenario (parallel-sided path), output parameters were substantially unchanged, however the run-time of the simulation would rapidly become intractable. Bicycles arrive stochastically according to Poisson-distributed arrival intervals (based on a set overall rate for the run) and at random lateral positions in the pathway at rates of 100–5000 bicycles per hour (in steps of 100 bicycles per hour). To reduce the effect of stochastic noise, 25 runs per parameter combination were undertaken.

All bicycles are assumed equal in terms of parameters, the exception being their individually desired speed which is drawn from a Normal distribution with the specified parameters (Table 3.2).

A graphical interface (Figure 3.16) was coded to allow visual calibration and verification that the model was behaving as expected. Outputs from runs were recorded as to the average speed of bicycles and miscellaneous other data such as instances of crashing (where crash conditions are defined as an intersection of the bicycle envelope with another bicycle envelope, boundary or off-path area), distance between bicycles, random seed values, run IDs, etc.

Given the overall number of runs, the simulations were completed on the Iridis4 High Performance Computing cluster (at the University of Southampton), however the model is generally capable of running in acceptable time on mid-range (or better) desktop hardware, for reasonable ranges of parameters.

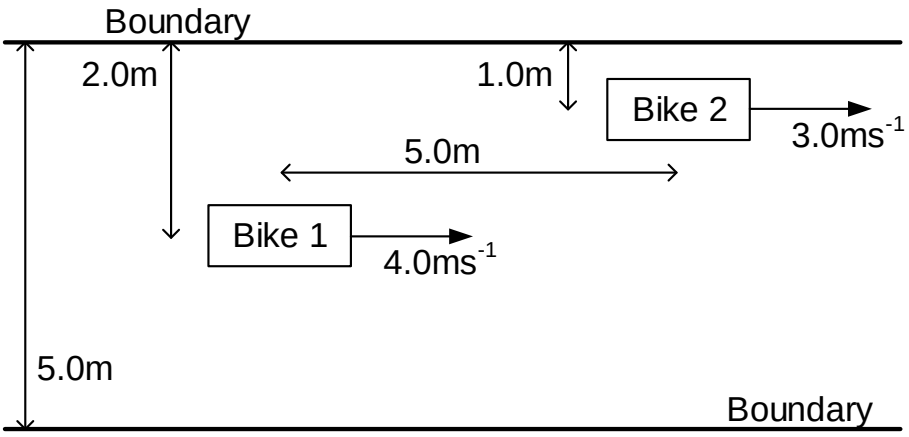

Figure 3.9: General arrangement of bicycles in worked example (not to scale)

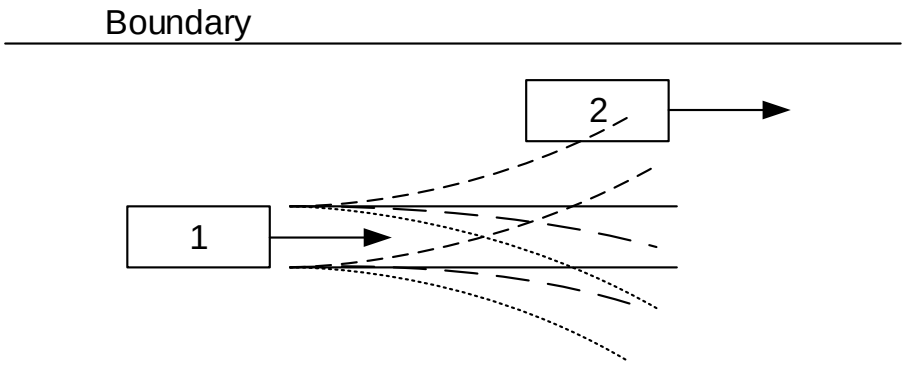

Figure 3.10: Illustrative path choices of Bicycle 1 (not to scale; non-exhaustive; southern boundary omitted for clarity)

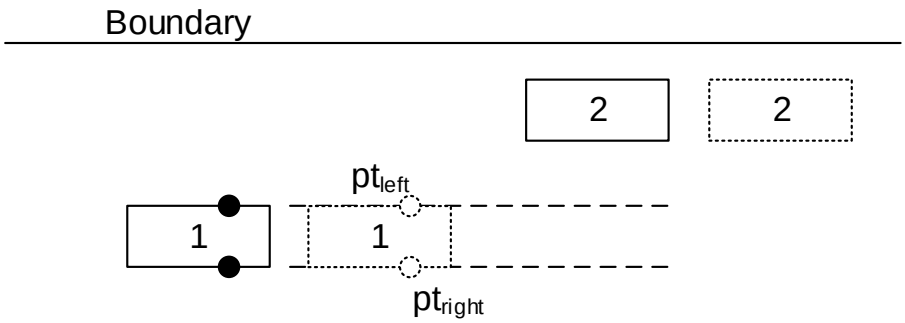

Figure 3.11: Illustrative pt_{left} and pt_{right} projection from Bicycle 1 (not to scale)

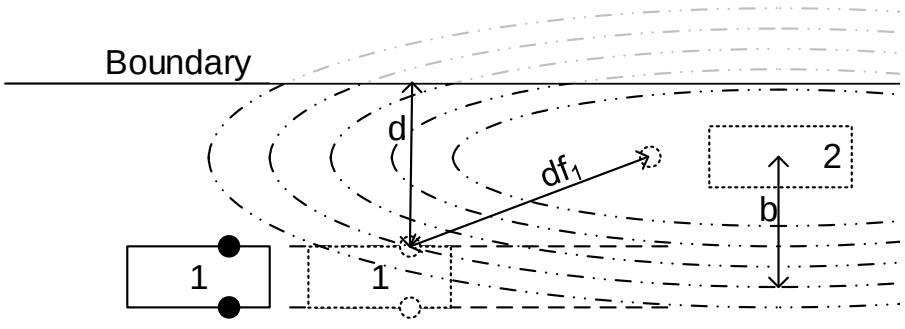

Figure 3.12: Illustrative repulsive field projection by Bicycle 2 (not to scale)

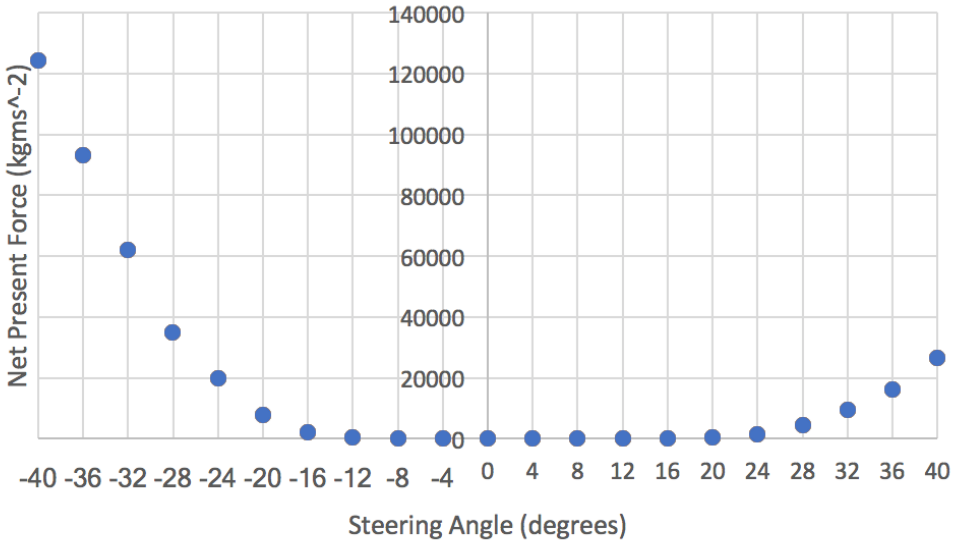

Figure 3.13: Net present repulsive force across potential trajectories for Bicycle 1

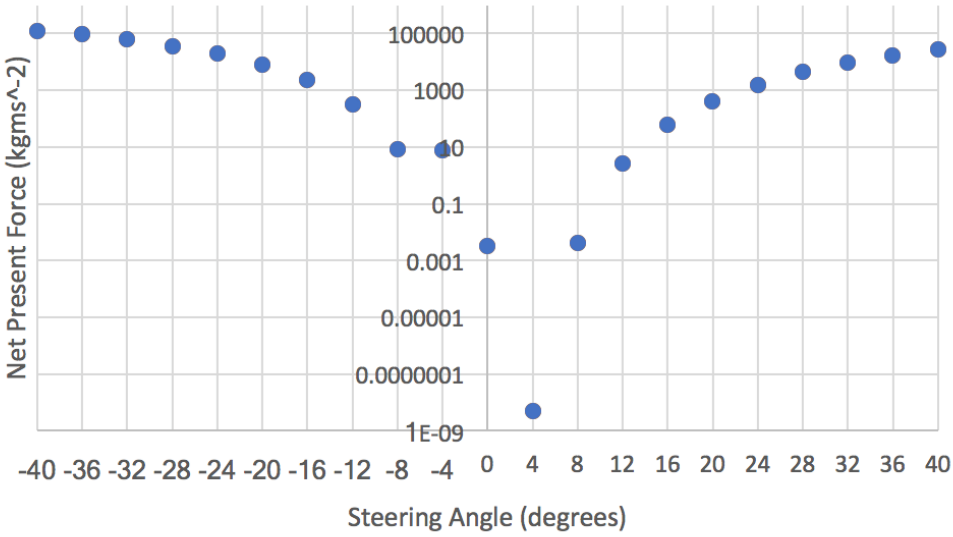

Figure 3.14: Net present repulsive force across potential trajectories for Bicycle 1 (log scale)

Figure 3.15: Trajectories of bicycles following worked example (figures in m; not to scale)

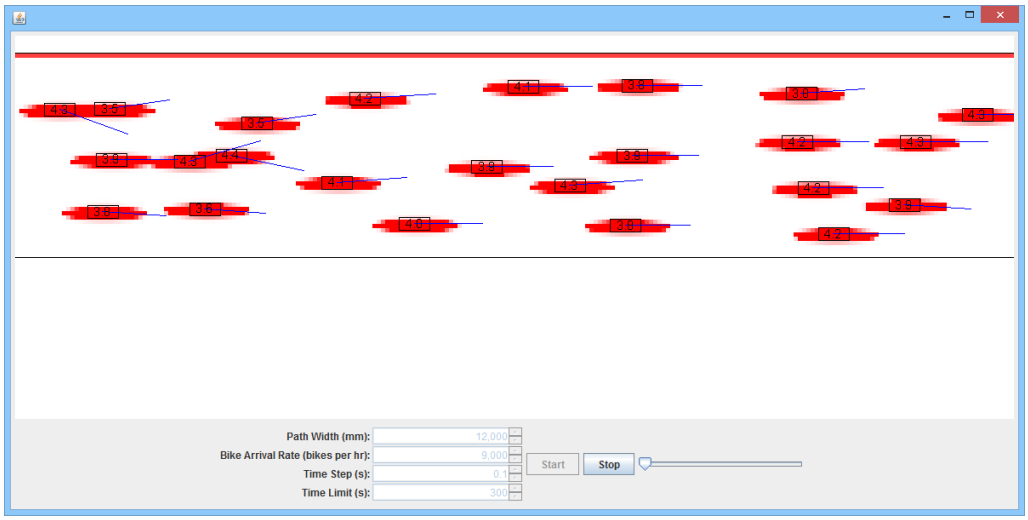

Figure 3.16: Graphical interface of model implementation (indicated parameters run for illustrative purposes)

3.6 Simulation Outcomes

The two parameter variables – path width and bicycle arrival rate – are essentially covariant if expressed in terms of flow density. For the purpose of the following, these are therefore combined into a bicycles per metre width per hour (bpmph) measure. Two sets of the parameter sweeps discussed above (Section 3.5) were completed: one with all agents set such that they could not vary their speed (i.e. $v(t)$ was fixed to $v_{desired}$ for each agent), and the other as described above. All other parameters remained fixed. The former scenario therefore uses the assumption from Botma (1995), whereas the latter does not. Should the assumption be valid, observation of a similar output from both sets across a range of input parameters would be expected. The result of this comparison is shown in Figure 3.17.

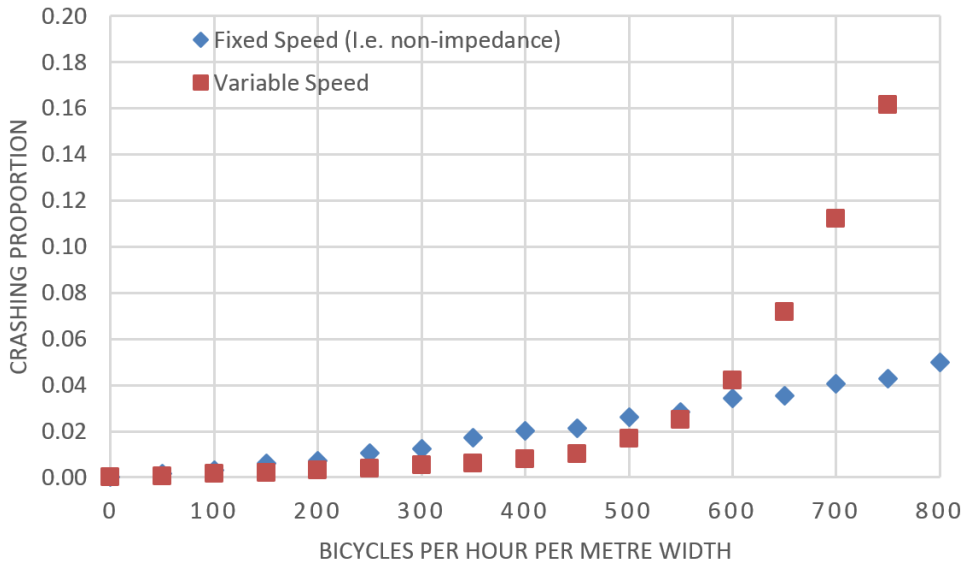

Figure 3.17: Crash proportions across a range of arrival rates for fixed and variable speed simulation runs

For fixed speed operation, the output indicates a proportion of agents are experiencing a 'crashing' state across essentially all arrival rates with proportions steadily increasing. The implication of this being that the interaction of bicycles is potentially non-negligible, even at the low end of flow rates.

For variable speed operation, a fundamentally different outcome is observed. In this case, crash proportions are approximately zero until an inflection point is reached at approximately 550bpmph. A state change has occurred and the proportion of agents in a crashing state escalates rapidly (much more so than for fixed-speed operation). This corresponds with the visual observation of the simulation that congestion has occurred. The inflection point therefore indicates a capacity limit to the arrangement. Rather than the flow degrading slowly with increasing arrival, the flow degrades sharply. When this happens

the shockwave propagates through the flow with dissipation only occurring once inflow to the area of breakdown has dropped sufficiently. This corresponds with equivalent observations of highway and pedestrian traffic.

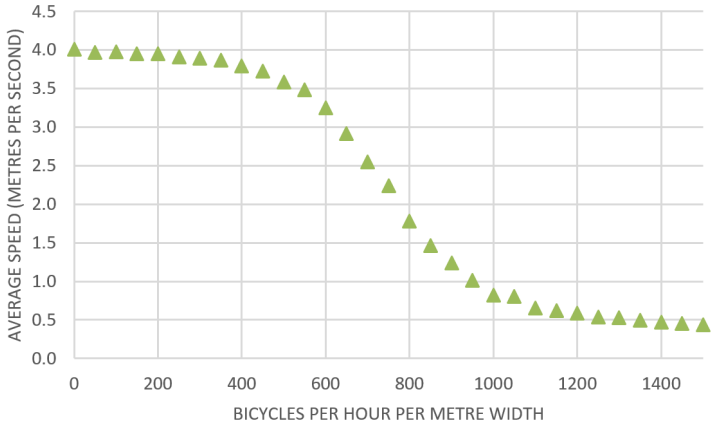

Figure 3.18: Average bicycle speeds across runs undertaken at given inflows

Exploring this further and plotting the average speed against this measure, Figure 3.18 is produced. This clearly demonstrates outputs in keeping with established three-phase traffic theory. Below inflow rates of approximately 500bpmph, speeds are high and close to the cyclists' desired speed. Once densities increase, and up to an inflow of approximately 1000bpmph, flow breakdown is increasingly likely. Beyond 1000bpmph, flow breakdown occurs consistently with average speeds being low. Between these two values, some runs will complete without flow breakdown, others will experience limited breakdown which resolves, and the remainder will experience flow breakdown which does not resolve. The balance between these three states depends on the inflow rate which, as it is stochastic in nature, means all three phases can potentially exist for a given parameter combination in different locations on the same path.

Breakdown tends to occur in the entry area, as inflow is not stopped if the inflow area is not clear. In essence, the bicycles are arriving from a notionally infinite width space into a fixed width one and consequently, a de facto bottleneck. Breakdown at the entry in this manner is paralleled in pedestrian behaviour, such as crowding around a door or at a corridor entrance (Hoogendoorn & Daamen, 2005).

The results of each parameter combination's set of runs is consequently bimodally distributed between those runs were speeds were high (i.e. no breakdown occurred) and where speeds were very low (i.e. breakdown was rapid and ongoing) as flow breakdown often takes substantial time to resolve (as in reality). The resulting plot (Figure 3.18) averages out that behavioural detail and thus it must be underlined that for a parameter combination such as at 750bpmph, few bicycles will be found travelling at 2.0ms^{-1} at any given time step — instead, this is an average of runs comprised mainly of slow/stopped

bicycles and those travelling essentially unimpeded. The only variable between those runs is the stochastic nature of arrivals and therefore the occurrence of flow breakdown due to the 'noise' in the system, and in the absence of any exogenous factors, parallels real highway traffic behaviour. Similarly, no bicycle will be travelling at a speed less than their minimum but greater than zero (in a given time step), so indicated average speeds beyond 1000bpmph are results of the aggregation of stopped bicycles with those travelling at a speed greater than, in this case, 0.92ms^{-1}.

Given the inputs and their limited literature/empirical basis, it is important to note that the exact numbers observed are unimportant. However, the inputs to the simulation are based on first-principles and the limited literature that exists, and yield simulation outcomes in the range of values which are expected from that literature (Table 2.1). At the low end Transport for London (2014a) specifies a capacity in the range of 133–363bpmph; at the high end, Navin (1994) computes a maximum capacity of 4000bpmph.

Most importantly however is the fundamental difference that is observed by the inclusion of speed-changing behaviour. Speed-changing behaviour is explicity ignored as an assumption in Botma (1995) but is shown here to result in a fundamental shift in outcomes within the range of values which a practitioner is likely to be concerned with.

3.7 Conclusions and Future Work

The lack of appropriate micro-scale bicycle modelling tools is a potential barrier to the assessment and quantification of schemes of various types. Such a barrier has been experienced by pedestrian scheme designers (and highway designers) in the past and the potential exists to utilise the basis of some successful pedestrian tools to produce models which are suitable for the modelling of bicycles. In addition, such bicycle models would have the obvious extension for interconnection with established pedestrian models and thus the modelling of more complex shared-space arrangements than is possible with current tools; such as by consideration alongside compatible pedestrian–vehicle models to create combined cycle–vehicle–pedestrian models (e.g. Anvari et al., 2014; Pascucci et al., 2015; Schönauer et al., 2012).

No literature exists establishing the behavioural traits of interacting cyclists and the limited literature that does exist with regard to capacity is in disagreement by orders of magnitude (Table 2.1). Similarly, empirical data of the scale necessary to robustly establish these parameters are not widely available and are expensive to obtain. This is further highlighted by the parameters used for this model and collated in Table 3.2 which almost universally are required to calibrate the model but which are only informed by a limited amount of literature. A clear next step of work from this project is for practitioners to validate the theoretically-justified parameters and outcomes in the context of empirical data.

Work to deliver the Highway Capacity Manual took the work of Botma (1995) and folded those numbers into the fundamental Level of Service measure often used by highway designers. In the process, the caveats attached to any assumption in their production have been lost. However, it is shown here though a modified implementation of the Social Force Model (as first proposed by Helbing & Molnár, 1995) that the core assumption that the Botma (1995) paper is founded upon, results in a fundamental step-change in outputs within the range of values that capacity is considered.

There remains a wide range of research required in order to quantify cycle infrastructure. If there remains an aspiration to produce a robust modelling framework for cyclists (as, for example, the current wave of 'shared-space' research indicates is the case), then the underlying numerical basis must be similarly robust. Large-scale empirical data is required over a range of facilities, and similarly, large-scale studies of individual cycle behaviour are required over a range of demographics and nationalities. These do not currently exist in any form in the literature as they are resource-intensive to collect, however, if cycle modelling is to become as robust as motor vehicle modelling, then they are a necessity. The next chapter (Chapter 4) moves on from the issue of modelling cycle flow to consider the larger issue of large-scale empirical data, and how it might be collected.

Chapter 4

Methodologies for the Collection of Empirical Data

The previous chapter highlighted that even to create a relatively simple model of bicycle movement, a range of parameters are required; and Chapter 2 illustrated the lack of literature to inform those. The standing requirement then is a need to collect both a wider range and a larger volume of data than is currently available. The following chapter considers the high-level methodological routes to achieving this. However first, a brief recap of this document's progress (thus far) against its aims and objectives is presented.

4.1 Progress Against Aims and Objectives

Section 1.3 summarised the aims and objectives of this document. Recalling and considering each in turn:

4.1.1 Aims Already Addressed

4.1.1.1 Aim 1: Demonstrate that the state-of-the-art of the quantitative cycle infrastructure capacity literature is insufficient to meet practitioner needs

Chapter 2 has reviewed and identified the current state-of-the-art with regard to the quantitative capacity of cycle infrastructure. Consideration has been made of both the academic and practitioner literature and where data exists at all, there is a wide disagreement of the capacity of even the simplest infrastructure across two orders of magnitude. Consequently it is concluded that the current literature is not fit for purpose and that it is necessary to return to first principles. This aim has therefore been achieved.

4.1.1.2 Aim 2: Identify the appropriateness of quantitative modelling techniques as applied to cycle infrastructure

Chapter 3 revised the first principles of traffic flow theory, used the literature to reinforce the wider applicability of these to non-motor vehicle modes such as pedestrians and cyclists, and considered the current state of cycle modelling and those aspects from other modelling methods which could be adopted for use in a cycle model.

Once the basis for modelling was established, Section 3.3 then moved on, using the Social Force Model defined for pedestrians by Helbing & Molnár (1995) as a foundation, and developed a cycle simulation model. The model requires a range of parameters to be used in calculation and it is clear (summarised in Table 3.2) that the literature backing for the majority of the parameters is limited, and that others require empirical calibration.

The developed model is used (Section 3.6) to test the underpinning assumption in Botma (1995) which has been indirectly incorporated in the Highway Capacity Manual (Transportation Research Board, 2010) despite the caveats noted therein relating to the assumption that bicycles do not impede one another, something which is intuitively not the case and which contradicts the fundamental principles of traffic flow theory (Section 3.1).

The model outputs indicate a substantive qualitative difference in outcome in the absence of bicycle interaction through speed change than is the case if bicycles can change their speed as would be expected in reality. Furthermore, the results of the model with the inclusion of interactions correspond to expected traffic flow theory principles (e.g. Figure 3.18 versus the speed-flow component of Figure 3.1). However, the quantitative outputs of the model are

a function of those inputs and as noted, these are lacking in literature basis. Consequently, whilst there is clearly practical scope for the use of simulation tools for the consideration of cyclist circumstances, the underlying availability of data is not at a level necessary to inform them.

This aim is therefore achieved in so far as the testing performed indicates that the current state-of-the-art is not yet sufficient to be used in any new generation of simulation tools, but that there is clear value for the aspiration to do so.

4.1.2 Aims to be Addressed

Given that presented in this document up to this point, the lack of availability of basic cyclist parameters (e.g. speed distributions, acceleration profiles, etc.) is a clear barrier to the development (and validation) of new tools to assist the cycle infrastructure practitioner. Further, the lack of data and understanding relating to how cyclists behave – both individually and collectively – also undermines the ability to develop those tools, as does the general lack of empirical data pertaining to cyclists.

Aim 3 of this document is to "Improve understanding of methodologies for data collection pertaining to cyclists", therefore the remainder of this chapter will consider the potential data collection methodologies which can be used to address the identified shortfall in the literature.

This chapter will also frame the structure of the remainder of this document which captures Aim 4: "Develop and apply those methodologies to determine if there is scope for their further development and/or use", and its sub-objectives which are addressed throughout the remainder.

4.2 Methodologies

Chapter 2 of this document demonstrates the limited nature of the quantitative literature with regard to cycle infrastructure. Chapter 3 then went on to establish the current state of the cycle modelling literature and developed a model to test one of the core assumptions underpinning a key piece of the quantitative literature – specifically, the non-interaction assumption in Botma (1995) – starting at first principles, and found that accepting/rejecting the assumption results in fundamental step-changes in outcomes in exactly the ranges of values that practitioners in high cycle mode share jurisdictions are most likely to be concerned with. In both the literature and the simulation model developed, there is a lack of real empirical data and experimentally-validated underlying knowledge relating to the behaviour of cyclists.

Given this lack of data, it is reasonable to seek methods to address this. There are three possible routes to consider:

1. Gather data on the cyclist (in real or experimentally-controlled situations) by direct instrumentation;

2. Gather data on the cyclist (in real or experimentally-controlled situations) by remote observation;

3. Gather data on the cyclist (in experimentally-controlled situations) by substituting reality for a simulated reality.

Noted here as a demonstrative case study, in 2013, Transport for London commissioned TRL Ltd. to perform some off-road trials regarding cycle-friendly (or 'Dutch-style') roundabout layouts for use in London. TfL could not do these trials on real junctions as the layouts are unproven in the UK, the cost and disruption involved in the changes would be enormous and the safety of the design would have to be demonstrated to the Department for Transport before they could be rolled-out (Greater London Authority, 2013). TRL constructed a full-scale roundabout (on their private site in Crowthorne, Berkshire, UK) with a number of different marking and layout configurations and performed a range of tests; the results are presented in Yor et al. (2015). This project, its scale and associated cost, is demonstrative of the current methodology by which such options could be optioneered and tested in a controlled circumstance.

The data collected in the TRL study was video-based observation of priority violation (to establish if the layout was working as intended), self-report questionnaires from participants, and focus groups. They did not directly collect data relating to the speed of cyclists or their paths taken; to do so would have required the deployment of physical measuring equipment (such as pneumatic tube sensors or hand-held speed guns) or the use of personnel either for live observation or off-line review of the video. In any case, the cost involved

would have escalated further. Additionally, note the sunk costs, even before users are involved, have included the cost of building an entire roundabout (albeit in part from temporary materials).

There is clearly then a standing need for more cost-effective methodologies to capture pertinent cyclist data. The remainder of this section considers the three categories of methodology noted above and sets out a way forward for this project.

4.2.1 Direct Instrumentation

Direct instrumentation can provide practical data on the cyclist. Parkin & Rotheram (2010) used 'on-board' GPS equipment to derive values for speeds, accelerations and assess the effect of gradients on those. Parkin & Meyers (2010) used a bicycle-mounted camera to collect data on the lateral passing distance of motor vehicles with regard to the bicycle on UK roads; Llorca et al. (2015) and Garcia et al. (2015) used instrumented bicycles – fitted with laser rangefinders, GPS trackers and multiple video cameras – to collect similar data for motor vehicles and other cyclists in Spain; and Mehta et al. (2015) used ultrasonic range sensors fitted to a bicycle to similar effect. Vansteenkiste et al. (2013) and Vansteenkiste, Van Hamme, et al. (2014) both focus upon the riders' eye view and how this varies with speed/path width and with speed/curve radius, respectively; Dozza & Fernandez (2014) used instrumentation to collect data for kinematic movement parameters and frequency analysis.

These examples are noted as they are representative of both the strengths and weaknesses of bicycle instrumentation. For each participant in those studies, or each observed interaction, data is collected which is robust in measurement and can be supported by a range of parameters (depending on the particular instrumentation used). Instrumented bicycles (and riders) can be taken out onto the public highway to gather real data (e.g. Llorca et al., 2015; Parkin & Rotheram, 2010), or alternatively can be placed in off-road experimentally-controlled situations to gather data instead (e.g. Vansteenkiste et al., 2013; Vansteenkiste, Van Hamme, et al., 2014).

However, all these examples are studies that involve a number of participants in the single digits or low tens order of magnitude and usually of an available demographic (typically university students) which may not be representative of the population at-large. Furthermore, the conspicuous fitting of a bicycle with instrumentation (if not also the rider, such as with head-mounted eye-tracking equipment) is likely to generate a substantial experimenter effect such that the data collected may not actually represent real behaviour. If one wishes to achieve experimental control, then studies must be performed off-road in contrived settings which in themselves may generate unrealistic behaviour in perhaps all but the most extravagant situations.

Ultimately, whilst there may be scope for the exploration of bicycle instru-

mentation, especially given the recent growth in smartphone ownership and the suite of sensors they contain, the main path to collecting more data by these methods is simply to go and do so, building larger data sets and sample sizes, albeit that this is not necessarily cost-effective.

Consequently, direct bicycle instrumentation methodologies are out of the scope of this project which is seeking to focus upon the new development of the underdeveloped methodologies (i.e. remote observation and simulated reality) to enable more effective data collection, not simply wider application of existing methods.

4.2.2 Remote observation

Remote observation moves the equipment to measure the required parameters from the cyclist/bicycle combination, to a remote observer or observation point. A variety of equipment exists for this purpose: from video, radar and thermal detection, to on/in-road sensors (such as pneumatic tubes or induction loops), manual counting and other human observation. These methods are well-used and long-standing for motor vehicle observations. They also translate directly from 'real-life' to experimental deployments, though as with vehicle instrumentation, off-road experimental deployments may have limited transferability to real on-road circumstances unless cost and effort is expended on a scale which is broadly impractical. Bespoke observation surveys and regular counts are a common feature of the transport and development planning processes and live (motor vehicle) traffic information is now a 'stock feature' on most smartphones. This illustrates the raw scalability of remote observation of vehicles, something which cannot possibly be matched at any reasonable cost with individual vehicle instrumentation.

However, when it comes to remote observation of cyclists and cycle infrastructure operation, the options are substantially more limited. Surveys – whether the observers are deployed on-site (with or without sensing equipment) or video data which is scrutinised after the fact – are expensive and time-consuming, and consequently can only be used in targeted circumstances.

Fixed highway equipment has limited application to bicycle detection. Loops installed in the road need to be reliably maintained such that they pick up cyclists and are relatively expensive to install. Temporary pneumatic tube detectors may not be sensitive enough to pick up a cyclist and must be replaced regularly; and automatic number plate recognition (ANPR) and other automated radar/camera detectors either cannot or do not pick up cyclists.

In contrast to the maturity of motor-vehicle sensing networks, the first digital-display cycle counter in the UK was installed in only 2013 (Laker, 2014). Physically sending out observers (for example to use hand-held speed guns to detect speeds) is also costly and has the added problem of an experimenter effect, which may affect the quality of the data recorded.

Specialist proprietary equipment does exist that is designed to (also) capture cyclists but this is still relatively immature when compared to equipment to remotely sense motor vehicles. The Flir Thermicam device range (FLIR Intelligent Transportation Systems, 2016) – marketed as "the world's first integrated thermal traffic sensor" – operates by proprietary thermal image analysis to distinguish vehicle from cyclist (and pedestrian). However, the system operates based on 'detection zones' in the observed image area, whilst more flexible than detection systems such as induction loops or radar sensors and capable of providing speed and flow data for vehicle traffic, is more limited with regards cyclists.

At the leading edge of cyclist detection, machine vision and learning techniques in use in autonomous vehicles currently in development, are not yet developed to the point at which they can be used reliably (and verifiably) in urban traffic (e.g. Levin, 2016).

In summary then, the current (limited) methods in hand for the remote observation of cyclists are inadequate, costly (or of sufficient cost that they are simply not used), or unreliable, and yet the ability to observe cyclists remains a potentially valuable source of information; both for the practitioner and the modeller. Focusing on the more promising, video-based systems provide the scope and capture of data which is required, the standing issue being able to convert that visual information to useful data, in a timely manner. As noted, machine vision techniques are a promising field of research as they can be applied to stereoscopic (i.e. three-dimensional) imaging and can cope with movement of the camera (for example, if it is fixed to a moving vehicle); as such this is a field of active development with regard to autonomous vehicles.

However, in the push to advanced machine vision imaging techniques, the immediate needs of the practitioner are not being addressed. Video surveys are already in use in a wide range of projects as a documentary survey technique; however, analysis of the video usually entails a human to observe the video to make counts, or perform more complex operations like timing movement to establish speed of motion. In the scale of project for which a video survey might be used, the cost of the staff member/surveyor is non-trivial and this cost expands should more complex data be required to be extracted from the video. This applies whether the video is collected from a bespoke/commissioned survey, or simply from the acquisition of existing video data sources, such as CCTV.

If therefore, a less 'human-intensive' method were available, the cost-effectiveness of video surveying would be improved, and the scope and extent of data available would increase. To that end, the component operations required to achieve this are well-established. For example, conversion of a known two-dimensional field of view (i.e. an image) to a three-dimensional plane (i.e. the ground) follows a mathematical process called a homography tranformation,

and established image filtering processes can isolate moving foreground objects from background objects. Indeed, the leading edge of research in that field relates to the efficiency of those computations and their application to complex domain problems such as automated driving, not the basic principles.

In this case, chaining those operations appropriately with simple physics calculations can then yield spatial, speed and acceleration data. Furthermore, by doing so in an automated, repeatable, non-proprietary, verifiable and deterministic manner (something which is not currently possible) – in that the video (and some spatial layout information) is provided as an input, and bicycle velocity data is output – both the utility of this methodology is demonstrated, and the process has immediate usability for the practitioner. Chapter 5 develops and applies such a methodology to a study involving real cyclists, demonstrating the scope for the wider use.

4.2.3 Simulated Reality

Simulators have a proven history of positive contribution to user training. By placing an operator in an appropriate simulator, the operator can be exposed to real and contrived scenarios which enable them to learn and trial different strategies of behaviour such that they are capable of proficiently handling real circumstances. Simulators have long been used to train astronauts (Woodling et al., 1973), pilots (Airbus SAS, 2016) and drivers (OKTAL, 2016). A simulator also allows the (repeatable) testing of scenarios that would not be possible in reality (owing to cost, practicality or safety).

Simulators are also of great value in research. In many respects, they share a lot of properties of instrumented vehicles discussed above (Section 4.2.1), but by replacing reality with a virtual re-creation, circumstances and parameters can be controlled and situations can be experimentally repeated. In fact, this is their greatest strength in comparison to the methods outlined above, as high-fidelity experimental recreations of real on-road circumstances are possible in a way that is simply not practicable with the other methods.

Cycle simulators are extremely rare in the literature and elsewhere. A strand of studies using a bicycle simulator at the University of Iowa appear in Grechkin et al. (2012); Plumert et al. (2004, 2011). The simulator therein consists of a bicycle on a fixed turbo-trainer in a three-walled (each 3.0m by 2.4m) projection room. All three papers focus exclusively upon the use of the simulator for such for the study of child (and to a lesser extent adult) on-bike motor-traffic crossing behaviour. However no information on the methodology for the detection of the bicycle's state is given; and no details are provided as to the design development process for the simulator itself. The only information available with regard to any of the design aspects of their simulator are some citations of other (limited) research relating to the 'in-house' virtual environment tools used in the facility (e.g. Cremer et al., 1997; Willemsen et al.,

2003). These sources do not provide a basis upon which to develop a simulator.

A very recent publication captures the validation of a virtual reality bicycle simulator at Monash University, Australia (O'Hern et al., 2017). This simulator consisted of a virtual reality headset for the user who uses "instrumented road-bicycle that has its rear wheel mounted to a bicycle trainer". No information relating to the instrumentation or trainer is provided; nor is any information related to the simulator architecture, content or design methodology.

One alternative to the academic literature was found on a personal blog of a hobbyist. Yan (2016) presents the development of a basic virtual reality bicycle simulation. The simulator utilises a rear-wheel rotation count detector connected to a PC running a 3D simulation, which is then delivered to a smartphone mounted in a Google Cardboard headset to provide immersive visualisation. The system is set up to move along a predetemined path at a rate dictated by the computed speed of the bicycle. Consequently, the only control the user has within the simulation is the rate of forward motion. As a result, a system in this form would not be of use for more complex simulations, however this source has the benefit of substantive detail with regard to the hardware interfaces utilised.

Finally, some products are available commercially that apply elements of bicycle simulation. A number of companies such as Bkool (Bkool, 2016) and CycleOps (CycleOps, 2016) offer commercial products that attach to the user's bicycle and, by using a complementary display surface such as a television (these products are mainly targeted at the home/enthusiast market), provide a virtual environment through which the user can control an avatar cyclist and 'race' with local 'AI', other human opponents (on similar devices connected across the internet), or themselves. The mechanisms of operation are generally similar to the 'hobbyist' example detailed above, though some systems incorporate some element of lateral steering into the simulation in a proprietary manner (e.g. Tacx B.V., 2016e). In all cases, the commonality is that they are focussed on the interest and enjoyment of the participant in a training context, as opposed to reflecting realistic and reproducible behaviours in a highway-based environment (as is the desire in this circumstance).

Otherwise, cycle simulators are not an explored area. Furthermore, in none of the above cases is detail available to allow a simulator of sufficient fidelity to be created and then used for replicable and repeatable experimentation. Alternatively, if one were simply to try to use one of the aforementioned proprietary systems, the experimenter would be limited (or possibly unable) to define the environment in the simulation beyond those options provided by the manufacturer. In addition, the experimenter would remain compelled to validate behaviour in the given simulator against reality if it were wanted to be used to observe simulated 'real' behaviour. The proprietary nature of such systems precludes a detailed examination of the inner workings of such a simulator and

additionally, reduces or removes the ability to make calibration or experimental adjustments.

Consequently, if (per Aim 4b) the value of simulators is to be established, there first remains a need to draw together simulator background material from the more advanced simulator categories (e.g. driving) to inform and define a specification for a bicycle simulator. There is then a need to openly define a replicable simulation software environment and hardware interface for use with a bicycle to meet that specification. Finally, there is a need to test such a simulator to ensure that it is fit for purpose. These items will be therefore be returned to and addressed later in this document in Chapters 6, 7 and 8.

4.3 Going Forward

To recap, this chapter sought to address Aim 3 of this document (Section 1.3) which is to "Improve understanding of methodologies for data collection pertaining to cyclists". The potential methodologies being: individual bicycle instrumentation, remote observation, and simulation of reality. The chapter then covered the state-of-the-art for each of these methodological categories, establishing the scope for the chapters that follow.

Of those methodologies, direct instrumentation of bicycles is to be excluded from further consideration for the purposes of this project. If the interactions and behaviours of cyclists on a large scale are to be better understood, the sheer volume of bicycles that would need to be instrumented, and/or the sheer level of instrumentation required for each bicycle, is simply not scalable. Extraneous factors such as the behaviour of other road users are difficult to capture systematically and the equipment required to do so is so invasive as to impose a substantial experimenter effect on the cyclist and other road users. Notwithstanding the objective value of instrumented bicycles, this lack of practically-valid scalability is reasonable grounds for the exclusion of direct instrumentation from this project.

This therefore leaves two remaining avenues for consideration: remote observation and simulated reality. Simulated reality will be considered in later chapters (Chapters 6, 7 and 8); the following chapter (Chapter 5) will focus upon remote observation. The reader is also directed to Figure 1.1 for a graphical illustration of this structure. By exploring these aspects, this document serves to expand the pool of methodologies available to both the academic and the practitioner for the collection of cyclist data of more detail and with a greater ability to experimentally replicate and verify.

Chapter 5

Acquiring Empirical Data by Remote Observation

5.1 Introduction

The TRL study discussed in the previous chapter (Section 4.2) was, in effect, a microcosm of the state-of-the-art for experimental off-road data collection. Excluding the interview groups and questionnaires (which cover subjective data we are not immediately concerned with here), the focus was mainly on video which was then analysed by humans post-study. The fact that video captures a large range of information which can be reviewed off-line is of great value. This is one of the reasons why transport professionals will often choose video surveys for their projects, as sometimes client (or project) requirements change and one does not wish to have to go out and do a fresh survey.

This was also the motivation for the work completed in this chapter. A study with real cyclists was organised to collect data pertinent to their behaviour such as speed, acceleration etc. The range and volume of data required would have been prohibitively time-consuming to analyse manually, therefore an alternative was sought.

In mid-2014 when the work was planned, there were no commercially-available packages that could analyse video as required by this project in an automated manner. It was therefore necessary to develop a process to do so.

5.2 Computer Vision Applied to Bicycle Observations

Observing the movement of a bicycle on video in an automated manner is not something which has been captured in the literature. Video observations of cyclists that are then analysed manually is not uncommon (within the broader limited category of cycle observations in the literature at all) but analysis is completed manually (e.g. Z. Li et al., 2012).

As has been established in Chapters 2, 3 and 4, there is a standing need for empirical data relating to cyclists. Video provides a good candidate process for capturing data pertaining to cyclists, however at the time this work was completed (2014), analysis of video was time consuming, and thus costly. Indeed any video may capture cyclists and there are a huge number of CCTV cameras already distributed across the road network for routine monitoring (not to mention the ability to deploy specific equipment for the purpose). If methods then existed whereby the video data could be analysed automatically, then the cost (and consequently the main barrier to the collation of more information) could be reduced.

A wide range of applicable well-known processes and algorithms exist in the field of computer vision. In fact, automated vehicles currently coming to market depend fundamentally upon them. However, at the time this work was completed, these had not then been applied in an automated way to cyclists. Essentially however, a bicycle is simply a moving object on video which is a process that has been widely explored. It was therefore the case that if this could be successfully achieved, then the cost and difficulty of obtaining cyclist parameters from existing and future video sources could be markedly reduced. The core operations of video analysis are a relatively mature field in computer science and the reader is directed to a text such as L. G. Shapiro & Stockman (2001) for a comprehensive overview.

The obvious current course of action would be simply to use a software library such as OpenCV (OpenCV, 2014) to apply all the relevant algorithms. However, the OpenCV Foundation received a significant investment from Intel and DARPA in late 2014, after this chapter's work was completed (OpenCV Foundation, 2014), which greatly improved the OpenCV library both in quality and scope of code implementation. Some aspects, in particular the implementation of background subtraction (necessary to isolate a foreground moving object from the background field-of-view; e.g. Kaewtrakulpong & Bowden, 2001), were not well implemented prior to that time, hence the methodology developed below. Additionally, and regardless of the use (or not) of a library, an understanding of the underlying process was critical for the quality of the research to be delivered.

Computer vision and machine learning are rapidly evolving fields and cur-

rent (as of early 2017) practice has overtaken the work presented here. However in any case, the use of a library or proprietary product does not avoid the need to logically construct a toolchain of the individual sub-processes (to some degree) for the use case here. Doing so openly as detailed below also has the added benefit of reproducibility and transparency, which is not the case should one use a proprietary product.

Machine vision applied to cyclists remains an evolving field. For example, Jin et al. (2015) asserts that video observations of on-road cycles were analysed "Using video-processing technology", but no detail of this is provided, and the observation area was augmented with white line markings to enable their analysis. Alternative interventions could include having the participants wear particular items, or fitting bicycles with marks; e.g. Zhang et al. (2013). Notably, the process defined in this chapter did not require an intervention on the cyclists.

In terms of proprietary products, the Luxriot Video Analytics Package (A&H Software House Inc., 2017) allows the capture of cyclist data from video by virtue of classification (e.g. object of given size range is a bicycle) of moving objects across screenlines or detection zones in offline video (or live video when combined at further cost with other Luxriot packages). Ling et al. (2017) used the Luxriot package to determine cyclist path and speed information related to a tram-track crossing point.

An alternative proprietary product has been developed by FLIR (FLIR Intelligent Transportation Systems, 2016) and, mentioned in mentioned in Section 4.2.2, is a thermal-imaging camera that utilises in-image detection zones to establish cycle counts and flows. However, these products are not capable (at time of writing) of providing complex information about the cyclists such as speed or specific paths (detection zones are in the order of at least a metre in given dimension). Classification is achieved by zonal presence in the image, thus should a cyclist not be in the given cycle zone of the image, or the zone be encroached upon by alternative users (not uncommon for highway infrastructure), there will be an impact on the data collected. It should also be reiterated that the use of this product demands the specific installation of the FLIR Thermicam equipment on-site.

Were the work described in this chapter completed at the time of compilation of this document, then it is probable that (should budget allow) one or more of the commercial products may have been selected. The progress that has occurred in the intervening period being substantial. However, even if that were the case, the need to be able to reliably and repeatably perform the analysis required would still be hindered by the closed and proprietary nature of the commercial products. As delivered, the process detailed in this chapter is deterministic and reproducible.

The remainder of this chapter therefore presents an analytical methodol-

ogy applied to a small-scale study to derive cyclist parameters automatically from observed video. Such procedures have the potential to greatly increase the availability of cyclist parameter data from existing and future video data sources. The majority of this chapter was published in Procedia Computer Science (Osowski & Waterson, 2015) through the 4th International Workshop on Agent-based Mobility, Traffic and Transportation Models, Methodologies and Applications (ABMTRANS) in June 2015. The work was also presented at said conference.

5.3 Practical Event and Participation

An experiment was designed to establish procedures which can be applied more widely to derive empirical values, which for example, can then inform the calibration of the agent-based bicycle model utilised in Chapter 3. Whilst direct 'on-the-day' measurements could have been undertaken for this small-scale study, this does not scale effectively to wider-scale data collection. The focus of measurement was instead on video observation. The purpose being to develop techniques after-the-event that could be used to derive the required data and perform the necessary analysis, without the need for direct interaction with the cyclist.

6 cyclists (existing regular cyclists; postgraduate students; 5 male, 1 female; ages 24–32) were recruited for the event (University of Southampton ERGO Ethics ID 12379) which lasted approximately 2h30. A closed car-park, at the National Oceanography Centre (NOC) in Southampton, was laid out with a circuit in a 'figure-8' arrangement (Figure 5.1; overall route length approx. 520m). Edges were delineated by coloured rope on the asphalt surface or by the existing car park kerb infrastructure, and supplemented by coloured training cones. The layout incorporated straight and curved sections of various radii (3–10m) and widths (2–6m).

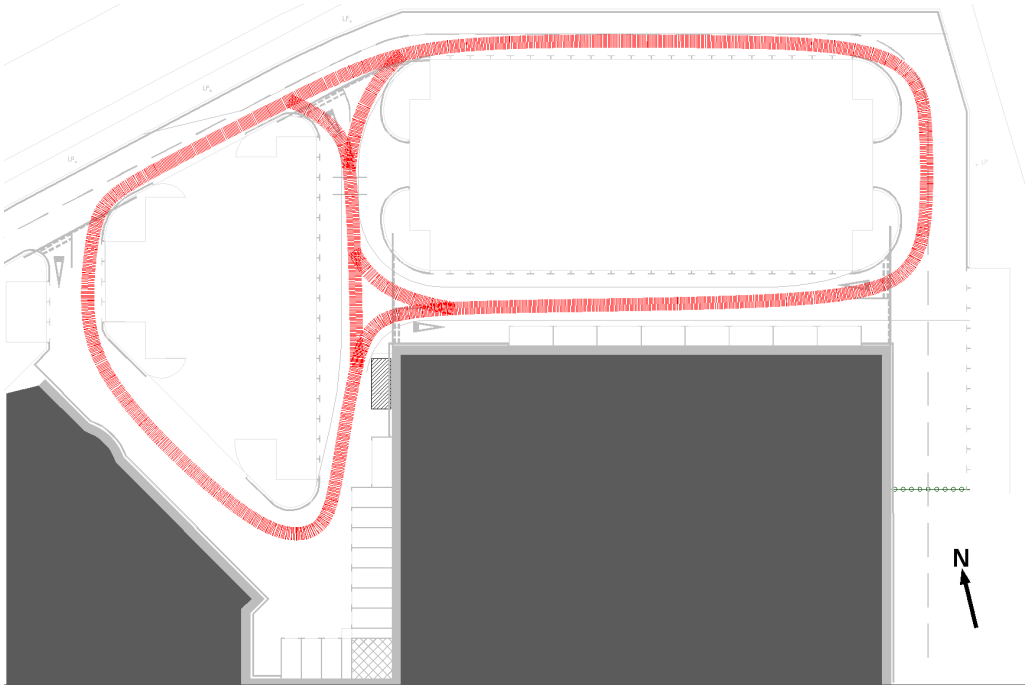

Figure 5.1: NOC General Arrangement plan and indicative route availability

Participants each cycled around the course individually on routeing of their choice for 5 minutes to familiarise themselves with the layout and to collect data related to free-flow and uninhibited travel. Thereafter, the participants all took to the course together for various arrangements of individual and group

cycling. For brevity and with regard to necessary focus on the analytical process, only the individual familiarisation laps are discussed here.

5.3.1 Data Collected

5 video cameras were installed in the adjacent building at approximately 10m height and provided visual coverage of the majority of the laid course (Figure 5.2). The following demonstrative analysis is limited to one of the high-level cameras (Camera 1).

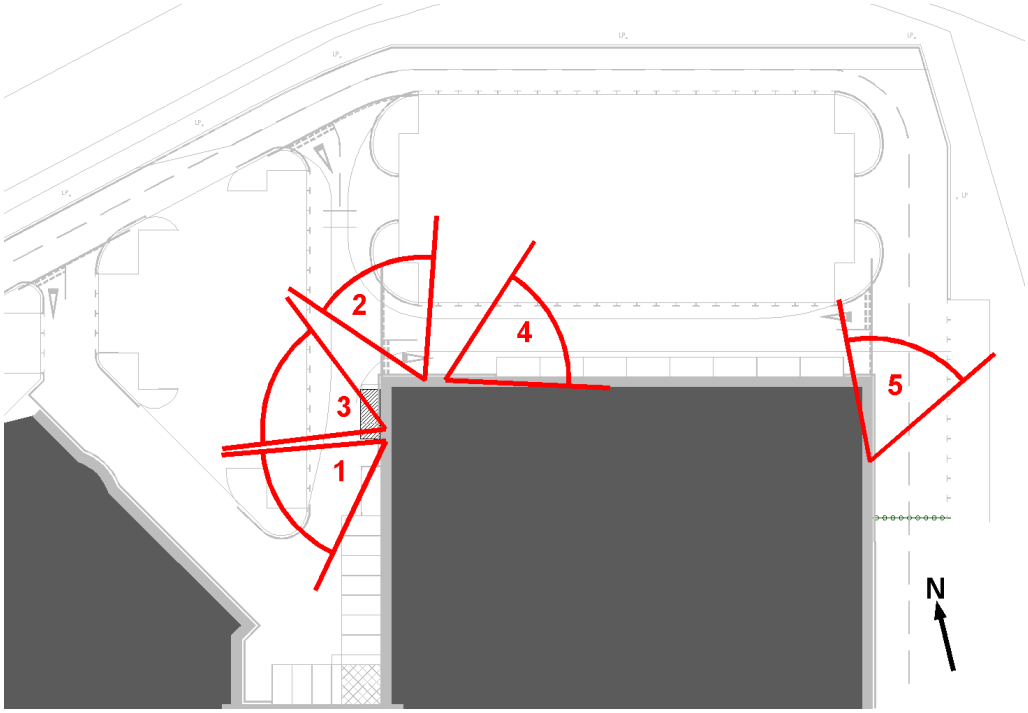

Figure 5.2: NOC General Arrangement plan with indicative high-level camera fields-of-view

5.4 Video Analysis

In order to derive the data desired, a procedure is required which isolates the moving object in the camera field-of-view (in this case the bicycle) and converts this into true spatial and temporal coordinates. Processes which can be used for each of the individual steps required, are well-studied in the field of computer vision with a variety of algorithms, libraries and implementations both possible and available. The work presented here is focussed on the combination of these methods into a 'toolchain' that yields the outcome of cyclist data required. The choice of specific methodology/algorithm within each stage of the process is ultimately a matter for the implementer. Indeed, as technology, software and hardware develop into the future, this is rightly the case. However the overall order of operations and the type of processes required, is defined by this work and is as follows:

Video decomposition Video data files are first broken down into their constituent image frames; for each video (in this case, approximately 3 hours) this corresponds to approximately 270000 frames (videos are recorded at 25 frames per second). FFmpeg libraries (FFmpeg Project, 2014) are used for this process. Subsequent analysis is then performed on a frame-by-frame basis.

Image preparation Frame data is resampled to compensate for equipment differences in pixel aspect ratio and to reduce the image resolution to manageable proportions (generally, 1440×1080 images were downscaled to 720×405 resolution). The image is then converted from RGB values to greyscale (using the ITU HDTV conversion standard (ITU, 1990)) with pixels masked out (set to 0; i.e. black) according to manually-defined areas. This masking was applied to remove unnecessary computation (i.e. areas substantially off the course) and remove those areas of the image that would reduce the effectiveness of the process (e.g. window reflections, tree movement, etc.).

Background subtraction The process of separating foreground and background was implemented broadly as Snowdon (2008). A 'rolling-window' of historic pixel values (over 5sec; i.e. 125 frames) is retained for each pixel between frames. For the frame being processed, the current pixel is compared to the modal value of the pixel throughout the length of the window. Pixel values range from 0 to 255. The modal value is however not a simple mode of the actual pixel values, but is in fact the mode of a composite of Gaussians ($\bar{x}$ = pixel value, $\sigma = 2.0$), one for each frame's pixel value. This process compensates for noise between frames where simple mode values may produce a complex tightly-clustered multi-modal histogram leading to incorrect conclusions as to the real modal value.

However, a side-effect of this is that the algorithm is sensitive to physical movement of the camera. This was not an issue here as the cameras were fixed on tripods, but would be a notable limitation for mobile/vehicle-mounted cameras, or those on less stable mountings.

'Blob' operations A (two-pass) connected-component labelling algorithm (L. Shapiro & Stockman, 2000) was applied to the foreground pixels and each component was then enclosed in a rectangle with a pixel padding (3px) around the component. Rectangles with a dimension below a minimum size (9px) are discarded as noise. Following this, rectangles that overlap are 'unioned' and redefined as a single larger rectangle (Figure 5.3a). This allows the linking of multiple parts of the same real object which have not already been linked by pixels into the same connected-component (e.g. a bicycle's wheel and the rider may not be identified as one contiguous foreground component but should be combined for the purpose of analysis). However, also note that this is not limited to parts of the same object and thus where a second moving object is close to, or partially occluded by the first, it may be combined into a larger single object rectangle. This is not an issue here as the density of cyclists was low and the angle of the camera high, but would be required to be addressed where density of cyclists were high or occluding objects were present in the field of view.

(a) Foreground object rectangle enclosure (Participant 1 at Camera 1)

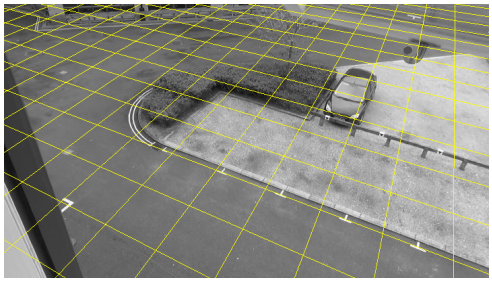

(b) Camera 1 image overlaid with computed homography ground plane (2m grid)

Figure 5.3: Demonstration images

This process yields a list of frame-indexed rectangle objects which are then analysed as follows:

Centroid linking Each rectangle object has a centroid (geometric centre). If a centroid was found in the previous frame within a given radius (10px) then this is considered to be part of the same path. If no nearby centroid was found, then a new path is created. Whilst naïve, this method is effective for sparse foreground objects at reasonable frame rates. Upon completion of this process, paths shorter than a specified length (25 frames; i.e. 1 second) are removed as these are likely to be noise.

Homography translation The use of a homography matrix allows the transformation of points between two different two-dimensional planes in three-dimensional space. In this case, a homography matrix is computed for each camera (using the OpenCV library, (OpenCV, 2014)) to convert two-dimensional points in the plane of the camera image (i.e. pixel coordinates) to two-dimensional points in the plane of the ground covered by the image (and vice versa; e.g. Figure 5.3b). Ground coordinates were manually matched to the image using CAD drawings of the car park site. Path vertices were thus able to be converted from image to real coordinates by use of this matrix.

The output of this process is participant paths located in both (frame-indexed) time and space; the spatial arrangement of which is plotted in Figure 5.4. A demonstrative exploration and discussion of the data follows below. For a more comprehensive analysis, the reader is directed to Chapter 8, where the data from this study is considered alongside the data collected on the bicycle simulator tested in that chapter.

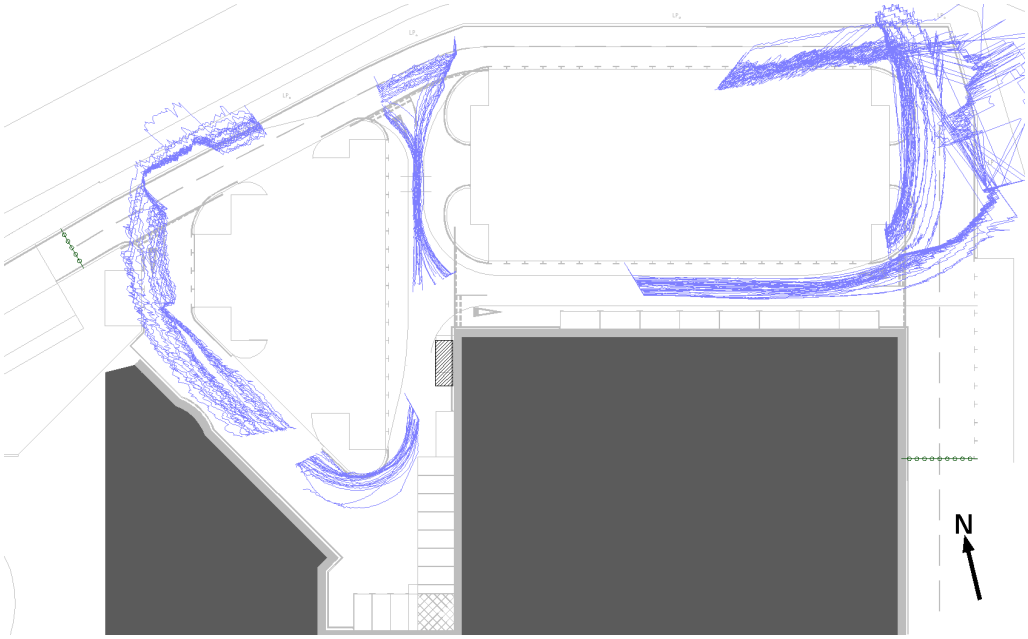

Figure 5.4: Plot of all participants' paths taken in Real Life situation

5.5 Analysis of Video-derived Spatiotemporal Data

Participant paths located in real spatial and temporal coordinates allows, as an example, the analysis that follows. This process is of interest as it yields broadly precise information without any need to interfere with the cyclist. Whilst this experiment was performed under controlled circumstances, subject to issues regards the quality of non-experimental data, the process is applicable to real cyclists (and indeed other road users) and therefore provides the opportunity for wide-scale collection of new data from existing datasets and data collection streams. Such data is currently not widely available or requires relatively costly intervention with individual cyclists (e.g. Parkin & Rotheram, 2010; Vansteenkiste, Zeuwts, et al., 2014), usually resulting in small sample sizes.

The following demonstrates the usefulness of the data derived from the analysis of the experiment's video data.

5.5.1 Spatial Distribution

Superimposition of paths onto a base plan of the course allows a visual inspection of the distribution of paths resulting from the analytical procedure above for the whole NOC experimental area (Figure 5.4), and for individual parts thereof (Figure 5.5a).

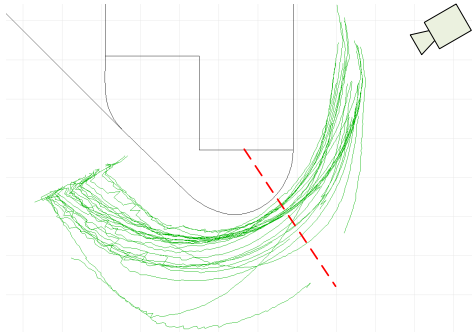

(a) Paths followed by participants on their individual laps (2m grid) around fixed kerbed area with indicative camera location. Cordon line indicated.

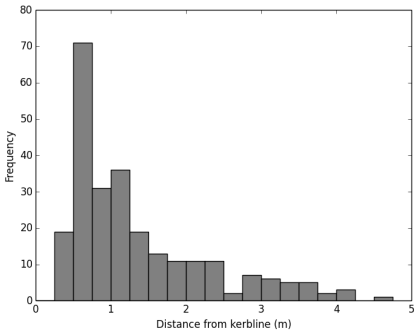

(b) Frequency distribution of paths across section through course taken by all participants throughout the event, inclusive of group laps.

Figure 5.5: Analysis of path distributions at Camera 1

Visual inspection of paths can, at best, provide qualitative insights as to the cyclist's behaviour. What is instead desired is quantitative data.

Of relevance to the development of a valid agent-based model considered in Chapter 3, is the spatial distribution of cyclists relative to fixed obstacles, and the spatial distribution relative to one another. A quantitative understanding of these distributions could (for example) allow the fitting of a probability function and by extension, a force profile for the given object. Consider therefore

Figure 5.5.

Using Chapter 3 as a frame of context, the original Social Force Model assumes force profiles with a negative exponential distribution (given the model is based in particle and field theory). Such force profiles imply that in infinite space, an agent would be repelled at whatever distance away (from a boundary or another agent) it was. Practically, this may be a suitably accurate modelling approximation, but the distribution of agents that results may be better approximated by a more truncated functional form.

Figure 5.5b illustrates the distribution of cyclists across a section outward from the radius of the bend throughout the experimental period. Whilst the cyclists obviously experience a repulsion such that they did not contact the foliage on the interior of the radius, beyond that point the repulsion (if any) would seem to be of rapidly diminishing importance and perhaps more rapid than would be suggested by an exponential repulsive force. However, further analysis is required before conclusions can be drawn, owing to the limited sample size here, and the occlusion that occurs when multiple cyclists are travelling as a group.

That said, where the homography transformation has been properly calibrated, and where camera angles are appropriate (such as the data presented here; e.g. Figure 5.5a), it is clear that the spatial (and temporal) data which has been collected is of a good quality. Figure 5.4 however shows data from all the cameras collected and illustrates the sensitivity of this process to inaccuracies and imprecision in the homography transformation process and the loss of both that develops as a result of this and of shallow camera angles. It should however be restated that those errors are a function of the data collection set-up and calibration, and not an inherent property of this process.

5.5.2 Participant Speed Profiles

Spatially and temporally-located paths allow the derivation of speed information (i.e. distance moved per unit time). Cyclist average speeds are generally considered to fall in the range 4.0ms^{-1} (CROW, 2007) to 6.0ms^{-1} (Parkin & Rotheram, 2010) but sources are limited.

Once smoothed to reduce the effect of sampling noise, Figure 5.6 shows the relatively consistent speeds selected by Participant 1 as they pass through the camera's field of view. Direction of travel is included for reference as the increased noise in the data measured further from the camera is apparent. Indicated time gaps are those periods of time where the participant was elsewhere in the course than within the field-of-view of this camera.

Whilst some smoothing is necessary (both the raw and smoothed data is shown in Figure 5.6), resulting data is generally in the range expected with computed speeds average around 5ms^{-1} ($\bar{x} = 5.25$ms^{-1}, $\sigma = 0.61$ms^{-1}, $n = 568$). The increased noise in the data measured further from the camera is

visible in the raw data but following smoothing is much less extreme and could be further improved by more effective smoothing methods, such as a Kalman filter.

The possession of speed data allows some more interesting exploration of the data to be considered. Figure 5.5b illustrated the distribution of distance from the inner kerb of the corner covered by Camera 1. If the speed information for each of those crossings is derived and plotted against distance from the inner kerb, then the result is Figure 5.7. A reasonable hypothesis is that speed chosen to travel around the corner may be correlated with the distance chosen.

Although visually, the data seems to indicate a limited trend of decreasing speed with distance from the inner kerb, with an $r = -0.294$, this data indicates that the two are not significantly linearly correlated at the 5% significance level ($p = 0.129$; $n = 28$). It is of course possible that a larger data set might indicate that the relationship is significant, or that it may be substantially non-linear, however this is difficult to establish from this small data set.

Noise aside, the value of this process is demonstrated by the outputs as the speed of a cyclist has been measured here without any specialist equipment or any direct interaction with the cyclist. The relevance of this data for establishing individual speed profiles, and more developed parameters such as the relationship between speed and distance from a corner has also been demonstrated. Future work can therefore expand the scope of data collected and refine the 'noisiness' of the resulting data.

5.5.3 Participant Acceleration Profiles

In the same way that the rate of position change provides speed data, the rate of change of speed data allows the determination of acceleration (and deceleration) parameters. Literature acceleration figures for cyclists are even harder to find than those for speed – a range $+1.0\text{ms}^{-2}$ to -1.5ms^{-2} is indicated in CROW (2007) versus a maximum observed acceleration of $+0.71\text{ms}^{-2}$ in Parkin & Rotheram (2010) – and lack data reflecting the likely changes in acceleration behaviour with regard to current speed. For the development of robust microsimulation models, an understanding of cyclists' acceleration behaviour will ultimately be necessary.

Figure 5.8 shows the acceleration profile derived from the speed data for Participant 1 at Camera 1. This data is substantially more 'noisy' (given it is a function of the speed data), however smoothing produces acceleration values around 0.0ms^{-2} ($\bar{x} = 0.09\text{ms}^{-2}$, $\sigma = 3.67\text{ms}^{-2}$, $n = 568$) indicating a relatively constant speed choice, which is broadly corroborated by the video observations of the participant and the relatively slow underlying changes in speeds shown in Figure 5.8. Other sections of the course were involved in interventions that would have deliberate substantial speed changes (and thus non-trivial accelerations), but these are not covered here.

Albeit an identified problem of lesser magnitude with the computed speed data, the noise in the measurements of acceleration are a major impediment to their practical use. The noise inherent in this acceleration data yields values multiple orders of magnitude greater than those indicated in the literature noted above, making conclusions difficult to draw. Further work should seek to apply more intelligent filtering to the location observations (e.g. a Kalman filter) such that unnecessary noise due to the video observation does not become amplified in the subsequent analysis. Also relevant, more precise location information (such as if the video recording and analysis were performed at a higher image resolution) would likely have the effect of reducing this computed variability.

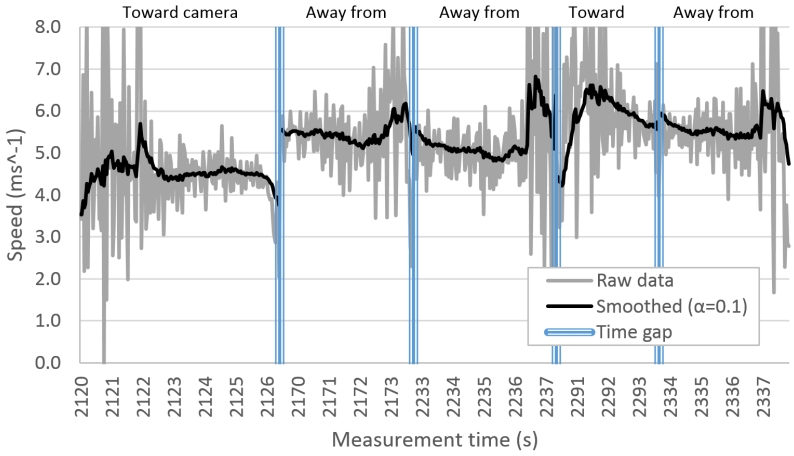

Figure 5.6: Speed profile derived from video analysis of Participant 1 on Camera 1

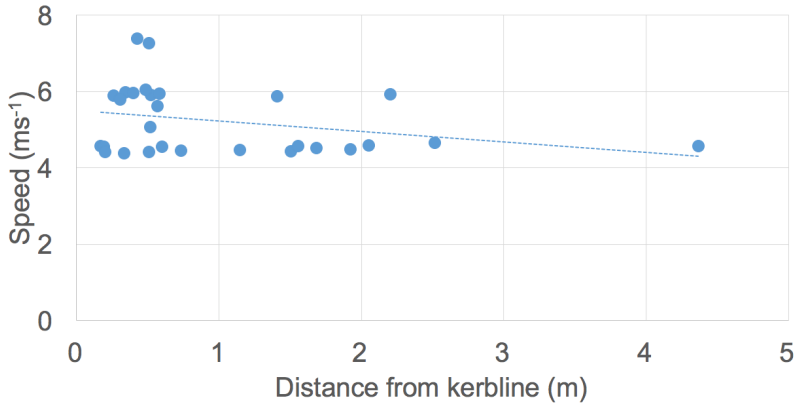

Figure 5.7: Plot of all participants' individual laps' cordon crossing speed vs. distance from kerm

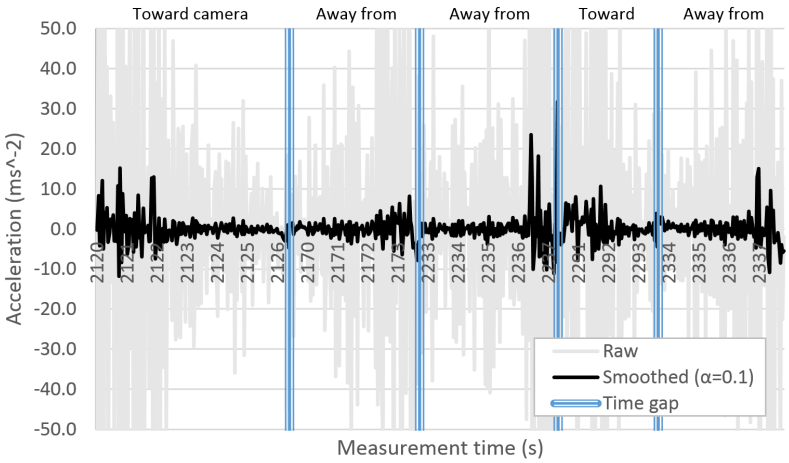

Figure 5.8: Acceleration profile derived from video analysis of Participant 1 on Camera 1

5.6 Implications

5.6.1 Outcomes

The methodology in this chapter presents a process for the extraction of key parameters for the cyclist from simple video imagery. The methodology is widely applicable not only to bicycles but to any object that moves through a fixed-camera's field-of-view; common examples in transport being pedestrians and other road vehicles. Consequently, the process provides an opportunity to massively increase the pool of transport data collected whilst simultaneously reducing the costs for doing so. Given that collection of empirical data is a relatively costly process, this is to be welcomed.

5.6.2 Limitations

As it stands, there are some technical limitations with this methodology. However, the concept is proven and the taking forward of this work to address the shortcomings, is now easily specified, something which would not have been the case previously (the process in a proprietary product being a matter of commercial sensitivity for example). The following points of note will assist in the specification of such a programme.

The primary limitation of this methodology is the run time. A run of this methodology for a given video takes in the order of 50–100× real-time. For offline analysis this is essentially unimportant as one can simply start the process and walk away. However, this precludes its use in its current form for the analysis of real-time video. Whilst there is clearly a huge scope for usage otherwise, this would be a desirable end goal. Sufficient general hardware resource is unlikely to be available to address this overhead in the next decade (if one assumes a 'Moore's Law' doubling of compute power every 18 months), so the algorithms must be refined. The current bottleneck is the background subtraction and so this should be the primary focus of an efficiency effort.

The other major difficulty (again a function of the background subtraction algorithm) is the sensitivity of the process to movements of the camera and consequently the entire field of view. As a background is established on a pixel-by-pixel basis, it does not require much movement to cause dramatic shifts in background values for a pixel, especially around sharp contrasts in the image (e.g. kerblines). This may not be a huge issue for fixed cameras but given that most cameras mounted 'in the wild' are atop poles, a wobble of the support pole might be enough to 'upset' the background subtraction algorithm and cause fixed objects to be identified as objects that have moved. A digital image stabilisation algorithm could compensate for such motion prior to any background subtraction, however given the background subtraction algorithm itself must be addressed in an effort at improving computational efficiency, this should also be borne in mind.

An issue which did not arise in this study but is worth noting, is that this imaging was performed with visual cameras in the daytime. There is no reason the same could not be done with infrared-band images (as used, for example, by the FLIR Thermicam equipment) and consequently at night time. However, both methodologies would be sensitive to precipitation – indeed this is one of the primary methods for satellite imaging cloud cover and rainfall – and thus this is a potential limitation to this methodology, though potentially not an insurmountable one (e.g. Garg & Nayar, 2007).

The most obvious improvement to be made, and one which can be done within the scope of the existing methodology and without further development work, is to use multiple cameras to cover the same field-of-view. This would allow the systematic errors in homography to be reduced, improving spatial accuracy and reducing 'noise' in the measured values. Furthermore, whilst also producing more accurate data, an improved precision to the computed data will also reduce the noise present in the derived data (speed and acceleration) and reduce the attendant need for post-processing.

5.6.3 Conclusions

The new process established here has taken proven concepts in image recognition and applied them in a practically-useful way to get precise empirical data from cyclists solely from video observation. This presents huge scope for the expansion of the scale of data which can be collected from existing and new video sources for cyclists and other modes, and does so in an open manner not available from proprietary products. It also substantially reduces costs associated with collecting such data, something which is highly attractive to the practitioner.

In Section 4.2.2, the need for such tools was highlighted. As a result of this established proof-of-concept, there are now clear, tightly-scoped improvements to be taken from this chapter to improve the quality of the results, reduce the computational cost in their acquisition – though it is reiterated that the cost of computer time is a *fraction* of that of a human to perform the same task – and extend the circumstances to which the methodology can be applied. That said, data such as speeds and spatial placements, not to mention simple counts, are all possible with acceptable accuracy from the work presented. Therefore, even in its current form, the work presented can still be applied to real data to derive the kinds of parameters that practitioners require to inform design.

As discussed in Section 5.2, practical techniques have advanced since the work in this chapter was completed and thus practitioners would be wise to consider the potential for the available proprietary products at their time of use, however, the methodology presented in this chapter is openly-defined at each stage and therefore remains of use both as an alternative and/or an augmentation for commercial and/or proprietary products.

This study illustrated some of the challenges in producing controlled and replicable empirical situations. As noted in the previous chapter (Chapter 4), controlling and replicating experimental situations is one of the core strengths of simulators. Consequently, the next chapter (Chapter 6) moves on to explore the potential for a bicycle simulator.

Chapter 6

Bicycle Simulator Design Literature and Specification

Chapter 2 established the lack of availability of empirical cyclist data and the lack of fundamental understanding of cyclist behaviour. Chapter 4 then framed the potential methodologies for collecting such data, with Chapter 5 moving on to present a new methodology for remotely collecting cyclist data, using automated analysis of video. The wider application of this methodology will yield a more comprehensive gamut of behavioural parameters, but is potentially limited by the ability of the practitioner to experimentally control the situation to which the cyclist is exposed; a particular problem if one wishes to experiment on the public highway.

This is where the use of a simulator could assist. Simulators provide various practitioners with a range of benefits. They allow training of users in controlled, replicable, and safe circumstances where failures can be accommodated without risk to life or equipment (e.g. Schwebel et al., 2016), and failures or difficult situations can be identified and addressed. Simulators also allow the testing of things difficult or impossible to do in reality such as scenario testing in air crash investigations (e.g. Bundeshaus Nord, 2001), pre-flight training for space missions (e.g. Woodling et al., 1973), testing of driver impairment due to alcohol or drugs (e.g. Helland et al., 2016), and experimental replication of repeatable and controllable scenarios.

In this use case, a simulator: allows the set up of experimental testing and replication in a controlled laboratory condition; would provide a capability to ex ante test schemes and scheme options; and would enable off-line quantification of real circumstances in a simulator context. More generally, it provides a greater potential to further interrogate and understand *why* cyclists behave as they do and therefore support data showing *how* cyclists behave (such as that collected by the methodology established in Chapter 5).

Chapter 4 (Section 4.2.3) discussed the limited literature that involves the use of bicycle simulators. In all cases, the design of the simulator systems used are under-defined or simply provided (in part or entirely) as a given, under-

mining the ability to replicate the given studies and deliver new ones. What is therefore needed before a simulator can be constructed and used, is an examination of the background for how one might create a simulator, and a defined specification such that one can then be constructed. In line with Aim 4b: "Identify the scope for, design, construct and test an immersive simulator as a data collection method for bicycle parameters", this chapter therefore explores the design literature, expanding to the general simulator literature in the absence of cycle-specific literature, and presents a specification for a bicycle simulator. This specification will then be used to construct and test a simulator in the chapters to follow.

6.1 Simulator Design

6.1.1 Emulating real life behaviour

Objective fidelity (the replication of the physical characteristics of the environment and vehicle) and perceptual fidelity (the degree to which the user behaves and feels as they would in reality), or more simply the level of realism, is considered to be an important determinant of the effectiveness of a given simulator (Groot et al., 2011). Considering both driving and aircraft simulators (to take two well-studied examples), these range from simple interactive simulations one uses whilst sat at a desk, to 'full immersion' purpose-built pieces of equipment. The more premium of these simulators are constructed on motion platforms which provide the user with a real perception of motion as well as realistic vehicle controls and a wide field-of-view coverage.

The placement of the user into a vehicle cockpit goes a long way to dealing with the issues of objective fidelity. If the interaction of the user is by the same means as they would interact with that vehicle in reality (i.e. through using a flight yoke or turning a steering wheel) and their view of the world is the same as in reality (i.e. three-dimensional and covering the available field-of-view afforded by the vehicle), then a level of 'mundane realism' (the extent to which an experiment is similar to everyday life; Blascovich et al., 2002) and an objective fidelity can be said to have been achieved.

Perceptual fidelity is more difficult to quantify because to do so would depend on an understanding of real-life behaviour, its motivations and the ability to compare the two. Where real data is available, such as those provided by an aircraft flight data recorder or an accelerometer sensor placed in a car, then such comparisons can be made. Simply, the fidelity can be numerically validated by replicating a known situation in the simulator, comparing the measured behaviour in the simulator with the data from reality, and finally establishing if the two are significantly different from one another.

That said, replicating real behaviour doesn't necessarily mean the simulator is providing fidelity; if the underlying simulation or the interactions/perceptions thereof are not suitably realistic, then the behaviours measured in the simulator may reflect artefacts of the control system (e.g. see the discussion in Tydeman, 2004, with regard to flight simulators' use in accident investigation). The need then for the control system to become 'transparent' (the objective fidelity) and the behaviour to be realistic (the perceptual fidelity), are clearly closely interwoven. If a simulator user believes themselves to be present in the simulation, if their actions and interactions are the same as would be performed in reality, and their engagement with the equipment (both use and perception) is as would be the case in reality, then the simulator may be sufficient to provide useful data.

6.1.2 Automatic brain actions

Amongst the many skills of the human species, the bipedal walking ability, is perhaps one of the most defining. Walking requires little in the way of concious effort and less so the biomechanical nuances required to do so. This is highlighted in circumstances of injury rehabilitation, as a key component of the process is the concious processing of the biomechanical movement patterns, or more commonly in novel situations, such as stepping onto a slippery surface (Malone & Bastian, 2010). The ability to walk without substantial cognitive effort frees us to do other things at the same time (e.g. socialise, hunt, navigate).

These (semi-)automatic skills can be deeply learned and situationally nuanced. The 'broken escalator effect' – the strange sensation and imbalance one experiences when stepping on to an escalator which is switched off – is widely experienced and has been shown to be quickly learned (Reynolds & Bronstein, 2003). This type of effect only then becoming apparent when the unusual circumstance of the broken escalator is experienced.

Similarly, the skill of driving also requires varying levels of driver cognitive involvement depending on the familiarity with the vehicle, situation and other factors external to the vehicle. Anecdotally, this can be observed by the ability of a driver to hold a conversation whilst driving, but possibly trailing off into silence when approaching a junction or needing to make a manoeuvre where cognitive load increases. In a perhaps extreme demonstration, 'highway hypnosis' is a widely-reported phenomenon wherein a driver can find themselves having driven a great distance, properly and inclusive of necessary manoeuvring, and yet have no recollection of how they arrived there (Cerezuela et al., 2004; Williams, 1963).

More generally, this type of effect could be referred to as 'affordance'; a term coined by Gibson (1979). An affordance is "a specific combination of the properties of [an object's] substance and its surfaces taken with reference to an animal". It is for example, the reason why humans 'know' to pull a vertical door handle, push a horizontal one, or hold a cup by the handle (Extra Credits, 2014a). Indeed, some of Gibson's early work was based in observations with regard to the psychology of driver behaviour and was an early proponent of Lewin's field theory (Lewin, 1951, which is a basis for the SFM used in Chapter 3) with regard to vehicle motion (Gibson & Crooks, 1938).

The human ability to perform complex behaviours without substantial cognitive load by virtue of being placed in a given situation can be leveraged here. It could reasonably be argued that if one is put on a bicycle (providing an objective fidelity and an opportunity for a behavioural affordance in terms of interaction), and feels sufficiently engaged in a simulation (to maximise the potential for immersion; see below), that they would be likely to behave in a manner in which they would ordinarily (i.e. perceptual fidelity). Some element of 'buy-in' from a participant can be assumed as they will have chosen to partic-

ipate in a given study and thus be open to experiencing it as the participant will have chosen to suspend disbelief and cross 'the magic circle' (Extra Credits, 2014b). Amongst other things, this is a matter of informed consent in recruitment.

However, building and maintaining an appropriate level of 'engagement' is a more complex issue and can be encouraged by the set up of the simulator and experiment. Engagement is variously described as immersion or presence, and is discussed now.

6.1.3 Immersion and Presence

Immersion (as a concept) is widely understood in common parlance, yet its technical definition is not agreed. Literature relating to immersion overlaps with concepts in a range of aspects including a feeling of "being there" in a given situation in the media (referred to as "spatial presence", Weibel et al., 2008), or of feeling as if you are really in a situation with another human ("social presence", Weibel et al., 2008), losing track of the world around you ("flow", Csikszentmihalyi, 1990), cognitive absorption/engagement (Brockmyer et al., 2009) and combinations of all of the above. However, even within those aspects, definitions are not agreed and are variously a matter for ongoing debate (Jennett et al., 2008).

Indeed, whilst immersion is often used as a byword for some of these principles, there are clear differences. Jennett et al. (2008) note an example of the ability of a person to become immersed in a game such as Tetris, without believing you are in a world surrounded by falling blocks. Or alternatively, feeling like you are present in a virtual reality world, without simultaneously experiencing a lost sense of time (e.g. if performing a boring task). Jennett et al. (2008) also show data that indicates that a key component of video game immersion is an emotional charge; whereas Csikszentmihalyi (1990) claims that flow involves a serene mindset (with regard to one's "psychic energy"), a term more commonly referred to as being 'in the zone' (e.g. Shainberg, 1989).

Other examples of being immersed include: the 'falling away' ('invisibility') of interaction controls such that the user feels like the user interface controls are an extension of themselves, and an emotional involvement with the simulated universe. For example, player bonding with the 'Weighted Companion Cube' in the computer game Portal (Valve Corporation, 2007): The developers' deliberately engineered the manipulation of the player through in-game dialogue and a heart texture decal applied to the cube, and finally to the player subsequently being forced to 'euthanise' said cube (Elliott, 2008).

Part of the difficulty of defining immersion surrounds the fact that it is so difficult to define clearly. In computer gaming and other modern media, immersion is often a key design goal, but concepts such as the suspension of disbelief (a term coined with reference to the theatre and usually ascribed to Samuel

Taylor Coleridge; e.g. Christou, 2014), are long standing prior to the modern era. In addition to the difficulty of scientifically defining immersion, is the difficulty in measuring it. Jennett et al. (2008) present an attempt to build an immersion questionnaire, where the focus is on consistent reporting of a subjective experience. They show that "people can reliably reflect on their own immersion in a single question". However, note this is an entirely subjective measure, and as such, no objective measures have been found in the literature.

Kwon et al. (2013) demonstrated that participants in a job interview virtual reality (VR) simulation felt anxious at most levels of realism but that the higher the level of immersion, the more anxious the participants felt. This was the case even though the participants were not actually in a job interview, highlighting the fact that merely being placed in a given environment is perhaps enough for the brain to respond on a subconscious and automatic level. Portman et al. (2015) describes the potential value of the 'theatre' of virtual reality in achieving stakeholder buy-in to architectural schemes through immersing them in an environment which would not have the same impact if simply presenting static 3D images or paper-based plans/elevations.

Weibel et al. (2008) found that participants report higher presence in situations where they were playing against a human opponent, and also report higher flow and enjoyment; underlining the value of social presence as a tool for building immersion. Ivory & Kalyanaraman (2007) found that increased realism in the virtual graphical environment enhances feelings of player immersion and both self-reported and physiological arousal; but that this was not significantly related to the level of violence in the media, indicating that high levels of engagement may not depend on the specifics of the subject matter.

Nacke et al. (2010) found that audio addition to a game impacted positively on player perception and a range of factors including flow, immersion and tension. Addition of music also had a similar effect on flow and immersion in the absence of other sound but little or a negative effect in the presence of other situationally-appropriate sounds. However, this result is likely a result of the type of music used as appropriate music is long known as a tool for building audience immersion in visual media including video games (Extra Credits, 2012). That said, given the general lack of music in real cycling scenarios, it is unlikely that adding music would be of great benefit here. Conversely, an ambient audio track (e.g. general urban noise) may be of benefit in increasing immersion and should be included.

Unfortunately, whilst immersion is difficult to quantify and build, its destruction can be swift and easily done. Unrealistic or unintuitive behaviours that are not narratively explainable or internally consistent, and of particular relevance to the digital field, difficult or clumsy controls and interface can all result in a lack of (or the collapse of) immersion (e.g. Smith, 2016; Warr, 2016).

In summary, measuring immersion is subjective, but participants 'know it

when they see it'. Immersion can be described as a fusion of a range of experiences and yet does not require all those aspects to be felt. Ensuring that the developed simulator maximises immersion, and that no obvious content or interface aspects serve to break immersion, is therefore a key requirement. One obvious way to maximise the potential for immersion is by way of immersive visualisation techniques.

6.1.4 Immersive Visualisation

An important part of creating the immersive experience is the method of visualisation. Hou et al. (2012) showed that participants experience a greater sense of physical and self presence in front of a larger screen. This is supported by a background of literature for a range of media types, referred to therein. Immersion aside, a user will not be able to react to observed events in a simulation without an appropriate means of having seen them. Broadly there are three main types of visualisation methods available: standard video display units (VDU), a larger area screen-projection, and a head-mounted display unit (HMD).

A VDU is commonly a flat-screen LCD or LED display of the type most usually used on a desktop computer. Larger screens are often available by use of television displays which are generally the same underlying display technology (though with the inclusion of tuners and other television-specific hardware). Projections made to a screen are available in different forms but are commonly PC-connected digital LCD projectors arranged such they project onto a flat planar material screen of neutral colour. These range from the scale of a small meeting room to a full-scale cinema screen. HMD devices of various forms are available (and have long been used in a research context; e.g. see review in Sharples et al., 2008) but generally consist of a LCD/LED screen mounted a few centimetres from the eyes with in-built lenses to enable to user to focus upon it despite its close proximity. A slightly different image is presented to each eye and the user's brain forms a three-dimensional image ('stereoscopic 3D'). This methodology differs from the other display types where a 3D rendered image is always perceived two-dimensionally.

At time of writing, the HMD device market is in a period of rapid development in the consumer space with a number of new devices coming to market – for example, the Oculus Rift (Oculus VR LLC, 2016b), the HTC Vive (HTC Corporation, 2016) and the Sony Playstation VR (Sony, 2016). However, how best to utilise these devices is still not clear: "The focus so far has been largely on how well the technology works, rather than how we'll use it. As a result, many companies are still figuring what works and what doesn't." (Lane, 2016).

The use of HMDs provides some core benefits:

- User immersion is maximised by excluding the ability for the user to see anything except the simulation.

- Data can be collected relating to the user's head orientation.

However, there are some disbenefits:

- Peripheral vision is limited (in current hardware) and thus users must turn their head to observe areas outside of their forward view cone (approximately 110 degrees in the Oculus Rift device).
- Users can suffer from simulator/motion sickness. See Section 6.2.
- Users cannot see their own body.

Clearly, these benefits are desirable in this scenario. In particular, a more 'immersed' user may be more likely to behave as they would in reality, forgetting (to some extent) that they are on a simulator and resulting in perceptual fidelity for the simulator. On that basis, it is a timely opportunity to explore the use of VR in the development of this simulator. However, for comparison it would be prudent to use alongside a more conventional VDU screen, given the novelty of consumer HMDs.

Encouragingly, Morel et al. (2015) found that there was no significant behaviour difference between VR in a HMD and real life for their balance assessment experiment. However, they did find that reactions were overstated and delayed in the VR compared to real life and theorised that this may have related to a consistent underestimation of perceived distance in VR. They also confirmed work of Horlings et al. (2009) which found that VR caused an order of magnitude increase in postural sway similar to that of the participants' eyes being closed. However, this may be linked to the limited field-of-view of the HMDs (in the order of 30 degrees) used in those studies. Chiarovano et al. (2015) used an Oculus Rift DK2 HMD (the second, and final, development version prior to the consumer release of the Oculus Rift) which has a 110 degree field-of-view and found no such significant behaviour.

Alshaer et al. (2017) also found that having a self-avatar (a virtual representation of the participant's body) in VR had significant effects on correct judgement of width distances. Specifically, they were concerned with doorway/ obstacle gap width acceptance by wheelchair users.

The disbenefits, whilst potentially not inconsequential (a lack of peripheral vision may be a issue that becomes apparent in use), are likely to be outweighed by these benefits and to some extent can be minimised by good simulation design. Simulator sickness however, is a common and recurring thread in the literature (not simply relating to VR HMDs) and so must be properly considered.

6.2 Simulator Sickness

6.2.1 Causation and Incidence

Simulator sickness (SS) is distinct from motion sickness in that it is sickness that occurs as a result of a motion that one sees but does not feel, with motion sickness instead commonly occurring as a result of motion that one feels but does not see. Alternatively called Visually Induced Motion Sickness (Kennedy et al., 2010) and occasionally referred to, by some authors, as 'simulator adaptation syndrome', SS is a common feature in the literature regarding simulator studies and has become more pertinent in mainstream knowledge with the increasing exposure of consumer-VR technology.

Motion sickness has been well-known since antiquity when humans took to the sea and began to ride animals for transport (Golding, 2006). The propensity for motion sickness is commonly accepted to be an evolutionary response to the brain's perception that the mismatch between motion felt (by the inner-ear vestibular system) but not seen, is a potential threat due to ingested neurotoxins and appropriately dealt with through a vomiting response (Brooks et al., 2010; Golding, 2006; Reason & Brand, 1975). This is supported by the observations that species down to the level of fish (e.g. fish being transported can become sea sick; Golding, 2006; Reason & Brand, 1975) can suffer from motion sickness, and also supported by the finding that persons without properly functioning inner-ears ("labyrinthine defectives") do not get motion sick (Bos et al., 2008).

Alternative hypotheses for the causes of motion sickness exist, such as (non-exhaustively): postural instability (Riccio & Stoffregen, 1991), where the sufferer cannot (or has not) maintained the stability of their posture in the presence of a perturbing situation; eye-movement, and their stimulation of the vagus nerve (Ebenholtz, 1992; and see Brooks et al., 2010 for summary); and activation of a vestibular-cardiovascular reflex (see Golding, 2006, for summary). However, while the cause(s) of motion sickness are still very much undecided, the balance of evidence is thought to favour the 'toxin detector' hypothesis discussed above (Golding, 2006).

Discussed in Chapter 4, O'Hern et al. (2017) incorporated VR equipment into a cycle simulator and reported a drop out due to SS of 2 participants out of 30 total (7%). No information is provided however as to the nature of this drop out (e.g. type of symptoms), nor is any information provided as to reports of any SS symptoms in participants that did complete the study.

In the literature regarding driving simulators, simulator sickness is commonly reported. Helland et al. (2016) found some 32% of participants (6 of $n = 19$) in their driving simulator study interrupted or withdrew due to SS symptoms, and this was a group of participants that had already been screened (using the Apfel scale for post-operative vomiting; Apfel et al., 1998) in an at-

tempt to exclude those most likely to suffer from SS. As an aside, whilst Apfel et al. (1998) finds that a probability of post-operative vomiting (after opioid-supplemented inhalational anaesthesia for ear/nose/throat surgery) can be predicted by a patient history risk score, the procedure has not been validated for simulator sickness; a point mentioned in Helland et al. (2016).

Similarly, Brooks et al. (2010) performed a driving simulator confirmation meta-study (from three other studies) in which 17% of participants (19 of $n = 114$) withdrew due to SS. In all of these cases in the other studies mentioned here, those who did not withdraw were not necessarily symptom free.

A study by Allen et al. (2006) summarised in Classen et al. (2011) noted drop-out rates due to SS of 14% and 37% in participants grouped by age to young (ages 21–50, $n = 51$) and old (ages 70–90, $n = 67$) respectively. Classen et al. (2011) more widely noted the need for "well-designed [randomised controlled trials] to clearly show the impact of underlying factors on the occurrence of [simulator sickness]." and also noted surprise at the lack of such given that driving simulators are "perfect for conducting studies under repeatable and controlled conditions."

More widely than driving simulators, in a study involving locomotion in a virtual environment, the unscreened control group in Stanney et al. (2003) experienced a drop-out rate due to SS of 77.5%; the "streamlined" control group (with a more limited number of degree-of-freedom of motion) experienced a 22.5% drop-out rate and across the entire study, 12.9% (142 of $n = 1102$) dropped out due to SS. In a VR latency study performed by St. Pierre et al. (2015) which was static save for head movement, 9.2% (11 of $n = 120$) of participants dropped out.

As noted earlier, various HMDs and VR set ups have been used over the last few decades in a range of academic and industrial contexts. However, consumer-grade HMD VR devices have only very recently come to market and as such wide-scale use outside of a laboratory setting is a new and poorly-studied situation. In particular, this wide-scale use is sufficiently novel that it has not seen substantial (if any) research filter through to the academic literature that captures issues of SS. King (2016) notes that "up to half of all VR users will feel sick, dizzy or headachey when using VR," and anecdotal reporting of VR-related sickness symptoms is common (e.g. Extra Credits, 2014c).

As a final note, the cycle simulator system that features in Grechkin et al. (2012) and Plumert et al. (2004, 2011) does so with only a single mention of simulator sickness, specifically in Grechkin et al. (2012) which states: "During the familiarization session, participants were instructed to notify the experimenter if they experienced any simulator sickness." As there is no further mention, one can only assume no participants did, though as has been illustrated by the above, this seems unlikely.

6.2.1.1 Quantifying Simulator Sickness

To quantify SS, Kennedy et al. (1993) developed the Simulator Sickness Questionnaire (SSQ) which "is the gold standard to measure SS." (Classen et al., 2011). Derived from the wider symptom sets in the then best practice of Motion Sickness Questionnaires (MSQ; which contain symptom sets varied between approximately 25–30 symptoms depending on the particular MSQ chosen), Kennedy et al. excluded: those symptoms which were found to occur with $< 1\%$ frequency (e.g. vomiting) and thus display limited statistical value; those symptoms which showed no changes in frequency/severity in simulators; and those symptoms which might give misleading indications (e.g. boredom). The final list of 16 symptoms is as follows:

- General discomfort
- Fatigue
- Headache
- Eyestrain
- Difficulty focusing
- Increased salivation
- Sweating
- Nausea
- Difficulty concentrating
- Fullness of head
- Blurred vision
- Dizzy (eyes open)
- Dizzy (eyes closed)
- Vertigo
- Stomach awareness
- Burping

Based on statistical analysis of data from 10 different simulators, three symptom clusters were identified: nausea, oculomotor and disorientation. Scoring for each category is on an integer rating of 0–3 where each score represents: "not at all", "slightly", "moderately" or "severely", respectively. Scores for each symptom in each symptom cluster are then factored by a binary factor based on whether the symptom belongs to that cluster, summed for the cluster and then factored by a constant (different for each symptom cluster) to reach a final score for the symptom cluster. Note that some factors belong to more than one category: "Blurred vision" for example, belongs to both the oculomotor and disorientation symptom clusters. Factoring is performed "to produce

scales with similar variabilities on which values can be more readily compared" (Kennedy et al., 1993). A total score is reached by summing each symptom cluster, summing those scores, then factoring by a constant. An absence of symptoms results in a zero score and the availability of subscales provides a practitioner with the ability to focus on what might be wrong with a simulator in a more focussed way.

In Kennedy et al. (1993) approximately half of the responses (of $n = 3691$) to the survey administered in a number of different flight simulators, reported a range of SS symptoms at a range of severities. The distribution of reported scores from the calibration sample that approximately half of participants experience no symptoms with the remainder experiencing a range from mild to severe. The 75th percentile value was suggested as a comparison value as it is "essentially the midpoint of the part of the population that was adversely affected by the exposure".

It seems likely that even in a best case scenario, SS is an issue that will need to be considered with this simulator. It would therefore be prudent to consider the current best practice so as to minimise the potential for SS.

6.2.2 Best Practice

Given their substantial investment in the success of VR, Oculus (manufacturer of the consumer Rift HMD) have compiled a best practice guide for developers (Oculus VR LLC, 2016a). They cite a number of factors to minimise the likelihood of SS (list quoted from said guide):

Acceleration Minimize the size and frequency of accelerations;

Degree of control Don't take control away from the user;

Duration of simulator use Allow and encourage users to take breaks;

Altitude Avoid filling the field of view with the ground;

Binocular disparity Some find viewing stereoscopic images uncomfortable;

Field-of-View Reducing the amount of visual field covered by the virtual environment may also reduce [dis]comfort;

Latency Minimize it; lags/dropped frames are uncomfortable in VR;

Distortion correction Use Oculus VR's distortion shaders;

Flicker Do not display flashing images or fine repeating textures; and

Experience Experience with VR makes you resistant to simulator sickness (which makes developers inappropriate test subjects)."

Each point is briefly discussed by the Best Practice Guide and a limited number of literature sources are cited therein where relevant. As noted in Oculus' introduction, "The exact causes of simulator sickness (and in fact all forms of motion sickness) are still being researched[(See Section 6.2.1 for a discussion.)]. Simulator sickness has a complex etiology of factors that are sufficient but not necessary for inducing discomfort, and maximizing user comfort in the VR experience requires addressing them all."

Based upon the choice to explore the use of a VR HMD in this project, and for lack of any other agreed and consolidated best practice with regards to the minimisation of SS, these items will be addressed as follows below.

6.2.2.1 Acceleration

As discussed in Section 6.2.1, the most widely accepted theory of the cause of SS (Brooks et al., 2010; Reason & Brand, 1975) comes from the mismatch between what one sees and the motion one feels (primarily through the inner-ear vestibular system). One cannot necessarily feel motion at a constant velocity but instead is able to sense changes in that velocity (i.e. accelerations) and so the Oculus Best Practice Guidance suggests that "An instantaneous burst of acceleration is more comfortable than an extended, gradual acceleration to the same movement velocity."

In addition to the issue of translational acceleration, rotational acceleration must be considered. To a large extent, rotational accelerations should be minimal as the participant remains upright. However, some aspects, such as navigating around corners, are unavoidable as they are a key part of the simulation. Unfortunately, the slow adjustment of speed is also a key function of a bicycle simulator. As an illustrative example, an instantaneous acceleration from a stop to 'cruising speed' would likely break immersion entirely, wrench control from the participant and be completely unrealistic. It is unfortunately the case that the 'best practice' with regard to acceleration cannot be followed.

6.2.2.2 Degree of control

The Oculus Best Practice Guidance states that "Taking control of the camera away from the user or causing it to move in ways not initiated by the user can lead to simulator sickness" and that it is therefore important to "not decouple the user's movements from the camera's movements in the virtual environment."

Despite this, there are potentially circumstances where reducing the degree of control of the user may actually be beneficial in reducing SS. The Stanney et al. (2003) study discussed in Section 6.2.1 above, found that a reduced degree of control (from 6 to 3 degrees-of-freedom) actually reduced the reported incidences of SS.

In any case, save for the movements within the environment caused by the movement of the user given they are on a vehicle (if you turn a bicycle, then you turn with it), there would be no need to remove visual control of movement from the user.

6.2.2.3 Duration of simulator use

Stanney et al. (2003) found that reported SS scores (at 15 minute intervals) increased in VR with length of exposure up to the studied 60 minutes. However, substantial reporting of symptoms was found even at the 15 minute session length (total SSQ score at 15min some 62% of the total at 60min). Therefore, whilst utilised scenarios should be limited in duration if possible, a scenario yielding an inherently high incidence of SS is likely to manifest in any practical exposure length.

6.2.2.4 Altitude

The more ground plane that is in view, the more that the effect of movement is imposed on the participants' visual field and the greater the likely increased feelings of vection (the illusion of self-motion in the opposite direction based on the movement of the surroundings So et al., 2001). Given that a simulation from the cyclist's-eye-view would be approximately 1.75m above ground level, filling a substantial part of the field-of-view with ground texture is unfortunately unavoidable.

A study by So et al. (2001) of the relationship between SS and speed of movement (both translation and rotation) which was constant for each participant and showed that reported induced SS scores increase rapidly with increased speed in the range 3–10ms^{-1} and level off thereafter through speeds of up to 60ms^{-1}. Speeds below 3.3ms^{-1} or above 59.2ms^{-1} were not tested. The visual field in said study was close to ground level in a simulated urban environment (as opposed to being high in the air or space, for example). These results indicate that the range of values likely to be experienced in the simulation (in the order of 5ms^{-1}) are not likely to induce the highest levels of SS in of themselves.

However, it should be noted that "vection is considered to be an important (actually desired) element of creating *presence* within the simulated environment. [their emphasis]" (Kennedy et al., 2010), and therefore may be a worthy trade-off to achieve higher levels of immersion.

6.2.2.5 Binocular disparity

It is possible that participants simply do not react well to the stereoscopic 3D aspect of the VR simulation; the issue arising because the eyes are focusing (with lenses) on a screen at one depth whilst viewing an image with potentially

a different depth of field. To some extent this is unavoidable however it may be possible to isolate this effect should a person react particularly strongly to the stereoscopic effect as this should translate into the specific eye-related symptom set in the SSQ. The best practices with regards to ensuring interpupillary distance is properly translated into the camera location in scene, is a function of the libraries used to enable the use of the Oculus Rift HMD. Furthermore, this would not be the case for a screen-based situation where the 3D images presented to the participant are non-stereoscopic.

6.2.2.6 Field-of-View

Similar to the issues of personal-incompatibility with stereoscopic 3D, large field-of-view coverage is a (key) function of the use of a VR HMD. The Oculus Rift HMD has a field-of-view of approximately 110 degrees. This is also not easily adjustable (save for artificially imposing (in software) a 'tunnel vision' on the rendered images sent to the eyes) and will not be explored here. Screen-based field-of-view (if one is to avoid distortion) is also a function of the hardware and its placement, though the size of the screen proposed and its placement is likely to occupy a substantially smaller part of the participant field-of-view than that of the HMD (approx. 30 degrees, versus 110 degrees).

6.2.2.7 Latency

Latency is a time delay in a given system. This may occur owing to the processing overhead of the simulation, the time it takes the sensors/accelerometers to take and provide measurements to the system, the time it takes the system to respond to these changes, and the time it takes to present those changes back to the user. It is important in a simulator because substantial delays will be difficult for the user to engage realistically with. However, it is of high importance with HMD VR devices as delayed in-simulation head motion versus head motion in reality (sometimes referred to as 'motion-to-photon latency') has been found to be significantly related with SS (St. Pierre et al., 2015), although earlier work referred to critically therein (Draper et al., 2001), did not find any significant effect on SS.

The simulation environment should be calibrated and optimised at design time to ensure that the a refresh rate of at least the native level of the VR HMD is achieved (90Hz). Since variable latency has been found correlated with higher levels of SS (St. Pierre et al., 2015) than constant latency, the control system, which operates with different latency paths than the HMD, will be set up to ensure that latency is constant, even if this results in a slightly higher level of latency than is theoretically achievable.

6.2.2.8 Distortion correction

An issue arising throughout the image display process is that from the use of flat screens at close proximities, given the differing distances from the eye to the centre of the screen compared to the more distant edges. This is particularly an issue for the use of HMDs given their proximity to the eye. As a consumer product, the Oculus Rift HMD has officially-distributed software libraries for interfacing with the HMD which make the necessary corrections. This is not applicable in screen-based situations for all but the largest (or closest) screens, the scale of which are not proposed here. The hardware set up (e.g. Section 7.1.2) precludes placing a screen within around 2m of the participant and large projection screens are not proposed in this study. Though, if those are used then the projectors generally ship with adjustment settings such as keystone adjustment.

6.2.2.9 Flicker

The simulation content proposed does not involve any obvious sources of flickering or flashing. The HMD unit screens operate at 90Hz which is likely to be imperceptible in terms of flicker. All but the most premium of VDU equipment typically operates at lower refresh rates, e.g. 60Hz. These are usually fixed functions of the hardware however new systems such as 'G-Sync' and 'FreeSync' are recent to market and adapt the refresh rate of the display to the output of the particular manufacturer's display card hardware. The Best Practice Guidance states that "Flicker plays a significant role in the oculomotor component of simulator sickness." and therefore if this is an issue, it would manifest in that specific set of symptoms. If this is the case then it can be reviewed post-study or flagged for future consideration.

6.2.2.10 Experience

The Best Practice Guidance suggests that "new users should not be thrown immediately into intense game experiences; you should begin them with more sedate, slower-paced interactions that ease them into the game." A 'learning' scenario should therefore be provided for the users of the simulator, to allow them to accustom themselves to the controls etc.

6.2.2.11 Summary

Where possible, established best practices (such as they are) have been adopted or designed in to the proposed specification detailed in Section 6.3 below. However, in some key areas – such as acceleration – this is not possible as it would undermine the primary purpose of the simulator.

Though undesirable and potentially unpleasant for the participant, ultimately the occurrence of SS can be considered a 'study overhead' in that its

implication is simply to demand a larger sample size to achieve any given data requirement. Any study is likely to suffer from a level of participant withdrawal; this would simply be a higher rate.

However, SS is not a binary condition and an amount of participants will experience SS symptoms and not feel the need to withdraw. Furthermore, simply excluding participants with *any* symptoms is not practical as this may include *all* the participants. For data to be useful then, the behaviour of participants (or more particularly the data collected from them) must not substantially change as a result of SS symptoms. Fortunately, some limited literature exists related to this issue.

6.2.3 Effect on Data

It is intuitive that people behave differently when they do not feel well. However, of issue here is whether that potentially different behaviour will yield fundamentally different outcomes in the simulation. Clearly, such an assumption can only truly be validated by a comprehensive and large-scale trial.

In the absence of such, Classen et al. (2011) performed a comprehensive literature review of the driving simulator literature in an effort to establish the basic prevalence of SS and to attempt to classify factors that may increase incidence of reported SS. They found that despite a number of the available studies mentioning SS, none were "Class 1" in that they were either not properly blinded/randomised, had low numbers of participants, or not assessed using the established "gold standard" SSQ procedure. They note: "This finding was surprising because driving simulators are perfect for conducting studies under repeatable and controlled conditions. A need exists to determine, through RCTs, the causes of SS."

Subsequently, Helland et al. (2016) found "that there is no significant influence of simulator sickness on the important driving impairment measure [standard deviation of lateral position].". Their particular concern was to do with the use of a simulator to study impaired driving under the influence of alcohol and to frame this within the body of literature (e.g. Muttray et al., 2013) that has established standard deviation of lateral position as a robust parameter of driving impairment under the influence of drugs. They found significant evidence between simulator sickness, and chosen mean speed and steering reversal frequency; their interpretation being "that simulator sickness primarily causes the subjects to drive more slowly and avoid unnecessary steering wheel reversals in an attempt to ease symptoms" and that this was "most pronounced in sober subjects". The measures of speed and steering reversal frequency were recorded to be significantly negatively correlated with reported levels of SS, however it should be noted that these two parameters are themselves not independent as evidenced by the lack of a significant correlation between steering reversal per distance driven and speed. Additionally, no significant effects were

found between SS and standard deviation of chosen speed, brake or accelerator pressures per distance driven, or other steering wheel movement measures (specifically speed, movement per distance driven, and as mentioned, reversals per distance driven).

More generally, it may not be ethically possible to establish the extent to which the experience of SS may introduce bias into the data. Of issue: if one is seeking to deliver a study in which participants are required to give full and informed consent, and in which it is impossible to provide the participant with objective warning (given that pre-knowledge of the potential of an experience of SS may contribute to reported SS), it is perhaps then not possible that *any* such study could deliberately and specifically be ethically administered, properly randomised/controlled for SS factors, or recruited for.

For the purpose here, given that this literature indicates that the outcome measure of the variability of lateral placement is unaffected, and lacking in any other literature indicating the validity of measured behaviours versus reality, it is reasonable to take forward a working assumption that spatial/lateral placement is likely independent of the level of SS reported by a participant, if any. A comparison of reported SS, and chosen speed and steering reversals will be considered based on the outcome data from the study. Participants will be informed that they may potentially experience SS and will be free to withdraw at any time.

6.3 Specification for a Simulation Interface

The movement of a bicycle in normal use can be reasonably simplified to (constrained) motion of a point object on a two-dimensional plane. Furthermore, that motion can be readily explained as the result of the combination of the action of the cyclist pedalling/braking (i.e. forward motion and control of the speed of such) and of the cyclist turning the handlebars to steer (i.e. so the forward path of motion takes the path of an arc, rather than a straight line). Whilst the reality of a cyclist's motion is obviously more complex – for example, a cyclist will 'lean into' a sharp turn or could pick up their bicycle and move it laterally when stopped – a simplification to these two control measures is a reasonable starting assumption. Such an assumption has the benefit of reducing the need for motion capture to two degrees of freedom (i.e. the forwards/backwards axis, and yaw).

One might therefore expect that if the experience of being on a bicycle can be embedded into a simulation, then the user may feel sufficiently immersed such that they behave in the simulation as they would in reality; i.e. their thought processes fall back to the 'base actions' required of a cyclist (see Section 6.1.2). In this particular use case, the fidelity of in-simulator versus in-reality behaviour is not only desirable, it is the goal (see discussion in Section 6.1.1). The required simulation interface is therefore as follows:

1. All user interaction through the use of a standard bicycle (in a fixed indoor position) to maximise objective fidelity (Section 6.1.1);
2. Capture of forward/backwards axis motion, utilised to provide speed information;
3. Capture of steering direction to provide turn information;
4. Live visualisation of the simulation to the 'rider' so as to achieve a high level of immersion and response to their chosen actions (Section 6.1.4); and
5. Visualisation equipment should be calibrated to each user's requirements (e.g. Section 6.2.2.5).

In addition to these required items, some points are desirable but not absolute requirements:

6. Users of the simulator should be able to use their own bicycle. This maximises comfort, removes the need to adjust to a new bicycle and avoids presenting unnecessary barriers to immersion (Section 6.1.3).
7. A 'learning' scenario should be provided to allow participants to become accustomed to the controls and visualisation method for the purpose of participant comfort (and to reduce likelihood of SS; Section 6.2.2.10).

8. Appropriate audio should be included to increase chances of high levels of immersion (Section 6.1.3).

9. Latency of interface operation and visual rendering should be minimised to improve usability of the interface and reduce probability of SS (Section 6.2.2.7).

10. Exposure sessions should be limited in duration so as to reduce probability of SS (Section 6.2.2.3).

Now that the literature base has been established, and a specification has been compiled, this document moves on to the design and creation of a simulator.

Chapter 7

Simulator Construction and Implementation

7.1 Simulator Construction

Taking the specification for a bicycle simulator established in the previous chapter (Section 6.3), this chapter details the implementation of a simulator to deliver that specification.

7.1.1 Speed Detection

For a static bicycle to provide a facsimile of a real riding experience, resistance needs to be applied to the rider's pedalling. This can be applied directly to the pedal crank as is commonly the case for an integrated stationary exercise bicycle (e.g. Technogym, 2016) and usually takes the form of electromagnetic resistance. Alternatively, resistance can be applied to the driven (i.e. rear) wheel as is the methodology with rear-wheel attached 'turbo-trainer' devices (e.g. Tacx B.V., 2016a), or by simple 'roller' devices (the cycling equivalent of a 'rolling road'/chassis dynamometer; e.g. Tacx B.V., 2016d). However, none of these solutions are specifically appropriate for the use case here. It was therefore necessary for this author to develop a bespoke solution.

As a starting point, an 'off-the-shelf' cycle trainer unit was chosen. Such a trainer provides the potential for the use of any compatible (i.e. the rear hub is Quick Release 'QR' compatible) standard bicycle (possibly including the bicycle belonging to a given participant; Specification Item 6), does not require the learning process required to safely use a roller unit, and does not create safety issues surrounding the use of a full-immersion virtual reality HMD, such as not being able to maintain balance on the bicycle. A Tacx T2780 Bushido Smart Trainer (Tacx B.V., 2016a) was chosen (Figure 7.1); as this was self-contained, provided variable resistance (so potentially providing for the future ability to explore the effects of gradients), and connects to other equipment wirelessly by use of the ANT+ open-access protocol (Dynastream Innovations, 2016, analo-

gous to Bluetooth but optimised for ultra-low power devices such as sensors). Alternative real-wheel trainers would also be suitable, provided they meet the requirements specified here.

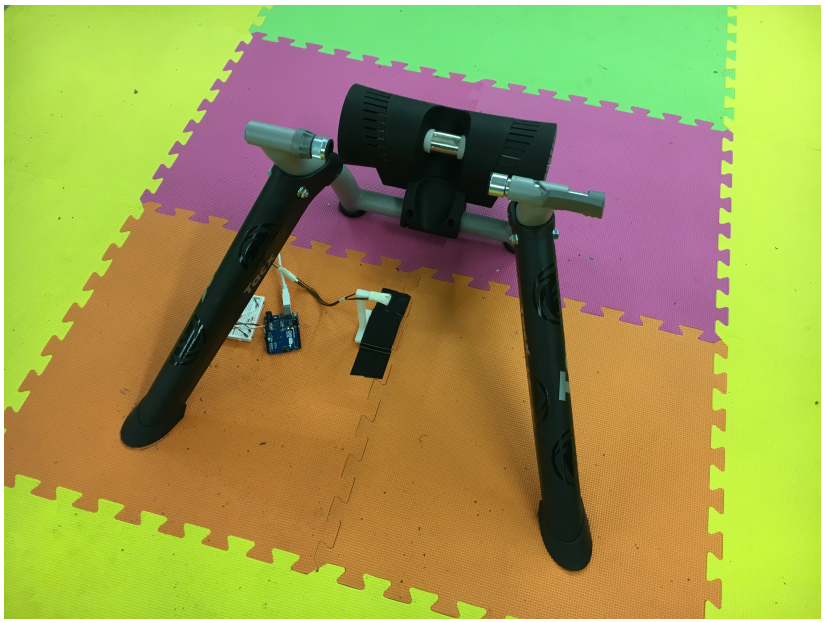

Figure 7.1: Image of the Tacx Bushido T2780 Smart Trainer

The trainer uses the ANT+ network protocol to broadcast speed readings from the bicycle wheel at a rate of 4Hz and, as libraries for this transmission protocol are open-access, can be integrated into other software with much greater ease than a closed/proprietary unit. A standard PC with a USB ANT+ sensor is used to receive the broadcast data.

Ideally, the chosen rear-wheel trainer device would provide live and accurate speed information to the simulation. However, during testing, an issue arose whereby at speeds below approximately 8kmh^{-1} (2.2ms^{-1}) the reporting resolution of the trainer device falls such that as actual speeds approach zero, the device continues to report the last speed it considered to have been valid (usually a speed value in the range 4–7kmh^{-1} (1.1–1.9ms^{-1}). Then, at even slower speeds below approximately 4kmh^{-1} (1.1ms^{-1}) and indeed at an actual stop, the device continues to broadcast speed data as if motion were still in progress. A grey zone is shown in the live graphs in the turbo-trainer's official iOS app (Tacx B.V., 2016c; Figure 7.2), which would indicate that this is a 'feature' of the hardware and not a set-up error on the part of this author.

Given the primary design case of rear-wheel trainer devices is not for low speeds, it is likely that alternative products may also output false data in a similar manner. Existing virtual cycling solutions are targeted at higher speed riding (such as simulated racing), and therefore may simply ignore the erroneous data at the low end and/or: either heuristically assume a speed of zero below a given 'cut-off' value, or respond to some proprietary component of the broadcast signal which is not documented or available through the ANT+ pro-

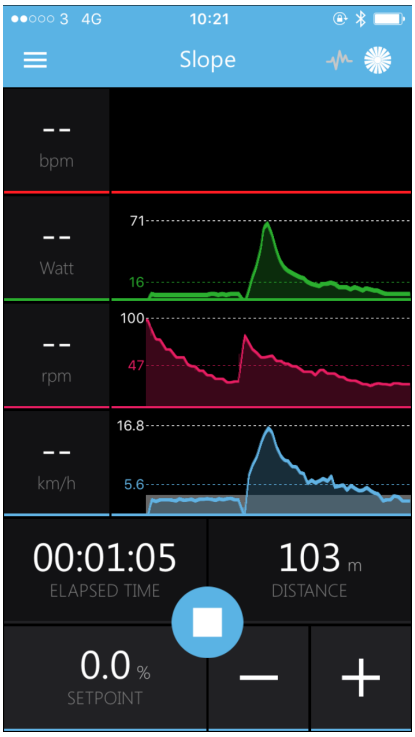

Figure 7.2: Screen capture of the Tacx App in operation (Tacx B.V., 2016c). Note the grey zone indicated on the lowermost (speed) graph.

tocol, or simply be left as something the user simply has to live with. This particular trainer uses a combined ANT+/Bluetooth broadcast; it is possible that the Bluetooth component of the signal carries additional proprietary information. However, in this case (urban utility cycling), these issues arise within the range of speeds that are desired for interacting with the simulation so an alternative needed to be implemented.

As an alternative, consideration was made of the (broadly universal) methodology used by basic cycle computers to detect the speed of the moving bicycle. Most of these cycle computer devices utilise a magnet attached to a wheel spoke and a Hall Field-Effect Transistor (HFET) placed such that the magnet passes close to the HFET once per wheel revolution. The pulse generated by the passing magnet can be recorded; with known dimensions of the wheel, distance can be calculated; and, given a known time, speed can be inferred.

Whilst simple to implement here, attaching a magnet to participants' bicycles may be undesirable (Specification Item 6), and a detection resolution of a single wheel revolution (in the order of 2m travel distance) may not provide a sufficient quality of data at very low speeds. Simply attaching more magnets (such that there are multiple pulses per revolution) moves further into the potentially unwelcome territory and may start to impact the gyroscopic balance of the wheel at highest speeds. As an alternative then, a simple infrared (IR) reflective photomicrosensor (a paired infrared-emitting LED and phototransis-

tor that is set up to detect a reflection/interruption at close range to the IR emitted by the LED; Farnell, 2016) was set in a fixed stand and held approximately 5mm from the passing spokes near the rim of the rear wheel as illustrated in Figure 7.3. This provides a resolution down to the individual spoke (in the order of 5cm travel distance) but does not require a physical intervention on the bicycle itself and is broadly insensitive to the specific wheel used (provided differing numbers of spokes are accounted for).

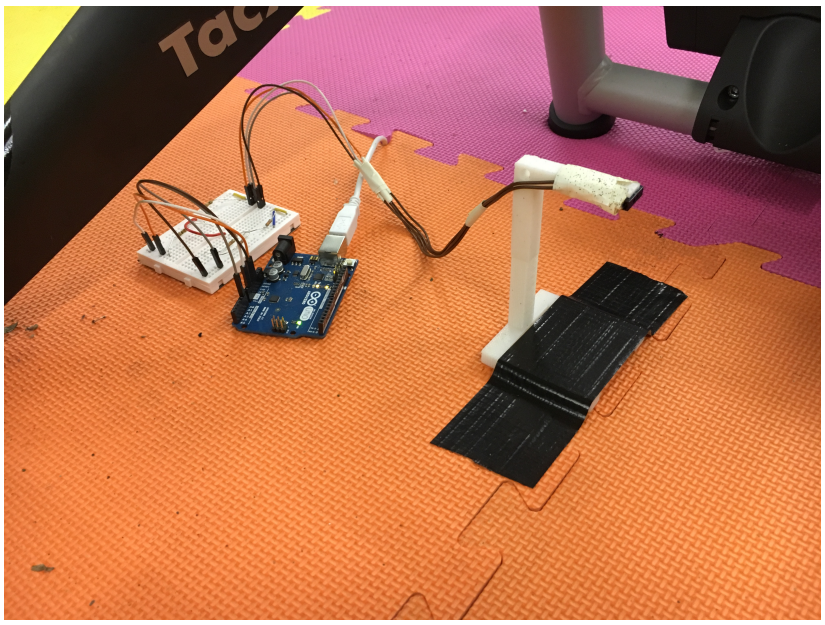

Figure 7.3: Photo of IR sensor mounted on 3D printed support with accompanying control unit

The IR sensor was connected to an Arduino Uno (Arduino, 2016) unit (also shown in Figure 7.3) and a script created that registered the pulses generated by the passing of the spokes as the wheel rotated (Figure 7.4). The passing of the tyre valve does not give a detectable IR return as the valve is dark rubber and therefore absorbs the majority of the IR incident upon it. The presence (or not) of reflective material in the field-of-view of the IR sensor – based on a simple threshold value for sensed level of IR – is then posted over a USB serial interface to the simulation PC. This methodology has the benefit of providing a much higher spatiotemporal resolution (down to an approx. 5cm or 300Hz) than the rear-wheel trainer was capable of (no better than around 1m or 4Hz), albeit at the need of further equipment.

Whilst the methodology is sound, the microcomputer used was not capable of sufficient monitoring resolution (see Section 7.2.1.2); specifically, the limit of the Arduino Uno used is approximately 1kHz which is insufficient to capture the rapidity of spoke movement at all but the lowest speeds. Bridging adjacent spokes with IR-reflective tape (visible in yellow in Figure 7.4) provided a reliability over a much wider range of speed values (including the full range of interest for simulation). When set up this way, the sensor provides reliable speed

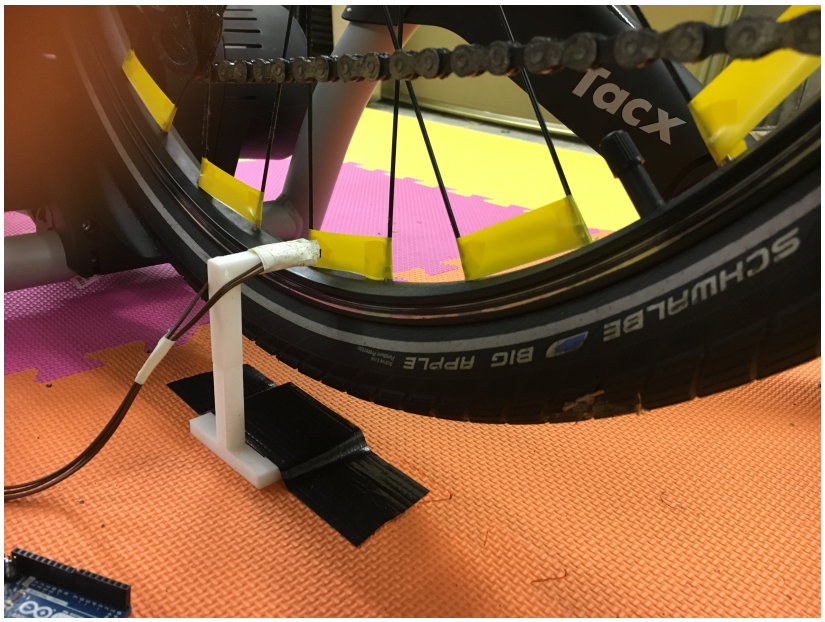

Figure 7.4: Photo of IR sensor positioned in proximity to passing spokes

values in the range 0.2–10.0ms^{-1} (0.7–36.0kmh^{-1}).

At the high end, once the speed of wheel rotation reaches an equivalent travel speed of more than approximately 10ms^{-1} (36kph^{-1}), and even with the bridged spokes, the utilised sensor again does not have sufficient resolution to continue to work reliably (owing again to the stroboscopic effect resulting from the microcomputer sampling resolution). For the purpose of this study, this speed is beyond the top end of concerned values, however there is clearly scope for transitioning between the high resolution IR speed sensor and the lower resolution rear-wheel trainer (which does support such speeds) across the range of values which both operate reliably. Outputs from the two sensors agree well across a range 2.5–10.0ms^{-1} (9.0–36.0kmh^{-1}). Bridging the spokes on a participant's bicycle with tape is perhaps undesirable (Specification Item 6), but not a complex or irreversible operation. That said, this equipment set up used a single bicycle (appropriately adjusted) between all participants.

It should also be noted that if an alternative such as HFET pulsing had been used, then the result would have been the same. Also, adjustments of the sensor position such that the spoke movement is slower (i.e. closer to the hub) would reduce the ability to accurately track the wheel rotation owing to the standard arrangement of interleaving wheel spokes at the hub.

7.1.2 Steering Detection

Steering (and thus yaw) of a bicycle is delivered by the direct adjustment of the handlebars by the user. Detection of this angle is therefore achieved through the specific detection of the handlebar setting or an appropriate proxy. Due to the standard construction of a front wheel fork, detection of any of: the angle

of the handlebars, angle of the fork, angle of the the wheel, or angle of the tyre; provide an equivalent measure of the handlebar setting.

An elaborate purpose-built system could replace the entire front fork assembly with a dedicated hardware control device equivalent to using a steering wheel controller in a driving simulation, or a flight yoke in a fight simulator. However this would entail modifications to the bicycles used, such that the bicycle's handlebar assembly did not interfere with the controller, or vice versa. The most obvious alternative is the same type of methodology used for bicycle instrumentation in the literature (e.g. Miah et al., 2016), which requires the fitting and calibration of the bicycle with a suite of sensors. However, again this requires a substantial intervention on a given bicycle.

A much better alternative in this scenario – and indeed the methodology taken by most products on the market in this field (e.g. Tacx B.V., 2016e) – is to detect the orientation of the wheel by resting it in a support device which rotates with the tyre as it is turned about the vertical axis. As noted, commercial products are available, however none that were able to be used in a non-proprietary context were found. As such, a purpose-built equivalent of such a device was instead implemented by the author.

The basic arrangement of such a system consists of:

- A support for the wheel such that the bicycle does not fall from the support and such that the support rotates with the wheel about the vertical axis allowing the detection of intended steering angle.
- A base that contains the support section and detection/connection equipment.
- A bearing to allow smooth and unimpeded rotation of the movable support section.
- A method to detect the wheel's orientation.

A front wheel support was designed and 3D-printed in ABS plastic with a recess to accommodate a tyre (Figure 7.5a). A base section was also designed and the two sections each support a single thrust bearing at the interface to ensure smooth motion (Figure 7.5b). Because a bicycle tyre does not rotate about the 'world' vertical but in fact along an axis parallel to the axis of the head tube (usually around 20 degrees to 'world' vertical), the base section was designed to an angle of 20 degrees off horizontal (i.e. perpendicular to the axis of wheel steering rotation). It should be noted that some off-the-shelf systems do not make this accommodation, instead sitting under the contact point of the tyre, however these systems do not tend allow more than a limited rotation in either direction (a bicycle wheel actually sweeps through a torus as it rotates owing to the curvature of the front fork). This is insufficient for this purpose where sharper turns are to be expected (e.g. in the process of urban navigation), but is more appropriate for the racing contexts in which those devices are

designed to be utilised (where contextually, a sharp turn would lead to a fall from the bicycle). An abrasive surface was affixed to the bottom of the base to ensure there was no slip potential.

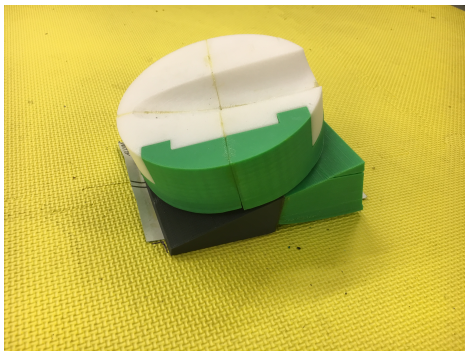

(a) As used

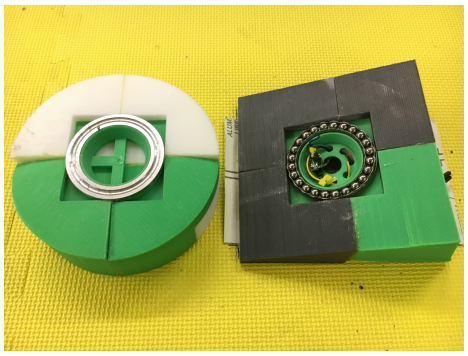

(b) Taken apart to view internal bearing (NB. electronics visible relate to abortive rotary encoder sensing and were not used).

Figure 7.5: Front wheel steering support

The next step is the detection of rotation along the rotation axis of the wheel support unit. Some systems use a proprietary potentiometer (or similar) orientated about the rotation axis. However, a potentiometer with sufficient robustness and sensitivity was unable to be found within budgetary constraints. The wear profile of this use case is such that the area of the potentiometer about the forward direction is likely to receive a large amount of wear and would also need to be sensitive to small corrective movements. As noted above, off-the-shelf systems can target (or produce) components focussed on the narrow range of motion expected in high speed scenarios (e.g. by using gearing); this was not explored here.

As an alternative to a mechanical detection, an optical rotary encoder and code wheel can be used together to detect rotation. An incremental optical rotary encoder transforms the rotation of a code wheel though its field of view into electrical impulses which can be used to infer direction and rate of rotation. The majority of rotary encoders operate using an infrared emitter aimed at (and placed near to) a detector, with the code wheel passing between the two. The code wheel is usually simply a disc of opaque material with regular gaps or slits that rotate through the field of view. The rate and direction of this movement is detected by the rotary encoder and translated into a digital output. These types of devices are often used in industrial contexts to measure shaft rotation, are common on mechanical computer mice, and are also preferable when the range of motion is not desired to be limited. They are also desirable here as they do not suffer from wear. The components are contactless, so instead all the mechanical wear is borne by the thrust bearing, which is designed for the task and is much more robust.

An incremental (a type dedicated to the rate and scale of rotation of the

shaft which is better suited for this scenario where motion substantially less than full rotation is concerned) optical rotary encoder (specifically Broadcom item AEDB-9140-A13: RS Components, 2016) was fitted to the base of the wheel support unit and an appropriate code wheel was fitted to the moveable wheel support in such a way that it rotated within the field of view of the rotary encoder. The encoder was connected to an Arduino Uno (Arduino, 2016) unit which tracks the relative change in position on the code wheel using the signal returned from the encoder. The encoder utilises quadrature difference (waves that are 90 degrees out of phase) between two different waveforms to count either up or down depending on the direction of turn.

However in testing, it transpired that this method is error-prone for small and sharply-made rotations that entail a switch of direction. Specifically, a substantial drift in the 'zero' point developed from which difference from forward direction would be computed. Given that for most forward motion, very small sharp adjustment rotations are made, this methodology was abandoned. Code wheels and optical encoders with a higher level of precision are available, but not within the budgetary constraints afforded this project.

Despite this, for the purpose of supporting the wheel and enabling smooth rotation, the wheel support unit remains otherwise viable and is retained. Driven by the demonstrated capabilities of using video in Chapter 5, the author instead chose to observe wheel orientation by remote detection from the unit by the use of a PC-connected USB webcam. This camera is set forward of and above the bicycle, and looks back down upon it such that the centre of the field-of-view is set upon the centreline of the tyre. The surrounding floor area is of a highly contrasting material and the PC runs a software service, written by this author, which observes the live image and looks for the area that is 'not floor'. As bicycle tyres are generally dark and the floor is light in the image, a binary low-pass filter of the image pixel values is sufficient for the purpose. The software seeks the points along the edge of the assumed tyre furthest from the centreline in the image and therefore has captured the leading edge of the wheel. Finally, trigonometry is applied to yield a heading angle for the tyre from this information. The software then provides this information as a network service that can be read by the simulation which runs independently (See Section 7.2.1).

7.1.3 Visualisation

The benefits of virtual reality (VR) and head-mounted devices (HMDs) was explored in Section 6.1.4. Given the use of a 3D rendering engine (see Section 7.2.2 below), set up of visualisation is simply a matter of connecting the relevant equipment to the controlling PC and setting it up such that it does not interfere with the user. In the case of VR, this means routing the cables such they do not risk being tangled about the user, and in the case of the

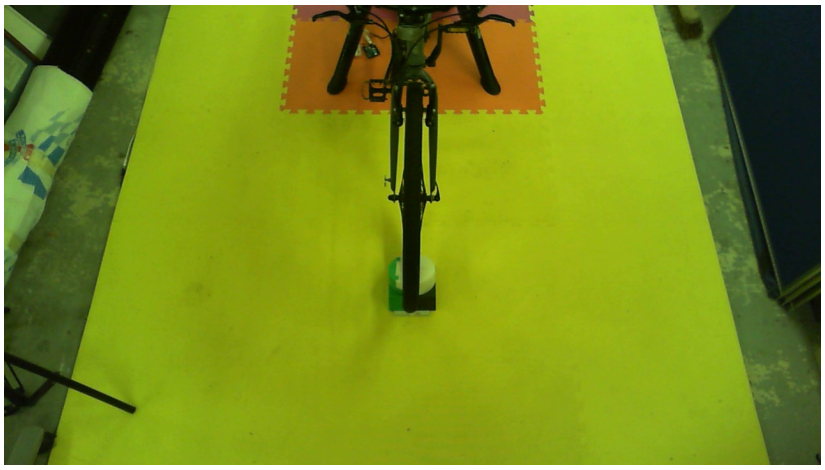

Figure 7.6: Camera image for observation of steering angle

screen, that it is in the centre of the forward field-of-view of the user. Figure 7.7 demonstrates the spatial set up for the screen-based situation.

Figure 7.7: Equipment as set up for the screen-based situation

7.2 Simulation Design

7.2.1 Hardware interfaces

The simulation interface constructed thus far collects relevant sensor data from the participants' actions. The next step is to take this raw sensor output and convert it to real measurements, then to post those real valued data points to the simulation.

This separates concerns such that the simulation does not have to expend overhead on interpretation of raw sensor input. It also ensures that if a sensor were to change or be replaced, the simulation would not have to be changed, and vice versa. The following therefore details the software side of the sensor systems. Note that although some of the hardware components were 'off-the-shelf', the software parts of the sensing systems were purpose-built for this simulator by this author.

7.2.1.1 Software Structure

Interface to the hardware components applied to the bicycle is through a single self-contained piece of software written by this author. The following goes into some detail; readers without an interest in coding issues, may wish to skip to Section 7.2.2.

Inputs are through the connected camera and microcomputer; and over two network ports (one each for speed and direction). Each of the two network ports is managed by a thread instance of a `SpokeSensorServer` class (for speed) and a `SteeringDetectorServer` class (for steering), each of which extends a `LocalhostServer` class.

`LocalhostServer` is an abstract class which wraps a thread which attaches to a network port and (until terminated) listens for client attempts to connect to that port (see Section 7.2.2 below). Then when a successful connection is made, sends that client a small bundle of information based on the overridden `GetServerMessage()` method specified in the implementing class.

`SpokeSensorServer` extends `LocalhostServer` with regard to speed detection, specifically through the movement of the spokes. The class spawns a thread to: read information from the USB-connected microcomputer over a serial interface; and make the necessary calculations to determine speed (Section 7.2.1.2 below). The class also overrides `GetServerMessage()` to return a short ASCII string containing that calculated speed information, which is then transmitted over the network port by this `LocalhostServer` thread.

`SteeringDetectorServer` extends `LocalhostServer` with regard to steering direction, measured by visual observation of the angle of the front wheel. The class spawns a thread to read information from the USB-connected camera which is attached to the PC and is set to observe the bicycle's front wheel, calculate the extent and direction to which the wheel is turned (Sec-

tion 7.2.1.3 below), and expose that as a data property. The class overrides `GetServerMessage()` to return a short ASCII string containing that calculated direction information, which is then transmitted over the network port by this `LocalhostServer` thread.

7.2.1.2 Speed Detection Algorithm

An infrared microsensor is set up close to (but clear of; ca. 5mm) the spokes of the rear wheel of the supported bicycle (Section 7.1.1). The sensor is fixed to a simple ABS 3D-printed stand and wired through an appropriate resistor to limit current to the IR emitter diode, and though an appropriate (in this case, 10K) pulldown resistor to stabilise the output of the sensor (Figure 7.8). The microcomputer provides the power for the emitter (and sensor), a common ground, and reads the analogue value from the detector on one of the microcomputer's analogue input pins.

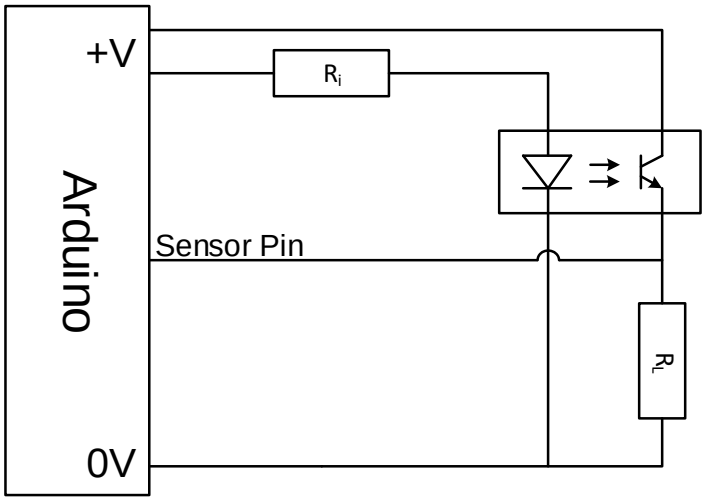

Figure 7.8: Circuit diagram for speed sensor hardware

When there is no reflection of the emitted IR light, then the sensor returns a low value (it is essentially reading the ambient IR level in the experimental space), and when the emitted IR is reflected from an object, then the sensor returns a high value. Absolute values depend on the nature of the reflecting surface and the specific choice of current limiting resistor (R_i in Figure 7.8).

The microcomputer was set to apply a simple cut off value against the input detected on the analogue pin. If the input is over this value, then detection is assumed; if below, then no detection is assumed. The state of detection (or not) is codified as a '0' or a '1' and then transmitted over a serial interface to the connected PC.

Once received by the PC (specifically by `SpokeSensorServer` which listens to a virtual serial port on the USB interface), a running window of the last 3

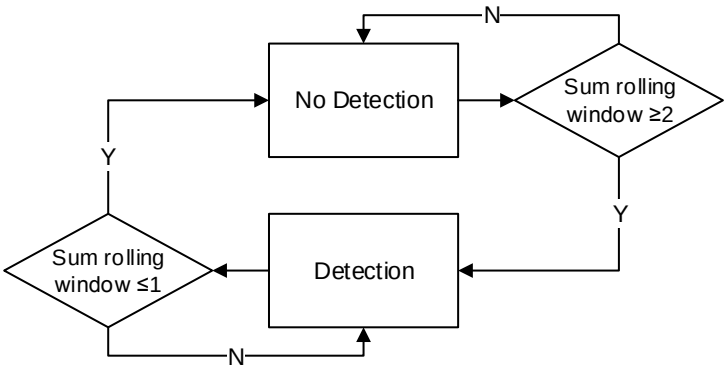

Figure 7.9: State diagram for speed sensor hardware

detection states is maintained. A state machine is then maintained as to the current state of the detector; Figure 7.9 summarises this graphically. In short, the total value of the states within the window are considered. If they sum to 2 or more, and the current state is 'no detection', then a state change happens to a 'detection' state. Conversely, in a 'detection' state, if the detections sum to 1 or less, then the state changes to 'no detection'. This state machine has the property of being robust to small numbers of erroneous detections (or non-detections) which might otherwise be considered a state change if the raw feed were considered. A wider window can be used, however this comes at a tradeoff to the frequency with which changes can be detected.

The time interval since the last transition from 'detection' to 'no detection' and from 'no detection' to 'detection' is recorded, averaged and converted to speed based on the known circumferential distance between two spokes:

$$speed = \frac{(\frac{2}{numSpokes}) \times \pi \times wheelDiameter}{detectionInterval} \tag{7.1}$$

Note that Equation 7.1 uses the term $\frac{2}{numSpokes}$ because the reflective tape used between each pair of spokes has, in effect, halved the number of spokes that can be observed per rotation of the wheel.

Finally, an upper limit on $detectionInterval$ is applied and if the interval exceeds that period, speed is assumed to be zero. An interval of 1s therefore provides reliable speed data down to approximately 0.2ms^{-1} (0.7kmh^{-1}), however does have the consequence that an absolute stop takes that period of time to register. This can be conspicuous at sudden stops from high speeds. If a particular scenario demands it, then the interval can be shortened, with the effect being a loss of speed resolution at extremely slow speeds.

Finally, the computed speed value is exposed as the property which the `LocalhostServer`'s thread and methods deliver to the network port.

7.2.1.3 Steering Detection: Actual to Effective

An off-the-shelf USB webcam was set up above and ahead of the bicycle and an on-screen overlay used to align the field of view such that the horizontal centre was directly forward of the front wheel, and that the vertical centre was set to the centreline of the wheel (Figure 7.10).

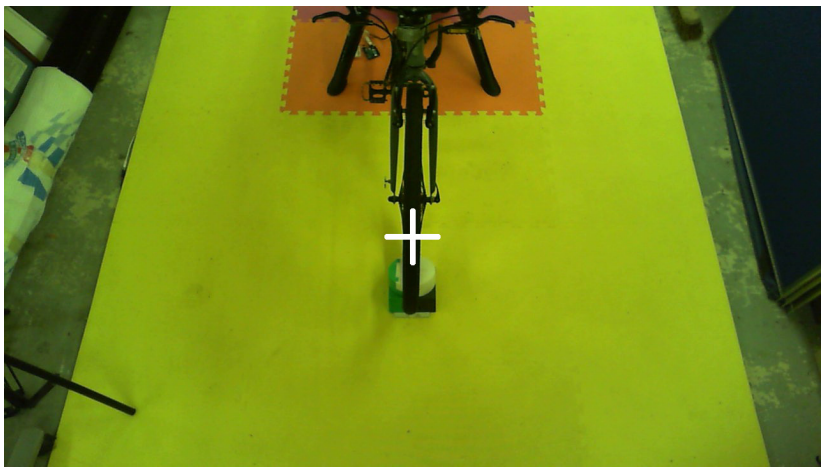

Figure 7.10: Field of view of steering camera with reticule indicated in white

This arrangement takes advantage of the fact that as a user steers to one side, the leading edge of the tyre in the field of view of the camera, moves in a predictable manner; namely, as the hypotenuse of a triangle drawn from the axis-of-rotation of the wheel (Figure 7.11) to the tyre edge (i.e. the radius of the wheel) and a triangle base from the axis to the camera. The observed tyre edge furthest from the 'zero-axis' is the leading edge and provides direction (left or right). Calibration for tyre dimension – both the known radius of the wheel (i.e. the hypotenuse) and a pixel equivalent width (i.e. d in Figure 7.11) for the given set up – allows basic trigonometry to be applied to calculate θ as illustrated in Figure 7.11.

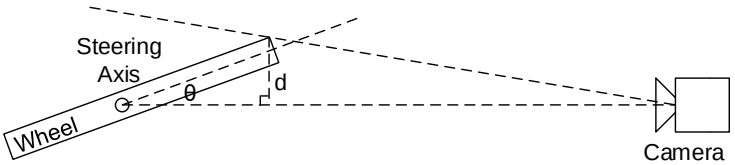

Figure 7.11: Diagram of camera observation methodology (not to scale)

The resulting value for θ is therefore a reasonable approximation of the true direction of the front wheel which is then exposed as a property which the `LocalhostServer`'s threads and methods deliver to the network port.

Unfortunately, getting from a real steering angle to an effective steering angle is a non-trivial operation. As elucidated by Cossalter (2006) with regard

to motorcycles – though the overwhelming majority of the content is directly transferable – the bicycle is acted upon by 15 different forces, yet can only be controlled by two adjustments (speed and steering). As demonstrated by Yin & Yin (2007), an attempt to mathematically define the bicycle-rider dynamic unit is anything but simple. The two rigid-body interactions of a bicycle (the front fork assembly, and the remainder of the frame and tail) are complex, the bicycle itself also being an inverted pendulum. The demonstrable stability of a bicycle in motion remains an open question in the literature (Jones, 2006).

A common simplification of the equivalent for four-wheeled vehicles is the Ackermann simplification (e.g. Hamme et al., 2015) which, amongst other things, ignores vehicle lean and tyre slip. Miah et al. (2016) made a similar simplification for their bicycle dynamic model, however the use of a bicycle demands consideration of bicycle roll. They therefore take equations from Yi et al. (2006) such as:

$$tan(\beta) = \frac{tan(\phi) \times sin(\eta)}{cos(\theta)} \tag{7.2}$$

where β is the effective steering angle, ϕ is the actual steering angle, η is the front fork angle and θ is the frame roll angle.

Note that, if the bicycle roll is constantly zero, as in this case, and one considers the front fork angle constant (assuming, for example, the bicycle has no suspension), then the equation essentially simplifies to an effective angle being equal to the steered angle multiplied by some variable (with a limited numerical range of the range of angles concerned). For a usual front fork angle of approximately 70 degrees, the equivalent factor from actual to effective is approximately 0.9.

Given that the roll angle being zero is a function of the hardware rather than a choice by the user, then even using these simplified equations may not be realistic. In early testing, a simple constant-factor damping (a factor of 0.5) of the instantaneous steering angle was found to be appropriate; i.e. a steering angle of 20 degrees would result in an effective steering angle of 10 degrees.

7.2.2 Simulation Engine

An off-the-shelf computer simulation engine was selected to underpin the system. The Unity engine (Unity Technologies, 2016b) provides a complete and (mostly) platform-agnostic 3D rendering engine and has 'out-of-the-box' support for virtual reality hardware complete with some 11 years of development history, commercial and community support. This is preferable over a bespoke solution as it allowed the focus on the actual simulation and its use, rather than the underlying development process itself. Unity features such as the physics engine can be replaced (where necessary) to suit this particular use case, whilst features such as texture mapping, lighting and post-processing can

be leveraged to make a more immersive simulation. Scripting APIs which use the .NET platform, provide the ability to build the necessary components to deliver this simulator by augmenting the engine in a co-operative manner.

7.2.2.1 Using the Unity Engine for bespoke simulation code

This section briefly summarises the scripting operation of the Unity engine. Those without a programming background (or interest in same) may wish to skip to the next section (Section 7.3). Whilst comprehension of the following does not demand a pre-requisite knowledge of the Unity engine, a reimplementation by the reader would require a working familiarity. To that end, a comprehensive range of learning tools, tutorials and documentation, are noted as available from Unity at `http://www.unity3d.com/learn/`.

The basic unit of object in Unity is the `GameObject` class type (in the `UnityEngine` namespace) which technically extends `System.Object` objects in the Mono.Net hierarchy (the open-source implementation of the C♯ framework used by Unity). In any given deployment, the engine uses one or more 'scene(s)', each of which contains a single 'scene graph'. `GameObject`s are either attached directly to the root node of the scene, or to one another, forming a more complex scene hierarchy.

Each `GameObject` maintains a `Transform` (comprising a 3D location vector, rotation quaternion and a 3D scaling vector) and none, one or more `Component`s which can include a range of engine features such as `MeshRenderer`s (components responsible for getting its attached mesh(s) into the graphics rendering pipeline), shaders (components responsible for rendering textures to screen and interacting with the lighting calculations), `Collider`s (used by the physics engine to test collisions between objects), and scripts. It is by attaching scripts to `GameObject`s that behaviours can be assigned and interactive scenes can be constructed.

Any script which is attached to a `GameObject` must inherit from the `UnityEngine.MonoBehaviour` class which provides the relevant interfaces to the entity's `GameObject` properties and which allow the scene to be run. The result of this arrangement is that simply by use of some 'boilerplate' syntax for access (specifically the pseudo-overriding of methods and some 'under-the-hood' reflection), the engine provides all the mechanisms necessary to enable a complex agent-based simulation to be built and visualised. All scripts attached to `GameObject`s have their methods (if they have them) iterated in the same defined order; Figure 7.12 summarises.

As illustrated by the scale of the flow chart in Figure 7.12, there are many methods which can be overridden to produce desired behaviour. The key methods used in this project for scripts to enable the bicycle hardware to interact with scenes are `Start()`, `FixedUpdate()` and `OnDestroy()`:

`Start()` Called once per script instance and after the script is enabled. Script

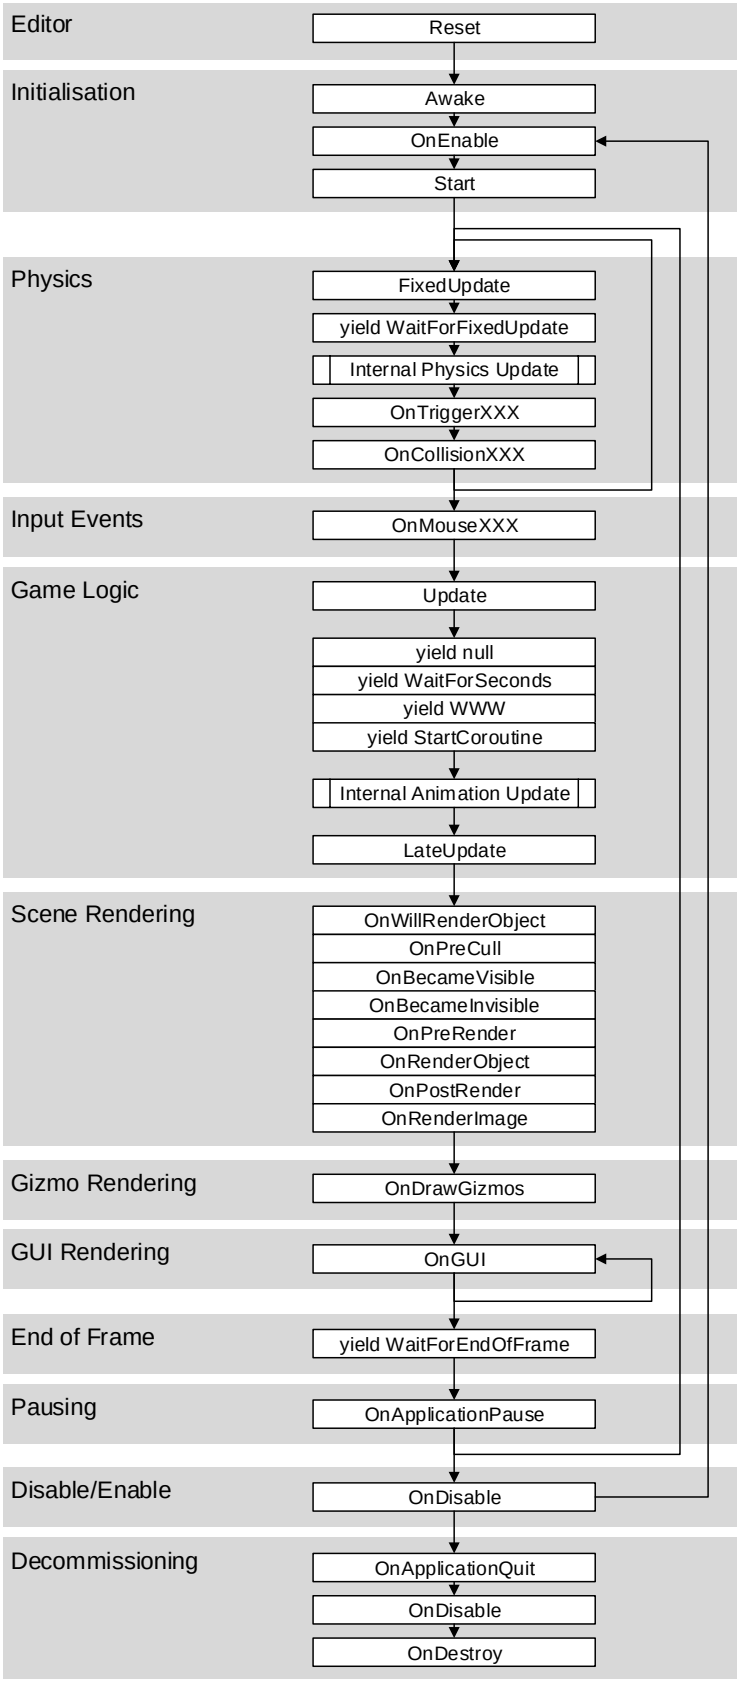

Figure 7.12: Execution order for `MonoBehaviour` implementing objects. (Reproduced from Unity Technologies, 2016a)

constructors are not exposed for `MonoBehaviour` instances so the `Start()` method is used for this purpose.

`FixedUpdate()` Called once per physics engine iteration (default is 50Hz but here because of the higher refresh rate requirements, a rate of 150Hz was used), on a fixed time basis (as opposed to `Update()` which is called once per frame rendered). Updates to the transforms of the hardware agent (Section 7.2.2.2) are therefore taken in this method which is independent of the frame rate being delivered to the HMD or VDU.

`OnDestroy()` Called once a `GameObject` has been set for disposal. This is either directly (e.g. to remove the object from the active scene) or because the scene has been closed. This method provides a way to ensure that the object can clean up after itself. Often not needed, this is used in this project because some of the hardware agent scripts (Section 7.2.2.2) create their own background threads, and these are not managed by the Unity engine. Note that this is not an object destructor.

7.2.2.2 Interfacing code

Given the above operating structure, this author implemented scripting for interfacing with the hardware as illustrated in Figure 7.13.

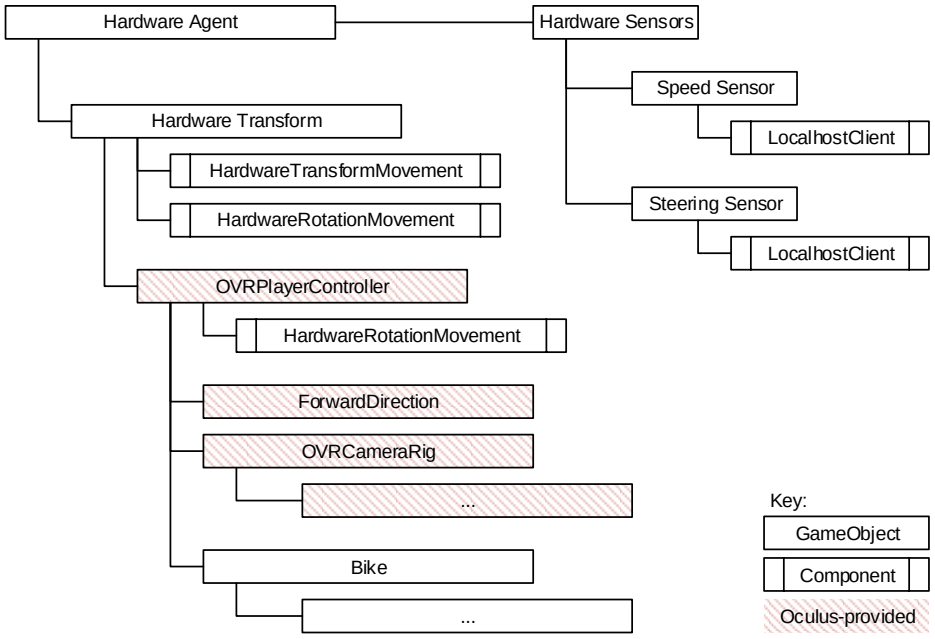

Figure 7.13: Structural software diagram for the Hardware Agent `GameObject`.

Three different scripts were created to enable all the hardware interfacing to Unity:

`HardwareTransformMovement` This script finds all instances of `LocalhostClient` in its `Start()` method and isolates the one responsible for speed data. A `FixedUpdate()` method is implemented that uses the most recent value of this speed data and moves the position of the transform of the `GameObject` to which it is attached, forward by the distance calculated from that speed and the known (constant) time since the last movement was applied.

`HardwareRotationMovement` This script finds all instances of `LocalhostClient` in its `Start()` method and isolates the ones responsible for steering and speed data. A `FixedUpdate()` method is implemented that uses the most recent value of this steering data and applies a rotation to the transform of the `GameObject` to which it is attached, factored by: the distance that the `GameObject` will have moved in the known (constant) time since the last rotation was applied; the circumference of the circle that would have been inscribed by such a steering manoeuvre; the proportion of that circumference that would have been inscribed, given that distance travelled; and finally, the actual angle through which the agent would have rotated in the period of time concerned given those calculations. In testing this was found to be far too sharp a response (Section 7.2.1.3), so the steering was damped by a factor of 0.5 which was found to be appropriate (although see Sections 8.4.2.1 and 8.4.3.2 later).

`LocalhostClient` This script spawns a thread (in `Start()`) which listens on a given network port, sets up a TCP connection and receives a unit of data. This repeats continuously (including a small delay so as not to monopolise the system resources) until the tread is terminated by a flag set in the `OnDestroy()` method. The received data is then parsed and stored as a property which is available to the rest of the simulation. Two instances of this class are used, each listening on a different network port; one port (and thus `GameObject`) receiving speed data and the other receiving steering data. The ports used here are local loopback ports so all the software runs on the same PC; however, this set up allows future expansion where the computer physically connected to the hardware need not be the same computer running the simulation itself.

All the relevant `GameObject`s are attached to a single parent `GameObject` node, as are instances of the `HardwareTransformMovement` and `HardwareRotationMovement` classes, which ensures that they all move collectively by sharing a common root `Transform`.

The `OVRPlayerController GameObject` is a 'prefab' (one or more `GameObject`s saved as a single asset outside of the scene and which can be reused as a template and be placed as an instance or instances in that or any other scene) provided by Oculus as a drag-and-drop way of integrat-

ing the Oculus Rift VR equipment into a scene. Amongst other things, it contains `GameObject`s with `Camera` components attached, to provide visual inputs for the HMD. The `OVRPlayerController` instance also has a `HardwareRotationMovement` script attached despite being a child node of the overall Hardware Agent, to ensure that the 'head' of the `OVRPlayerController` turns when the agent does. Otherwise, head rotation is (by design) independent of body rotation.

Finally, a bicycle mesh is attached to the Hardware Agent `GameObject` such that if the user looks downwards, they see a properly orientated bicycle underneath them, rather than nothing at all. However, no body avatar was implemented.

7.3 Design Conclusion and Progress

The previous chapter (Chapter 6) presented the literature basis for the development of a simulator (Sections 6.1 and 6.2) and developed a specification against which a bicycle simulator could be built (Section 6.3). Consequently, following the development and construction of the necessary and appropriate hardware and software components presented in this chapter, there is now a bespoke simulator in hand which is ready for testing with participants. Of particular value, this document has detailed both the design specification process and the implementation, which allows the reimplementation of this set-up by others; something not currently possible from the published literature.

A range of design barriers have been overcome and addressed as necessary. Off-the-shelf and purpose-built components have been developed and brought together by this author, and a combined hardware and software system built which is capable of detecting the state of a standard bicycle (Specification Item 1) in terms of effective speed (Specification Item 2) and steering direction (Specification Item 3). The system simulates a 3D environment which is agnostic to the visualisation method used provided such can be connected to a standard PC. This ensures a live and interactive visualisation can be presented to the participant (Specification Item 4) and that the calibration of the visualisation is a function of that hardware (and can be calibrated as such; Specification Item 5) and does not depend on the underlying simulation. In achieving these 5 specification items (Section 6.3), the core requirements for the simulator have been met.

In addition, with the limited exception of applying (removable) reflective tape to the bicycle spokes, there is no fundamental barrier to participants' use of their own bicycle (Specification Item 6). Further, low overhead hardware has been created which provides the implementer with an ability to minimise latency or choose latency as a trade-off where this has measurement effects (Specification Item 9). These items, combined with the implementation in the following chapter (Chapter 8) of the other 'desirable' points in the simulator specification defined in Section 6.3, ensures a complete achievement of the entire specification.

The next chapter therefore takes this simulator and tests its efficacy with human participants.

Chapter 8

Simulator Testing

Thus far, this document has showcased the lack of robust quantitative literature relating to cycle infrastructure (Chapter 2) and the case for the need to properly quantify the same. Chapter 4 then went on establish the three obvious avenues for collecting the type, experimental rigour and scale of data that is required to provide the sorts of data sets required in order to begin to build a canon of empirically-backed literature and practitioner tools. Sending participants out on instrumented bicycles was discounted therein as both not being appropriately scalable and not allowing the experimenter to capture the full experimental circumstance without substantial experimenter effect; both of which are key (though not necessarily mutually inclusive) requirements to develop the scale of empirical data needed to address that identified lack of robust quantitative cycle literature.

However, whilst Chapter 5 presented a highly-scalable methodology for collecting cyclist data, it could be potentially limited in some aspects by an inability to apply experimental controls or by not physically having access to the given circumstance that one wishes to test. A simulator would serve to address some of these limitations and so the previous chapters set out the case for this, the literature basis, a specification, and the development process for both hardware and software.

This chapter now takes the simulator delivered at the end of the previous chapter, and tests it with human participants. This chapter establishes what aspects were successful, those where improvements need to be made, and frames the outcomes in terms of comparative data from reality. Specifically, that data collected in Chapter 5, and using the same participants from that study to ensure direct comparability.

8.1 Simulator Test Set-up

8.1.1 Scenarios

At this stage of development, the primary concern is validating the operation of the simulator both in terms of objective fidelity and perceptual fidelity (see Section 6.1.1). Also, as discussed in Section 7.2.2, simulation scenarios are prepared in the form of Unity 'scenes'. Finally, in Section 6.3 a need to provide a 'learning' scenario was noted in an effort to allow the participants to adjust to the controls and to reduce the potential for SS.

Two scenarios are therefore proposed:

"Driving School" A scenario with a minimum of diversity and requirements, this scenario will present the participant with the control scheme in a featureless flat plane. An area will be set out with road cones in various configurations (e.g. straights, curves of varying radius, slalom) which will allow the participant to navigate around or through them to build their confidence. By avoiding the inclusion of any other features, the participant will focus on basics of the control, comfort and operation of the simulator. No data will be analysed for this scenario so as to further increase the comfort of the participant and avoid their fear that they might be "doing something wrong".

National Oceanography Centre Car Park As the desire here is to replicate real behaviour and a selection of empirical data pertaining to real behaviour was collected in Chapter 5, a recreation of the NOC will be created. This scenario will replicate the layout as explored by the participants in reality, with focus on creating a high-quality replication as more realistic virtual environments are reported to result in increased presence for the user. Where possible, 'photorealistic' textures will be used for surfaces and objects, and details such as audio will be included to further increase the chances of immersion. Section 6.1.3 considered the literature basis for these inclusions.

The following sections detail the key parts of each of the two scenarios:

8.1.1.1 "Driving School"

A plane of 500m × 500m is textured with asphalt. This extent is provided such that when in the centre, the scene is perceived to be endless in all directions. In the centre of the plane, a 100m × 100m area is enclosed with red/white textured meshes of water-filled barriers. Within that area, red/white cones were laid out as illustrated in Figure 8.1.

The participant starts in the position nearest the camera in Figure 8.1 and is then free to explore the area at their own pace (Figure 8.2). The starting

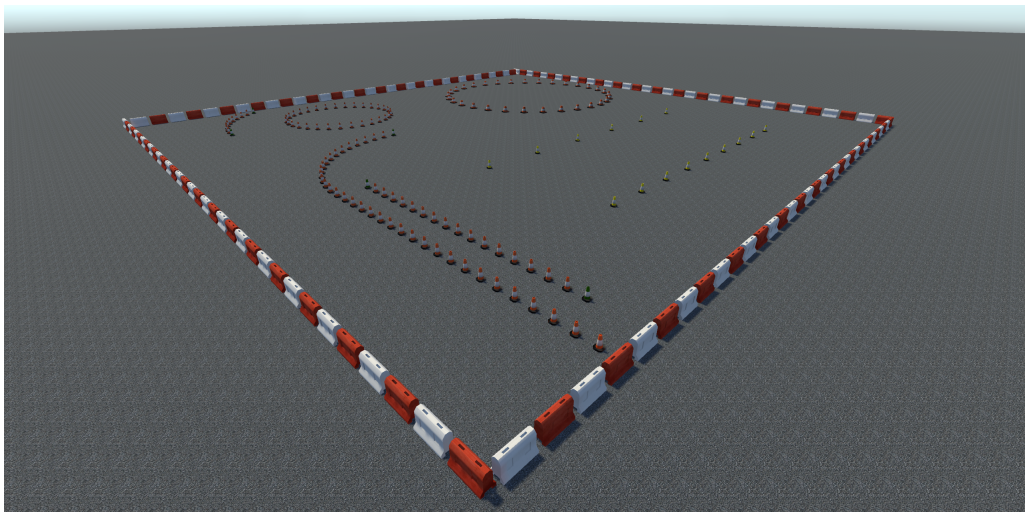

Figure 8.1: General Arrangement in-simulation of ‘Driving School’

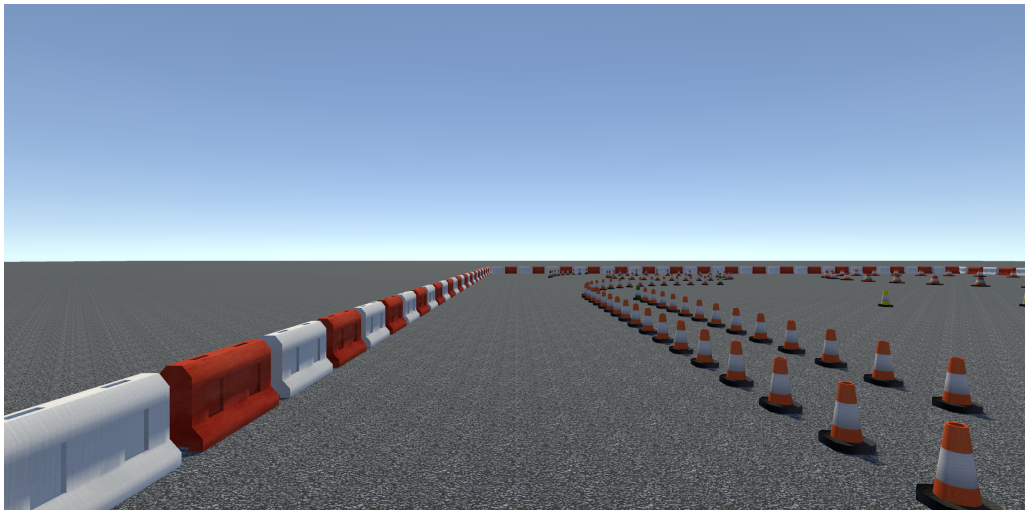

Figure 8.2: Screenshot of starting view in ‘Driving School’ simulation

section is straight and so this encourages the participant to explore the forwards start/stop of the simulator before moving on to the steering. However, the colliders for the cones are deactivated, as are those for the waterfilled barriers. And so, if the participant wishes to immediately move off in a different direction, then they are free to do so.

8.1.1.2 NOC Car Park

Buildings for the site and surrounding area are imported from the Ordnance Survey Buildings Data Set (Ordnance Survey, 2016). Buildings were then normalised in altitude to match the ground plane as whilst there is some minor variation in ground elevation in the data set, the site is essentially flat as it is built on reclaimed land. Normalising building heights ensures that all buildings appear at ground level and do not 'float' which would be visually odd and would likely break immersion. The exterior of the NOC building was textured with a light-coloured brick, as per reality. However, detailed textures including windows and other exterior features were not used owing to the resources required to do so. Kerbed areas within the NOC car park were extracted from CAD drawings provided by NOC Estates & Facilities (see Figure 5.1), extruded to standard kerb height and imported to the scene. The result is the environment pictured in Figure 8.3.

Figure 8.3: General Arrangement in-simulation of NOC

The ground plane was mostly textured with asphalt, and those areas which in reality are gravel, with a gravel texture. Solid 'fence' textures were placed around the perimeter in place of the real fences. These fences were then obscured with a selection of tree objects (with a randomisation applied to the extent of the scales and orientations so as to avoid an impression of 'cut-and-paste' objects, potentially breaking immersion), as is the case in reality. The kerbed areas had foliage added to provide an equivalence to the shrubbery present in real life, and those areas also had individual tree objects placed in

a position and scale to match reality. A light wind zone was added to the scene to provide some movement to the trees/foliage and add realism.

Figure 8.4: Screenshot of NOC simulation showing ground and markings, hedges and foliage, and marked way

Ground markings to match reality were constructed on the CAD drawing base then imported to the scene as a layer ordered above the ground plane (Figure 8.4). Similarly, blue lines representing the rope lines used were added, as were small cone objects to supplement (per reality).

To further increase the potential for immersion, a 10 minute ambient audio track (featuring ambient background traffic noise, seagulls, birds and a passing aircraft) was recorded in Southampton City Centre which provided a good representation of that which would have been experienced in the original study.

Finally, various objects were added to the scene to further replicate reality: a shipping container, vehicles (all the same SUV-type so as to avoid spending resource time replicating various vehicle models) in places where vehicles were parked in reality, and wooden barrier separators in the parking areas (Figure 8.5).

Participants all started the scenario at the same place and with the same orientation (Figure 8.6).

8.1.2 Experimental Set-up

Following the design, construction and preliminary testing associated with the development presented thus far, a small pilot study was organised to test the simulator on actual cyclists that were not involved in the design process for the simulator.

The group of cyclists from the practical study presented in Chapter 5 was re-approached and 5 of the 6 participants re-recruited (still postgraduate students; 4 male, 1 female; ages 25–33) for this study. This study consisted of:

- Familiarisation with the bicycle simulator equipment;

- Familiarisation with the VR HMD for those sessions using VR;
- Introduction to and undirected exploration of the "Driving School" scenario to ensure familiarity with the operation of the bicycle in the virtual space (Section 8.1.1);
- Cycle-based circulation of the NOC car park in virtual space, analogous to the October 2014 study (Chapter 5);
- Administration of a Simulator Sickness Questionnaire; and
- Semi-structured debrief: Participants were directed within the scope of a free-form discussion, to consider certain aspects of the study (e.g. perceived level of immersion, latency in the control system, etc.) if they were not already doing so.

In the screen-based situation, the use of the VR HMD was replaced with a large (107cm) computer connected LED TV screen display at approximately 1.5–2.0m in front of the participant (visible in Figure 7.7). In that situation, the HMD familiarisation was not necessary; all other procedure remained unchanged.

Owing to the experimental space and equipment available, switching between VR and screen-based operation was physically non-trivial. As such, the order of exposure was not randomised amongst the participants. However, given the presence of the training scenario and the small sample size, it is unlikely that a learning effect could have been isolated. In any case, the interval between the two sessions ranged from 5–28 days ($\bar{x} = 13$ days).

8.1.3 Research Questions

As established in Chapter 4, the primary purpose of the creation of a simulator is to gather data on the cyclist in experimentally controlled conditions. Then, as further established in Chapter 6 – specifically Section 6.1.1 – this can be further broken down to a desire for objective fidelity (that the circumstance *feels* like riding a bicycle), that a participant consequently feels immersed in the simulation and thus that the simulator achieves a perceptual fidelity (that the participant behaves as if they were riding a bicycle). As such, three broad research questions arise:

1. Does the simulator feel, to the participant, like they are actually riding a bicycle?
2. Do participants feel immersed in the simulation? Does the use of VR increase the feeling of immersion?
3. Does the participant behave as they would (or did) in reality?

Whilst not highly specific, these research questions are reflective of the exploratory nature of this project at this stage.

That said, the literature basis presented in the previous chapter (Chapter 6) presents the opportunity to capture some additional and more specific research questions which can be tested with regard to simulator sickness:

4. Do participants experience Simulator Sickness?
5. Does the reported level of Simulator Sickness increase with speed of movement in the simulation? Per So et al. (2001), see Section 6.2.2.1.
6. Does the recorded steering reversal frequency decrease with increased reported level of Simulator Sickness? Per Muttray et al. (2013), see Section 6.2.3.

8.1.4 Data Collected

A set of data relating to:

- Simulation time;
- Rendering frame rate;
- HMD 3-axis orientation (fixed zero for the Screen scenario);
- Bicycle 3-axis position in the simulation;
- Bicycle 3-axis orientation in the simulation;
- Recorded speed (from the network service); and
- Recorded steering angle (from the network service).

were collected from the participants for subsequent exploratory analysis. Whilst 3-axis data were recorded, it should be noted that the bicycle was fixed in position in the y-axis (vertical) and fixed in orientation in the x and z axes. Data points were collected at each physics time step of the Unity engine (Section 7.2.2), at a rate set to 150Hz.

With the exception of the high level desire to assess the quality of the simulation and interactivity, the specifics of the subsequent analysis (e.g. consideration of speed in Section 8.3.1, paths through the corner presented in Sections 8.3.2.2) was not detailed to the participants so as to avoid influencing their behaviour. Informed and written consent was obtained at the start of each session in accordance with the approved ethics: University of Southampton Ethics ERGO ID: 18464 for the study.

The analysis that follows considers the NOC scenario so that it can be framed in the setting of the previous study (Chapter 5). All statistical tests are performed at the 5% significance level. Prior to a consideration of the quantitative outputs from the experiment, it is useful to consider the larger scale qualitative results.

Figure 8.5: Screenshot of NOC simulation showing container, vehicle, and wood barriers

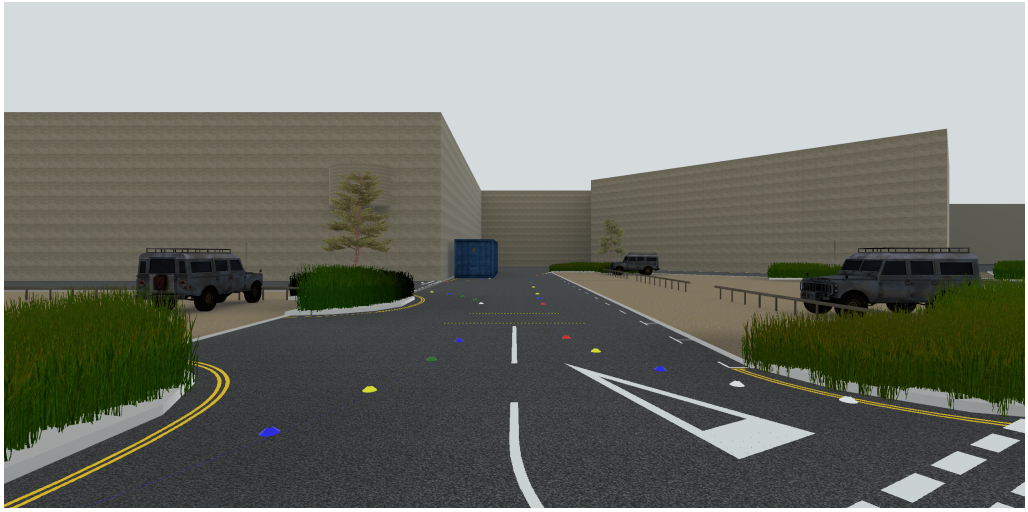

Figure 8.6: Screenshot of starting view in NOC simulation

8.2 Qualitative Outcomes

The primary source of qualitative data from the study was in the form of verbal semi-structured discussion. Participants were free to verbalise their thoughts as they arose during the study and these were documented without interrogatory response (so as not to distract the participant) as they arose.

Then, following completion of the study, whilst the participant was seated off the simulator (but after the SSQ had been administered), comments that had been made were repeated back to the participant with an appropriate request for elucidation if any. For example: "You said during the simulation that you felt the simulation is 'not photorealistic but your brain takes it for real'. Could you expand on what you meant by that? Was there anything you felt was particularly unrealistic, or realistic?"

Such questioning continued until noted material was exhausted. Finally, the following topics were prompted for further thoughts (if any):

- Level of immersion;
- Feelings of realism in simulation;
- Feelings of realism in control equipment (pedalling and steering); and
- Any other thoughts.

It should be noted that whilst the tone of the discussion was maintained as deliberately informal, care was taken to ask neutral and open questions so as to ensure the participant did not provide responses that they thought the experimenter wanted to hear and that responses were properly representative of the participants' thoughts.

8.2.1 Immersion

Participants in the VR situation universally reported that their feeling of immersion in the simulation was high. All but one participant had no substantial experience of virtual reality and in that single case, was earlier developmental equipment, not a release-grade consumer device. Participants seemed, at the least, surprised (if not openly impressed) at the visual quality of the recreation: "Wow, you've recreated that well [...it is] not photorealistic but your brain takes it for real."

The lack of obvious visual cues to break immersion allowed the participants to become mentally involved in the scenario, as was intended. Indeed, such was the level of immersion that actions that one would avoid in real life, such as cycling through a solid object, elicited a visceral reaction: "Riding through things was a psychically [sic] uncomfortable experience." and "[...] you felt like you were going to hit a cone [in the Driving School scenario] if you went [in]to it."

Not all participants qualified their immersion so strongly. One participant stated that it "mostly feels like I am riding a bike [... but] you don't forget [reality] because it doesn't look real."

In the screen-based situation, reported immersion was much more variable. Some participants did not feel immersed in the experience at all: "No [I did not feel immersed], it felt like I was playing a video game. It wasn't a particularly immersive experience.", "I was very aware I was in a garage.". Whereas others were able to become sufficiently immersed that they felt involved in the scenario: "It's immersive enough that I feel like I have to be careful about how [I was] cycling. I feel like if I were to crash into something, it would be bad."

The VR situation saw one participant withdraw after only 1m30 due to SS and therefore that participant did not experience the VR scenario at the NOC. In the screen-based situation, this participant reported: "the first three minutes felt like a long time, the last seven flew." Losing track of time is a common 'symptom' of immersion/presence (Brockmyer et al., 2009).

The NOC scenario featured an ambient audio track to in an effort to further increase the potential for immersion. No participants noticed the presence of this audio, however, the real environment of the simulator was adjacent to some exterior building works which may have served to overwhelm the subtlety of the audio track. Similarly, details such as seagull calls in the audio, may have been interpreted as being exogenous to the simulator, when they were in fact endogenous.

In summary, whilst immersion was universally reported to be high in the VR situation, the reported immersion levels in the screen-based situation were much more variable. To some extent, the full field-of-view nature of the VR works to force a level of immersion on the participant, whereas the screen-based experience is likely to be more voluntary. In general, where immersion was achieved, obvious breaks of immersion were minimal and non-obvious, and mostly related to issues between the participant and bicycle, or were exogenous to the experimental set-up. Improvements to the simulator interface, and a less 'rough-and-ready' experimental space, would both help to address these aspects.

8.2.2 Steering Function and Latency

Whilst participants universally agreed that the steering did not feel like it had any lag in it, all participants struggled to some degree. It was universally agreed that the steering felt over-sensitive ("very very sharp", "the amount you had to turn the bike [to steer a certain angle] seemed remarkably small", "[The steering] feels very sensitive; I only have to turn it a very little.").

One participant adapted – i.e. was able to control the steering in a smooth and definite manner – very quickly to the steering but this participant was the one that was terminated due to simulator sickness after less than 1m30s. An-

other participant also adapted quickly however it transpired upon debrief that this participant has ceased to be a regular cyclist between the October 2014 study and this (July 2016), and has past experience of assisting the development of a human–vehicle control interface (although that case was a motor vehicle steering assembly). Given the small sample, it is difficult to separate these out as merely coincidences, however a third participant also adapted quickly despite this participant being a regular utility and recreational cyclist. The remaining two participants were observed to continue to operate the steering in a somewhat 'jerky'/overcompensatory manner throughout.

One participant was able to elaborate on one aspect of the steering that was unrealistic; that being, the lack of increasing resistance to turning as speeds increase: "If you're going fast, there was no resistance [...] If it were a real bike, you would expect the bike to fight back." In reality, when one is cycling at speed, the rotation of the mass of the front wheel generates a gyroscopic effect which, amongst other things, provides an effective resistance to the turn being applied by the rider. The simulator provides no such resistance, being as easy to turn at no speed at all as it is at a high simulated speed. A more complex steering support mechanism could theoretically apply such a resistance, or an effective resistance could be perhaps simulated by simply damping the handlebar–steering sensitivity as simulated speed increases, though this may only be effective if the rider could not see their arms.

A further damping of the steering sensitivity may be sufficient to address the user complaints of over-sensitivity. Whilst this should be considered in consort with the literature considered in Section 6.1.4 which indicates that body motions are amplified in VR, it should also be noted that the complaints of over-sensitivity persisted into the screen-based situation, albeit that the impact on actual steering outcomes seemed less extreme (see Section 8.3.2 below).

More of an issue than the steering itself however, is the intertwined issue of leaning (i.e. roll rotation); discussed later (Section 8.2.4).

8.2.3 Speed/Pedal Function and Latency

Generally participants felt there was a close agreement between their pedalling action and the response of the simulation. The potential for a lag related to a sudden stop (see Section 7.2.1.2) was not discovered by the participants as extreme/sudden stops were not a feature of the simulation scenarios presented.

The bicycle was set to a medium gear (38/18 gearing ratio) and the participants were instructed to not change gear as, amongst other issues, they could not see the handlebars while wearing the VR HMD, and so would not know what changes they were making. One participant elaborated on this constraint: "You pedal a lot but don't move the wheels much. It doesn't accelerate fast, but [it does behave] naturally. [...] When I cycle, I switch gears a lot, so that was weird." Another participant summarised their feelings on the pedalling in-

terface as: "[I] felt like I was pedalling reasonably consistently and the bike was being consistent."

No participant appeared to explore freewheeling – continuing to ride forward under the rider's existing momentum without pedalling; equivalent to 'coasting' in a motor vehicle – in any detail, though two participants noted this explicitly afterwards. To a great extent, the lack of freewheeling is a function of the trainer used in that the resistance arrangement does not freewheel in a realistic manner, coming to a stop relatively rapidly after pedalling is suspended. However one participant normally rides a fixed-gear bicycle, so actually found this constraint improved their immersion: "[The use of the rear-wheel trainer] was natural, you stop pedalling and you stop."

8.2.4 Roll Rotation

As the bicycle was held in a fixed trainer, there was no ability for the rider to lean into turns. This was almost universally a substantial issue for the participants, with a number of them appearing to come close to falling from the bicycle whilst wearing the VR HMD. Bear in mind the HMD fully obscures their vision and the support for the bicycle adds approximately 5cm of clearance from the ground which potentially puts the participant out of the range of being able to put a foot on the ground to stabilise themselves. After a few minutes, most seemed to adapt to not leaning (either consciously or subconsciously) to some extent, however often appeared to revert to a latent desire to do so when inadvertently cycling into an object or performing some other action which distracted them from thinking about the operation.

One participant (of those that struggled the most with steering) stated: "I would always physically lean whilst turning but this [bicycle] doesn't, so this gives you a slightly odd sensation that you're about to throw yourself off the bike which is slightly disconcerting." Another noted that at slow speeds: "If it were a real bike, you would have fallen. [...My] brain says 'panic' when you're slow because you're going to fall over, [but] only when stopping, not so much of an issue at high speeds — except when crashing of course." A third: "This is weird; I feel like I'm learning to ride a bicycle again."

Another participant stated: "[It] felt like you should lean, but obviously you don't. That wasn't much of a problem." However, this should be framed in the context of being the participant that necessitated termination after 1m30s. It is eminently possible that this *is* in fact a problem, and that the problem manifests as simulator sickness.

A common issue is a difficulty for the participants to articulate the effect of their experience of the effect of the lack of physical roll: "[I was] expecting to be tilting and I just wasn't. I feel like I was trying to overcompensate for something but that wasn't happening in reality." The specifics of the 'something' could not be further elucidated, despite further prompting to do so. More gen-

erally, a number of the participants came to a similar conclusion with regards to dealing with the issue of the lack of ability to apply a roll motion. “I was having to fight what my body wanted to do. It was so much like riding a real bike that it didn’t quite respond how I was expecting it to”.

The necessity of the rider to lean into corners is an integral part of maintaining balance on a bicycle. However, the dynamics of a bicycle is itself is a complex system and this reasoning removes the agency of the rider as, additional to the need to balance, the gyroscopic effect of the front wheel (when turned) results in a force that encourages steering further into that direction of lean. In the absence of any lean to begin with, the act of leaning in of itself can cause a steering action, thus necessitating the leaning action that caused it. A simple demonstration of this can be witnessed by watching a cyclist who is able to cycle comprehensively without holding on to the handlebars. The question then generated by this act of agency is, in the case of ‘usual’ riding: Does the rider lean in order to turn, or does the rider turn and thus need to lean?

In reality, the two actions may be inseparable. At the very least, although leaning whilst using this simulator is physically possible, the lack of steering response and therefore forces experienced in reality versus those expected, generates obvious vestibular confusion. It is potentially the case that this is a contributory factor to the incidences of simulator sickness discussed below (Section 8.2.5).

8.2.5 Simulator Sickness

After each session, a SS questionnaire (SSQ) was administered in accordance with the procedure in Kennedy et al. (1993) – i.e. postexposure only and results excluded for those who were “sick” or in anything “other than their usual state of health”. Consent forms contained a health declaration as, given the set up, already being unwell was grounds for exclusion from the study. Consequently, all SSQ returns were valid for inclusion. No other pre-screening was performed.

It should also be restated (see full discussion in Section 6.2.1.1) that despite the SSQ scores being the ‘gold standard’ for quantifying SS, they are based on ordered categorical subjective scoring and that scores are not linear in terms of symptom severity – hence this consideration of the SS’s placement in the Qualitative Results section of this document – or even realistically inter-comparable between individuals. However, in the aggregate, they can provide information which can be informative.

In 60% of cases here, VR sessions were terminated due to simulation sickness (SS) on the part of the participant. Furthermore, as noted earlier, one of those participants succumbed to SS in under 1m30s. The remaining two participants completed a full 10 minutes in the NOC scenario.

Interestingly, one of the completed participants reported not experiencing

SS at all, yet responded to the SSQ with a symptom set (and scoring) that agrees well with the other participants, indeed the reported total score is actually higher than some participants that were terminated due to SS. Given that the participant could not have known the responses from the others, it is safe to assume that they were answering truthfully but perhaps did not consider those symptoms to be SS. Although objectively, it should be noted that the participant responded "moderately" to the query as to the status of the nausea symptom so it is not known what experience would have been necessary for them to say they felt unwell.

Administration of the SSQ highlighted non-trivial sickness scores in both this participant and the other that completed the full session. This perhaps illustrates the value of the SSQ in elucidating the complex symptom intermix of SS and potentially highlights a mismatch between simple high-level self-reporting and the more detailed reflection upon symptoms that the administration of a SSQ invokes in the participant. This contrasts somewhat with Helland et al. (2016) where they considered a simple single high-level self-reported score to be valid.

More generally, in a number of cases the participants were surprised to be suffering from SS as they did not consider themselves usually susceptible to motion sickness, perhaps inadvertently highlighting the different causation methodology for SS versus motion sickness (Section 6.2.1).

In the screen-based situation, all five participants completed a full 10 minutes in the NOC scenario. Informally, none reported feeling unwell and formal SSQ scores generally agree with this; although one participant reported a non-trivial score in the oculomotor symptom cluster.

Table 8.1: Summary Simulator Sickness Questionnaire Scores

Situation	VR ($n = 4$)	Screen-based ($n = 5$)
Nausea mean	66.8	7.6
Oculomotor mean	12.1	6.1
Disorientation mean	33.4	5.6
Total Score mean	41.1	7.5

Results of the SSQs are summarised in Table 8.1. Symptoms across the individuals are focused within the nausea and disorientation symptom clusters in the VR situation and generally low in the screen-based situation. This clustering, being similar across the participants, would seem to indicate a common underlying causation and is in line with the expected symptoms of SS due to vestibular causes (i.e. the simulator itself), as opposed to the potential for oculomotor symptoms derived specifically from the use of the HMD. Other data supports the assertion that the HMD was operating optimally, specifically that it operated at a consistent 90Hz refresh rate, as designed. For reference, the native refresh rate of the screen used was 60Hz and the simulation therefore also

operated optimally on this hardware.

By maintaining all other settings the same and changing only the visualisation method, one can compare 'like-for-like' in terms of immersion and SS, considering the visualisation method to be the experimental factor. Such experimental targetting of possible SS factors is something that is noted as lacking in the literature (Classen et al., 2011; Kennedy et al., 2010). Whilst literature such as Bos et al. (2008) would suggest that 'vertical mismatch' is sufficient to cause SS, and here the symptom cluster does appear to be similar between the situations (indicating that reported SS may be independent of the visualisation method), reported SS is seemingly linked to degree of immersion and a vertical mismatch should be found in both situations as the equipment (and the lack of ability to roll) is unchanged. However, a lack of quantitative measure for immersion and the small sample here, does not allow this to be quantified as a result of this study.

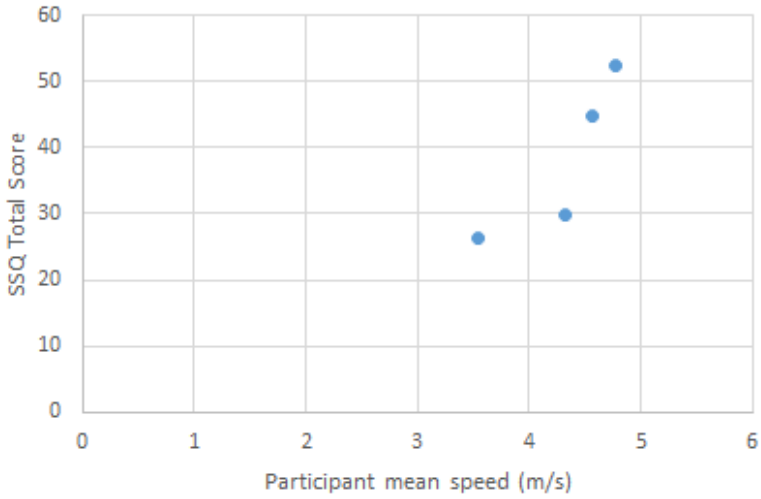

Figure 8.7: Virtual reality participant simulator sickness questionnaire total scores versus mean speed

Section 6.2.2.4 noted research by So et al. (2001) that indicated a positive relationship between SSQ scores and speed of motion through an immersive scene up to approximately 8–10ms^{-1}. Figure 8.7 presents the data from this study (in the VR situation) and whilst a similar relationship appears visually to be the case, due to the small sample size, no conclusions of significance can be drawn with regard a possible linear correlation for this data at the 5% significance level ($r = 0.871$, $T(2) = 2.519$, not significant, one-tailed $p = 0.064$).

8.2.6 Qualitative Conclusions

Generally, users of the simulator experienced high levels of immersion with the VR and lower levels of immersion with the screen visualisation. This would seem to support the intuitive assumption that more immersive visualisation

leads to greater immersion for the participant. It would further support a case for pursuing such comprehensive visualisation techniques in the future.

This however must be contrasted against the incidence of simulator sickness which was substantial and near universal in the VR situation. Clearly, a simulator that makes all participants ill is undesirable and whilst some level of SS can be accommodated by simply recruiting larger participant groups, comprehensive SS is impractical. A potential link to the inability to roll about the longitudinal axis may be present, although this study was not set up to identify it, and may be worthy of future exploration, given that this is a fundamental activity on a bicycle. If the (not insubstantial) issues surrounding the prevalence of SS can be addressed, there is clearly a wide scope for the use of this simulator for qualitative purposes.

Interfaces with the simulation were considered to be generally good, if not perhaps a little too effective in some cases (i.e. the "sharp" steering). Adjustments can be made to future uses of the simulator to account for this, and thus should not preclude further deployment of the equipment. That said, the simulator arrangement is arguably in a state which is objectively difficult (though not necessarily impossible) to improve further: "I don't know if there's much you can do to improve as you're fighting against biology. It's kind of fundamental."; "[The movement of the bike was an] oddly smooth ride [as compared to a road surface]. I don't know if that would have made me feel more sick."; "[I] experience[d] a slight sense of nausea which peaks when I ride through things." (NB. There was no collision detection active in any of the scenarios); and "There [was] no wind feedback."

By contrast, reducing the effort to provide realism is not an objective improvement which was even identified by the participants: "It would be interesting to know if the physical sensation of resistance in steering with speed were actual or if it is more just the less manoeuvrability you have got? But then, the more of that stuff you do, the less it is riding a bike and the more it is [a game] controller."

With the above in mind, the following moves on to consider the quantitative data collected from the simulator, and considers it alongside the real-life data collected in Chapter 5.

8.3 Quantitative Results

The foregoing considered the qualitative outputs of the pilot simulator study. The nature of the data collected also allows more objective quantitative results to be presented and interrogated. The following presents a comprehensive consideration of that data.

8.3.1 Speed Selection

Figures 8.8a, 8.8b and 8.8c show histograms of a single participant's speed selected in each of the three situations: real-life, VR-based on simulator, and screen-based on simulator, respectively. For reference, Participant 1 in this chapter and in Chapter 5 are the same. Although Participant 1 struggled with the steering in the VR situation, the speed histograms are representative.

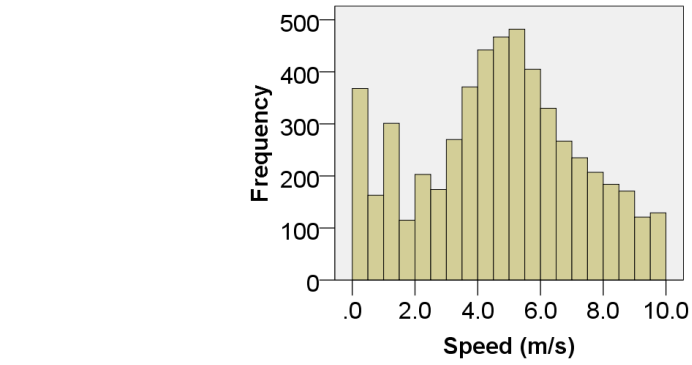

(a) Real life
(NB. truncated at 10ms^{-1})

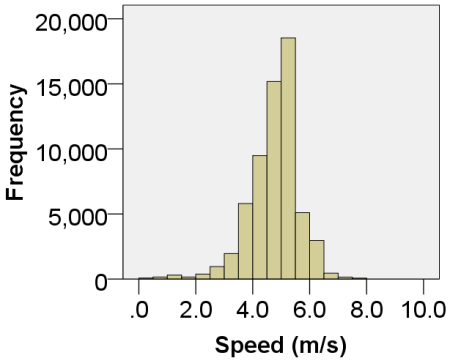

(b) Virtual reality

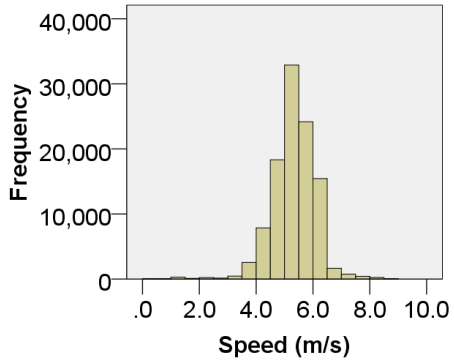

(c) Screen-based

Figure 8.8: Speed histograms for Participant 1

The real-life set (Figure 8.8a) includes a pair of peaks of data at and near zero owing to the methodology of calculating the speed; specifically, as speed is derived from participant position in the recorded video, the pixelation of the video causes a low-end quantisation of the calculated speed. Smoothing the raw data would remove these peaks as it should be noted that there are data points which have been truncated at the high end (e.g. in the tens, hundreds or some 800ms^{-1} for one point with another participant). As can be corroborated with

the original video evidence, at no point was any participant stopped or nearly stopped (nor travelling at 800ms^{-1}). The VR (Figure 8.8b) and screen-based (Figure 8.8c) data sets have had zero speed points excluded as these were: accurately reported; do not represent a valid component part of the selected-speed data sets in these situations; and occur at the start/end of the session prior to and after the participant was engaged with the simulation.

Table 8.2: Summary Speed Statistics

Situation	n	$\bar{x}$ (ms^{-1})	s (ms^{-1})	Median
Real-life	6	6.84	9.96	5.14
VR	4	4.41	0.83	4.58
Screen-based	5	5.23	0.69	5.37

Table 8.2 details some selected descriptive statistics for the speed data collected from the participants over the two studies/three sessions. Visually, the mean speed of participants is substantially higher in the real-life situation, however this (arithmetic) mean is skewed by the outlying high values resulting from the image analysis; this is evidenced by the arithmetic mean of the medians of the participants' individual arithmetic mean speeds being 5.14ms^{-1}, versus their weighted arithmetic mean of 6.84ms^{-1}. The impact of these outliers is also seen in the standard deviation value an order of magnitude higher than in the other two situations. As they are therefore more representative, the participants' median speeds were compared for the various combinations of the three situations using Mann-Whitney Rank-Sum tests:

Real-life vs VR $U = 1.71$, not significant, two-sided $p = 0.088$;

Real-life vs Screen-based $U = 1.83$, not significant, two-sided $p = 0.067$; and

VR vs Screen-based $U = 1.71$, not significant, two-sided $p = 0.086$;

In all combinations of situation, the median of the participants' individual mean speeds were not significantly different to one another at the 5% significance level. The implication of this result is that the selected speeds are broadly similar in the the three situations and can be considered comparable.

Helland et al. (2016) showed that participants tend to travel more slowly in their simulator (although they were concerned with motor vehicles) when they experienced higher levels of SS. Incidence of SS was higher in the VR-based situation (Section 8.2.5) and speed here was numerically lower (though not statistically significantly so; $p = 0.086$). Although the confounding factor of the physically different visualisation experience cannot of course be separated, the SS effect may be partly explanatory of these observations.

That said, two issues should be noted. Firstly, the sample of participants is small in each situation, limiting the statistical power of the tests (see Sec-

tion 9.2.2). Secondly, the selected speed is, to some extent a function of the turbo-trainer's resistance and settings. Given this is a commercial product in a competitive market, a substantial parallel experience between chosen speed on the trainer and chosen speed in reality is a specific design feature, and the product development process has explicitly sought to minimise the difference as perceived by the cyclist (Tacx B.V., 2016b).

8.3.2 Steering Selection

As no recording of steering was available in the real-life situation, this section presents data of interest in the screen-based and VR-based situations only. The data presented represents the *handlebar* steered angle, not the *effective* steered angle (although in the current implementation, this was simply factored).

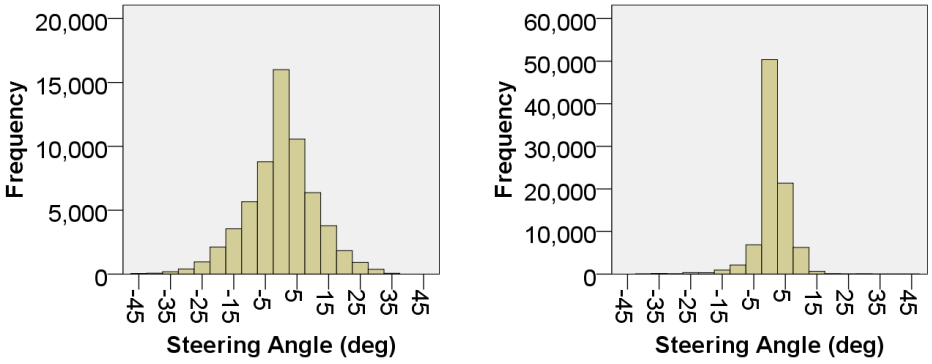

(a) Participant 1 (b) Participant 2

Figure 8.9: Steering histograms in VR situation

Figures 8.9a and 8.9b juxtapose histograms of chosen steering angle in the VR situation for a participant that reported struggling substantially with the steering (Figure 8.9a) and one who became accustomed to the steering more quickly (Figure 8.9b). There is a clear visual difference between the two participants with the participant that struggled producing a histogram with a significantly wider distribution (s.d. = 10.79 vs. s.d. = 5.32).

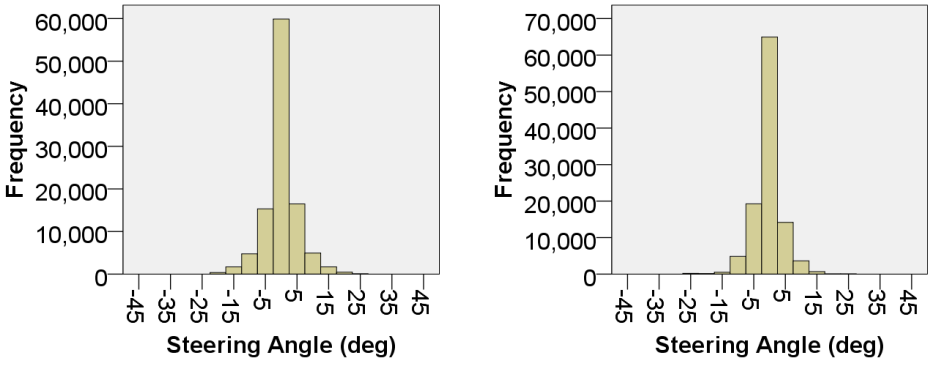

(a) Participant 1 (b) Participant 2

Figure 8.10: Steering histograms in Screen-based situation

Figures 8.10a and 8.10b present histograms of those same two participants in the screen-based situation. None of the participants reported (or were witnessed) struggling in the screen-based situation and this is observable here where the distribution looks visually similar for both participants, and similar to the distribution of those participants who did not struggle in the VR situation (e.g. Figure 8.9b).

Quantile-Quantile plots comparing these two participants in each situation are shown in Figure 8.11. These visually confirm the observations above. The QQ-plot of the participants in the VR situation (Figure 8.11a) shows data that deviates substantially from the 'equivalence' line. By comparison, the QQ-plot for the screen-based situation (Figure 8.11b) shows much closer agreement between the participants' behaviour at all but the most extreme steering values. However, given the large number of data points in both cases ($n_1 = 72955$ versus $n_2 = 95859$, and $n_1 = 106848$ versus $n_2 = 110010$ in the VR and screen-based situations respectively), an Independent Samples Kolmogorov-Smirnov test indicates significant evidence to reject the null hypothesis that the two participants' data are identically distributed ($p < 0.001$), albeit that this may be unimportant in the screen-based situation.

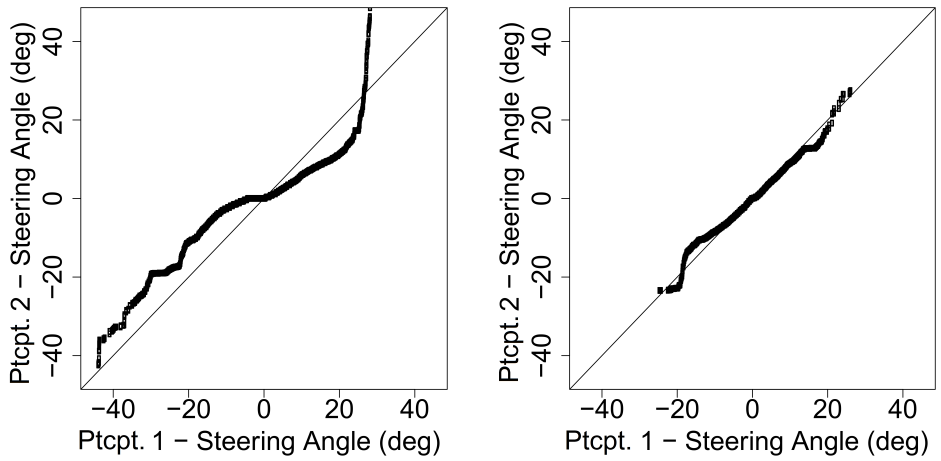

(a) VR-based situation (b) Screen-based situation

Figure 8.11: QQ-plots for Participant 1's steering selection

Table 8.3: Summary Steering Statistics

Situation	n	s (ms^{-1})	Mean IQR	Mean Range (deg)
VR	4	0.85	6.03	78.71
Screen-based	5	0.69	3.10	43.89

Table 8.3 displays some summary statistics for the steering data recorded for the participants. The median of the standard deviations of steering angle for the group in the VR situation is significantly different to that in the screen-based situation, at the 5% significance level (Mann-Whitney Rank-Sum

$U = 2.45$, $p = 0.014$). The medians of the ranges of steering angle is also significantly different (Mann-Whitney Rank-Sum $U = 2.45$, $p = 0.014$), however the medians of the IQRs of the participants (Mann-Whitney Rank-Sum $U = 1.71$, $p = 0.086$) are not significantly different. Indeed, the IQR of the participant shown in Figure 8.9a in the VR situation could reasonably be considered an outlier: This participant's IQR is an order of magnitude greater than the remaining participants in both the VR and screen-based situations.

This statistical breakdown highlights the main difference also observed visually in the participants. Generally, the control of the steering was more extreme in the VR with reports of over-steering, over-compensation, and sharp, sudden turns. Participants could not see their arms (owing to the full coverage of the HMD) and thus could only know their actual arm (and thus handlebar) positions though kinaesthetic sensation and its (lagged, in the worst case, 100ms) results in the simulated movement. Given most of the movement was on straight sections, the IQRs are statistically similar across the two situations but when turn-steering is applied, more extreme behaviour is observed. Such extreme turns were not witnessed in the screen-based situation, and this is supported by the smaller range of steering values observed across the participants.

Table 8.4: Summary Steering Reversal Statistics

Situation	n	$\bar{x}$ (Hz)	s (Hz)	Median (Hz)
VR	4	2.69	0.44	2.70
Screen-based	5	2.98	0.46	3.17

Steering reversal frequencies (calculated as the number of changes in the direction of steering movement recorded per participant per second) are shown in Table 8.4. The median of the frequencies in the screen-based situation is significantly greater than that in the VR-based situation (Mann-Whitney Rank-Sum $U = 1.96$, one-tailed $p = 0.025$), which even for this small sample, appears to confirm the outcome observed by Helland et al. (2016) who found that steering reversal is lower when SS scores are reported to be higher (Section 6.2.3). In this case, SS scores were higher in the VR situation, and steering reversal frequency was lower.

8.3.2.1 Path Selection

Figures 8.12 and 8.13 show spatial plots of the positions of the recorded paths in each of the two simulated situations: VR-based on simulator, and screen-based on simulator respectively. Compare with Figure 5.4 which plots the data collected for the real-life circumstance (processed as per Chapter 5). All three plots present data aggregated across all the participants.

Simple visual comparison of the three situations is sufficient to notice clear differences in the data. The real life situation's data (Figure 5.4) highlights

some of the systematic error inherent in the homography transformation process (see Section 5.4). Small errors in this transformation, coupled with shallow viewing angles for some of the cameras' coverage, result in plotted paths which clearly illustrate systematic error, and which (in some cases) transform to infeasible locations. For example, some would be located inside hedgerows.

By comparison, the screen-based (Figure 8.13) and VR-based (Figure 8.12) data show no obvious systematic error, based as they are, in the coordinate data taken directly from the simulation. The VR-based data clearly demonstrates a much greater visual variability in the paths chosen, given that a number of the participants struggled with controlling their progress through the simulation. The screen-based data show little in the way of erroneous paths taken through theoretically solid objects. In the VR situation, some participants passed through foliage or the container, however this was accurate, if not perhaps deliberate. There is one clear example (west) where a participant 'got lost' and had to circle around to come back to the circuit and another (south) where a participant simply chose to turn around.

8.3.2.2 Paths through a corner

Section 5.5.2 presented a limited analysis of a cross-section of the paths through a corner in the real life data set. That corner was selected as the camera view was from a high angle and thus the quality of the analysis was sufficiently robust to provide reasonably valid outputs. This section extends that analysis and performs a similar analysis on the screen-based and VR-based data sets.

The corner pictured from real life in Figure 8.14a and from the simulation in Figure 8.14b is of a circular profile and has a 3.05m radius. It is open on one side, has dense foliage to a height of approximately 1.0m on the inside, and is led into by straight approaches of at least 15m.

Visual comparison between the real life, VR-based and screen-based paths through the corner (Figures 8.15, 8.16 and 8.17, respectively), as with the whole of circuit comparison, demonstrates some clear differences. Albeit with substantially more visible variability in the VR, the real-life and VR situations (Figures 8.15 and 8.16) show a similar spatial arrangement with paths clustered towards the inside of the kerb on the apex. The screen-based situation on the other hand (Figure 8.17) shows a spatial arrangement of paths which do not cluster so closely to the apex.

Section 5.5.1 (specifically Figure 5.5) also presented a histogram showing a binned section of path distance from the inside kerb through the corner in the real-life data set, although this included data from the entire of the study including group riding. This chapter focuses on that data directly comparable to this simulator study, specifically the single rider data. Similar to the methodology in Section 5.5.1, a computer program was written that iterated through

time-adjacent data points in each set. If the given line segment intersected one of the defined cordons then the crossing was registered. Participant, time, direction and position on the cordon were recorded.

The profile in Figure 5.5 appears to show a distribution where cyclists attempt to stay close to the inside of the curve. A more complex series of sections was considered here spaced at approximately 5m intervals on the approaches or 20 degrees of turn around the corner itself. The following presents two representative cordons, namely Cordon D and Cordon G; both indicated on Figures 8.15, 8.16 and 8.17.

Considering the component sub-figures of Figure 8.18, no immediately obvious distributional differences are apparent between the various situations or directions, with the exception of the screen-based anti-clockwise situation (Figure 8.18f). Whereas in all the other cases, the majority of the data points are clustered within a few metres of the inner kerb, this case has all the data spread out much more substantially with the modal bin in the 5.0–6.0m distance and no data points within 2.0m of the inner kerb. This specific case is essentially a measure of the entry path to the corner in the screen-based situation.

Figure 8.19 shows a breakdown of the situations by direction for a cordon on the corner itself. Generally, the real-life (Figures 8.19a and 8.19b) and VR situations' (Figures 8.19c and 8.19d) data points are tightly clustered to the kerbside, as was identifiable in the spatial plots (Figures 8.15 and 8.16). However, the screen-based situation in the clockwise direction (Figure 8.19e) is clustered further from the kerb with zero data points falling in the 0.0–1.0m bin; and in the anti-clockwise direction, the modal bin is the 2.0–3.0m distance.

Combining the observations between Cordons D and G, it appears that entry placement is substantially different in the screen-based situation (in the anti-clockwise direction) and that this results in a wider path taken through the corner. It is not clear why this would be the case but perhaps may relate to the limited field-of-view with the screen, and the tight angle of the corner. Larger and more expansive screen(s) may alleviate this. In fact promisingly, the spatial placement in VR seems to be quite similar to that in the real-life situation, which adds further to the case for using highly-immersive techniques such as VR in this context.

8.3.2.3 Paths on a straight

As noted above, there are substantial systematic errors in the visual homography transformation, in particular when the viewing angle is shallow. In addition to the curve area images discussed above, a straight section adjacent to the NOC building was covered by a camera with a relatively steep angle of view and looking almost perpendicular to the cross-section of the marked pathway (Figure 8.20). Consequently, one can be reasonably confident that a section

through the lateral placement in this area, is likely to be relatively unaffected by systematic error in the axis we are concerned with, given the majority of systematic error on this view is along the depth-of-view axis.

Visual comparison between the real life, VR-based and screen-based paths on this straight section (Figures 8.21, 8.22 and 8.23, respectively) appears to show a substantially wider spread of paths in the VR situation compared to the Screen-based situation and further so compared to the real-life situation. For purpose of comparison, a cordon line (Q) is indicated on each of the three plots and was processed as per the above (Section 8.3.2.2; and per Section 5.5.1) by situation and direction.

Aggregated across all situations, the eastbound and westbound distribution of cordon crossing do not present any obvious fundamental visual difference and therefore, as distinct from the paths around a curve discussed above (Section 8.3.2.2), lateral placement at this cordon does not appear to obviously depend on direction of travel. The result of a Independent Samples Kolmogorov-Smirnov two-sided test does not show a significant difference in distribution of the cordon crossing point with regard to direction ($n_1 = 69$, $n_2 = 62$, $p = 0.537$). Participants tended to travel around the centre of the path (which was marked to 4.0m in width) and (almost) always remained within the marked way. The data point substantially outside the marked way occurred in the VR situation with one of the participants who was struggling with the steering control. Consequently, the following does not separate data in each situation by direction of travel.

Figure 8.25 shows the cumulative distributions of the cordon crossings broken down by situation. Again no obvious differences are apparent from a visual inspection and an inter-comparison of the distributions by Independent Sample Kolmogorov-Smirnov two-sided tests backs this up:

Real-life vs VR Not significant, $n_1 = 35$, $n_2 = 30$, $p = 0.761$;

Real-life vs Screen-based Not significant, $n_1 = 35$, $n_2 = 66$, $p = 0.516$; and

VR vs Screen-based Not significant, $n_1 = 30$, $n_2 = 66$, $p = 0.797$;

This result is however a desirable outcome with respect to participant behaviour, as it can be framed as further supporting evidence that – at least spatially – participant's behaviour did not change between the three different situations.

8.3.3 Look Angles

The use of a VR HMD in the VR-based situation provides the opportunity to capture data relating to head orientation. Whilst this is not technically equivalent to look direction as the participant can move their eyes to view within the HMD's field-of-view, it provides a reasonable approximation for all but the highest frequency eye movements.

View angles are reported by the HMD about the three axes of the HMD in the left-hand coordinate system used by the Unity engine. The result is three values: an X angle orientated about the lateral (pitch) axis where ahead is 0 degrees and angles increase in value as the participant looks downwards; a Y angle orientated about the vertical (yaw) where 0 degrees is orientated the same as the 'world coordinate' 0 degree direction (in this case, approximately south and an unfortunate outcome of the particular arrangement of CAD data utilised) and increases in angle to the participant's right; and a Z angle orientated about the longitudinal (roll) axis of motion where 0 degrees is upwards and reported angles increase as the HMD rolls to the left. X and Z angles can be used directly, Y-axis angles (as they are in world coordinates) must be corrected for the angle in which the bicycle was moving (i.e. transformed to participant coordinates); the residual 'delta Y' represents the angle in the vertical differenced from the current direction of travel where 0 is the direction of travel and angles increase to the right of the direction of travel.

Table 8.5: VR Look Angle Summary Statistics

Descriptive	X-Axis	Y-Axis	Z-Axis
n	286326	286140	286326
$\bar{x}$ (deg)	16.73	-0.14	-1.12
s (deg)	5.30	10.52	2.30
Median (deg)	15.48	0.19	-1.10
IQR (deg)	8.54	14.13	2.73
Range (deg)	68.08	106.46	30.28

X-axis rotation (Table 8.5) across all the participants tends to be slightly downward, usually between 10 to 20 degrees or more below horizontal. This is to be expected as the participant is seated on a bicycle so would naturally have some downward angle of view. Assuming an average height off the ground of 1.75m and a 15.5 degree downward viewing angle (the median for this data set), this computes to head-angle directed towards an area on the ground some 6.3m ahead; though of course the participant will probably be looking further ahead with the eyes, this is important as it is the neutral direction for look angle. With average speeds of 4.4ms^{-1} (Table 8.2), this is a distance approximately 1.4s travel time ahead. Assuming a negligible reaction time, this would equate to a maximum linear deceleration rate of -1.54ms^{-2}; a close match to the maximum deceleration rate (-1.5ms^{-2}) taken from CROW (2007) and used in Section 3.5 (Table 3.2). The implication therefore being that the rider generally directs their attention to a point in the road to which they could safely stop, if they were suddenly faced with the need to.

Z-axis rotation values (Table 8.5) are tightly distributed around the vertical, as would be expected given the fixed nature of the bicycle and the fact the participants remained upright throughout.

Y-axis rotation is more interesting. As would be expected, participants orientate their head toward the direction they are going (Figure 8.26); interquartile range however is only 14 degrees. Taking the same 'straight section' example considered in Section 8.3.2.3, Figure 8.27a shows the distribution of relative look direction for westbound riders (i.e. towards Camera 4), and Figure 8.27b the same data for the eastbound direction of travel. In both cases, the look angles are closely distributed about 0 degrees (i.e. ahead; $\bar{x} = 1.23$, $n = 11265$; and $\bar{x} = -2.36$, $n = 7801$; for west and eastbound respectively).

To frame the above in the context of a contrasting case, consider the corner to the northeast (cf. Figure 8.3). In particular, how far ahead around a curve might the participant be looking. If the paths which are travelling anticlockwise are isolated, the data presented in Figure 8.28 is retained. Here, the data remains clustered however this time it is clearly substantially to one side of a view ahead. Mean head angle is 11.0 degrees to the left. Whilst this is not specifically the angle the participant is actually looking – recall that the HMD tracks head orientation, not actual eye gaze direction – it confirms that the participants were likely looking further around the corner than was their current travel direction, as has been observed in reality (e.g. Vansteenkiste, Van Hamme, et al., 2014).

8.3.4 Quantitative Conclusions

Speed selection (as quantified by participant median speed) was not found to be significantly different across the different situations (real-life, virtual reality and screen-based). Though differences may be identified by larger sample sizes, the design of the turbo-trainer and its impact on the participant seems to have resulted in speeds which are objectively similar to those from reality. It is possible that the incidence of SS had some impact on recorded speeds in the VR, but this would require further study to confirm.

The ability (or not) of participants to control the steering to their satisfaction was noted as a qualitative issue (Section 8.2.2) and this is also identifiable in the data. Participants that struggled with the steering, display a wider spread of chosen steering angles; and in the VR, where participants had a greater tendency to over-steer, the range of measured motion was significantly greater. This may be ameliorated (to an extent) by the potential for the reduction in steering sensitivity, however literature indicates that exacerbation of movements (intended or otherwise) is a function of not being able to see one's own body and thus the use of VR itself. The less excessive measurements of movements in the screen-based situation would support this.

Path choice in the VR was found to be more erratic than in the screen-based situation, as might be inferred from the more extreme steering responses. However in the aggregate, spatial placement in the VR agreed well with the observed paths in reality, implying a close agreement in this important out-

turn data type. Spatial placement in the screen-based situation was similar but noticeably different on the sharp corner near Camera 1 (Figure 5.2) where paths were distinctly further from the apex of the curve. This may perhaps be a function of the lack of ability to look around the corner whilst travelling (as one can in reality or in VR), as head angle information from the HMD shows that participants tend to turn their head up to around 20 degrees (further to their current turn extent) as they navigate a turn; something which was not possible with the screen-based arrangement. More extensive screen equipment may provide the ability to present visual information further around the curve and thus eliminate this effect, however this was not able to be tested here.

On straight sections of the circuit, the lateral distribution of path choices was not substantially different, regardless of situation. As a general conclusion then, and save where identified above, the choice of path appears to be relatively consistent regardless of the situation chosen. This, and the other quantitative results presented in this section, indicate a relatively good numerical agreement between the simulator and reality, and lend further credence to a conclusion of a successfully designed and implemented simulator.

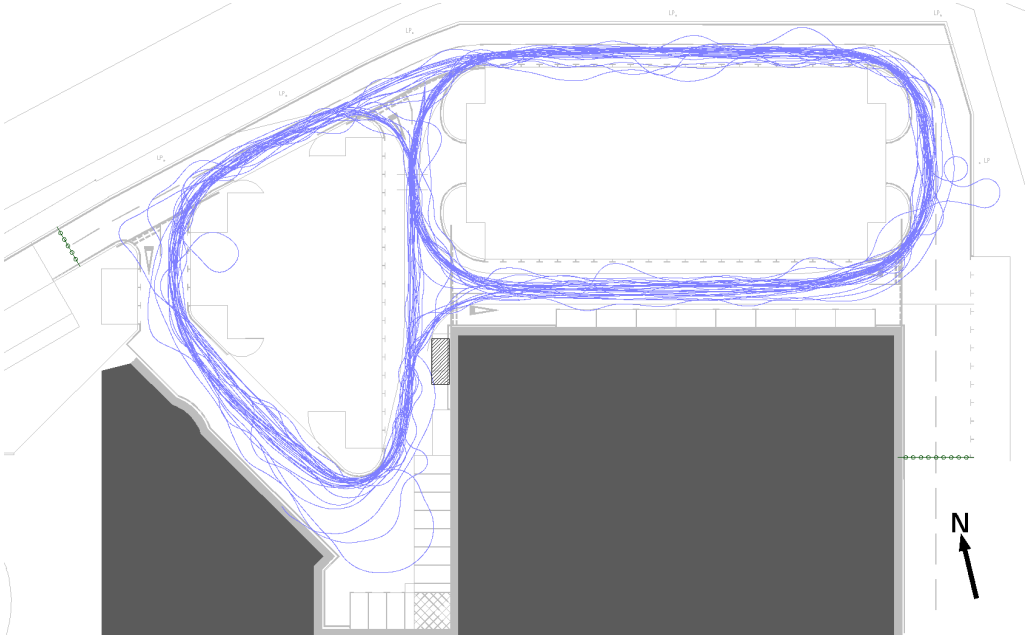

Figure 8.12: Plot of all participants' paths taken in VR situation

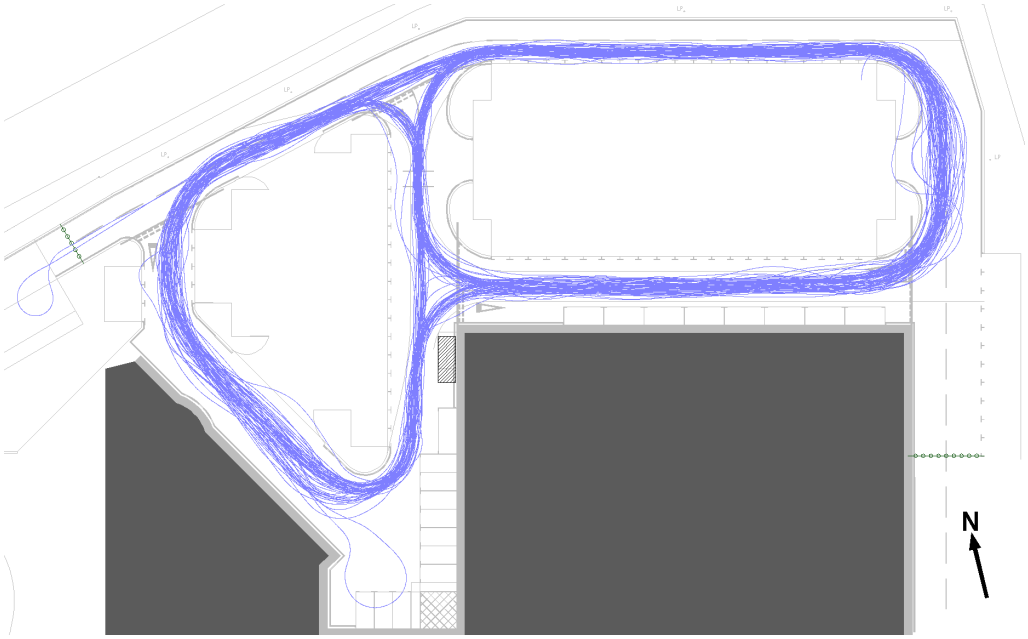

Figure 8.13: Plot of all participants' paths taken in Screen-based situation

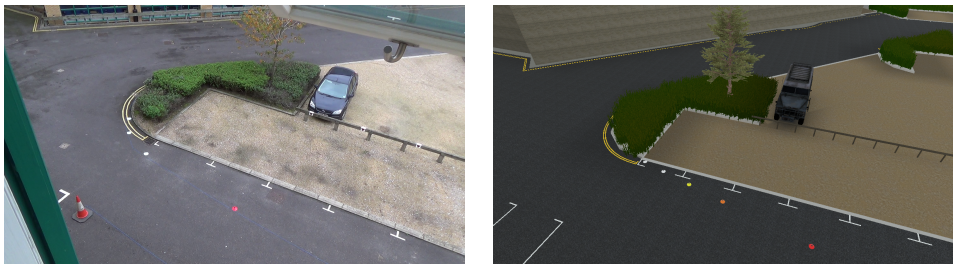

(a) Corner in reality (b) Corner in simulation

Figure 8.14: Images captured of corner

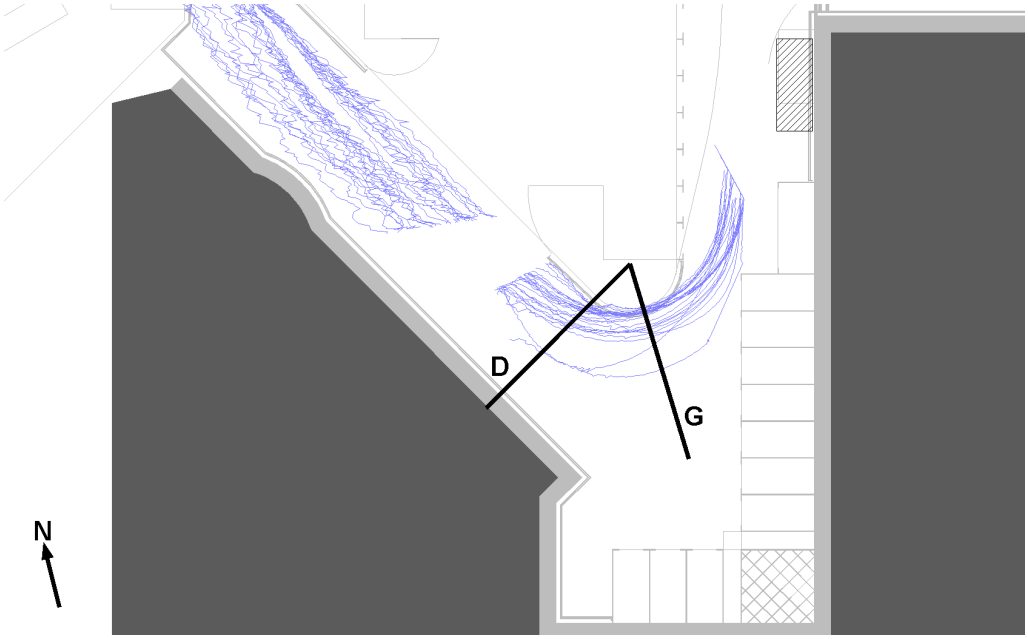

Figure 8.15: Plot of all participants' paths taken at Camera 1 corner in Real Life situation (Cordons D and G shown)

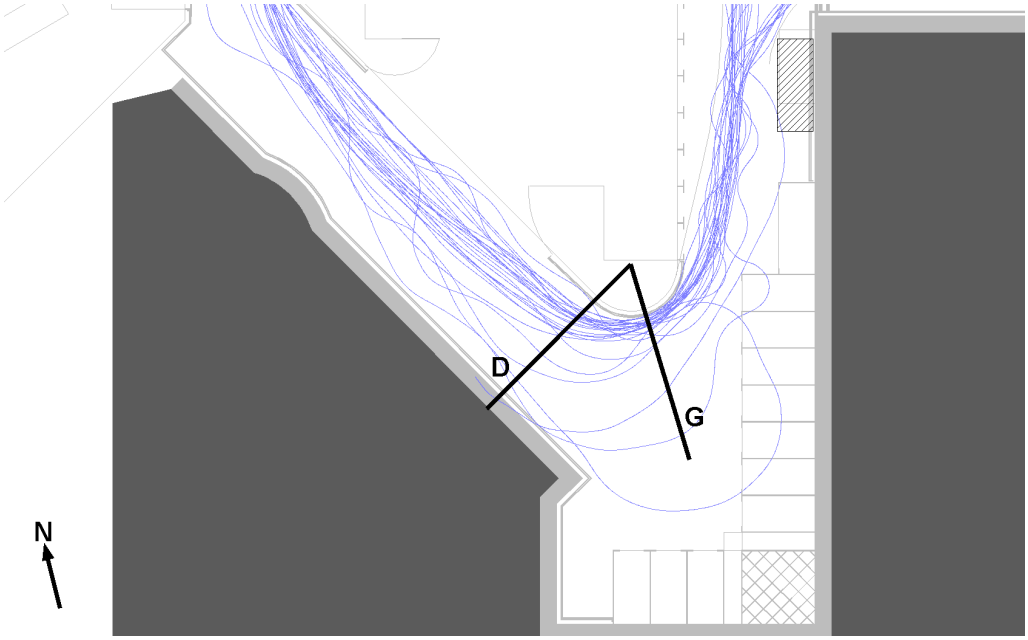

Figure 8.16: Plot of all participants' paths taken at Camera 1 corner in VR situation (Cordons D and G shown)

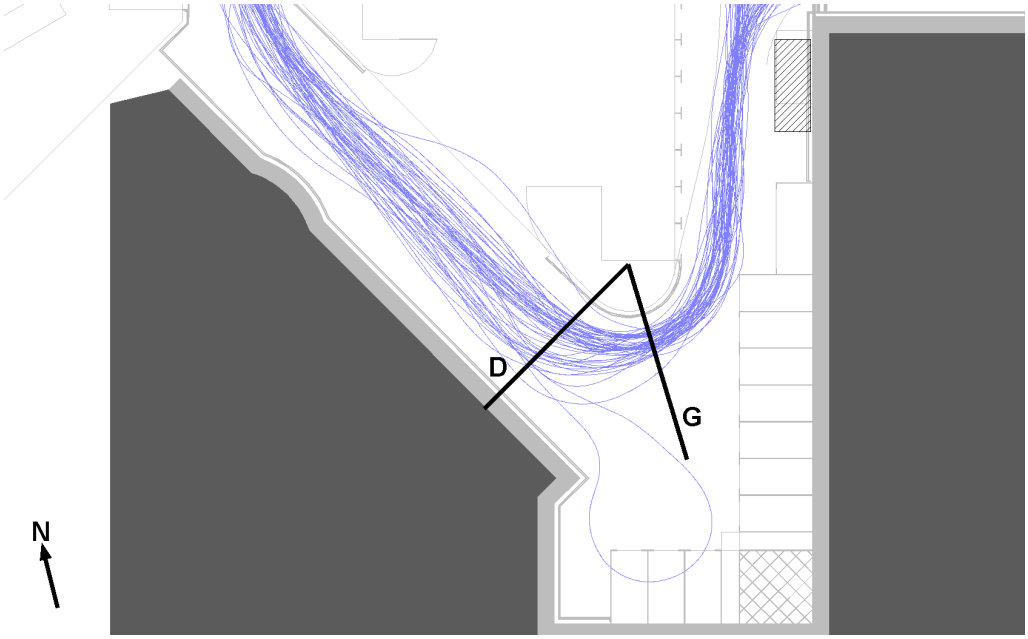

Figure 8.17: Plot of all participants' paths taken at Camera 1 corner in Screen-based situation (Cordons D and G shown)

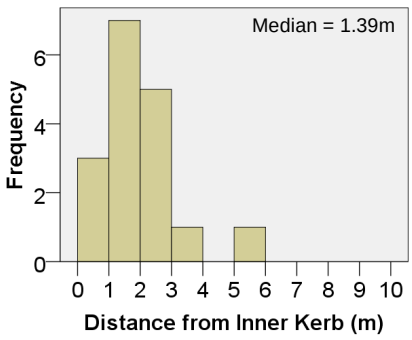

(a) Real-life (clockwise)

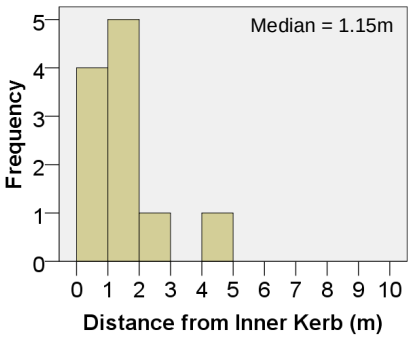

(b) Real-life (anti-clockwise)

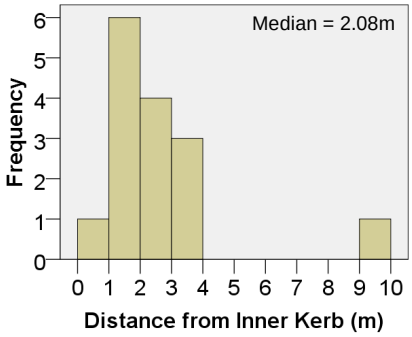

(c) VR situation (clockwise)

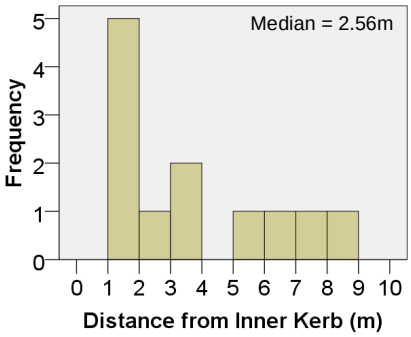

(d) VR situation (anti-clockwise)

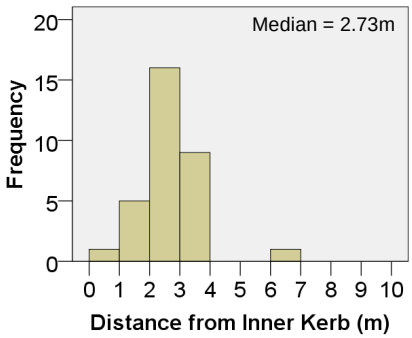

(e) Screen-based situation (clockwise)

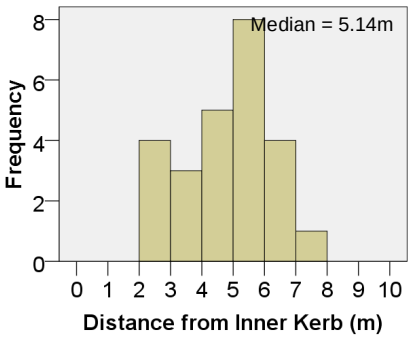

(f) Screen-based situ. (anti-clckws.)

Figure 8.18: Histogram of all participants' crossing positions of Cordon D

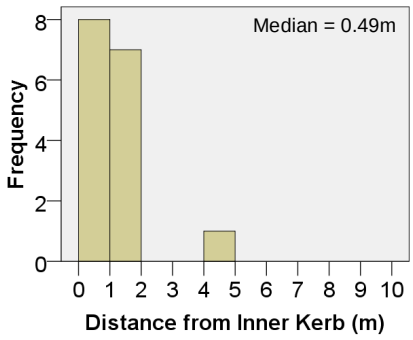

(a) Real-life (clockwise)

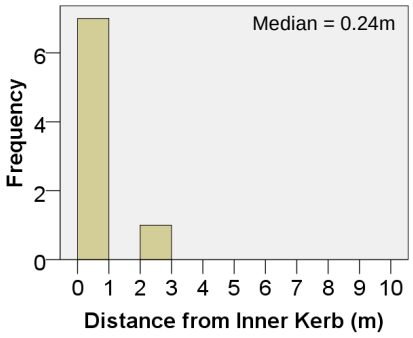

(b) Real-life (anti-clockwise)

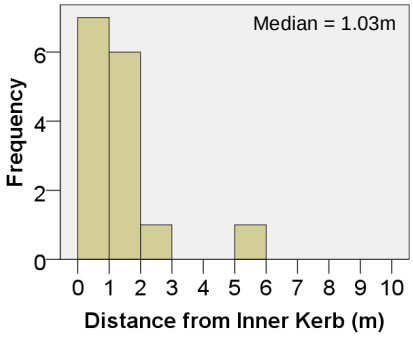

(c) VR situation (clockwise)

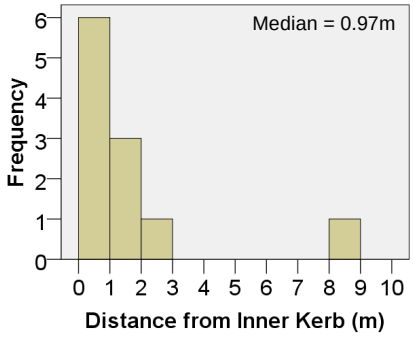

(d) VR situation (anti-clockwise)

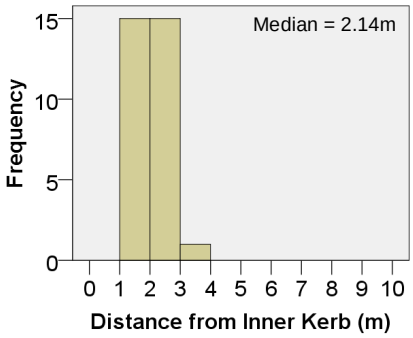

(e) Screen-based situation (clockwise)

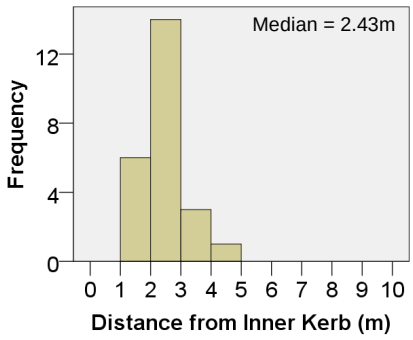

(f) Screen-based situ. (anti-clckws.)

Figure 8.19: Histogram of all participants' crossing positions of Cordon G

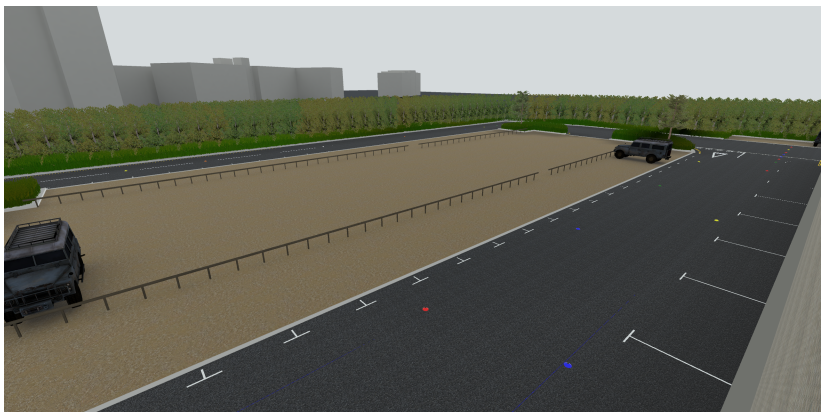

Figure 8.20: Screenshot of Camera 4 field-of-view (straight section to lower right)

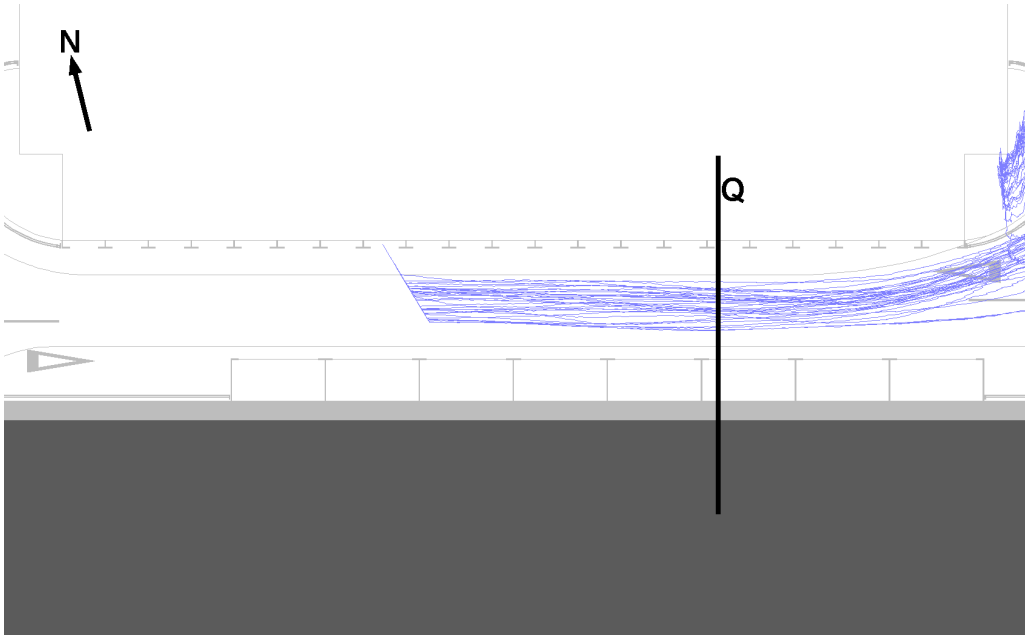

Figure 8.21: Plot of all participants' paths taken at Camera 4 straight in Real Life situation (Cordon Q indicated)

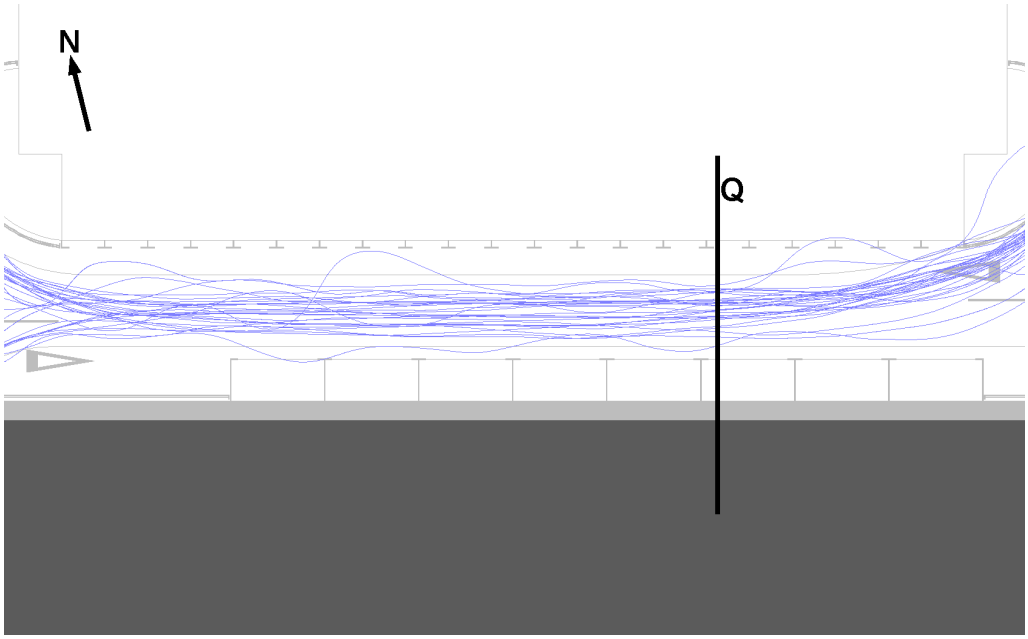

Figure 8.22: Plot of all participants' paths taken at Camera 4 straight in VR situation (Cordon Q indicated)

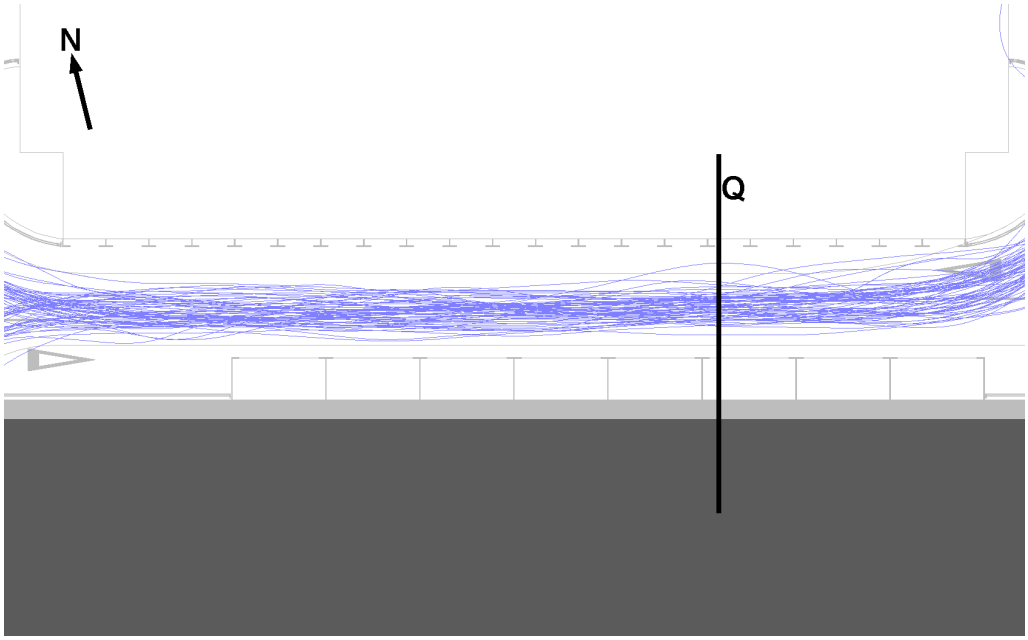

Figure 8.23: Plot of all participants' paths taken at Camera 4 straight in Screen-based situation (Cordon Q indicated)

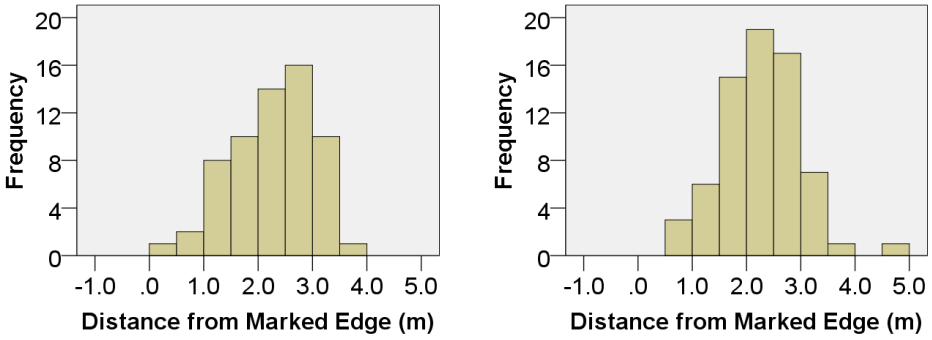

(a) Eastbound Crossing (b) Westbound Crossing

Figure 8.24: Histograms of Crossing Points for Cordon Q (Distance from southern edge of marked path)

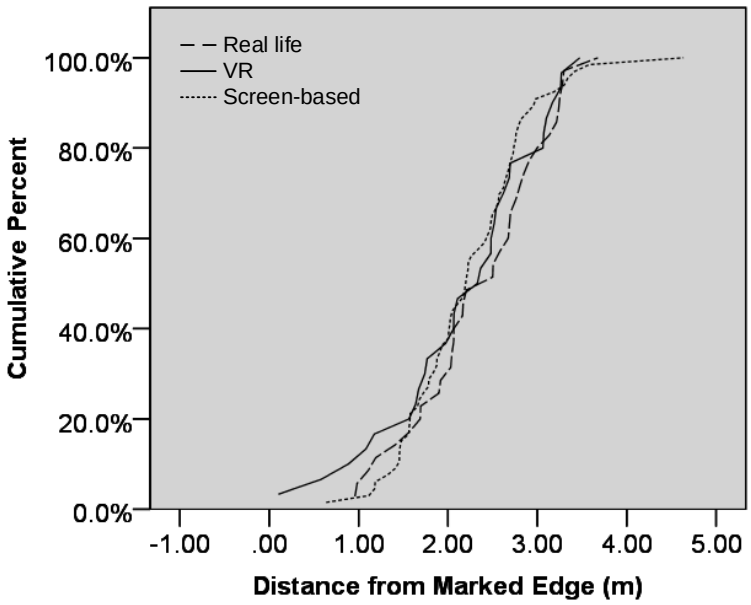

Figure 8.25: Ogive plot of of all participants' crossing position of Cordon Q

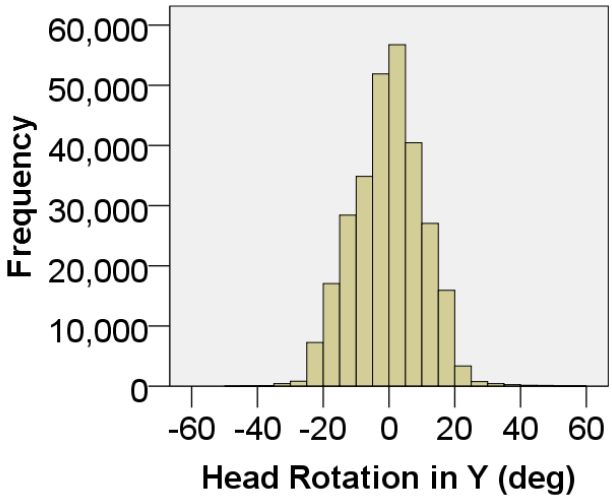

Figure 8.26: Histogram of all participants' 'delta-Y' head rotation

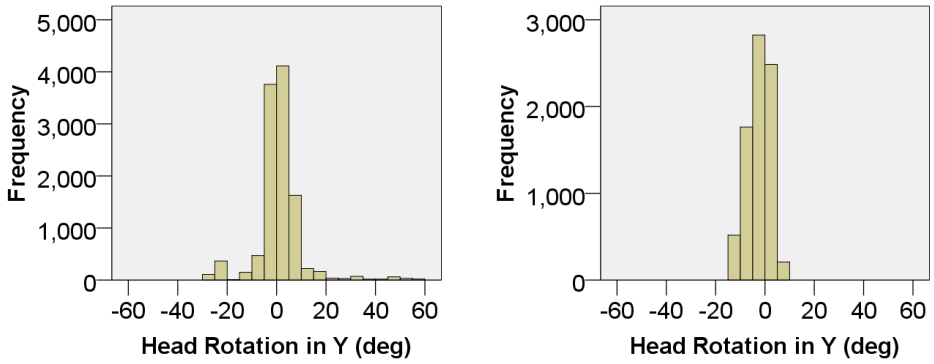

(a) Heading west (b) Heading east

Figure 8.27: Histogram of all participants' look angle on the straight section adjacent to Camera 4

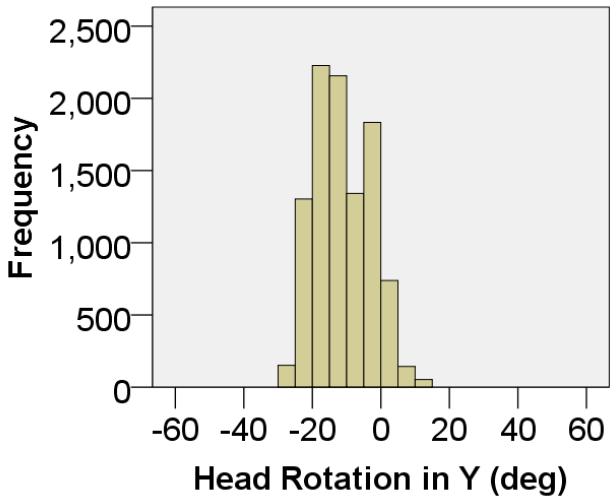

Figure 8.28: Histogram of all anticlockwise participants' 'delta-Y' head rotation on northeast corner

8.4 Conclusions

This chapter (and the previous Chapter 6) has presented the literature basis, design, construction and testing of a bicycle simulator. This simulator (subject to some improvements identified) has been broadly successful and can now be taken forward and used to begin to develop a canon of understanding of cyclist behaviour in controlled and replicable circumstances.

That in itself is no mean feat, however this document has not presented a simulator as a 'fait accompli' but instead has led the reader through the entire development process. This is important as if – as is the case in much of the literature – the efficacy of the simulator is simply presented as a given, the reader's confidence in the validity of the results is rightly reduced.

More generally, and whilst used here to develop a bicycle simulator, (relevant parts of) the process detailed could well be applied to other simulators. Furthermore, the reader has been alerted to those aspects of the simulator which were difficult to achieve and those aspects which require further refinement. Again, this is of value as it provides a much more robust starting point for others (for example, a specification and detailed replicable procedures), and makes clear that the design and construction process itself is not a simple turnkey exercise, as is the usual style in the literature with regards to simulators of any type.

The last part of this chapter recalls the research questions identified early in this chapter (Section 8.1.3), answers them as is now possible with the data in hand and draws together the key limitations of the set up in its current form.

8.4.1 Research Questions

Three main research questions were identified in Section 8.1.3. First, did the participant feel like they were actually riding a bicycle (RQ 1)? Secondly, did participants feel immersed in the simulation and was this perception increased by the use of VR (RQ 2)? And finally, did the participants behave as they did in reality (RQ 3)?

8.4.1.1 RQ 1: Does the simulator feel, to the participant, like they are actually riding a bicycle?

Feedback from the participants consistently agreed that the use of the simulator felt similar to, but not exactly the same as, riding a real bicycle. The virtue of the use of an actual bicycle almost certainly in large part reinforced this, something which would perhaps have been reinforced further if participants had used their own bicycles. The response to pedalling and the turbo-trainer was universally reported to feel realistic, albeit that the participants did not explore freewheeling. The responsiveness of the steering on the other hand was

reported as excessive and detracted from participants' feelings of realistic bicycle behaviour.

Conclusion Generally yes; subject to steering improvements required.

8.4.1.2 RQ 2: Do participants feel immersed in the simulation? Does the use of VR increase the feeling of immersion?

Immersion was universally reported as high in the VR situation, and was variable (but subjectively lower) in the screen-based situation. Attention to detail with regards to the layout of the virtual environment was well received and where immersion was broken, this generally related to the issues identified with the steering. Immersion in VR was sufficiently high that behaviour that would have caused injury in reality (e.g. collision with a solid object) elicited visceral reactions in the participants.

Conclusion Yes in VR; variably on screen. VR distinctly increases reported immersion.

8.4.1.3 RQ 3: Does the participant behave as they would (or did) in reality?

Quantitative observations indicate that participants tend to behave in a manner not distinctly different from reality on the simulator. Speeds were not found to be significantly different to reality on the simulator and spatial placements were generally found to be similar to those in reality. Where spatial behaviour differed in the simulator, the likely cause is the lack of expanse of the field-of-view of the screen-based set up; this can be addressed in future improvements to the simulator.

Conclusion Generally yes in VR, subject to greater variability in behaviours related to the steering control issues identified. Generally yes with screen visual interface subject to some differences at extreme angle turns which would be ameliorated with more expansive screen(s).

8.4.1.4 RQ 4: Do participants experience Simulator Sickness?

In addition to the key research questions, a series of questions developed from the literature with regard to simulator sickness. As it transpired, simulator sickness was a major issue in the use of VR, despite the application of best practices (where possible). Even those participants that did not consider themselves to have SS, reported high SS scores when the SS questionnaire was administered.

Conclusion Yes for all participants in VR, to varying degrees, with a drop-out rate of 60%. No substantive SS was reported in the screen-based situation.

8.4.1.5 RQ 5: Does the reported level of Simulator Sickness increase with speed of movement in the simulation?

Using the SS scores to test some of the observations in the literature, plotting participants' reported total SS scores against their speed produce a compelling plot indicating an increase of reported score with speed. However given the small number of data points, no statistically significant relationship can be concluded (despite a r-score of 0.871).

Conclusion Based on the small sample ($n = 4$), there is compelling visual evidence, however no statistically significant evidence.

8.4.1.6 RQ 6: Does the recorded steering reversal frequency decrease with increased reported level of Simulator Sickness?

The median steering reversal frequency was found to be significantly lower in the VR situation, and reported SS scores were substantially higher in that situation. However, without having controlled for the effect of the different visual interface between the screen-based and VR situations, this confounding effect cannot be excluded.

Conclusion This is indicated based on this data but should be accepted with some trepidation based on the potential effect of confounding factor(s).

8.4.2 Limitations

A number of individual limitations that exist with the current state of the simulator can be identified at this point. The issues surrounding simulator sickness are of a larger scale and are addressed later in Section 8.4.3.

It is important to highlight these issues here, not simply to demonstrate awareness of them but to draw them together to assist future practitioners in this field. All too often, literature presents the development of new and groundbreaking research as some sort of pseudo-magical achievement in which all requirements were met, and all issues are either opportunities or were supposedly intended to be the case from the start. As an example, the strand of work discussed in Section 4.2.3 that connects across multiple studies (Grechkin et al., 2012; Plumert et al., 2004, 2011) used a bicycle simulator and associated system which is presented as a complete and reliable system with no issues whatsoever surrounding it in terms of usability, participant comfort or behavioural fidelity. Whilst it is theoretically possible that this is the case, the reality is that this is unlikely and yet such a narrative serves only to create the impression that the outcomes of such work should be taken as given. Indeed, achieving publication may depend on such a narrative.

This author does not take that approach here. Delivering the simulator (and indeed the other parts of this document) has been, at times, a difficult

process that has required new, innovative work on the author's part to bring together existing and new systems to achieve the required objectives. Sometimes this has been successful, other times, less so. Detailing the development process of *any* simulator to the extent presented in this document is, in this author's opinion, itself an original contribution to the literature and certainly will serve to advance the knowledge of others who wish to create a new simulator. The key aspects which have been less successful are summarised here.

8.4.2.1 Steering Responsiveness

Across all the participants, the usual detail which appeared to be a barrier to both objective fidelity and immersion was the steering over-responsiveness. All participants reported that they felt the steering was too sensitive. This report was consistent between the VR and screen-based situations which implies the impact of not seeing one's own hands in VR can be discounted as the only cause. However note that the participants did not struggle with steering (as much) in the screen-based situation, even though they still considered it to be the case. Therefore it is possible the inability to see one's own hands exacerbated the situation.

The resultant steering was simply applied as 0.5 times the measured angle of the front wheel as a simplification to allow expedient testing. This was not changed even after it became an apparent issue in participant use, so as to ensure an experimental similarity between all the participants. Clearly, this factor needs to be revised down further. As discussed in Section 7.2.1.3, in reality the steering of a bicycle is actually highly complex with many degrees of freedom, which are then further complicated by the fixation of the bicycle in a vertical stand position. Future work should build on that established here.

8.4.2.2 Roll-axis Rotation

The bicycle apparatus was supported in a fixed stand turbo-trainer. This precludes the user of the bicycle from leaning to the side when turning, but also from the natural oscillation that occurs when the rider pedals (a frequency of around 1.2Hz; Dozza & Fernandez, 2014). Whilst this may have effects that manifest in simulator sickness (discussed later; Section 8.4.3), it also had some reported impact on breaking the immersion of the participant in the simulation.

Alternative trainer devices, such as 'rollers' (the bicycle equivalent of a 'rolling-road' chassis dynamometer; e.g. Tacx B.V., 2016d) may be beneficial in providing a more realistic facsimile for road-riding, but even a small attempt to turn the handlebars would likely result in a fall, not to mention the user learning curve. It is for those reasons that these were not used here.

Unfortunately, this leaves a conundrum as to how one can more accurately represent the roll rotation of the rider when both pedalling and making turning

movements. It is possible that there is no realistic way to address this issue and that this may be an inherent limitation of any bicycle simulator. Motion platforms, of the style used in the more elaborate full-motion simulators may provide an answer; however, with the standing lack of clear understanding of bicycle dynamics in the literature, it would likely be difficult to construct an arrangement which is both realistic and safe to the user, bearing in mind that the user does not have the benefit of being secured inside a vehicle cockpit as is the case for other types of vehicles' simulators. In any case this is, as a direct result of the work presented here, an identified direction for future work with regard to bicycle simulation and would likely also serve to improve the canon of literature with regards the general dynamics of bicycles.

8.4.2.3 Visual Field Coverage

By design, the use of a VR HMD provided for a large field-of-view to the participants (110 degrees); by contrast the screen set up used filled a substantially smaller field-of-view (some 30 degrees). An intermediate arrangement of using a larger screen (or screens) was not explored here.

Reported levels of simulator sickness were high with the VR (as is not an uncommon experience with VR) and were low with the screen. Conversely, immersion was usually high with the VR and variable with the screen arrangement. There therefore potentially exists a 'trade-off' zone between the two extremes whereby immersion can be maximised with increasingly larger screens that surround the participant more comprehensively, but that can be limited to the point at which simulator sickness becomes a more substantial issue (or indeed, the reverse could be true).

Given that it was found that there were no major behavioural differences due to the visualisation method and that where differences were identified, more extensive field-of-view would likely be a specific remedy, the behaviour of the participant should be relatively invariant to the extent of the visualisation in this 'trade-off' zone, but such modifications would allow a degree of 'tuning' of the immersion and simulator sickness response. In addition to being a desirable direction with which to extend this work, such a process would also allow specific testing of both simulator sickness responses and immersion perception, something which is an identified standing gap in the literature (e.g. Classen et al., 2011, also see Section 8.4.3 below).

8.4.3 Simulator Sickness

Simulator sickness has long been a feature in the literature and a wide body of work exists which notes and indeed quantifies the existence of SS. And yet, the cause of SS itself is still a matter of debate with up to four main actively competing causal hypotheses (Section 6.2). Furthermore, consideration of SS itself is almost always a by-product of the literature. With particular consideration

of driving simulators, a comprehensive literature review into SS by Classen et al. (2011) notes with surprise the lack of randomised controlled trials into SS, given that driving simulators are a best-case situation for experimentally testing SS in controlled circumstances.

Simulator sickness in the participants of this study was a substantial issue. Aside from the obvious source of vestibular mismatch – in so far as the participant sees motion they are not making in reality – and therefore the obvious potential for simulator sickness causation in line with the balance of evidence in the literature (Section 6.2.1), the simulator is also an active source of 'postural instability' in line with the motion sickness theory postulated by Riccio & Stoffregen (1991). Wherein, those authors note that passive restraint (such as being seat-belted into a seat) can reduce incidence of motion (and simulator) sickness, this example on a bicycle requires substantial postural stability given that the participant may fall from the bicycle relatively easily if they were to lose their balance. Whilst the two similar but separate causal methodologies cannot be separated in this study, it can be seen that whatever the actual aetiology of simulator sickness, some unavoidable aspects are a function (arguably a required feature; Section 6.2.2) of the simulator itself.

Given the relatively high incidence of SS, the work completed has highlighted this limitation which to a large extent, transcends this particular project which itself has served to highlight standing gaps in the knowledge. Consequently, it is considered in some detail here.

8.4.3.1 Procedures and Scores

In both the situations studied on the bicycle simulator (VR and screen-based) a simulator sickness questionnaire (SSQ) was delivered to all participants at completion/termination of their session. In the case of VR, all participants reported substantial SS scores and 60% of participant sessions were terminated due to same. Clearly, the sample of participants was small but there is little reason to believe that the high rate of SS is not representative of an underlying issue with the VR simulator. This conclusion is also reinforced by the practical absence of SS reported in the screen-based situation.

If therefore, the level of immersion in the VR is high but the level of SS is high, and the level of immersion on the screen-based simulator is variable, while the SS scores reported are low, then, is there a 'middle ground' which would achieve a desirable level of immersion but maintain low SS scores? A worthwhile future project would therefore be to repeat the study performed in this project with a more expansive screen set up (such as surrounding the participant with large projection screens). In effect then, the beginnings of a scale of immersive visualisation versus SS can be constructed as this set of studies would therefore be effectively controlling for the visualisation method. The project presented in this document has identified a potential specific value to

the surrounding projection set up, as without the completion of a VR-type arrangement, there would be no rationale as to why a new person approaching this problem might not simply start from scratch with VR again.

With regard the SS scores it should be noted that some element of the scoring captures 'symptoms' which are a result of being in an enclosed space and cycling for 10 minutes; sweating being the prime example. Sweating scores contribute to the nausea category but have not been excluded. The studies also did not screen for prior susceptibility to SS in large part because there is no established method for doing so. Pre-exposure scores were not captured for the rationale established in Kennedy et al. (1993): that being, the use of pre-exposure questionnaires has a poor reproducibility and does not add to the quality of data observed. Although medical consent questions did screen for pre-existing conditions (in line with Kennedy's methodology), no exclusions were necessary. Consequently, the high rate of SS may not be unusual in comparison to other simulators (where some attempt to screen participants is made). It is possible that pro-active screening would have reduced rates of reported SS and, given that VR sickness is observed in about half of users (King, 2016), and that drop-out rates in the literature involving screened participants are in the order of some tens of percent (Section 6.2.1), pro-active screening may have removed sufficient number of participants to not enable the comparison between the same participants in the different situations. By establishing this base-line as has been done here, future improvements can be compared to it, rather than needing to be framed in isolation.

8.4.3.2 Equipment set up

The floor of the experimental space was laid with 1cm foam padding for the safety of the participants of the simulator studies. Whilst the equipment set up and weight meant that the bicycle equipment was stable, the stability was less than it would have been on a hard floor. Added to this consideration is the findings in Chiarovano et al. (2015) and Horlings et al. (2009) that participants in VR exhibit significantly (and substantially) more postural sway than those not in a VR environment – similar in magnitude to a person with their eyes closed – and that this sway is amplified on a foam flooring (albeit that the foams used in these studies were less dense and in the order of 6–10cm thick, the effects may well be similar). It is possible that the safety set up of the equipment actually exacerbated the problem it was seeking to protect from. An improvement to the test set up would be to ensure contact (or indeed fix) the bicycle rig to the hard floor underneath the foam area. Not only would this be a safety improvement (although it is stressed that no one fell from the bicycle), it would potentially reduce the extremity of the imbalance that participants felt in VR. This may also reduce SS reported scores if there is a link between this 'vertical mismatch' (or indeed 'postural stability') and SS experienced. Bos et

al. (2008) in particular indicates that that vestibular-mismatch causation of SS can be isolated to mismatch in the forces about the longitudinal axis, which would be particularly the case here. Again, this study provides a reasoned and robust baseline against which future improvements can be judged.

The VR equipment used (Oculus Rift HMD) had a field-of-view of 110 degrees; the screen-based set up less (approximately 30 degrees). The presence (or not) of this peripheral view may be of importance when simulating real road-based environments with other users. In particular, the ability to notice other road users in the peripheral vision may be critical to safely navigating those environments. This does not perhaps substantially impact the data collected in this specific study given there were no other road users present in the simulation, but, as noted, should be kept in mind for any future work that builds on this project.

No participants explored speeds outside of the capability of the detection rig applied the bicycle (0–10ms^{-1}). As discussed in Section 7.1.1, the reported speeds from the rear-wheel trainer were disregarded – as the trainer did not properly report speeds below approximately 6kmh^{-1} (1.67ms^{-1}) – in favour of a spoke-passing measurement. In testing, it was found that the particular microcomputer device used could not sample robustly or rapidly enough (the equipment could only sample at less than 1kHz) given the variable background IR levels in the test space and so alternate spokes were bridged with electrical tape to provide a suitably IR-reflective surface. This meant of course that an intervention on the specific bicycle used was required for the purpose of this test which was not ideal. Future work should seek to ensure a microcomputer with a higher sample frequency (in the order of 1MHz), smoothly hand-off the data capture to the trainer as speeds increase, or both, to ensure that a set up can be used which does not require any adjustments to the bicycle; although it should be noted the tape used was temporary and could have been set up in a few minutes or less.

8.4.3.3 Experimental procedure

Owing to the difficulties of switching the equipment set up between the VR and screen-based situations, all participants experienced the VR situation first, and the screen-based situation second. It is possible that there was therefore some learning or habituation between the two situations. Ideally, the order would have been randomised to control for this however with the small number of participants, it would be difficult to detect all but the most extreme learning effects without undertaking a larger study. To some extent, the use of the Driving School scenario ensures that a learning effect for using the equipment is captured *before* the main NOC scenario is experienced in each case. In any case, a larger study should control for this potential learning effect, by randomising order and thus controlling for the situation.

However, with specific regard to the Driving School scenario, all participants were exposed to this prior to the NOC Car Park scenario. The purpose of the Driving School scenario was to accustom the participants to the operation of the bicycle and steering, and ensure their comfort before they were exposed to the experimental circumstance. However, the Driving School was laid out such that there was a strong focus on turning and steering control of the bicycle (by design). This may have inadvertently exacerbated the resulting SS in the participants, especially if it was the experience of turning that was a major contributor to the vestibular mismatch which is a likely cause (or at least a contributor) to the feelings of SS. Indeed, the participant that dropped out of the VR study within 1m30 did so in the Driving School scenario and a number of participants noted that they felt more comfortable in the NOC scenario; although they did not specifically note a link to their feelings (or not) of SS.

As noted in Section 6.2.3, it is not clear from the literature as to whether feelings of SS result in differing resultant behaviour in the simulation; the results here are similarly unclear. There are clear differences between the outcomes in the VR and screen-based situations, but these may result from the immersion differences, rather than the SS; although the two may also be covariant. A larger study may provide the scale of data required (and the resulting statistical power necessary) to control for other factors and establish if changes in reported SS are correlated with changes in behavioural outcomes. This remains a noted gap in the literature and the wider need to establish this impact is further underscored by the work completed here.

8.4.4 Final Conclusions

As stated at the start of Section 8.4, this chapter (and indeed this entire document) has broadened the scope of methodologies for the collection of cyclist data both on an individual basis and with an increasing ability to study the more complex behaviours in an experimentally-controllable manner.

In its current form, the simulator designed, presented and tested here can be taken forward as is and can be a source of valuable and new data. Moreover, with the further development of those clearly identified aspects which are less than ideal (see above), the simulator can be further improved such that it can be used in a wider and more reliable manner. This document has sought to present the totality of the development process both to highlight the difficulty of the task to the reader, but also to provide a pedagogical source for those who may wish to develop a simulator themselves. In addition, the simulator developed here can also be used more widely, to investigate gaps in the knowledge relating to simulator sickness and immersion. Aim 4b of this document: "Identify the scope for, design, construct and test an immersive simulator as a data collection method for bicycle parameters", has therefore been successfully achieved.

The final chapter of this document to follow, draws back together the strands of these methodologies for collecting cycling data and highlights avenues for future research that the work presented in this document has enabled.

Chapter 9

Conclusions

9.1 Recapitulation

This document has presented a constellation of work completed over the period of this PhD. First, following the introduction and motivations of Chapter 1, a comprehensive review of the quantitative cycling literature is presented in Chapter 2. The lack of any practical grounding for the limited literature is explored as the literature in hand is shown to be generally questionable when one attempts to apply quantities to it.

Exploring the implications of some of this literature in more detail, the next chapter (Chapter 3) returns to traffic first principles and establishes the state-fo-the-art with regard to cycle modelling. Using this as a 'stepping-off point', the development of an agent-based simulation model based on the Social Force Model (originally developed by Helbing & Molnár, 1995) is presented. This simulation was used to test one of the key underpinning assumptions in Botma (1995) which has been indirectly incorporated in the Highway Capacity Manual, and thus into industry practice. Using the simulation to test the effect of a non-interaction assumption whilst controlling for every other factor, a fundamental qualitative difference between scenarios with/without the assumption was found. This discovery indicates that a capacity measure based on this assumption may not be correct, and more particularly, that the relationship between the quantitative level of service and flow rate experienced by the cyclist may be fundamentally different than that assumed thus far in the literature which depends critically on the emergent interactions of the cyclists, rather than simply the properties of the individual cyclist.

However, the real quantitative data necessary to reinforce this is expensive to obtain and therefore there is not a wide availability of empirical data to validate the model in Chapter 3. This is an issue much wider than this author's work; indeed, it is a key part of the reason that large-scale data for cyclists does not exist. Chapter 4 therefore considered the three main methodologies for capturing real cyclist behaviour. Bicycle instrumentation was excluded from further consideration as it does not scale well or provide a sufficient level of de-

tail (free from serious experimenter effect) to be a best use of effort in building the scale and detail of empirical data required to inform quantitative literature and quantitatively-led practice. The remaining methodologies – remote observation (such as by camera), and use of a simulator – provide an opportunity to address these issues; remote observation being inherently scalable, and simulators providing an ability to experimentally control circumstance, capture comprehensive data, and test scenarios which one does not have access to in reality.

Chapter 5 therefore starts to address this potential for data collection by presenting a methodology designed by the author to capture cyclist parameters (such as speed and acceleration) from video data remote to the cyclist. Amongst other benefits, this methodology would not entail having to directly interact with the cyclist (such as by fitting their bicycle with detection equipment) and would scale from a few cyclists (such as in the pilot study) to any number of cyclists. The methodology is agnostic to the video source so could be applied to existing CCTV capture and historic footage already held; subject to some limitations on quality, image stability, angle etc. (see Section 5.6.2). This is valuable because there is a wide infrastructure of CCTV equipment deployed on the road system and video survey data is already routinely collected by practitioners. The methodology is also agnostic to mode type and therefore is as easily applicable to motor vehicles – both as a new method, or a supplement to existing methodologies – or pedestrians, as it is to cyclists. This therefore provides an ability to, for example, derive speed data from CCTV or video survey without the need to install (or have previously deployed) speed detection equipment. Such capabilities represent real and valuable steps forward in practical methodology and are detailed such that they can be modified and improved as required by practitioners; something not possible with more recent proprietary products.

Chapter 4 also made the strategic case for the use of a simulator to collect detailed empirical data. Indeed, one of the main problems with collection of empirical data, in particular with regard to cycling where there are potentially many more factors to consider than for simpler movement such as that for motor vehicles or pedestrians (see Chapter 3), is that it can be difficult to get data from analogous circumstances, data which is not affected by the changing circumstances around the cyclists, and data from circumstances you can experimentally control. Chapters 6, 7 and 8 therefore present the background, development (of particular note, this is the first time the development process for a simulator itself has been presented in such detail in the literature) and testing of an immersive bicycle simulator (respectively).

The simulator creates similar opportunities afforded by flight simulators and driving simulators in that it allows the testing of situations in an experimentally-controlled and robust manner, the testing of situations which

would be impossible to test in reality (such as accident scenarios), and the ex ante testing of innovative scheme designs at a fraction of the cost that physical construction (even off-road facsimiles such as Yor et al., 2015) would entail. The design and testing of the simulator resulted in general success but also provides a solid basis for improvements and the development of a future work programme to take it forward. In particular, by establishing a key set of outcomes which were successful, and by highlighting the scale of the issue of simulator sickness in this novel circumstance, a programme for testing, controlling for and improving those aspects which were less successful is developed.

9.1.1 Original Contribution

As a consequence both of the novel development of this simulator, and of the development of an automated methodology for the collection of cyclist parameters from video, this author has established a new and comprehensive set of methods with which the practitioner can now approach the standing issue of the lack of empirically-based cyclist data. Prior to this work, if the practitioner wished to test a scheme they were designing then they would be forced to build a facsimile (e.g. Yor et al., 2015) or simply implement the scheme without the backing of any quantitative tools with respect to the primary users (e.g. Transport for London, 2015b). Now, as a direct result of the work completed here (and only possible because of it), a commercial opportunity exists for a practitioner to be able to create a simulation of the proposal and place actual users into that simulated (and replicable) environment, exploring the users' behaviour in the process. More widely, the simulator provides a foundation for the comprehensive non-commercial research required in the field of cyclist behaviour.

For existing schemes already in place, the methodology created and presented in this document can be applied to new or existing video data collected and analysed to yield simple counts or more complex data such as spatial usage, speed data, and (with suitable quality video) even more complex derived data such as acceleration/deceleration behaviour. Such data is greatly lacking in the literature and was an expensive human-necessary operation to collect. This author's work enables analysis to be performed automatically and at a fraction of the cost (both in reduced human-time and avoiding the need for expensive commercial software packages), something which was not previously possible.

Combined, these achievements represent a real step forward in the ability of the practitioner to assess, quantify and understand cyclist behaviour. Moreover, these two main tranches of work cover the needs for data collection for schemes from ex ante design through ex post implementation, and cover scales from the the detail level of individual cyclists to many cyclists. Also presented in this document is a comprehensive exploration of the quantitative literature

and an establishment of its obvious (and less obvious) invalidities for anything over than the most trivial number of cyclists.

Despite the delivery of these successful outcomes, there are opportunities to learn from the difficulties and to scope future work informed by both those and the accomplishments. The next section highlights the wider limitations and that scope for the future work that results.

9.2 Limitations and Scope for Future Work

This section considers the impact of some of the more general limitations that this project has illustrated and frames those relevant items of future work. Difficulties such as the lack of academic (or practitioner) literature considering cycle infrastructure quantitatively has been a running thread – and indeed a large part of the motivation (Section 1.2) – throughout this document, and will not be restated here. However, there are a few key themes which underlie the limitations which specifically emerge from this work.

9.2.1 The Need for Refined Implementations

A common theme across all the parts of this document is the developed need for the refinement of the implementation of the various models, analyses and simulations. The existence of the work in this project serves as a valid base from which to start each of these work streams. In the absence of this work, the basic process would need to be established and validated with real data, something which is done throughout this document.

In short, the refinements universally necessary are to improve the implementation of the various proven concepts to reduce their computational complexity; and improve them such that they run closer to (or ideally better than) real-time (i.e. 1s in simulation time is computed in 1s of real time or less). This is quite a high bar in many cases and commercial simulation packages – such as Vissim and Aimsun, to choose two popular examples – rarely run that fast for anything other than the most trivial networks. However, the high computational complexity means that as the scale of these implementations grow, the run time rapidly escalates to an unreasonably long scale. If the simulation or analysis cannot be run in reasonable time, then it is of limited practical use. Given the starting point of a absence of tools, this document has sought to prove concepts and ensure validity; future work should focus on the refinement to form practical (and in some cases, commercial) products.

9.2.1.1 Refinement of the SFM

Chapter 3 presented a simulation model based on a modified version of the social force model. The SFM was selected because it fit the requirements to demonstrate the validity of the non-interaction assumption in Botma (1995), it operates in two continuous dimensions, is well understood owing to its wide use in pedestrian microsimulation packages, and replicates known cyclist behaviour such as virtual lane formation. The model demonstrated that the non-interaction assumption is not likely to be valid and thus has been suitably useful for the purposes described.

This said, wider use of the model may not be easily possible with the implementation described as it may not scale to practical use. In particular,

the algorithms deployed have a high computational complexity (in the order of $O(n^2)$) and thus the simulations rapidly become intractable on standard hardware once the number of cyclist agents increases beyond the order of 50. Should future work wish to continue to refine, expand or calibrate the model with real data, then that chapter has provided a robust starting point to do so, but refinement to the algorithm such as to maintain tractability (e.g. switching to a nearest neighbour consideration of other agents) will be necessary.

Implementation aside, the underlying model appears to be an appropriate model for the circumstance, with the utilised literature values providing reasonable outputs in the range of values one might expect given the limited literature in hand and the qualitative impact of the non-interaction assumption being clearly demonstrable. The value of the simulation can only then be improved by calibration with real data, and future work should focus upon this.

9.2.1.2 Refining real-life data collection

As best evidenced by Figure 5.4, there is a substantial systematic error present in some parts of the real-life post-analysis data captured in Chapter 5. This could be easily remedied by multiple camera views which could correct for errors in the homography in the single camera view. Therefore, whilst the real data is that was collected is limited in this document, the process is sound and able to be rolled out almost 'as is' to any future work.

The main barrier to such a wider use is the rate at which processing of images is undertaken. To some extent some optimisations can be made by reducing image sizes or cutting the frame rate of the video from which the images are taken, however, the basic volume of pixels that need to be considered results in a run time substantially slower than real-time. Better than real-time operation is desirable because the process can then be run on live images and so this represents a key future avenue for this body of work and should be pursued. Indeed, such future work will benefit greatly from being built upon the process established in Chapter 5.

9.2.2 'Small-n' sample

Both practical studies presented in this document (Chapters 5 and 6/7/8) were 'small-n' studies in that they involved a small number of participants (6 and 5 respectively). In both cases this was because each was a study to establish proof-of-concept, and to develop experimental and analytical methods to support future, larger-scale studies.

In the case of the real-life study in Chapter 5, the study indicated that subject to some minor data capture adjustments (such as multiple cameras), the methodology was sound. Difficulties organising such an event on a larger scale emerge from managing the number of participants that would be desired and the cost of doing so. The reasoned conclusion to running such an event is to

capitalise on existing events which fit the requirements. Such events involve closure of roads to motor traffic, turnout potential in the scale of tens of thousands, and sufficient extent that suitable sites for capturing video data are numerous. Additionally, the management and welfare coverage is dealt with by the organisers and the recruitment of participants (such as it is) is covered by the sponsors. This means that the scientific focus can instead be on the capture and analysis of data, not of the management required to organise a large-scale public event. Capturing data from such an event was not taken forward within this project but given the work undertaken here, capturing and analysing such data would now simply be a 'handle-turning' exercise and would not require further development of methods required to do so. Future work should therefore apply the methods developed herein to the wider scale of which they are capable.

Chapter 8 also involved a small number of participants. This was by design as the development of the equipment was still novel and thus organising and exposing a large number of participants would have been foolhardy had it turned out that there were key parts of the simulator which did not work. Furthermore, utilising the same participants as the study in Chapter 5 brings benefits of direct comparability on a per-person basis. Despite the small number of participants, as it turned out, the study was broadly a success with those aspects that were universally agreed to not be right (such as the sensitivity of the steering) being items which could be adjusted for subsequent studies. These adjustments were not made here to ensure that all participants experienced the same experimental conditions as it was possible that individuals would have a variable perception of such items. The clear next step is now to perform larger studies in light of adjustments made following feedback from the first study. Such studies would provide the validation required for the equipment with a statistically useful number of participants, and would allow a focus on addressing those aspects of the experience which may have contributed to the incidence of simulator sickness observed (discussed in Section 8.4.3) and on other relevant outputs.

As noted with regard the real-life study, the work completed throughout this project has been a fundamental necessity to enable such work; future work can now focus on refinement and the practical use of the equipment, as opposed to needing to develop the principles and the interfaces required from scratch, as would have been the case otherwise.

9.2.3 Simulator Sickness

The impact of simulator sickness is discussed in detail in Section 8.4.3, and will not be restated here. The issue has been brought up again to re-highlight the general issue in the field which is that the cause(s) of simulator sickness (and indeed motion sickness) are not agreed, in many ways are poorly understood,

and are generally not well studied.

The work presented in this document has drawn together relevant literature relating to simulator sickness and found it lacking. A major recommendation to take forward is a systematic study of simulator sickness on a larger scale. Amongst other benefits, the simulator developed in this document can assist such studies by being a experimentally-controllable platform for their completion.

9.3 Future Work

Chapter 2 demonstrated the lack of robust quantitative literature surrounding bicycle infrastructure. Chapter 3 that followed, reinforced that observation by demonstrating the probable invalidity of one of the few pieces of quantitative literature that does exist. As such there is a huge scope for quantitative work relating to cyclists. The following notes a few of the specific parts that this work has enabled:

9.3.1 Social Force Model for Bicycles

The social force model (and similar) has been widely used for transport modelling for many years. As such there is a strong understanding in the literature pertaining to other individual modes but this does not extend to inter-modal models, which are still in their infancy. Consequently, the cycle model that is presented in Chapter 3 has the potential to become part of the body of work which may lead to a comprehensive model for all transport modes; e.g. for use in a 'shared-space' situation. Work taken forward should focus primarily on empirical validation (on a large scale) of the operation of the model, and computational refinement of the algorithmic processes so as to ensure the simulation (or simulations it may inform) remain tractable.

9.3.2 Remote Observation of Cyclists

A methodology to collect large scale data was developed in Chapter 5. The concept is proven and consequently future work can now focus on the practical use of the methodology to collect wide-scale data (for any mode as, it is again noted, the methodology is mode-agnostic). Whilst the state-of-the-art has progressed in the time between the completion of that work, and the current time of writing (early 2014 through early 2017), the methodology presented remains of value both as an open and defined (as opposed an obfuscated proprietary process) alternative or augmentation to current proprietary tools.

Considering the methodology presented, as with the SFM implementation noted above, runtime can be lengthy and therefore refinement can be targetted on the reduction of runtime with an eventual goal of better-than-realtime run speeds on desktop hardware. This would enable the methodology to be run on real-time video streams such as CCTV. This would have the capability of revolutionising the collection of traffic data both in terms of the scope of data which can be collected, and the cost of doing so.

However, even without that refinement, the cost of computer runtime is substantially lower than the equivalent cost of a human to perform the same task; and becoming cheaper all the time, whilst for humans generally the reverse is true. Additionally, the process developed is essentially deterministic and thus not prone to the same sorts of errors that are common with human

operators. In this author's professional experience, even a simple traffic count can yield data pertaining to the same traffic with differences in the order of 10% between analysts. Off-line analysis of existing and current data sets can therefore be more widely explored by the developed process, as is. Data which could inform necessary future work, such as the determination of accurate PCU values for bicycles (Section 3.1.2), would now be much more readily achievable.

9.3.3 Bicycle Simulator

In its current state, the bicycle simulator developed in this project is to (at a minimum) Technology Readiness Level 3 (per the EU Horizon 2020 definitions; European Commission, 2015) – i.e. "experimental proof of concept". As has been noted, the simulator can be taken and applied to real and contrived circumstances in its current form, if one is prepared to accept the experimental 'overhead' of high rates of simulator sickness. However, the improvements identified in Section 8.4 (and above), such as steering improvements and exploration of intermediate-level immersive visualisation, would serve to improve the reliability of the simulator, may serve to reduce the incidence of simulator sickness, and would result in a 'better product'. Refinement should focus upon these items and potentially the exploration of the commercial (and research) value such a system provides and that this work has created.

More widely, one can also consider the novel use cases that this simulator allows, such as the linking of multiple instances of the simulator on a local network (or the internet) and allowing multiple participants to cycle in the *same* environment simultaneously. Explorations of the interactions of cyclists with one another is a field of work which is entirely unexplored, and certainly unexplored in any controlled circumstance such as this. Again, such new areas of research can now only be explored owing to the work invested in this project and the successes delivered.

9.4 Final Words

The forgoing chapters, and this chapter, have brought together a comprehensive literature review with regard to the quantitative aspects of cycle infrastructure delivery and have highlighted how tenuous that literature is. Methodologies were developed which now allow the wide-scale collection of transport data for cyclists (and indeed other modes) in a cost effective and replicable way that was not previously possible. Chapters 6, 7 and 8 presented the design, construction and testing of a simulator which allows the collection of cyclist data on a more targetted basis both in contrived and experimentally-controlled circumstances.

In Section 1.3, four core aims/objectives were presented for this project:

1. Demonstrate that the state-of-the-art of the quantitative cycle infrastructure capacity literature is insufficient to meet practitioner needs;
2. Identify the appropriateness of quantitative modelling techniques as applied to cycle infrastructure;
3. Improve understanding of methodologies for data collection pertaining to cyclists; and
4. Develop and apply those methodologies to determine if there is scope for their further development and/or use. In particular:
   (a) Identify the scope for the use of video observation as a data collection method for bicycle parameters.
   (b) Identify the scope for, design, construct and test an immersive simulator as a data collection method for bicycle parameters.

The literature review presented in Chapter 2 illustrated the current state of the literature with regard to the the quantitative capacity of cycle infrastructure, and found it lacking (Objective 1). This is important not solely simply for good practice, but because the scale of cyclists that world cities (such as London) will have to contend with is of a magnitude not previously addressed (also see Section 1.2). As the current state of the literature is not up to the required standard, Chapter 3 moved on to establish the state-of-the-art of cycle modelling (Objective 2) and framed that in the wider foundation of fundamental traffic flow theory. A model based on the SFM and the limited literature was created and demonstrated that the non-interaction assumption for cyclists (from Botma, 1995) which has filtered into practitioner literature, is not robust for anything other than the most trivial levels of flow.

Chapters 2 and 3 also demonstrated that the literature basis is not quantitatively sufficient to inform current and future practice. Chapter 4 therefore focused on the broad methodologies which could be used to gather further and wider empirical data pertaining to cyclists (thereby achieving Objective 3).

This chapter highlighted, amongst other items, the difficulties in scaling empirical data collection for cyclists and whilst individually valuable, the inherently non-scalable nature of some methods, such as bicycle instrumentation. The remaining methodologies taken forward: remote video observation and image analysis (Objective 4a), and the bicycle simulator (Objective 4b), have produced tools which can collect a range of parameters relating to cyclists (and indeed other modes in the case of the video analysis presented in Chapter 5) and for situations to which one could not possibly hope to access with instrumentation in the case of the simulator (i.e. delivering Objective 4).

With final regard to the project's motivation, one cannot perhaps simply return to industry with all the answers required to implement quantitatively-robust cycle infrastructure. But, if the author, reader, or indeed another practitioner, now finds themselves requiring an ability to establish the capacity of a given segment of cycle infrastructure, then they can turn to the simulation model created in this project (Chapter 3), and the future work this has laid the groundwork for. If that person requires a way to collect spatial (and related) information for one or many road users, then they can turn to the methodology established in Chapter 5. And, if that person requires a more comprehensive way to collect data on the individual in experimental circumstances, perhaps with regard to a scheme they have yet to deliver, then Chapters 6 and 8 provide a simulator system they can use.

Individually, each of these things provides useful tools and processes to meet the imminent need of the practitioner. However, if as a field, we desire a comprehensive understanding of the properties and behaviour of cyclists, then a combination of all of these items will be needed. This project has built the foundation necessary to achieve that understanding.

References

4ChordsNoNet. (2016). *Please Slow Down on CS3 – Head on Crash.* Retrieved from [2016-10-16]https://www.youtube.com/watch?v=5RH5HBq5hOg

A&H Software House Inc. (2017). *Luxriot VCA Module.* Retrieved from http://www.luxriot.com/product/luxriot-vca-module/

Airbus SAS. (2016). *Airbus Flight Training.* Retrieved from [2016-09-19]http://www.airbus.com/support/training/flight/

Allen, R. W., Park, G., Fiorentino, D., Rosenthal, T. J., & Cook, M. L. (2006). Analysis of simulator sickness as a function of age and gender. In *Driving simulator conference 2006 europe, paris.*

Alshaer, A., Regenbrecht, H., & O'Hare, D. (2017). Immersion factors affecting perception and behaviour in a virtual reality power wheelchair simulator. *Applied Ergonomics*, *58*, 1–12. doi: 10.1016/j.apergo.2016.05.003

American Association of State Highway and Transportation Officials. (1999). *Guide for the Development of Bicycle Facilities* (4th ed.). Washington, D.C.: American Association of State Highway and Transportation Officials.

Amigo, I. (2016). *When Cyclists Oppose Bike Lanes.* Retrieved from [2016-09-06]https://nextcity.org/daily/entry/cyclists-oppose-bike-lane-plan-madrid

Anvari, B., Bell, M. G., Angeloudis, P., & Ochieng, W. Y. (2014). Long-range Collision Avoidance for Shared Space Simulation based on Social Forces. *Transportation Research Procedia*, *2*(0), 318–326. doi: 10.1016/j.trpro.2014.09.023

Anvari, B., Bell, M. G. H., Sivakumar, A., & Ochieng, W. Y. (2015). Modelling shared space users via rule-based social force model. *Transportation Research Part C: Emerging Technologies*, *51*, 83–103. doi: 10.1016/j.trc.2014.10.012

Apfel, C. C., Greim, C. A., Haubitz, I., Goepfert, C., Usadel, J., Sefrin, P., & Roewer, N. (1998). A risk score to predict the probability of postoperative vomiting in adults. *Acta Anaesthesiologica Scandinavica*, *42*(5), 495–501. doi: 10.1111/j.1399-6576.1998.tb05157.x

Arasan, V. T., & Koshy, R. Z. (2005). Methodology for Modeling Highly Heterogeneous Traffic Flow. *Journal of Transportation Engineering*, *131*(7), 544–551.

Arduino. (2016). *Arduino Board Uno.* Retrieved from [2016-10-01]https://www.arduino.cc/en/Main/ArduinoBoardUno

Atkins. (2013). *SATURN.* Epsom, Surrey, UK: Atkins-ITS Transport. Retrieved from http://www.saturnsoftware.co.uk/

Baker, L. (2016). *DfT must allow investment in transport schemes to improve air quality - MPs.* Retrieved from https://www.transportxtra.com/publications/local-transport-today/news/49932/dft-must-allow-investment-in-transport-schemes-to-improve-air-quality--mps

Barcelo, J. (2010). *Fundamentals of Traffic Simulation* (Vol. 145). doi: 10.1007/978-1-4419-6142-6

Barker, J. B., Biehler, A. D., Brown, L. L., Clark, W. A. V., & Ekern, D. S. (2008). *NCHRP Report 616: Multimodal Level of Service Analysis for Urban Streets* (Tech. Rep.). Washington, D.C.: Transportation Research Board.

Bkool. (2016). *Bkool Simulator.* Retrieved from [2016-10-18]http://www.bkool.com/en-GB/cycling-simulator

Blascovich, J., Loomis, J., Beall, A. C., Swinth, K. R., Hoyt, C. L., & Bailenson, J. N. (2002). TARGET ARTICLE: Immersive Virtual Environment Technology as a Methodological Tool for Social Psychology. *Psychological Inquiry*, *13*(2), 103–124. doi: 10.1207/S15327965PLI1302_01

Bos, J. E., Bles, W., & Groen, E. L. (2008). A theory on visually induced motion sickness. *Displays*, *29*(2), 47–57.

Botma, H. (1995). Method to Determine Level of Service for Bicycle Paths and Pedestrian-Bicycle Paths. *Transportation Research Record*, *1502*, 38–44.

Botma, H., & Papendrecht, J. H. (1991). Traffic Operation of Bicycle Traffic. *Transportation Research Record*, *1320*, 65—-72.

Brackstone, M., Waterson, B., & McDonald, M. (2009). Determinants of following headway in congested traffic. *Transportation Research Part F: Traffic Psychology and Behaviour*, *12*(2), 131–142. doi: 10.1016/j.trf.2008.09.003

British Cycling. (2014). *Time to #ChooseCycling: British Cycling's vision for how Britain can become a true cycling nation* (Tech. Rep.). British Cycling.

Brockmyer, J. H., Fox, C. M., Curtiss, K. A., Mcbroom, E., Burkhart, K. M., & Pidruzny, J. N. (2009). Journal of Experimental Social Psychology The

development of the Game Engagement Questionnaire : A measure of engagement in video game-playing. *Journal of Experimental Social Psychology*, *45*(4), 624–634. doi: 10.1016/j.jesp.2009.02.016

Brooks, J. O., Goodenough, R. R., Crisler, M. C., Klein, N. D., Alley, R. L., Koon, B. L., ... Wills, R. F. (2010). Simulator sickness during driving simulation studies. *Accident Analysis and Prevention*, *42*(3), 788–796. doi: 10.1016/j.aap.2009.04.013

Bundeshaus Nord. (2001). *Final Report No . 1793 by the Aircraft Accident Investigation Bureau* (Tech. Rep. No. November).

Cambridge Cycling Campaign. (2014). *Making Space for Cycling: A guide for new developments and street renewals* (Tech. Rep.). Cambridge, UK: Cambridge Cycling Campaign.

Carrignon, D. (2009). Assessment of the impact of cyclists on heterogeneous traffic. *Traffic Engineering and Control*(July), 323–325.

Cerezuela, G. P., Tejero, P., Chóliz, M., Chisvert, M., & Monteagudo, M. J. (2004). Wertheim's hypothesis on 'highway hypnosis': Empirical evidence from a study on motorway and conventional road driving. *Accident Analysis and Prevention*, *36*(6), 1045–1054.

Chiarovano, E., de Waele, C., MacDougall, H. G., Rogers, S. J., Burgess, A. M., & Curthoys, I. S. (2015). Maintaining Balance when Looking at a Virtual Reality Three-Dimensional Display of a Field of Moving Dots or at a Virtual Reality Scene. *Frontiers in neurology*, *6*(July), 164. doi: 10.3389/fneur.2015.00164

Christou, G. (2014). The interplay between immersion and appeal in video games. *Computers in Human Behavior*, *32*, 92–100. doi: 10.1016/j.chb.2013.11.018

CIA. (2016a). *CIA World Factbook: Netherlands.* Retrieved from [2016-08-26]`https://www.cia.gov/library/publications/the-world-factbook/geos/nl.html`

CIA. (2016b). *CIA World Factbook: United Kingdon.* Retrieved from [2016-08-26]`https://www.cia.gov/library/publications/the-world-factbook/geos/uk.html`

City of Copenhagen. (2011a). *Copenhagen: City of Cyclists - Bicycle Account 2010* (Tech. Rep.). Copenhagen, Denmark: City of Copenhagen: The Technical and Environmental Administration.

City of Copenhagen. (2011b). *Good, Better, Best: The city of Copenhagen's Bicycle Strategy 2011-2025.* Copenhagen, Denmark: City of Copenhagen: The Technical and Environmental Administration.

City of Copenhagen. (2013a). *Copenhagen City of Cyclists - Bicycle Account 2012* (Tech. Rep.). Copenhagen, Denmark: City of Copenhagen: The Technical and Environmental Administration.

City of Copenhagen. (2013b). *Micro simulation of cyclists in peak hour traffic* (Tech. Rep. No. February). Copenhagen, Denmark: COWI.

City of Davis Bicycle Advisory Commission. (2009). *City of Davis Bicycle Plan* (Tech. Rep.). Retrieved from `http://cityofdavis.org/bicycles/pdfs/Bike-Plan-2009.pdf`

City of Los Angeles Department of City Planning. (2011). *2010 Bicycle Plan: Technical Design Handbook.* Los Angeles, CA, USA. Retrieved from `http://planning.lacity.org/cwd/gnlpln/transelt/NewBikePlan/Txt/LACITYBICYCLEPLANTDH.pdf`

Classen, S., Bewernitz, M., & Shechtman, O. (2011). Driving simulator sickness: An evidence-based review of the literature. *American Journal of Occupational Therapy*, *65*(2), 179–188.

Cossalter, V. (2006). *Motorcycle Dynamics* (2nd ed.). Author.

Cremer, J., Kearney, J., & Willemsen, P. (1997). Directable Behavior Models for Virtual Driving Scenarios. *Transactions of the Society for Computer Simulation*, *14*, 87—-96.

CROW. (1993). *Record 10: Sign up for the bike – Design manual for a cycle-friendly infrastructure.* Ede, Netherlands: Author.

CROW. (1998). *Recommendations for traffic provisions in built-up areas: ASVV* (4th ed.). Ede, Netherlands: Author.

CROW. (2007). *Record 25: Design Manual for Bicycle Traffic.* Ede, Netherlands: Author.

Csikszentmihalyi, M. (1990). *Flow.* HarperPerennial.

CycleOps. (2016). *CycleOps Virtual Training.* Retrieved from [2016-10-18]`https://www.cycleops.com/virtualtraining/overview`

Cycling Embassy of Denmark. (2012). *Collection of Cycle Concepts 2012.*

Cycling Embassy of Great Britain. (2014). *Sustrans' 'Cycle Design Handbook'.* Retrieved from `http://www.cycling-embassy.org.uk/news/2014/05/02/sustrans?-?cycle-design-handbook?`

De Leuw, Cather & Company. (1972). *Bicycle Circulation and Safety Study* (Tech. Rep.). Davis, CA: City of Davis.

Dekoster, J., & Schollaert, U. (1999). *Cycling: the way ahead for towns and cities.*

Department for Transport. (2004). *Design Manual for Roads and Bridges: Volume 13: COBA 11 User Manual.* London, UK: Department for Transport.

Department for Transport. (2005). *Design Manual for Roads and Bridges: TA90/05: The Geometric Design of Pedestrian, Cycle and Equestrian Routes* (Vol. 6) (No. February 2005). London, UK: Department for Transport. Retrieved from `http://www.dft.gov.uk/ha/standards/dmrb/vol6/section3/ta9005.pdf`

Department for Transport. (2006). *Traffic Advisory Leaflet 1/06: General Principles of Traffic Control by Light Signals: Part 3 of 4* (Tech. Rep. No. March). London: Department for Transport.

Department for Transport. (2007). *Manual for Streets.* London: Thomas Telford Publishing.

Department for Transport. (2008). *Local Transport Note 2/08: Cycle Infrastructure Design* (Tech. Rep. No. October).

Department for Transport. (2012). *Local Transport Note 1/12: Shared Use Routes for Pedestrians and Cyclists* (Tech. Rep. No. September).

Department for Transport. (2014). *Value for Money Assessment for Cycling Grants* (Tech. Rep. No. August). London, UK: Department for Transport.

Department for Transport. (2016). *Cycle Traffic and the Strategic Road Network.* Highways England. Retrieved from `http://www.standardsforhighways.co.uk/ha/standards/ians/pdfs/ian195.pdf`

Dozza, M., & Fernandez, A. (2014). Understanding bicycle dynamics and cyclist behavior from naturalistic field data (November 2012). *IEEE Transactions on Intelligent Transportation Systems*, *15*(1), 376–384.

Draper, M. H., Viirre, E. S., Furness, T. A., & Gawron, V. J. (2001). Effects of Image Scale and System Time Delay on Simulator Sickness within Head-Coupled Virtual Environments. *Human Factors*, *43*(1), 129—-146.

Duives, D. C., Daamen, W., & Hoogendoorn, S. P. (2015). Quantification of the level of crowdedness for pedestrian movements. *Physica A: Statistical Mechanics and its Applications*(xxxx). doi: 10.1016/j.physa.2014.11.054

Dynastream Innovations. (2016). *What is Ant+.* Retrieved from [2016-10-01]`https://www.thisisant.com/consumer/ant-101/what-is-ant/`

Ebenholtz, S. M. (1992). Motion Sickness and Oculomotor Systems in Virtual Environments. *Presence: Teleoperators and Virtual Environments*, *1*(3), 302–305. doi: 10.1162/pres.1992.1.3.302

Elliott, S. (2008). *Beyond the Box: Orange Box Afterthoughts and The Future of Valve.* Retrieved from [2016-08-28]http://www.1up.com/features/beyond-the-box

European Commission. (2015). Technology readiness levels (TRL). *HORIZON 2020 WORK PROGRAMME 2014-2015 General Annexes, Extract from Part 19 - Commission Decision C*(2014), 4995.

Extra Credits. (2012). *Extra Credits: Video Game Music.* Retrieved from [2016-09-14]https://www.youtube.com/watch?v=CKgHrz_Wv6o

Extra Credits. (2014a). *Affordances - How Design Teaches Us Without Words - Extra Credits.* Retrieved from [2016-09-14]https://youtu.be/QCSXEKHL6fc

Extra Credits. (2014b). *The Magic Circle - How Games Transport Us to New Worlds - Extra Credits.* Retrieved from [2016-09-14]https://www.youtube.com/watch?v=qZ-EY9gTsgU

Extra Credits. (2014c). *Simulation Sickness - Causes and Cures for Game Headaches - Extra Credits.* Retrieved from [2016-09-26]https://www.youtube.com/watch?v=CvimYs7tnRM

Farnell. (2016). *Omron Electronic Components EESY51 Photomicrosensor, Reflective, 3.5mm.* Retrieved from [2016-10-01]http://uk.farnell.com/omron-electronic-components/eesy171/photomicrosensor-reflect-3-5mm/dp/2301879

Federal Highway Administration. (2015). *Separated Bike Lane Planning and Design Guide* (No. May). U.S. Department of Transportation. Retrieved from http://www.fhwa.dot.gov/environment/bicycle_pedestrian/publications/separated_bikelane_pdg/page00.cfm

FFmpeg Project. (2014). *FFmpeg.* FFmpeg Project. Retrieved from https://www.ffmpeg.org/

FLIR Intelligent Transportation Systems. (2016). *ThermiCam.* Retrieved from http://www.flirmedia.com/MMC/CVS/Traffic/IT_0013_EN.pdf

Fruin, J. J. (1971). *Pedestrian Planning and Design* (1st ed.). New York, NY: Metropolitan Association of Urban Designers and Environmental Planners Inc.

Garcia, A., Gomez, F. A., Llorca, C., & Angel-Domenech, A. (2015). Effect of width and boundary conditions on meeting maneuvers on two-way separated cycle tracks. *Accident Analysis and Prevention*, *78*, 127–137. doi: 10.1016/j.aap.2015.02.019

Garg, K., & Nayar, S. K. (2007). Vision and rain. *International Journal of Computer Vision*, *75*(1), 3–27.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception.* Taylor & Francis. Retrieved from https://books.google.co.uk/books?id=8BSLBQAAQBAJ

Gibson, J. J., & Crooks, L. E. (1938). A Theoretical Field-Analysis of Automobile-Driving. *The American Journal of Psychology*, *51*(3), 453. doi: 10.2307/1416145

Gillett, F. (2016). *Hair-raising moment cyclists are thrown off bikes in central London cycle-superhighway pile-up.* London, UK. Retrieved from http://www.standard.co.uk/news/london/hairraising-moment-cyclists-are-thrown-off-bikes-in-central-london-pileup-crash-a3367621.html

Golbuff, L., & Aldred, R. (2011). *Cycling Policy in the UK A historical and thematic overview* (Tech. Rep.). London, UK: Sustainable Mobilities Research Group.

Golding, J. F. (2006). Motion sickness susceptibility. *Autonomic Neuroscience: Basic and Clinical*, *129*(1-2), 67–76.

Gould, G., & Karner, A. (2010). Modeling Bicycle Facility Operation. *Transportation Research Record: Journal of the Transportation Research Board*, *2140*, 157–164.

Gowri, A., Venkatesan, K., & Sivanandan, R. (2009). Object-oriented methodology for intersection simulation model under heterogeneous traffic conditions. *Advances in Engineering Software*, *40*(10), 1000–1010. doi: 10.1016/j.advengsoft.2009.03.015

Greater London Authority. (2013). *The Mayor's Vision for Cycling in London: An Olympic Legacy for all Londoners* (Tech. Rep.). London, UK.

Grechkin, T. Y., Chihak, B. J., Cremer, J. F., Kearney, J. K., & Plumert, J. M. (2012). Perceiving and acting on complex affordances: how children and adults bicycle across two lanes of opposing traffic. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(1), 23–36.

Groot, S. D., Winter, J. C. F. D., Mulder, M., Division, S., & Wieringa, P. A. (2011). Nonvestibular Motion Cueing in a Fixed-Base Driving Simulator : Effects on Driver Braking and Cornering Performance. *Presence: Teleoperators and Virtual Environments*, *20*(2), 117–142.

Halcrow. (2010). *Pedestrian Modelling - Simulating movements to effectively manage flows.* London, UK.

Hamme, D. V., Goeman, W., Veelaert, P., & Philips, W. (2015). Robust monocular visual odometry for road vehicles using uncertain perspective projection. *EURASIP Journal on Image and Video Processing*, *2015*(1), 10. doi: 10.1186/s13640-015-0065-6

Helbing, D., Buzna, L., Johansson, A., & Werner, T. (2005). Self-Organized Pedestrian Crowd Dynamics: Experiments, Simulations, and Design Solutions. *Transportation Science*, *39*(1), 1–24. doi: 10.1287/trsc.1040.0108

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, *407*(6803), 487–90. doi: 10.1038/35035023

Helbing, D., & Johansson, A. (2013a). Pedestrian, Crowd and Evacuation Dynamics. In R. Meyers (Ed.), *Encyclopedia of complexity and systems science.* Berlin: Springer-Verlag. Retrieved from `http://www.springerreference.com/pdf/61/60554.pdf`

Helbing, D., & Johansson, A. (2013b). Pedestrian, Crowd, and Evacuation Dynamics.

Helbing, D., & Molnár, P. (1995). Social force model for pedestrian dynamics. *Physical Review E*, *51*(5), 4282–4286.

Helbing, D., Molnár, P., Farkas, I. J., & Bolay, K. (2001). Self-organizing pedestrian movement. *Environment and Planning B: Planning and Design*, *28*, 361–383.

Helland, A., Lydersen, S., Lervåg, L.-E., Jenssen, G. D., Mørland, J., & Slørdal, L. (2016). Driving simulator sickness: Impact on driving performance, influence of blood alcohol concentration, and effect of repeated simulator exposures. *Accident Analysis & Prevention*, *94*, 180–187. doi: 10.1016/j.aap.2016.05.008

Holroyd, E. M. (1963). Effect of Motorcycles and Pedal Cycles on Saturation Flow at Traffic Signals. *Roads and Road Construction*, *41*(490), 315–316.

Homburger, W. S. (1976). *Capacity of bus routes, and of pedestrian and bicycle facilities* (Tech. Rep.). Berkeley, CA: University of California, Berkeley.

Hoogendoorn, S. P., & Daamen, W. (2005). Pedestrian Behavior at Bottlenecks. *Transportation Science*, *39*(2), 147–159. doi: 10.1287/trsc.1040.0102

Horlings, C. G. C., Carpenter, M. G., Küng, U. M., Honegger, F., Wiederhold, B., & Allum, J. H. J. (2009). Influence of virtual reality on postural stability during movements of quiet stance. *Neuroscience Letters*, *451*(3), 227–231.

Hou, J., Nam, Y., Peng, W., & Lee, K. M. (2012). Effects of screen size, viewing angle, and players' immersion tendencies on game experience. *Computers in Human Behavior*, *28*(2), 617–623. doi: 10.1016/j.chb.2011.11.007

HTC Corporation. (2016). *HTC Vive.* Retrieved from [2016-10-18]`https://www.vive.com/uk/`

Hummer, J., Patten, R., Toole, R., Schneider, R., Green, J., Highes, R., & Fain, S. (2006). *Evaluation of Safety, Design, and Operation of Shared-Use Paths* (Tech. Rep. No. July). Retrieved from http://ntis.library.gatech.edu/handle/123456789/4787

Institute for Transport Studies. (2007). *DRACULA.* Leeds, UK: University of Leeds. Retrieved from http://www.its.leeds.ac.uk/software/dracula/

ITU. (1990). *Basic Parameter Values for the HDTV Standard for the Studio and for International Programme Exchange.* Geneva, Switzerland.

Ivory, J. D., & Kalyanaraman, S. (2007). The effects of technological advancement and violent content in video games on players' feelings of presence, involvement, physiological arousal, and aggression. *Journal of Communication*, *57*(3), 532–555.

Jarrett, J., Woodcock, J., Griffiths, U. K., Chalabi, Z., Edwards, P., Roberts, I., & Haines, A. (2012). Effect of increasing active travel in urban England and Wales on costs to the National Health Service. *The Lancet*, *379*(9832), 2198–2205.

JCT Consultancy. (2011). *LinSig.* Lincoln, UK: JCT Consultancy. Retrieved from http://www.jctconsultancy.co.uk/Software/LinSigV3/linsigv3.php

Jennett, C., Cox, A. L., Cairns, P., Dhoparee, S., Epps, A., Tijs, T., & Walton, A. (2008). Measuring and defining the experience of immersion in games. *International Journal of Human Computer Studies*, *66*(9), 641–661.

Jin, S., Qu, X., Zhou, D., Xu, C., Ma, D., & Wang, D. (2015). Estimating cycleway capacity and bicycle equivalent unit for electric bicycles. *Transportation Research Part A: Policy and Practice*, *77*, 225–248. doi: 10.1016/j.tra.2015.04.013

Jones, D. E. H. (2006). The stability of the bicycle. *Physics Today*(April 1970), 51–56.

Kaewtrakulpong, P., & Bowden, R. (2001). An Improved Adaptive Background Mixture Model for Real-time Tracking with Shadow Detection. In *Proceedings of the 2nd european workshop on advanced video based surveillance systems* (pp. 1–5). Retrieved from http://personal.ee.surrey.ac.uk/Personal/R.Bowden/publications/avbs01/avbs01.pdf

Kelly, G. (2016). *Horrific cycling crash shows danger of overtaking on 'at capacity' superhighways.* London, UK. Retrieved from http://www.telegraph.co.uk/men/recreational-cycling/horrific-cycling-crash-shows-danger-of-overtaking-on-capacity-su/

Kennedy, R. S., Drexler, J., & Kennedy, R. C. (2010). Research in visually induced motion sickness. *Applied Ergonomics*, *41*(4), 494–503. doi: 10.1016/j.apergo.2009.11.006

Kennedy, R. S., Lane, N. E., Berbaum, K. S., & Lilienthal, M. G. (1993). *Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness* (Vol. 3).

Kimber, R., McDonald, M., & Hounsell, N. B. (1986). *RR67: The prediction of saturation flows for road junctions controlled by traffic signals* (Tech. Rep.). Transport and Road Research Laboratory.

Kimber, R., Semmens, M. C., & Shewey, P. J. H. (1982). Saturation flows at traffic signal junctions: Studies on test track and public roads. In *Institute of electrical engineers conference on road traffic signalling (conference publication no.1 207)* (pp. 1—-4).

King, T. (2016). VR Sickness. *Custom PC*(157), 12.

Klotz Associates. (2007). *Street Smarts Task Force Bicycle Facilities Toolbox.* Austin, TX, USA.

Kwon, J. H., Powell, J., & Chalmers, A. (2013). How level of realism influences anxiety in virtual reality environments for a job interview. *International Journal of Human Computer Studies*, *71*(10), 978–987. doi: 10.1016/j.ijhcs.2013.07.003

Laker, L. (2014). *Hackney cycle counter to clock its 500,000th cyclist.* Retrieved from [2016-09-19]`http://www.cyclingweekly.co.uk/news/hackney-cycle-counter-to-clock-its-500000th-cyclist-19353`

Landis, B. W., Vattikuti, V. R., & Brannick, M. T. (1997). Real-Time Human Perceptions Toward a Bicycle Level of Service. *Transportation Research Record*, *1578*, 119–126.

Lane, R. (2016). Developing for VR. *Custom PC*(156), 94—-95.

Legion Ltd. (2013). *Legion.* London, UK: Legion Ltd. Retrieved from `http://www.legion.com/`

Levin, S. (2016). *Uber admits to self-driving car 'problem' in bike lanes as safety concerns mount.* San Francisco, CA, USA. Retrieved from `https://www.theguardian.com/technology/2016/dec/19/uber-self-driving-cars-bike-lanes-safety-san-francisco`

Lewin, K. (1951). *Field theory in social science: selected theoretical papers.* New York, NY, USA: Harper & Row.

Li, F. (1995). Capacity and level of service for urban bicycle path in China. *Chinese Municipal Engineering*, *71*, 11–14.

Li, M., Shi, F., & Chen, D. (2011). Analyze bicycle-car mixed flow by social force model for collision risk evaluation. In *3rd international conference on road safety and simulation.* Indianapolis, IN, USA.

Li, Z., Wang, W., Liu, P., Bigham, J., & Ragland, D. R. (2012). Modeling Bicycle Passing Maneuvers on Multi-Lane Separated Bicycle Paths. *Journal of Transportation Engineering*(January).

Li, Z., Ye, M., Li, Z., & Du, M. (2015). Some Operational Features in Bicycle Traffic Flow. *Transportation Research Record: Journal of the Transportation Research Board*, *2520*, 18–24. doi: 10.3141/2520-03

Liang, X., Mao, B., & Xu, Q. (2012). Psychological-Physical Force Model for Bicycle Dynamics. *Journal of Transportation Systems Engineering and Information Technology*, *12*(2), 91–97. doi: 10.1016/S1570-6672(11)60197-9

L'Ile-de-France service de presse. (2011). *L'Ile-de-France adopte son plan vélo.* Paris, France. Retrieved from `http://www.iledefrance.fr/sites/default/files/medias/2013/04/documents/plan_velo.pdf`

Ling, Z., Cherry, C. R., & Dhakal, N. (2017). Factors influencing single-bicycle crashes at skewed railroad grade crossings. *Journal of Transport & Health*, 1–10. doi: 10.1016/j.jth.2017.01.004

Liu, X., Shen, D., & Ren, F. (1993). Operational Analysis of Bicycle Interchanges in Beijing, China. *Transportation Research Record*, *1396*, 18—–21.

Llorca, C., Angel-Domenech, A., Agustin-Gomez, F., & Garcia, A. (2015). Motor vehicles overtaking cyclists on two-lane rural roads: Analysis on speed and lateral clearance. *Safety Science*. doi: 10.1016/j.ssci.2015.11.005

Lydall, R. (2015). *Boris Johnson gives go-ahead for 'Crossrail' cycle superhighway through central London.* Retrieved from [2015-03-11]`http://www.standard.co.uk/news/transport/boris-johnson-gives-goahead-for-crossrail-cycle-superhighway-through-central-london-10004550.html`

Malone, L. A., & Bastian, A. J. (2010). Thinking About Walking: Effects of Conscious Correction Versus Distraction on Locomotor Adaptation. *Journal of Neurophysiology*, *103*(4), 1954–1962. doi: 10.1152/jn.00832.2009

McNally, M. G. (2007). The Four Step Model. In Hensher & Button (Eds.), *Handbook of transport modeling* (2nd ed., chap. 3). Pergammon.

Mehta, K., Mehran, B., & Hellinga, B. (2015). Evaluation of the Passing Behavior of Motorized Vehicles When Overtaking Bicycles on Urban Arterial Roadways. *Transportation Research Record: Journal of the Transportation Research Board*, *2520*, 8–17. doi: 10.3141/2520-02

Miah, S., Kaparias, I., Stirling, D. M., & Liatsis, P. (2016). 16-3977: Development and Testing of Prototype Instrumented Bicycle for Prevention of Cyclist Accidents. In *Transportation research board 95th annual meeting.*

Miller, R. (1976). *Width Requirements for Bikeways: A Level of Service Approach* (Master Thesis). University of California, Davis.

Ministry of Housing and Urban-Rural Development of China. (2012). *Code for Design of Urban Road Engineering: CJJ37-2012.*

Morel, M., Bideau, B., Lardy, J., & Kulpa, R. (2015). Advantages and limitations of virtual reality for balance assessment and rehabilitation. *Neurophysiologie clinique = Clinical neurophysiology*, *45*(4-5), 315–26. doi: 10.1016/j.neucli.2015.09.007

Muttray, A., Breitinger, A., Goetze, E., Schnupp, T., Geissler, B., Kaufmann, T., ... Letzel, S. (2013). Further development of a commercial driving simulation for research in occupational medicine. *International journal of occupational medicine and environmental health*, *26*(6), 949–65. doi: 10.2478/s13382-013-0164-5

Nacke, L. E., Grimshaw, M. N., & Lindley, C. A. (2010). More than a feeling: Measurement of sonic user experience and psychophysiology in a first-person shooter game. *Interacting with Computers*, *22*(5), 336–343. doi: 10.1016/j.intcom.2010.04.005

National Swedish Road Administration. (1977). *Swedish capacity manual.*

Navin, F. P. D. (1994). Bicycle Traffic Flow Characteristics: Experimental Results and Comparisons. *ITE Journal*(March), 31–36.

New York City Department of Transport. (2008). *Sustainable Streets: Improving Travel in a Thriving City* (Tech. Rep.). New York, New York, USA: New York City DOT. Retrieved from `http://www.nyc.gov/html/dot/downloads/pdf/stratplan_mobility.pdf`

New York City Department of Transportation. (2013). *Measuring the Street: New Metrics for 21st Century Streets* (Tech. Rep.). Retrieved from `http://www.nyc.gov/html/dot/downloads/pdf/2012-10-measuring-the-street.pdf`

New York State Department of Transportation. (2015). *Highway Design Manual: Chapter 17 Bicycle Facility Design.* Retrieved from `https://www.dot.ny.gov/divisions/engineering/design/dqab/hdm/chapter-17`

Oculus VR LLC. (2016a). *Oculus Best Practices: Simulator Sickness.* Retrieved from [2016-08-14]https://developer.oculus.com/documentation/intro-vr/latest/concepts/bp_app_simulator_sickness/

Oculus VR LLC. (2016b). *Oculus Rift.* Retrieved from [2016-10-18]https://www3.oculus.com/en-us/rift/

O'Hern, S., Oxley, J., & Stevenson, M. (2017). Validation of a bicycle simulator for road safety research. *Accident Analysis & Prevention*, *100*, 53–58. doi: 10.1016/j.aap.2017.01.002

OKTAL. (2016). *Automotive Simulators.* Retrieved from [2016-09-19]http://www.oktal.fr/en/automotive/range-of-simulators/range-of-simulators

OpenCV. (2014). *OpenCV Library.* Retrieved from http://www.opencv.org

OpenCV Foundation. (2014). *OpenCV Vision Challenge.* Retrieved from [2017-03-27]http://opencv.org/opencv-vision-challenge.html

Ordnance Survey. (2016). *Buildings Data Set.* EDINA. Retrieved from http://digimap.edina.ac.uk/

Oregon Department of Transportation. (2012). *2012 Oregon DOT Highway Design Manual: Chapter 12: Pedestrian and Bicycle.* Salem, OR, USA: Oregon DOT.

Osowski, C., & Waterson, B. (2015). Derivation of spatiotemporal data for cyclists (from video) to enable agent-based model calibration. *Procedia Computer Science*, *52*(0), 932–937. doi: 10.1016/j.procs.2015.05.168

Osowski, C., & Waterson, B. (2017). Establishing the validity of cycle path capacity assumptions in the Highway Capacity Manual. *International Journal of Sustainable Transportation*, *11*(6), 422–432. doi: 10.1080/15568318.2016.1266424

Parkin, J., & Meyers, C. (2010). The effect of cycle lanes on the proximity between motor traffic and cycle traffic. *Accident Analysis and Prevention*, *42*(1), 159–65. doi: 10.1016/j.aap.2009.07.018

Parkin, J., & Rotheram, J. (2010). Design speeds and acceleration characteristics of bicycle traffic for use in planning, design and appraisal. *Transport Policy*, *17*(5), 335–341. doi: 10.1016/j.tranpol.2010.03.001

Pascucci, F., Rinke, N., Schiermeyer, C., Friedrich, B., & Berkhahn, V. (2015). Modeling of shared space with multi-modal traffic using a multi-layer social force approach. *Transportation Research Procedia*, *10*(July), 316–326. doi: 10.1016/j.trpro.2015.09.081

Plumert, J. M., Kearney, J. K., & Cremer, J. F. (2004). Children's perception of gap affordances: Bicycling across traffic-filled intersections in an immersive virtual environment. *Child Development*, *75*(4), 1243–1253.

Plumert, J. M., Kearney, J. K., Cremer, J. F., Recker, K. M., & Strutt, J. (2011). Changes in children's perception-action tuning over short time scales: Bicycling across traffic-filled intersections in a virtual environment. *Journal of Experimental Child Psychology*, *108*(2), 322–337. doi: 10.1016/j.jecp.2010.07.005

Portman, M. E., Natapov, A., & Fisher-Gewirtzman, D. (2015). To go where no man has gone before: Virtual reality in architecture, landscape architecture and environmental planning. *Computers, Environment and Urban Systems*, *54*, 376–384. doi: 10.1016/j.compenvurbsys.2015.05.001

PTV. (n.d.). *How PTV Vissim is contributing to creating the world's best city for cyclists.* Karlsruhe, Germany: Author.

PTV. (2009). *PTV Vision: VISSIM – State-of-the-Art Micro-Simulation.* Retrieved from `http://www.ptvap.com/docs/VISSIM_AP_LowRes_opt.pdf`

PTV. (2013a). *Vissim.* Karlsruhe, Germany: Author. Retrieved from `http://vision-traffic.ptvgroup.com/en-uk/products/ptv-vissim/`

PTV. (2013b). *Visum.* Karlsruhe, Germany: Author. Retrieved from `http://vision-traffic.ptvgroup.com/en-uk/products/ptv-visum/`

PTV. (2013c). *Viswalk.* Karlsruhe, Germany: Author. Retrieved from `http://vision-traffic.ptvgroup.com/en-uk/products/ptv-viswalk/`

Reason, J. T., & Brand, J. J. (1975). *Motion sickness.* Academic Press.

Reid, C. (2016). *Cycling surge in London due to protected Cycle Superhighways, says TfL.* Retrieved from `http://www.bikebiz.com/news/read/cycling-surge-in-london-due-to-protected-cycle-superhighways-says-tfl/019677`

Resnick, M. (1994). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds.* Cambridge, MA, USA: MIT Press.

Reynolds, R. F., & Bronstein, A. M. (2003). The broken escalator phenomenon. *Experimental Brain Research*, *151*(3), 301–308. doi: 10.1007/s00221-003-1444-2

Riccio, G. E., & Stoffregen, T. A. (1991). *An ecological theory of motion sickness and postural instability* (Vol. 3) (No. 3).

Rouphail, N., Hummer, J., Milazzo II, J., & Allen, P. (1998). *Capacity Analysis of Pedestrian and Bicycle Facilities: Recommended Procedures for the "Bicycles" Chapter of the Highway Capacity Manual - FHWA–RD–98–108* (Tech. Rep.). McLean, VA, USA: Federal Highway Administration. Retrieved from https://www.fhwa.dot.gov/publications/research/safety/pedbike/98108/

RS Components. (2016). *Broadcom Incremental Encoder -0.5 7 V dc.* Retrieved from [2016-10-09]http://uk.rs-online.com/web/p/rotary-encoders/7967806/

Salter, R. J., & Hounsell, N. B. (1996). *Highway Traffic Analysis and Design* (3rd ed.; ‘, Ed.). Basingstoke, UK: Macmillan Press Ltd.

Schönauer, R., Stubenschrott, M., Huang, W., Rudloff, C., & Fellendorf, M. (2012). Modeling concepts for mixed traffic: Steps towards a microscopic simulation tool for shared space zones. *Transportation Research Record*, *2316*. Retrieved from http://trid.trb.org/view.aspx?id=1128890

Schwebel, D. C., Combs, T., Rodriguez, D., Severson, J., & Sisiopiku, V. (2016). Community-based pedestrian safety training in virtual reality: A pragmatic trial. *Accident Analysis and Prevention*, *86*, 9–15. doi: 10.1016/j.aap.2015.10.002

Seattle Department of Transportation. (2007). *Seattle Bicycle Master Plan.* Seattle, WA, USA.

Shainberg, L. (1989). *Finding ‘The Zone’.* Retrieved from http://www.nytimes.com/1989/04/09/magazine/finding-the-zone.html?pagewanted=all

Shapiro, L., & Stockman, G. (2000). Binary Image Analysis. In *Computer vision* (pp. 63—-106). Prentice-Hall. Retrieved from http://www.cse.msu.edu/~stockman/Book/2002/Chapters/ch3.pdf

Shapiro, L. G., & Stockman, G. C. (2001). *Computer Vision.* Prentice Hall. Retrieved from https://books.google.co.uk/books?id=FftDAQAAIAAJ

Sharples, S., Cobb, S., Moody, A., & Wilson, J. R. (2008). Virtual reality induced symptoms and effects (VRISE): Comparison of head mounted display (HMD), desktop and projection display systems. *Displays*, *29*(2), 58–69.

SIAS. (2000). *Paramics.* Edinburgh, UK: Author. Retrieved from http://www.sias.com/2013/sp/spproducts.htm

Smith, G. (2016). *Oculus Rift Guide: Everything You Need To Know Before You Consider Buying One.* Retrieved from [2016-09-14]https://www.rockpapershotgun.com/2016/03/28/oculus-rift-guide/

Snowdon, J. (2008). *Using gait to analyse outdoor scenes for automated visual surveillance* (Dissertation, University of Southampton). Retrieved from http://users.ecs.soton.ac.uk/jrs105/Report.pdf

So, R. H., Lo, W. T., & Ho, a. T. (2001). Effects of navigation speed on motion sickness caused by an immersive virtual environment. *Human factors*, *43*(3), 452–461.

Sony. (2016). *Playstation VR.* Retrieved from [2016-10-18]https://www.playstation.com/en-gb/explore/playstation-vr/

St. Pierre, M. E., Banerjee, S., Hoover, A. W., & Muth, E. R. (2015). The effects of 0.2 Hz varying latency with 20-100 ms varying amplitude on simulator sickness in a helmet mounted display. *Displays*, *36*, 1–8. doi: 10.1016/j.displa.2014.10.005

Stanney, K. M., Hale, K. S., Nahmens, I., & Kennedy, R. S. (2003). What to expect from immersive virtual environment exposure: influences of gender, body mass index, and past experience. *Human Factors*, *45*(3), 504–520. doi: 10.1518/hfes.45.3.504.27254

Still, G. K. (2000). *Crowd Dynamics* (PhD, University of Warwick). Retrieved from http://www.gkstill.com/CV/PhD/

Sustrans. (2014). *Sustrans Design Manual: Handbook for Cycle-Friendly Design* (No. April). Author.

Tacx B.V. (2016a). *Tacx Bushido Smart T2780.* Retrieved from [2016-10-01]https://www.tacx.com/en/products/trainers/bushido-smart

Tacx B.V. (2016b). *Tacx Bushido Smart T2780: Wireless Motor Brake.* Retrieved from [2016-10-09]https://www.tacx.com/en/products/trainers/bushido-smart#tab_2

Tacx B.V. (2016c). *Tacx Cycling App.* Apple iTunes Store. Retrieved from https://itunes.apple.com/gb/app/tacx-cycling-app/id665093316

Tacx B.V. (2016d). *Tacx Galaxia T1100.* Retrieved from [2016-10-01]https://www.tacx.com/en/products/trainers/galaxia

Tacx B.V. (2016e). *Tacx i-Genius Multiplayer Smart Trainer T2010.* Retrieved from [2016-10-01]https://www.tacx.com/en/products/trainers/i-genius-multiplayer

Technogym. (2016). *Technogym Forma Exercise Bike.* Retrieved from [2016-10-01]http://www.technogym.com/gb/bike-excite-forma.html

The Highways Agency. (2005). TA 91/05 Provision for Non-Motorised Users. , *5*(February 2005).

Tiwari, G., Fazio, J., & Pavitravas, S. (2000). Passenger car units for heterogeneous traffic using a modified density method. *Proceedings of fourth ...*, 246–257. Retrieved from http://onlinepubs.trb.org/onlinepubs/circulars/EC018/22_44.pdf

Transport for London. (2010a). *Business Case Development Manual* (No. October).

Transport for London. (2010b). *Traffic Modelling Guidelines v3* (Tech. Rep.).

Transport for London. (2010c). *Transport Assessment Best Practice Guidance* (Tech. Rep. No. April). Transport for London.

Transport for London. (2012). *Roads Task Force - Technical Note 4: How has cycling grown in London and how will it grow in future?* Retrieved from http://content.tfl.gov.uk/technical-note-04-how-has-cycling-grown-in-london.pdf

Transport for London. (2013). Central London Cycle Census: Technical Note. (October). Retrieved from https://tfl.gov.uk/cdn/static/cms/documents/cycle-census-technical-note.pdf

Transport for London. (2014a). *London Cycling Design Standards: Chapter 4: Cycle Lanes and Track.* London, UK: Transport for London. Retrieved from https://tfl.gov.uk/cdn/static/cms/documents/lcds-chapter4-cyclelanesandtracks.pdf

Transport for London. (2014b). *London Cycling Design Standards: Draft for consultation.* London, UK: Transport for London.

Transport for London. (2015a). *Have your say on a new segregated East-West Cycle Superhighway through central London.* Retrieved from [2015-03-11]https://consultations.tfl.gov.uk/cycling/eastwest

Transport for London. (2015b). *Item 7: Proposed Cycle Superhighways Schemes* (Tech. Rep. No. February). London, UK: Transport for London.

Transport for London. (2016). *Cycle Flows on the TfL Road Network.*

Transportation Research Board. (2010). *Highway Capacity Manual.* Washington, D.C.: National Research Council.

TRL Software. (2012). *ARCADY.* Crowthorne, Berkshire, UK: TRL Software. Retrieved from https://www.trlsoftware.co.uk/products/junction_signal_design/arcady

Twaddle, H., Schendzielorz, T., & Fakler, O. (2014). Bicycles in Urban Areas. *Transportation Research Record: Journal of the Transportation Research Board*, *2434*, 140–146. doi: 10.3141/2434-17

Tydeman, R. (2004). *The Use of Full Flight Simulators for Accident Investigation* (Tech. Rep.). Australian Society of Air Safety Investigators. Retrieved from http://asasi.org/papers/2004/Tydeman_FlightSimulators_ISASI04.pdf

Unity Technologies. (2016a). *Execution Order of Event Functions.* Retrieved from [2016-08-28]https://docs.unity3d.com/Manual/ExecutionOrder.html

Unity Technologies. (2016b). *Unity 3D.* Retrieved from https://unity3d.com/

Valve Corporation. (2007). *Portal.* Valve Corporation.

Van Mead, N. (2016). *Bike jams and unwritten rules: a day with Amsterdam's new 'bicycle mayor'.* Retrieved from https://www.theguardian.com/cities/2016/aug/11/cycling-amsterdam-bike-jams-bicycle-mayor-anna-luten

Vansteenkiste, P., Cardon, G., D'Hondt, E., Philippaerts, R., & Lenoir, M. (2013). The visual control of bicycle steering: The effects of speed and path width. *Accident; analysis and prevention*, *51*, 222–7. doi: 10.1016/j.aap.2012.11.025

Vansteenkiste, P., Van Hamme, D., Veelaert, P., Philippaerts, R., Cardon, G., & Lenoir, M. (2014). Cycling around a Curve: The Effect of Cycling Speed on Steering and Gaze Behavior. *PLoS ONE*, *9*(7), e102792. doi: 10.1371/journal.pone.0102792

Vansteenkiste, P., Zeuwts, L., Cardon, G., Philippaerts, R., & Lenoir, M. (2014). The implications of low quality bicycle paths on gaze behavior of cyclists: A field test. *Transportation Research Part F: Traffic Psychology and Behaviour*, *23*, 81–87. doi: 10.1016/j.trf.2013.12.019

Vasic, J., & Ruskin, H. J. (2012). Cellular automata simulation of traffic including cars and bicycles. *Physica A: Statistical Mechanics and its Applications*, *391*(8), 2720–2729. doi: 10.1016/j.physa.2011.12.018

Vejdirektoratet. (2012). *Grundlag for udformning af trafikarealer.* Retrieved from http://vejregler.lovportaler.dk/

Venkatesan, K., Gowri, a., & Sivanandan, R. (2008). Development of Microscopic Simulation Model for Heterogeneous Traffic Using Object Oriented Approach. *Transportmetrica*, *4*(3), 227–247.

Warr, P. (2016). *Wot I Think: Everybody's Gone To The Rapture.* Retrieved from [2016-09-14]https://www.rockpapershotgun.com/2016/04/15/everybodys-gone-to-the-rapture-pc-review/

Webster, F. V., & Cobbe, B. M. (1966). *Road Research Technical Paper No. 56: Traffic Signals* (Tech. Rep.). London, UK: Road Research Laboratory.

Wei, H., Huang, J., & Wang, J. (1997). Models for estimating traffic capacity on urban bicycle lanes. In *76th annual meeting of the transportation research board.*

Wei, H., Ren, F., & Liu, X. (1993). Research on the relationship between bicycle travelling state and bicycle road capacity. *Chinese Journal of Highway Transportation*, *6*(4), 60–64.

Weibel, D., Wissmath, B., Habegger, S., Steiner, Y., & Groner, R. (2008). Playing online games against computer- vs. human-controlled opponents: Effects on presence, flow, and enjoyment. *Computers in Human Behavior*, *24*(5), 2274–2291.

Willemsen, P., Kearney, J. K., & Wang, H. (2003). Ribbon networks for modeling navigable paths of autonomous agents in virtual urban environments. *Proceedings - IEEE Virtual Reality*, *2003-Janua*(October), 79–86.

Williams, G. W. (1963). Highway Hypnosis: An Hypothesis. *International Journal of Clinical and Experimental Hypnosis*, *11*(3), 143–151. doi: 10.1080/00207146308409239

Woodling, C. H., Faber, S., Van Bockel, J. J., Olasky, C. C., Williams, W. K., Mire, J. L. C., & Homer, J. R. (1973). *NASA TN D-7112: Apollo Experience Report: Simulation of Manned Space Flight for Crew Training* (Tech. Rep.). Washington, D.C., USA: National Aeronautics and Space Administration. Retrieved from `https://www.hq.nasa.gov/alsj/NASATND7112.pdf`

World Health Organization. (2011). *Health economic assessment tools (HEAT) for walking and for cycling: Methodology and user guide.* Retrieved from `http://www.euro.who.int/__data/assets/pdf_file/0003/155631/E96097rev.pdf`

Yan, P. (2016). *A VR Cycling Experience for $40.* Retrieved from [2016-01-26]`https://pauldyan.wordpress.com/2016/01/24/my-vr-bike/`

Yao, D., Zhang, Y., Li, L., Su, Y., Cheng, S., & Xu, W. (2009). Behavior modeling and simulation for conflicts in vehicles-bicycles mixed flow. *IEEE Intelligent Transportation Systems Magazine*, *1*(2), 25–30.

Yi, J., Song, D., Levandowski, A., & Jayasuriya, S. (2006). Trajectory Tracking and Balance Stabilization Control of Autonomous Motorcycles. In *Proceedings 2006 ieee international conference on robotics and automation, 2006. icra 2006.* (pp. 2583–2589).

Yin, S., & Yin, Y. (2007). Study on virtual force sensing and force display device for the interactive bicycle simulator. *Sensors and Actuators A: Physical*, *140*(1), 65–74. doi: 10.1016/j.sna.2007.06.018

Yor, I., Helman, S., & Vermaat, P. (2015). *RPN751 Off street trials of a Dutch-style roundabout* (Tech. Rep.). TRL Ltd. Retrieved from `http://www.trl.co.uk/media/839260/ppr751_dutch_roundabout_safety_v1.pdf`

Zhang, J., Mehner, W., Andresen, E., Holl, S., Boltes, M., Schadschneider, A., & Seyfried, A. (2013). Comparative Analysis of Pedestrian, Bicycle and Car Traffic Moving in Circuits. *Procedia - Social and Behavioral Sciences*, *104*, 1130–1138. doi: 10.1016/j.sbspro.2013.11.209